Probabilistic opinion pooling generalized Part two: The premise-based approach

Dietrich, Franz and List, Christian

Paris School of Economics, London School of Economics

May 2013

Probabilistic opinion pooling generalized
Part two: The premise-based approach*

Franz Dietrich
CNRS & UEA

Christian List
LSE

May 2013, revised July 2014

Abstract

How can different individuals' probability functions on a given σ-algebra of events be aggregated into a collective probability function? Classic approaches to this problem often 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'.

Keywords: Probabilistic opinion pooling, judgment aggregation, subjective probability, premise-based aggregation

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 aggregated into a corresponding 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

*Although both authors are jointly responsible for this paper and project, Christian List wishes to note that Franz Dietrich should be considered the lead author, to whom the credit for the present mathematical proofs is due. This paper is the second of two self-contained, but technically related companion papers inspired by binary judgment-aggregation theory. Both papers build on our earlier, unpublished paper 'Opinion pooling on general agendas' (September 2007).

under negation and countable disjunction, and by implication also 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 on it.[1] 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'). And 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 are 'constructible' by combining 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, in reasoning about a given set of events, we may conceptualize basic and non-basic events differently. We may conceptualize non-basic events as combinations of basic events, just as we form composite propositions by combining atomic propositions. Second, the way in which we assign probabilities to non-basic events is likely to differ from the way in which we assign probabilities to basic events. When we assign a probability to some non-basic event, say a conjunction or disjunction of basic events, this typically presupposes the assignment of probabilities to the underlying basic events. For example, the natural way to assign a probability to the disjunctive 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.[2] If this is correct, 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 to 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

[1]For a departure from the classical approach in which the 'agenda' of events for which probabilities are aggregated is a σ-algebra, see our companion paper Dietrich and List (2013b).

[2]The correlation question might be approached by looking for causal effects between, or common causes of, rain and heat. Of course, what we have described is a very styzlied method of probability assignment.

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, for instance, 'rain' and 'heat', then, in a first step, the collective probabilities for these two events are determined by aggregating 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 instead 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 treatment 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; see, for instance, 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 raise some questions. For example, although the distinction between 'basic' events and 'non-basic' events is arguably not *ad hoc*, there is no purely formal criterion for drawing that distinction.[3] However, there

[3]Indeed, one could construct ('derive') 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 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 a 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.

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 quite specifically, 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 specific 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 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). 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).

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, as soon as there are certain logical interconnections 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 precise nature of the unanimity-preservation requirement and in the kinds of interconnections that are assumed 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 those in our companion paper, though interpretationally different (Dietrich and List 2013b). 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 and McConway 1981; see also Wagner 1982 and 1985; Aczél, Ng and Wagner 1984;

Genest 1984a, Mongin 1995; and Chambers 2007).[4]

2 The framework

We consider a group of $n \geq 2$ individuals, labelled $i = 1, ..., n$, 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 $A^c := \Omega \backslash A$. The set of events to which probabilities are to be 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 $\Sigma = \{A, A^c, \Omega, \varnothing\}$, in which A might be the event that it will rain. 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 our companion paper (Dietrich and List 2013b), 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 critical 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. An expert committee may well be faced with an opinion pooling problem on the agenda $\Sigma = 2^\Omega$.

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

[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 and Genest, McConway and Schervish 1986), *multiplicative* pooling (Dietrich 2010), *supra-Bayesian* pooling (e.g., Morris 1974), and pooling of *ordinal* probabilities (Weymark 1997). A useful, though somewhat outdated, literature review is given in Genest and Zidek (1986).

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

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 $\cap_{A \in S} A$ is non-empty, and *inconsistent* otherwise; S *entails* an event B if the intersection of S is included in B (i.e., $\cap_{A \in S} A \subseteq B$).

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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$,

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

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 alone, we effectively 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 coherent, i.e., constitute a well-defined probability function.)

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 our companion paper, Dietrich and List 2013b). 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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$,

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

In the climate-change example of Section 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. Note, for example, that 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 probabilistic beliefs about other events, then we are led to regard those other beliefs as by-products of, or unrelated to, the individual's belief about the premise in

[5] These 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).

question. It then seems reasonable to treat those other beliefs as irrelevant to the question of what probability should be collectively assigned 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 requirement on premise-based opinion pooling 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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A) = 1$, then $P_{P_1,...,P_n}(A) = 1$.[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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A) = 1$, then $P_{P_1,...,P_n}(A) = 1$.

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]

[6]Equivalently, one can demand the preservation of the unanimous assignment of probability 0.

[7]In our companion paper (Dietrich and List 2013b), we make the opposite move of extend-

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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A|B) = 1$ (provided $P_i(B) \neq 0$), then $P_{P_1,...,P_n}(A|B) = 1$ (provided $P_{P_1,...,P_n}(B) \neq 0$).

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 $P_1, ..., P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A \backslash B) = 0$, then $P_{P_1,...,P_n}(A \backslash B) = 0$.

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 '$P_{P_1,...,P_n}(A|B) = 1$ (provided $P_{P_1,...,P_n}(B) \neq 0$)' is equivalent to '$P_{P_1,...,P_n}(B \backslash A) = 0$'. 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

ing consensus preservation to events outside the agenda. More precisely, we extend consensus preservation to events constructible from events in the agenda using the standard operations of conjunction (intersection), disjunction (union), or negation (complementation). In the present paper, there is no point in trying to extend consensus preservation to other events, since there are no events outside the agenda constructible from events in it (the agenda is a σ-algebra and is therefore closed under the relevant operations).

