

Probabilistic opinion pooling generalized

Part one: General agendas*

Franz Dietrich

Paris School of Economics & CNRS

Christian List

London School of Economics

May 2013, revised May 2016

Abstract

How can several individuals' probability assignments to some events be aggregated into a collective probability assignment? Classic results on this problem assume that the set of relevant events – the *agenda* – is a σ -algebra and is thus closed under disjunction (union) and conjunction (intersection). We drop this demanding assumption and explore probabilistic opinion pooling on general agendas. One might be interested in the probability of rain and that of an interest-rate increase, but not in the probability of rain *or* an interest-rate increase. We characterize linear pooling and neutral pooling for general agendas, with classic results as special cases for agendas that are σ -algebras. As an illustrative application, we also consider probabilistic preference aggregation. Finally, we unify our results with existing results on binary judgment aggregation and Arrovian preference aggregation. We show that the same kinds of axioms (independence and consensus preservation) have radically different implications for different aggregation problems: linearity for probability aggregation and dictatorship for binary judgment or preference aggregation.

Keywords: Probabilistic opinion pooling, judgment aggregation, subjective probability, probabilistic preferences, vague/fuzzy preferences, agenda characterizations, a unified perspective on aggregation

*We thank the referees and the editor for helpful and detailed comments. Although we are jointly responsible for this work, 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 first 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 Both papers build on our earlier, unpublished paper ‘Opinion pooling on general agendas’ (September 2007). Dietrich was supported by a Ludwig Lachmann Fellowship at the LSE and the French Agence Nationale de la Recherche (ANR-12-INEG-0006-01). List was supported by a Leverhulme Major Research Fellowship (MRF-2012-100) and a Harsanyi Fellowship at the Australian National University, Canberra.

1 Introduction

This paper addresses the problem of *probabilistic opinion pooling*. Suppose several individuals (e.g., decision makers or experts) each assign probabilities to some events. How can these individual probability assignments be aggregated into a collective probability assignment, while preserving probabilistic coherence? Although this problem has been extensively studied in statistics, economics, and philosophy, one standard assumption is seldom questioned: the set of events to which probabilities are assigned – the *agenda* – is a σ -algebra: it is closed under negation (complementation) and countable disjunction (union) of events. In practice, however, decision makers or expert panels may not be interested in such a rich set of events. They may be interested, for example, in the probability of a blizzard and the probability of an interest-rate increase, but not in the probability of a blizzard *or* an interest-rate increase. Of course, the assumption that the agenda is a σ -algebra is convenient: probability functions are defined on σ -algebras, and thus one can view probabilistic opinion pooling as the aggregation of probability functions. But convenience is no ultimate justification. Real-world expert committees typically do not assign probabilities to all events in a σ -algebra. Instead, they focus on a limited set of relevant events, which need not contain all disjunctions of its elements, let alone all disjunctions of countably infinite length.

There are two reasons why a disjunction of relevant events, or another logical combination, may not be relevant. *Either* we are not interested in the probability of such ‘artificial’ composite events. *Or* we (or the decision makers or experts) are unable to assign subjective probabilities to them. To see why it can be difficult to assign a subjective probability to a *logical combination* of ‘basic’ events – such as ‘a blizzard *or* an interest-rate increase’ – note that it is not enough to assign probabilities to the underlying basic events: various probabilistic dependencies also affect the probability of the composite event, and these may be the result of complex causal interconnections (such as the causal effects between basic events and their possible common causes).

We investigate probabilistic opinion pooling for general agendas, dropping the assumption of a σ -algebra. Thus any set of events that is closed under negation (complementation) can qualify as an agenda. The general notion of an agenda is imported from the theory of binary judgment aggregation (e.g., List and Pettit 2002, 2004; Pauly and van Hees 2006; Dietrich 2006; Dietrich and List 2007a, 2013; Nehring and Puppe 2010; Dokow and Holzman 2010; Dietrich and Mongin 2010). We impose two axiomatic requirements on probabilistic opinion pooling:

- (i) the familiar ‘independence’ requirement, according to which the collectively assigned probability for each event should depend only on the probabilities that the individuals assign to that event;
- (ii) the requirement that certain unanimous individual judgments should be preserved; we consider stronger and weaker variants of this requirement.

We prove two main results:

- For a large class of agendas – with σ -algebras as special cases – any opinion pooling function satisfying (i) and (ii) is *linear*: the collective probability of each event in the agenda is a weighted linear average of the individuals’ probabilities of that event, where the weights are the same for all events.
- For an even larger class of agendas, any opinion pooling function satisfying (i) and (ii) is *neutral*: the collective probability of each event in the agenda is some (possibly non-linear) function of the individuals’ probabilities of that event, *where the function is the same for all events*.