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

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

(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 corresponding to propositions in natural language, such as 'it will rain'. As argued in our companion paper (Dietrich and List 2013b), in such applications the classical assumption that the agenda is a σ-algebra (which we have retained here) is often unnatural.

However, 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 a particular period. Here, the set of worlds Ω could be $\mathbb{R}$, $\mathbb{Z}$, $\mathbb{N}$, or $\{0, 1, ..., m\}$, or it could be $\mathbb{R}^k$, $\mathbb{Z}^k$, $\mathbb{N}^k$, or $\{0, 1, ..., m\}^k$ (for natural numbers m and k). In such cases, the focus on the σ-algebra of events seems 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 $\Omega = \{0, 1, ..., m\}$);
- 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 summarize how such applications differ from the above mentioned applications involving events that correspond to 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, once there are certain interconnections between the premises in X, any pooling function satisfying independence on X and consensus preservation in one of the senses we have 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 called *neutral on* X ($\subseteq \Sigma$) if there exists some function $D : [0,1]^n \rightarrow [0,1]$ – the *local pooling criterion* for events in X – such that, for every profile $(P_1, ..., P_n) \in \mathcal{P}^n_\Sigma$, the collective probability of any event A in X is given by

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

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

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 our companion paper (Dietrich and List 2013b).

Theorem 1 (a) *For any non-nested (finite*[9]*) 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 that is not neutral on X.*

The result continues to hold if we weaken consensus preservation to conditional consensus preservation on premises:

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 that is not neutral 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 defined path-connectedness formally, we begin with a preliminary notion. Given 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.[10] 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.

[9] The finiteness assumptions in Theorems 1(a), 1(b), 2(a), 2(b), 3(a), 4(a), 4(b), 5(a), and 6(a) could each be replaced by the assumption that the σ-algebra generated by X is the agenda Σ (rather than a sub-σ-algebra of Σ), as is clear from our proofs. 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.

[10] A set S of events is *consistent* if $\cap_{C \in S} C$ is non-empty; S *entails* an event B if $\cap_{C \in S} C \subseteq B$.

We call the sub-agenda X *path-connected* if any two events $A, B \in X \backslash \{\varnothing, \Omega\}$ can beconnected by a path of conditional entailments, i.e., there exist events $A_1, ..., A_k \in X$ $(k \geq 1)$ such that $A = A_1 \vdash^* A_2 \vdash^* ... \vdash^* A_k = B$, 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 that is not neutral 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 $w_1, ..., w_n \geq 0$ with $w_1 + ... + w_n = 1$ such that, for every profile $(P_1, ..., P_n) \in \mathcal{P}_\Sigma^n$, the collective probability of any event A in X is given by

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

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 interconnected 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 that is not linear 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.[11] 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. For further discussion, see our companion paper (Dietrich and List 2013b).

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 that is not linear 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 that is not linear 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:

[11] 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 that is not linear on X.*[12]

7 Classical 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 and McConway 1981) as special cases. To see why this is the case, we 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:[13]

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

(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 $|X| > 4$, every pooling function $F : \mathcal{P}_X^n \to \mathcal{P}_X$ satisfying independence and consensus preservation is linear (and by implication neutral);*
(b) *if $|X| = 4$, there exists a pooling function $F : \mathcal{P}_X^n \to \mathcal{P}_X$ satisfying independence and consensus preservation that is not neutral (and by implication not linear).*

[12] This assumes 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 $k \in \{1, 2, 3, ...\}$ or is infinite.

[13] We require no restriction to a finite Σ, as observed in footnote 9.

8 References

Aczél, J., and C. Wagner (1980) "A characterization of weighted arithmetic means", *SIAM Journal on Algebraic and Discrete Methods* 1: 259-260

Aczél, J., C. T. Ng and C. Wagner (1984) "Aggregation Theorems for Allocation Problems", *SIAM Journal on Algebraic and Discrete Methods* 5: 1-8

Chambers, C. (2007) "An ordinal characterization of the linear opinion pool", *Economic Theory* 33(3): 457-474

Dietrich, F. (2006) "Judgment Aggregation: (Im)Possibility Theorems", *Journal of Economic Theory* 126(1): 286-298

Dietrich, F. (2010) "Bayesian group belief", *Social Choice and Welfare* 35(4): 595-626

Dietrich, F. and C. List (2007) "Arrow's theorem in judgment aggregation", *Social Choice and Welfare* 29(1): 19-33

Dietrich, F. and C. List (2013a) "Propositionwise judgment aggregation: the general case", *Social Choice and Welfare* 40: 1067-1095

Dietrich, F. and C. List (2013b) "Probabilistic opinion pooling generalized – Part one: general agendas", working paper

Dietrich, F. and P. Mongin (2010) "The premise-based approach to judgment aggregation", *Journal of Economic Theory* 145(2): 562-582

Genest, C. (1984a) "Pooling operators with the marginalization property", *Canadian Journal of Statistics* 12: 153-63

Genest, C. (1984b) "A characterization theorem for externally Bayesian groups", *Ann. Statist.* 12: 1100-1105

Genest, C., K. J. McConway and M. J. Schervish (1986) "Characterization of externally Bayesian pooling operators", *Ann. Statist.* 14: 487-501

Genest, C. and J. V. Zidek (1986) "Combining Probability Distributions: A Critique and Annotated Bibliography", *Statistical Science* 1 (1): 113-135

Genest, C. and C. Wagner (1984) "Further Evidence against Independence Preservation in Expert Judgement Synthesis", *Technical Report* 84-10, Dept. of Statistics and Actuarial Science, University of Waterloo

Kornhauser, L. A. and L.G. Sager (1986) "Unpacking the Court", *Yale Law Journal* 96(1): 82-117

Lehrer, K. and C. Wagner (1981) *Rational Consensus in Science and Society*, Dordrecht: Reidel

List, C. and P. Pettit (2002) "Aggregating sets of judgments: an impossibility result", *Economics and Philosophy* 18(1): 89-110

List, C. (2006) "The Discursive Dilemma and Public Reason", *Ethics* 116(2): 362-402

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

McConway (1981) "Marginalization and Linear Opinion Pools", *Journal of the American Statistical Association* 76: 410-14

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", *Journal of Economic Theory* 141(1): 100-113

Morris, P. A. (1974) "Decision analysis expert use", *Management Science* 20: 1233-41

Nehring, K. and C. Puppe (2010) "Abstract Arrovian Aggregation", *Journal of Economic Theory* 145(2): 467-494

Pettit, P. (2001) "Deliberative Democracy and the Discursive Dilemma", *Philosophical Issues* 11: 268-299

Priest, G. (2001) *An Introduction to Non-classical Logic*, Cambridge Univ. Press

Wagner, C. (1982) "Allocation, Lehrer Models, and the Consensus of Probabilities", *Theory and Decision* 14: 207-220

Wagner, C. (1985) "On the Formal Properties of Weighted Averaging as a Method of Aggregation", *Synthese* 62: 97-108

Weymark, J. (1997) "Aggregating Ordinal Probabilities on Finite Sets", *Journal of Economic Theory* 75: 407-432

A Proofs

This appendix contains all proofs. Since our results are mathematically related to those in the companion paper (Dietrich and List 2013b), some parts of the present results (notably, parts (a) of each theorem) will not be proved directly but reduced to results in the companion paper.

In Section A.1 we prove some lemmas which will help us translate between our results and those in the companion paper. Section A.2 contains the proof of Theorem 1, Section A.3 that of Theorems 2, and so on until Section A.7. Finally, Section A.8 contains the proof of Proposition 2.

A.1 The relationship to opinion pooling on general agendas

We now relate premise-based opinion pooling to opinion pooling on a general agenda as introduced in the companion paper. We start by generalizing the framework to arbitrary agendas. In general, an agenda is a non-empty set X of events of the form $A \subseteq \Omega$ which is closed under complementation (i.e., $A \in X \Rightarrow A^c \in X$). Interpretationally, X contains the events on which opinions are formed. It need not be closed under disjunction (union) or conjunction (intersection) of two events, and thus need not take the classical form of a σ-algebra. Examples are given in the companion paper.

Given an agenda X, an *opinion function* is a function $P : X \to [0, 1]$ which is coherent, i.e., extendible 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. Note that if X happens to be 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 $\mathcal{P}_X^n \to \mathcal{P}_X$ assigning to each profile $(P_1, ..., P_n)$ of individual opinion functions a collective opinion function, typically denoted $P_{P_1,...,P_n}$. 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.,

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

(and if, in addition, $\mathcal{P}_X = \{P|_X : P \in \mathcal{P}_\Sigma\}$, where this addition holds automatically whenever X is finite or $\sigma(X) = \Sigma$[14]). 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 the companion paper). The first two requirements are unrestricted versions of independence and consensus preservation:

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

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

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

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

Implicit consensus preservation. For all $A \in \sigma(X)$ and all $P_1, ..., P_n \in \mathcal{P}_X$, 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 $P_{P_1,...,P_n}$.

By contrast, individual i's *unrevealed* beliefs are any probabilistic beliefs which the agent privately holds relative to events in $\sigma(X) \backslash X$, and which are inaccessible based on the submitted opinion function P_i because different extensions of P_i to a probability function on $\sigma(X)$ 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 (possibly unrevealed) consensus is ever overruled.

Consensus compatibility. For all $A \in \sigma(X)$ and all $P_1, ..., P_n \in \mathcal{P}_X$, 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 $P_{P_1,...,P_n}$.

A final requirement pertains to *conditional* beliefs. Note that, based on an 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 $P_i(B) \neq 0$) is typically undefined, since typically $A \cap B \notin X$, so that $P_i(A \cap B)$ is undefined. Hence, if the agent happens to be certain of A given B, then this conditional certainty is typically 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 the companion paper for details).

Conditional consensus compatibility. For all $P_1, ..., P_n \in \mathcal{P}_X$, 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 $P_{P_1,...,P_n}$.

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 2 *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 proved in the companion paper:

Lemma 3 *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).

We finally note a simple sufficient condition for when a given pooling function induces a new one for a sub-agenda (see the companion paper for the simple proof):

Lemma 4 *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.*

A.2 Proof of Theorem 1

We draw on a measure-theoretic fact (proved in the companion paper):

Lemma 5 *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. (a) We reduce this part to the companion paper's Theorem 1(a). Let X be a non-nested finite sub-agenda of the σ-algebra agenda Σ. Suppose $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ is independent on X and (globally) consensus preserving. By independence on X (and finiteness of X), F induces a pooling function F' for agenda X (see Lemma 4). Now F' is independent and consensus-compatible (by Lemma 2), hence neutral by the companion paper's Theorem 1(a). So F is neutral on X (see Lemma 2).