We state three versions of each result, which differ in the nature of the unanimity-preservation requirement and in the class of agendas to which they apply. Our results generalize a classic characterization of linear pooling in the special case where the agenda is σ -algebra (Aczél and Wagner 1980 and McConway 1981).¹ For a σ -algebra, every neutral pooling function is automatically linear, so that neutrality and linearity are equivalent here (McConway 1981 and Wagner 1982).² As we will see, this fact does not carry over to general agendas: many agendas permit neutral but non-linear opinion pooling functions.

Some of our results apply even to agendas containing only logically independent events, such as ‘a blizzard’ and ‘an interest-rate increase’ (and their negations), but no disjunctions or conjunctions of these events. Such agendas are relevant in practical applications where the events in question are only *probabilistically* dependent (correlated), but not *logically* dependent. If the agenda is a σ -algebra, by contrast, it is replete with logical interconnections. By focusing on σ -algebras alone, the standard results on probabilistic opinion pooling have therefore excluded many realistic applications.

We also present a new illustrative application of probabilistic opinion pooling, namely to *probabilistic preference aggregation*. Here each individual assigns subjective probabilities to events of the form ‘ x is preferable than y ’ (or ‘ x is better than y ’), where x and y range over a given set of alternatives. These probability

¹Specifically, if the agenda is a σ -algebra (with more than four events), linear pooling functions are the only pooling functions which satisfy independence and preserve unanimous probabilistic judgments (Aczél and Wagner 1980, McConway 1981). Linearity and neutrality (the latter sometimes under the names *strong label neutrality* or *strong setwise function property*) are among the most widely studied properties of opinion pooling functions. Linear pooling goes back to Stone (1961) or even Laplace, and neutral pooling to McConway (1981) and Wagner (1982). For extensions of (or alternatives to) the classic characterization of linear pooling, see Wagner (1982, 1985), Aczél, Ng, and Wagner (1984), Genest (1984), Mongin (1995), and Chambers (2007). All these works retain the assumption that the agenda is a σ -algebra. Genest and Zidek (1986) and Clemen and Winkler (1999) provide surveys of the classic literature. For opinion pooling under asymmetric information, see Dietrich (2010). For the aggregation of qualitative rather than quantitative probabilities, see Weymark (1997). For a computational, non-axiomatic approach to the aggregation of *partial* probability assignments, where individuals do not assign probabilities to all events in the σ -algebra, see Osherson and Vardi (2006).

²This assumes that the σ -algebra contains more than four events.

assignments may be interpreted as beliefs about which preferences are the ‘correct’ ones (e.g., which correctly capture objective quality comparisons between the alternatives). Alternatively, they may be interpreted as vague or fuzzy preferences. We then seek to arrive at corresponding collective probability assignments.

Each of our linearity or neutrality results (with one exception) is logically tight: the linearity or neutrality conclusion follows *if and only if* the agenda falls into a relevant class. In other words, we characterize the agendas for which our axiomatic requirements lead to linear or neutral aggregation. We thereby adopt the state-of-the-art approach in binary judgment-aggregation theory, which is to characterize the agendas leading to certain possibilities or impossibilities of aggregation. This approach was introduced by Nehring and Puppe (2002) in related work on strategy-proof social choice and subsequently applied throughout binary judgment-aggregation theory. One of our contributions is to show how it can be applied in the area of probabilistic opinion pooling.

We conclude by comparing our results with their analogues in binary judgment-aggregation theory and in Arrovian preference aggregation theory. Interestingly, the conditions leading to linear pooling in probability aggregation correspond exactly to the conditions leading to a dictatorship of one individual in both binary judgment aggregation and Arrovian judgment aggregation. This yields a new unified perspective on several at first sight disparate aggregation problems.

2 The framework

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

The agenda. Let Ω be a non-empty set of possible *worlds* (or *states*). An *event* is a subset A of Ω ; its *complement* (‘negation’) is denoted $A^c := \Omega \setminus A$. The *agenda* is the set of events to which probabilities are assigned. Traditionally, the agenda has been assumed to be a σ -*algebra* (i.e., closed under complementation and countable union, and thereby also under countable intersection). Here, we drop that assumption. As already noted, we may exclude some events from the agenda, *either* because they are of no interest, *or* because no probability assignments are available for them. For example, the agenda may contain the events that global warming will continue, that interest rates will remain low, and that the UK will remain in the European Union, but not the disjunction of these events. Formally, we define an *agenda* as a non-empty set X of events which is closed under complementation, i.e., $A \in X \Rightarrow A^c \in X$. Examples are $X = \{A, A^c\}$ or $X = \{A, A^c, B, B^c\}$, where A and B may or may not be logically related.