(b) Consider a finite nested sub-agenda $X \neq \{\varnothing, \Omega\}$ of the σ-algebra agenda Σ. (As will become clear, the finiteness assumption 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 $(P_1, ..., P_n) \mapsto P_{P_1,...,P_n}$ for agenda Σ with the relevant properties. Without loss of generality, we may suppose that $\varnothing, \Omega \in X$. We start by establishing the following fact:

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

Indeed, suppose Theorem 1(b) holds in the special case. Let $\Sigma' := \sigma(X)$ $(\subseteq \Sigma)$. By assumption, there exists a pooling function $F' : \mathcal{P}_{\Sigma'}^n \to \mathcal{P}_{\Sigma'}$ with the relevant properties. Let $\mathcal{A}$ be the set of atoms of the (finite) σ-algebra Σ'. We define a pooling function $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$ as follows. Consider $P_1, ..., P_n \in \mathcal{P}_{\Sigma}$ and let $P' := F'(P_1|_{\Sigma'}, ..., P_n|_{\Sigma'})$. For each $A \in \mathcal{A}$ such that $P'(A) \neq 0$ there is an individual i_A such that $P_{i_A}(A) \neq 0$, since otherwise all individuals assign a probability of one to $\Omega \backslash A$ while $P'(\Omega \backslash A) \neq 1$, a contradiction as F' is consensus-preserving. By Lemma 5, P' can be extended to a probability function P on Σ. It is clear from that lemma's proof (given in the companion paper) that we may assume without loss of generality that

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

(In that proof, it suffices to choose the Q_A's appropriately, since each Q_A equals $P(\cdot|A)$ provided $P(A) \neq 0$.) We define $F(P_1, ..., P_n)$ to be this P. It remains to show that the just-defined pooling function F inherits all relevant properties from F'. This is obvious for the properties of independence on X and non-neutrality on X. To show that F is (globally) consensus-preserving, consider any $B \in \Sigma$ and any $P_1, ..., P_n \in \mathcal{P}_{\Sigma}$ such that $P_1(B) = \cdots = P_n(B) = 1$. We have to show that $P(B) = 1$, where $P := F(P_1, ..., P_n)$. 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, in turn, equals 1 because $P_{i_A}(B) = 1$. So

$$P(B) = \sum_{A \in \mathcal{A}: P(A) \neq 0} P(A) = 1.$$

This proves the claim.

By the previous 'Claim', we may assume without loss of generality that $\sigma(X) = \Sigma$. As X is nested, we may write it as $X = \{A, A^c : A \in X_+\}$ for a subset $X_+ \subseteq X$ which is linearly ordered by set-inclusion, and which contains both $\varnothing$ and Ω.

As an ingredient to 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 $(P_1, ..., P_n) \mapsto P_1$, with pooling criterion given by $D(t_1, ..., t_n) = t_1$.) As $X \neq \{\varnothing, \Omega\}$, there is a contingent event $A \in X$. As A is contingent, there are $P_1, ..., P_n \in \mathcal{P}_\Sigma$ 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 $D(1/2, ..., 1/2)$. As these probabilities sum to 1, it follows that

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

We now transform this on X neutral pooling function into a pooling function $(P_1, ..., P_n) \mapsto P_{P_1,...,P_n}$ 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 $x, y \in [0,1]$. (For instance, T could be defined by $T(x) = \min\{2x, 1\}$.)

Now consider any $P_1, ..., P_n \in \mathcal{P}_\Sigma$. We have to define the collective probability function $P_{P_1,...,P_n}$. We write P for the result of applying the neutral pooling function to $(P_1, ..., P_n)$. To anticipate, our definition will imply that

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

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

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

The restriction of $P_{P_1,...,P_n}$ 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 define

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

In particular,

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

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* in 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 $g(C_1 \cup C_2) = g(C_1) + g(C_2)$, as is easily checked using the additivity of P and the definition of g. To show σ-additivity, consider pairwise disjoint $C_1, C_2, ... \in \widetilde{X}$ such that $\cup_{k=1}^{\infty} C_k \in \widetilde{X}$. We have to show that

$$\delta_K := g(\cup_{k=1}^{\infty} C_k) - \sum_{k=1}^{K} g(C_k) \to 0 \text{ as } K \to \infty.$$

For all $K \in \{1, 2, ...\}$, note that the difference

$$(\cup_{k=1}^{\infty} C_k) \setminus \left(\cup_{k=1}^{K} C_k\right) = \cup_{k=K+1}^{\infty} C_k$$

may not belong to $\widetilde{X}$, but can be partitioned into a finite set $\mathcal{C}^K$ of events in $\widetilde{X}$. Clearly, $\mathcal{C}^K \cup \{C_1, ..., C_K\}$ partitions $\cup_{k=1}^{\infty} C_k$. By carefully inspecting the definition of g, one can see that

$$\delta_K = \sum_{C \in \mathcal{C}^K} g(C).$$

So, since $g(C) \leq KP(C)$ for each $C \in \widetilde{X}$ (by the definition of g and the property (iv) of T), we have

$$\delta_K \leq K \sum_{C \in \mathcal{C}^K} P(C) = KP(\cup_{k=K=1}^{\infty} C_k),$$

As $K \to \infty$, we have $P(\cup_{k=K=1}^{\infty} C_k) \to 0$ (by σ-additivity of P), and hence, $\delta_K \to 0$, as required.