An example of an agenda without conjunctions or disjunctions. Suppose each possible world is a vector of three binary characteristics. The first takes the value 1 if atmospheric CO₂ is above some threshold, and 0 otherwise. The second takes the value 1 if there is a mechanism to the effect that *if* atmospheric CO₂ is above that threshold, *then* Arctic summers are ice-free, and 0 otherwise. The third takes the value 1 if Arctic summers are ice-free, and 0 otherwise. Thus the set of possible worlds is the set of all triples of 0s and 1s, excluding the inconsistent triple in which the first and second characteristics are 1 and the third is 0, i.e., $\Omega = \{0, 1\}^3 \setminus \{(1, 1, 0)\}$. We now define an agenda X consisting of $A, A \rightarrow B, B$, and their complements, where A is the event of a positive first characteristic, $A \rightarrow B$ the event of a positive second characteristic, and B the event of a positive third characteristic. (We use the sentential notation ‘ $A \rightarrow B$ ’ for better readability; formally, each of A, B , and $A \rightarrow B$ are subsets of Ω .³) Although there are some logical connections between these events (in particular, A and $A \rightarrow B$ are inconsistent with B^c), the set X contains no conjunctions or disjunctions.

Probabilistic opinions. We begin with the notion of a *probability function*. The classical focus on agendas that are σ -algebras is motivated by the fact that such functions are defined on σ -algebras. Formally, a *probability function* on a σ -algebra Σ is a function $P : \Sigma \rightarrow [0, 1]$ such that $P(\Omega) = 1$ and P is σ -additive (i.e., $P(A_1 \cup A_2 \cup \dots) = P(A_1) + P(A_2) + \dots$ for every sequence of pairwise disjoint events $A_1, A_2, \dots \in \Sigma$). In the context of an arbitrary agenda X , we speak of ‘opinion functions’ rather than ‘probability functions’. Formally, an *opinion function* for an agenda X is a function $P : X \rightarrow [0, 1]$ which is *probabilistically coherent*, i.e., extendable to a probability function on the σ -algebra generated by X . This σ -algebra is denoted $\sigma(X)$ and defined as the smallest σ -algebra that includes X . It can be constructed by closing X under countable unions and complements.⁴ In our expert-committee example, we have $\sigma(X) = 2^\Omega$, and an opinion function cannot assign probability 1 to all of $A, A \rightarrow B$, and B^c . (This would not be extendable to a well-defined probability function on 2^Ω , given that $A \cap (A \rightarrow B) \cap B^c = \emptyset$.) We write \mathcal{P}_X to denote the set of all opinion functions for the agenda X . If X is a σ -algebra, \mathcal{P}_X is the set of all probability functions on it.

³Note that $A \rightarrow B$ (‘if A then B ’) is best interpreted as a non-material conditional, since its negation, unlike that of a material conditional, is consistent with the negation of its antecedent, A (i.e., $A^c \cap (A \rightarrow B)^c \neq \emptyset$). (A material conditional is always true when its antecedent is false.) The only assignment of truth-values to the events $A, A \rightarrow B$, and B that is ruled out is $(1, 1, 0)$. If we wanted to re-interpret \rightarrow as a material conditional, we would have to rule out in addition the truth-value assignments $(0, 0, 0)$, $(0, 0, 1)$, and $(1, 0, 1)$, which would make little sense in the present example. The event $A \rightarrow B$ would become $A^c \cup B (= (A \cap B^c)^c)$, and the agenda would no longer be free from conjunctions or disjunctions. However, the agenda would still not be a σ -algebra. For a discussion of non-material conditionals, see, e.g., Priest (2001).

⁴Whenever X contains A and B , then $\sigma(X)$ contains $A \cup B, (A \cup B)^c, (A \cup B)^c \cup B$, and so on. In some cases, *all* events may be constructible from events in X , so that $\sigma(X) = 2^\Omega$.

Opinion pooling. Given the agenda X , a combination of opinion functions across the n individuals, (P_1, \dots, P_n) , is called a *profile (of opinion functions)*. An *(opinion) pooling function* is a function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$, which assigns to each profile (P_1, \dots, P_n) a collective opinion function $P = F(P_1, \dots, P_n)$, also denoted P_{P_1, \dots, P_n} . For instance, P_{P_1, \dots, P_n} could be the arithmetic average $\frac{1}{n}P_1 + \dots + \frac{1}{n}P_n$.

Linearity and neutrality. A pooling function is *linear* if there exist real-valued weights $w_1, \dots, w_n \geq 0$ with $w_1 + \dots + w_n = 1$ such that, for every profile $(P_1, \dots, P_n) \in \mathcal{P}_X^n$,

$$P_{P_1, \dots, P_n}(A) = \sum_{i=1}^n w_i P_i(A) \text{ for all } A \in X.$$

If $w_i = 1$ for some ‘expert’ i , we obtain an *expert rule* given by $P_{P_1, \dots, P_n} = P_i$. More generally, a pooling function is *neutral* if there exists *some* function $D : [0, 1]^n \rightarrow [0, 1]$ such that, for every profile $(P_1, \dots, P_n) \in \mathcal{P}_X^n$,

$$P_{P_1, \dots, P_n}(A) = D(P_1(A), \dots, P_n(A)) \text{ for all } A \in X. \quad (1)$$

We call D the *local pooling criterion*. Since it does not depend on the event A , all events are treated equally (‘neutrality’). Linearity is the special case in which D is a weighted linear averaging criterion of the form $D(x) = \sum_{i=1}^n w_i x_i$ for all $x \in [0, 1]^n$. Note that, while every combination of weights $w_1, \dots, w_n \geq 0$ with sum-total 1 defines a proper linear pooling function (since linear averaging preserves probabilistic coherence), a given non-linear function $D : [0, 1]^n \rightarrow [0, 1]$ might not define a proper pooling function. Formula (1) might not yield a well-defined – i.e., probabilistically coherent – opinion function. We will show that whether there can be neutral but non-linear pooling functions depends on the agenda in question. If the agenda is a σ -algebra, the answer is known to be negative (assuming $|X| > 4$). However, we will also identify agendas for which the answer is positive.

Some logical terminology. An event A is *contingent* if it is neither the empty set \emptyset (impossible) nor the universal set Ω (necessary). A set S of events is *consistent* if its intersection $\bigcap_{A \in S} A$ is non-empty, and *inconsistent* otherwise. A set S of events *entails* another event B if the intersection of S is included in B (i.e., $\bigcap_{A \in S} A \subseteq B$).

Two kinds of applications. It is useful to distinguish between two kinds of applications of probabilistic opinion pooling. We may be interested in either of the following:

- (a) the probabilities of certain propositions expressed in natural language, such as ‘it will rain tomorrow’ or ‘the new legislation will be repealed’;

- (b) the distribution of some real-valued (or vector-valued) random variable, such as the number of insurance claims over a given period, or tomorrow's price of a given share, or the weight of a randomly picked potato from some farm.

Arguably, probabilistic opinion pooling on general agendas is more relevant to applications of type (a) than to applications of type (b). An application of type (a) typically gives rise to an agenda expressible in natural language which does not constitute a σ -algebra. It is then implausible to replace X with the σ -algebra $\sigma(X)$, many elements of which represent unduly complex combinations of other events. Further, even when $\sigma(X)$ is finite, it may be enormous. If X contains at least k logically independent events, then $\sigma(X)$ contains at least 2^{2^k} events, so its size grows double-exponentially in k .⁵ This suggests that, unless k is small, $\sigma(X)$ may be too large to serve as an agenda in practice. By contrast, an application of type (b) plausibly gives rise to an agenda that is a σ -algebra. Here, the decision makers may *need* a full probability distribution over the σ -algebra, and they may also be *able* to specify such a distribution. For instance, a market analyst estimating next month's distribution of Apple's share price might decide to specify a log-normal distribution. This, in turn, requires the specification of only two parameters: the mean and the variance of the exponential of the share price. We discuss opinion pooling problems of type (b) in a companion paper (Dietrich and List 2014), where they are one of our principal applications.

3 Axiomatic requirements on opinion pooling

We now introduce some requirements on opinion pooling functions.

3.1 The independence requirement

Our first requirement, familiar from the literature, says that the collective probability of each event in the agenda should depend only on the individual probabilities of *that* event. This requirement is sometimes also called the *weak setwise function property*.