Since g is σ-additive, and of course also σ-finite (i.e., Ω is a union of countably many events in $\widetilde{X}$ of finite g-measure, which is trivially true since $\Omega \in \widetilde{X}$), Caratheodory's Extension Theorem tells us that g extends uniquely to a measure on $\sigma(\widetilde{X}) = \sigma(X) = \Sigma$. Let $P_{P_1,...,P_n}$ be this extension. $P_{P_1,...,P_n}$ is indeed a *probability* measure since, firstly, g (and hence, $P_{P_1,...,P_n}$) is non-negative by the weak increasingness of T, and, secondly, $P_{P_1,...,P_n}(\Omega) = 1$ because $\Omega \in \widetilde{X}$ and $g(\Omega) = T(1) = 1$.

To complete the proof, we must show that the pooling function just defined, $(P_1, ..., P_n) \mapsto P_{P_1,...,P_n}$, is independent on X, (globally) consensus-preserving, but not neutral on X.

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

Non-neutrality on X. To show non-neutrality on X, it suffices to show that, for some $C \in X \backslash \{\Omega, \varnothing\}$, the pooling criteria D_C and D_{C^c} differ. (We require that

$C \notin \{\Omega, \varnothing\}$ to ensure that the criteria D_C and D_{C^c} are uniquely determined by the pooling function; the criteria D_Ω and $D_\varnothing$ are not uniquely determined and could be *chosen* to differ even if the pooling function were neutral on X.) This follows from the following argument. First, $X\backslash\{\Omega, \varnothing\} \neq \varnothing$ since $X \neq \{\varnothing, \Omega\}$ by assumption. So there is a pair $C, C^c \in X\backslash\{\Omega, \varnothing\}$. Without loss of generality, assume $C \in X_+$ and $C^c \in X\backslash X_+$. By the previous proof of independence on X, $D_C = T \circ D$ and $D_{C^c} = 1 - T \circ D(\mathbf{1} - \cdot)$. It follows that $D_C \neq D_{C^c}$, because $D_C(1/2, ..., 1/2) \neq D_{C^c}(1/2, ..., 1/2)$, as is clear from the fact that

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

where $T(1/2) \neq 1/2$.

Consensus preservation. Consider any $P_1, ..., P_n \in \mathcal{P}_\Sigma$ and any $A \in \Sigma$ such that each $P_i(A)$ is one. We have to show that $P_{P_1,...,P_n}(A) = 1$. Let P be the result of pooling $P_1, ..., P_n$ using the (at least on X) neutral pooling rule defined above. Since that pooling function is consensus-preserving, $P(A) = 1$. Note further that there is a $K > 0$ such that $P_{P_1,...,P_n}(B) \leq KP(B)$ for all $B \in \widetilde{X}$ (by property (iv) above). Since the (σ-additive and σ-finite) restrictions $P_{P_1,...,P_n}|_{\widetilde{X}}$ and $KP|_{\widetilde{X}}$ satisfy the inequality $P_{P_1,...,P_n}|_{\widetilde{X}} \leq KP|_{\widetilde{X}}$, their (by Caratheodory's Extension Theorem uniquely existing) measure extensions on $\sigma(\widetilde{X}) = \sigma(X) = \Sigma$ satisfy the analogous inequality. So, as these extensions are simply $P_{P_1,...,P_n}$ and KP, we have $P_{P_1,...,P_n} \leq KP$. In particular, $P_{P_1,...,P_n}(A^c) \leq KP(A^c) = K(1 - P(A)) = K(1 - 1) = 0$, whence $P_{P_1,...,P_n}(A) = 1$. ■

A.3 Proof of Theorem 2

(a) This part is reducible to the companion paper's Theorem 2(a) via Lemmas 4 and 2, in the same way in which we also reduced Theorem 1(a) to the companion paper's Theorem 1(a).

(b) This part follows immediately from Theorem 1(b), since (global) consensus preservation implies conditional consensus preservation on X by Proposition 1. ■

A.4 Proof of Theorem 3

(a) This part is reducible to the companion paper's Theorem 3(a) via Lemmas 4 and 2, just as Theorem 1(a) is reducible to the companion paper's Theorem 1(a), and Theorem 2(a) to the companion paper's Theorem 2(a).

(b) Now 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 start by proving that we may assume without loss of generality. that $\sigma(X) = \Sigma$.

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

Assume Theorem 3(b) holds for the special case and let $\Sigma' := \sigma(X)$ ($\subseteq \Sigma$). By assumption, there is a $F' : \mathcal{P}_{\Sigma'}^n \to \mathcal{P}_{\Sigma'}$ which, on X, is independent and consensus-preserving but not neutral. Consider a pooling function $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$ which, for any $P_1, ..., P_n \in \mathcal{P}_{\Sigma}$, generates a probability function in $\mathcal{P}_{\Sigma}$ which extends $F'(P_1|_{\Sigma'}, ..., P_n|_{\Sigma'})$ (where such an extension exists by Lemma 5 and the finiteness of Σ'). The so-defined pooling function F inherits all relevant properties from F': it is, on X, independent, consensus preserving, and non-neutral. This proves the claim.

Henceforth, 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., elements of $\tilde{\Sigma}\backslash\{\varnothing\}$ which are minimal with respect to set-inclusion). 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 $\bar{A} \in X\backslash\{\varnothing, \Omega\}$ be a (by assumption existing) event such that not for all $A \in X\backslash\{\varnothing, \Omega\}$ there is $\bar{A} \vdash\vdash^* A$, where "$\vdash\vdash^*$" denotes the transitive closure of $\vdash^*$. Consider any profile $(P_1, ..., P_n) \in \mathcal{P}_{\Sigma}^n$. To define the probability function $P_{P_1,...,P_n} \in \mathcal{P}_{\Sigma}$, 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

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

respectively. Let $P'_{P_1,...,P_n} \in \mathcal{P}_{\Sigma'}$ and $P''_{P_1,...,P_n} \in \mathcal{P}_{\Sigma''}$ be defined by

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

These two functions are indeed probability functions (on Σ' resp. Σ''), 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' is a superset of an atom of Σ' and hence of 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 $Y' \cup Y''$. Y is not a subset of any of Y' and Y'', because the latter sets 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 the claim.

We now extend the functions $P'_{P_1,...,P_n}$ and $P''_{P_1,...,P_n}$ 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 $\tilde{P}_{P_1,...,P_n}$ be the unique probability function on $\tilde{\Sigma}$ that behaves as follows on the atoms:

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

for all $A' \in \mathcal{A}(\Sigma')$ and all $A'' \in \mathcal{A}(\Sigma'')$. This function is indeed a probability function, because

$$\begin{aligned}\sum_{A \in \mathcal{A}(\tilde{\Sigma})} \tilde{P}_{P_1,...,P_n}(A) &= \sum_{A' \in \mathcal{A}(\Sigma'), A'' \in \mathcal{A}(\Sigma'')} P'_{P_1,...,P_n}(A')P''_{P_1,...,P_n}(A'') \\ &= \sum_{A' \in \mathcal{A}(\Sigma')} P'_{P_1,...,P_n}(A') \underbrace{\sum_{A'' \in \mathcal{A}(\Sigma'')} P''_{P_1,...,P_n}(A'')}_{=1} \\ &= 1.\end{aligned}$$

As one easily checks, restricting $\tilde{P}_{P_1,...,P_n}$ to Σ' (respectively Σ'') gives $P'_{P_1,...,P_n}$ (respectively $P''_{P_1,...,P_n}$), and so

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

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

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

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

Let C be as specified, and write $C_{P_1,...,P_n}$ for the event in question. As noted above, C takes the form $C = A' \cap A''$ with $A' \in \mathcal{A}(\Sigma')$ and $A'' \in \mathcal{A}(\Sigma'')$. By $P(C) > 0$ and (3), we have $\tilde{P}_{P_1,...,P_n}(A') > 0$ and $\tilde{P}_{P_1,...,P_n}(A'') > 0$. As $A' \in \mathcal{A}(\Sigma')$, 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 $\cap_{A \in X : \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1,...,P_n}$ can be written as $\cap_{A \in Y_{P_1,...,P_n}} A$, where the set

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

consists of exactly one member of each pair $A, A^c \in X \backslash (X' \cup X'')$. Thus $C_{P_1,...,P_n} = \cap_{A \in Y' \cup Y'' \cup Y_{P_1,...,P_n}} A$, where the set $Y' \cup Y'' \cup Y_{P_1,...,P_n}$ consists of exactly one member of each pair $A, A^c \in X$. So, as $\Sigma = \sigma(X)$, $C_{P_1,...,P_n}$ is either an atom or empty. Hence it suffices to show that $C_{P_1,...,P_n} \neq \varnothing$. 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 a $B \in Y \cap Y_{P_1,\dots,P_n}$ with $\bar{A} \vdash\vdash^* B$. Consider any $B' \in Y \backslash \{B\}$. 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 $B \in X'$, in contradiction as $B \in Y_{P_1,\dots,P_n}$. Further, as $\bar{A} \vdash\vdash^* B$ and $B \vdash^* (B')^c$, we have (ii) $\bar{A} \vdash\vdash^* (B')^c$. By (i) and (ii), letting $A := (B')^c$, the event $A_{P_1,\dots,P_n}$ ($\in \{A, A^c\}$) is well-defined. As $Y_{P_1,\dots,P_n}$ 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 we have $P_i(A) = 0$, and hence, $P_i(B') = 1$. Since this holds for all $B' \in Y \backslash \{B\}$, for all i we have $P_i(\cap_{B' \in Y} B') = P_i(B)$. Hence, as Y is inconsistent, for all i we have $P_i(B) = 0$. Thus $B_{P_1,\dots,P_n} = B^c$. So $B^c \in Y_{P_1,\dots,P_n}$, a contradiction since $B \in Y_{P_1,\dots,P_n}$.

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

Now we define $P_{P_1,\dots,P_n}$ as the unique function on Σ that assigns the following measure to the atoms of Σ. If an atom takes the form given in Claim 3, i.e., the form

$$B = C \cap (\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1,\dots,P_n})$$

where $C \in \mathcal{A}(\tilde{\Sigma})$ and $\tilde{P}_{P_1,\dots,P_n}(C) > 0$, then we define the atom's measure as

$$P_{P_1,\dots,P_n}(B) = \tilde{P}_{P_1,\dots,P_n}(C).$$

Any other atom has measure defined as zero.

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

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

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

There are two cases.

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

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

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

By definition of $P_{P_1,...,P_n}$, every atom of Σ that has positive probability is a subset of the event $\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1,...,P_n}$, and so this event has probability 1. It follows that, for all $A \in X$ such that $\bar{A} \vdash\vdash^* A$ and not $\bar{A} \vdash\vdash^* A^c$, we have $P_{P_1,...,P_n}(A_{P_1,...,P_n}) = 1$, and hence

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

This proves Claim 5.

By Claim 4, we have constructed a well-defined pooling function $(P_1, ..., P_n) \mapsto P_{P_1,...,P_n}$ for agenda Σ. By (4) and Claims 5 and 6, we know its behaviour on the entire sub-agenda X: the pooling function is independent on X and the local pooling criterion D_A of an event $A \in X$ is 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 everywhere except at $(0, ..., 0)$ the value 1) if $\bar{A} \vdash\vdash^* A$ but not $\bar{A} \vdash\vdash^* A^c$,
(iv) the non-linear criterion $1 - \hat{D}$ if not $\bar{A} \vdash\vdash^* A$ but $\bar{A} \vdash\vdash^* A^c$.

These pooling criteria also ensure unanimity preservation on X. To see that pooling is not neutral, 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.5 Proof of Theorem 4

(a) This part is reducible to the companion paper's Theorem 4(a) via Lemmas 4 and 2 (in the way in which we reduced Theorem 1(a) to the companion paper's Theorem 1(a)).

(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), since non-neutrality on X implies non-linearity on X. Now assume the other case, i.e., $|X \backslash \{\varnothing, \Omega\}| \leq 4$. We reduce the claim to the companion paper'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 3, F' is induced by a pooling function for agenda Σ which is independent on X, (globally) consensus-preserving, and not linear on X. ■

A.6 Proof of Theorem 5

(a) This part is reducible to the companion paper's Theorem 5(a) via Lemmas 4 and 2, again in the way in which (for instance) Theorem 1(a) is reducible to the companion paper's Theorem 1(a).

(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) and conditional-consensus-preserving, but not linear. We may assume without loss of generality that $\sigma(X) = \Sigma$, because the 'Claim' in the proof of Theorem 1(b) (proved using Lemma 5) holds analogously here as well.

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

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

Consider any $P_1, ..., P_n \in \mathcal{P}_\Sigma$. We have to define $P_{P_1,...,P_n}$. We write collective probabilities under the (at least on X) linear pooling function simply as

$$p(A) := P^{\text{lin}}_{P_1,...,P_n}(A) \text{ for all } A \in \Sigma,$$

and define

$$\begin{aligned} X_{\geq 1/2} &: = \{A \in X : p(A) \geq 1/2\} \\ X_{>1/2} &: = \{A \in X : p(A) > 1/2\} \\ X_{=1/2} &: = \{A \in X : p(A) = 1/2\}. \end{aligned}$$

Notice that for all $A \in X$ we have $A \in X_{>1/2} \Rightarrow A^c \notin X_{>1/2}$ and $A \in X_{=1/2} \Leftrightarrow A^c \in X_{=1/2}$.

(Although $p(A)$ and the sets $X_{\geq 1/2}, X_{>1/2}, X_{=1/2}$ depend on $P_1, ..., P_n$, our notation suppresses $P_1, ..., P_n$ for simplicity.)

To define $P_{P_1,...,P_n}$, we first need to prove two claims (which use that X is simple).

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 every pair $A, A^c \in X_{=1/2}$).

To show this, note first that $X_{=1/2}$ has of course 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 (with respect to set-inclusion), 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 by definition for $X^1_{=1/2}$, and holds for $X^2_{=1/2}$ too by the following argument. Let $A, B \in X^2_{=1/2}$ and suppose for a contradiction that $p(A \cap B) = 0$. 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 have

$$p(A) + p(B) + p(C) = p(A \cup B \cup C) \leq 1,$$

as p is a probability function. This is a contradiction, since $p(A) = p(B) = 1/2$ and $p(C) = p(A' \cap B') > 0$ (the latter as (i) holds for $X^1_{=1/2}$).

(ii) Suppose for a contradiction that some $X^j_{=1/2} \cup X_{>1/2}$ is inconsistent. Then (as X and hence $X^j_{=1/2} \cup X_{>1/2}$ is finite) there is a minimal inconsistent subset $Y \subseteq X^j_{=1/2} \cup X_{>1/2}$. As X is moreover simple, $|Y| \leq 2$, say $Y = \{A, B\}$. As $A \cap B = \varnothing$ and p is a probability function, we have

$$p(A) + p(B) = p(A \cup B) \leq 1.$$

So, as $p(A), p(B) \geq 1/2$, it must be that $p(A) = p(B) = 1/2$, i.e., that $A, B \in X^j_{=1/2}$. Hence, by (i), $p(A \cap B) > 0$, a contradiction since $A \cap B = \varnothing$.

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 $X = \{C^0_j, C^1_j : j = 1, ..., J\}$, where $J = |X|/2$ and where each pair C^0_j, C^1_j consists of an event and its complement. We may write Σ as the set of all unions of intersections of the form $C^{k_1}_1 \cap ... \cap C^{k_J}_J$, i.e., as

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

Recalling that Σ is the σ-algebra generated by X, the inclusion '$\supseteq$' in (7) is obvious, and the inclusion '$\subseteq$' holds because the right hand side of (7) includes X (as any $C^k_j \in X$ can be written as the union of all intersections $C^{k_1}_1 \cap ... \cap C^{k_J}_J$ for which $k_j = k$) and is a σ-algebra (as it is closed under taking unions and complements: just take the unions (respectively complements) of the corresponding sets $K \subseteq \{0,1\}^J$).

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

We are now in a position to define the function $P_{P_1,\ldots,P_n}$ on Σ. Since $\cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$ and $\cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C$ are non-empty by Claim 1, there exist worlds $\omega^1 \in \cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$ and $\omega^2 \in \cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C$, where we assume that $\omega^1 = \omega^2$ if $X_{=1/2} = \varnothing$, i.e., if $\cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C = \cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C = \cap_{C \in X_{>1/2}} C$. (Our notation for worlds again suppresses $P_1, \ldots, P_n$.) Let δ_{ω^1} and δ_{ω^2} be the corresponding Dirac measures on Σ, given for all $A \in \Sigma$ by $\delta_{\omega^j}(A) = 1$ if $\omega^j \in A$ and $\delta_{\omega^j}(A) = 0$ if $\omega^j \notin A$. We define

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

where ω^1, ω^2 of course depend on $P_1, \ldots, P_n$. (So $P_{P_1,\ldots,P_n}(A)$ is either 1 or 1/2 or 0, depending on whether $A \in \Sigma$ contains both, exactly one, or none of ω^1 and ω^2; further, $P_{P_1,\ldots,P_n} = \delta_\omega$ if $\omega^1 = \omega^2 = \omega$, i.e., if $X_{=1/2} = \varnothing$.)

As $P_{P_1,\ldots,P_n}$ is a convex combination of probability functions, $P_{P_1,\ldots,P_n}$ is indeed a probability function. The proof is completed by showing that the so-defined pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1,\ldots,P_n}$ has the desired properties, as shown in the next two claims.

Independence on X. We show that the pooling function is neutral on X (hence independent on X) with the pooling criterion D given in (6). To do so, consider any $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$ and any $A \in X$, and write $(t_1, \ldots, t_n) := (P_1(A), \ldots, P_n(A))$. We have to show that $P_{P_1,\ldots,P_n}(A) = D(t_1, \ldots, t_n)$. To do this, we consider three cases, and use p, $X_{>1/2}$, $X^1_{=1/2}, X^2_{=1/2}, \omega^1, \omega^2$ as defined above.

Case 1. $p(A) = D^{\text{lin}}(t_1, \ldots, t_n) < 1/2$. Then $D(t_1, \ldots, t_n) = 0$. So we must prove that $P_{P_1,\ldots,P_n}(A) = 0$, i.e., that A contains neither ω^1 nor ω^2. Assume for a contradiction that $\omega^1 \in A$ (the proof is analogous if we instead assume $\omega^2 \in A$). Then A includes the set $\cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$, as this set contains ω^1 and is (by Claim 2) an atom of Σ. So $A^c \cap [\cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C] = \varnothing$. Hence the set $\{A^c\} \cup X^1_{=1/2} \cup X_{>1/2}$ is inconsistent, so has a minimal inconsistent subset Y. Since X is simple, $|Y| \leq 2$. Y does not contain $\varnothing$, as A^c is non-empty (by $p(A^c) = 1 - p(A) > 1/2$) and as all $B \in X^1_{=1/2} \cup X_{>1/2}$ are non-empty (by

$p(B) \geq 1/2$). So $|Y| = 2$. Moreover, Y is not a subset of $X^1_{=1/2} \cup X_{>1/2}$, since this set is consistent by Claim 1. Hence $Y = \{A^c, B\}$ for some $B \in X^1_{=1/2} \cup X_{>1/2}$. As $A^c \cap B = \varnothing$ and as $p(A^c) = 1 - p(A) > 1/2$ and $p(B) \geq 1/2$, we have $p(A^c \cup B) = p(A^c) + p(B) > 1/2 + 1/2 = 1$, a contradiction.

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

Case 3. $p(A) = D^{\text{lin}}(t_1, ..., t_n) = 1/2$. Then $D(t_1, ..., t_n) = 1/2$. So we must prove that $P_{P_1,...,P_n}(A) = 1/2$, i.e., that A contains exactly one of ω^1 and ω^2. As $p(A) = 1/2$, exactly one of $X^1_{=1/2}$ and $X^2_{=1/2}$ contains A and the other one contains A^c, by Claim 1. Say $A \in X^1_{=1/2}$ and $A^c \in X^2_{=1/2}$ (the proof is analogous if instead $A \in X^2_{=1/2}$ and $A^c \in X^1_{=1/2}$). So $A \supseteq \cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$, and hence $\omega^1 \in A$. On the other hand, $\omega^2 \notin A$, 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. As $X \neq \{\varnothing, \Omega\}$, there is a contingent event $A \in X$, hence a probability function $P \in \mathcal{P}_\Sigma$ with $t := P(A) \notin \{0, 1/2, 1\}$. Now assume all individuals submit this P. If the pooling function were linear on X, the collective probability of A would again be t ($\notin \{0, 1/2, 1\}$), a contradiction since the collective probability is given by $D(t, ..., t)$ ($\in \{0, 1/2, 1\}$), as just shown.

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

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

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

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

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

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

(to see why $p(A \cup B)$, recall that $p(A \cup B) = P^{\text{lin}}_{P_1,\ldots,P_n}(A \cup B)$, where $P_i(A \cup B) = 1$ 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 $A \cup B$ contains both ω^1 and ω^2, whence $P_{P_1,\ldots,P_n}(A \cup B) = 1$, as desired. ■

A.7 Proof of Theorem 6

Proof of Theorem 6. (a) This part is reducible to the companion paper's Theorem 6(a) via Lemmas 4 and 2, once again in the way in which, for instance, Theorem 1(a) is reducible to the companion paper's Theorem 1(a).

(b) This part follows from Theorem 3(b) since non-neutrality on X implies non-linearity on X. ■

A.8 Proof of Proposition 2

Consider the σ-algebra agenda Σ, of which we assume that $|\Omega| > 2^3 = 8$, i.e., $|\Omega| \geq 2^4 = 16$. 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 Proposition 2 we have constructed a pooling function for agenda X which is neutral and consensus-preserving but not linear. By Lemma 3, this pooling function is induced by a pooling function for agenda Σ which, on X, is neutral and consensus-preserving but not linear. ■