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

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

One justification for independence is the Condorcetian idea that the collective view on any issue should depend only on individual views on that issue. This

⁵For instance, if X contains $k = 2$ logically independent events, say A and B , then X includes a partition \mathcal{A} of Ω into $2^k = 4$ non-empty events, namely $\mathcal{A} = \{A \cap B, A \cap B^c, A^c \cap B, A^c \cap B^c\}$, and hence X includes the set $\{\cup_{C \in \mathcal{C}} C : \mathcal{C} \subseteq \mathcal{A}\}$ containing $2^{2^k} = 16$ events.

reflects a *local*, rather than *holistic*, understanding of aggregation. (On a *holistic* understanding, the collective view on an issue may be influenced by individual views on other issues.) Independence, understood in this way, becomes less compelling if the agenda contains ‘artificial’ events, such as conjunctions of intuitively unrelated events, as in the case of a σ -algebra. It would be implausible, for instance, to disregard the individual probabilities assigned to ‘a blizzard’ and to ‘an interest-rate increase’ when determining the collective probability of the disjunction of these events. Here, however, we focus on general agendas, where the Condorcetian justification for independence is more plausible.

There are also two pragmatic justifications for independence; these apply even when the agenda is a σ -algebra. First, aggregating probabilities issue-by-issue is informationally and computationally less demanding than a holistic approach and thus easier to implement in practice. Second, independence prevents certain types of agenda manipulation – the attempt by an agenda setter to influence the collective probability assigned to some events by adding other events to, or removing them from, the agenda.⁶ Nonetheless, independence should not be accepted uncritically, since it is vulnerable to a number of well-known objections.⁷

3.2 The consensus-preservation requirement

Our next requirement says that if all individuals assign probability 1 (certainty) to an event in the agenda, then its collective probability should also be 1.

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

Like independence, this requirement is familiar from the literature, where it is sometimes expressed as a zero-probability preservation requirement. In the case of general agendas, we can also formulate several strengthened variants of the requirement, which extend it to other forms of consensus. Although these variants are not as compelling as their original precursor, they are still defensible in some cases. Moreover, when the agenda is a σ -algebra, they all collapse back into consensus preservation in its original form.

⁶When X is a σ -algebra, McConway (1981) shows that independence (his weak setwise function property) is equivalent to the *marginalization property*, which requires aggregation to commute with the operation of reducing the σ -algebra to some sub- σ -algebra $\Sigma^* \subseteq X$. A similar result holds for general agendas X .

⁷When the agenda is a σ -algebra, independence conflicts with the preservation of unanimously held judgments of probabilistic independence, assuming non-dictatorial aggregation (Genest and Wagner 1987; Bradley, Dietrich, and List 2014). Whether this objection also applies in the case of general agendas depends on the precise nature of the agenda. Another objection is that independence is not generally compatible with *external Bayesianity*, the requirement that aggregation commute with Bayesian updating of probabilities in light of new information.

To introduce the different extensions of consensus preservation, we begin by drawing a distinction between ‘explicitly revealed’, ‘implicitly revealed’, and ‘unrevealed’ beliefs:

- Individual i ’s *explicitly revealed beliefs* are the probabilities assigned to events in the agenda X by the opinion function P_i .
- Individual i ’s *implicitly revealed beliefs* are the probabilities assigned to any events in $\sigma(X)\setminus X$ by *every* probability function on $\sigma(X)$ extending the opinion function P_i ; we call such a probability function an *extension* of P_i and use the notation \bar{P}_i . These probabilities are ‘implied’ by the opinion function P_i . For instance, if P_i assigns probability 1 to an event A in the agenda X , this ‘implies’ an assignment of probability 1 to all events B outside the agenda that are of the form $B \supseteq A$.
- Individual i ’s *unrevealed beliefs* are probabilities for events in $\sigma(X)\setminus X$ that cannot be deduced from the opinion function P_i . These are only privately held. For instance, the opinion function P_i may admit extensions which assign probability 1 to an event B but may also admit extensions which assign a lower probability. Here, individual i ’s belief about B is unrevealed.

Consensus preservation in its original form concerns only explicitly revealed beliefs. The first strengthened variant extends the requirement to implicitly revealed beliefs. Let us say that an opinion function P on X *implies certainty of an event* A if $\bar{P}(A) = 1$ for *every* extension \bar{P} of P .

Implicit consensus preservation. For all $A \in \sigma(X)$ and all $P_1, \dots, P_n \in \mathcal{P}_X$, if, for all i , P_i implies certainty of A , then P_{P_1, \dots, P_n} also implies certainty of A .

This ensures that whenever all individuals either explicitly or implicitly assign probability 1 to some event, this is preserved at the collective level. Arguably, this requirement is almost as plausible as consensus preservation in its original form.

The second extension concerns unrevealed beliefs. Informally, it says that a unanimous assignment of probability 1 to some event should never be overruled, *even if it is unrevealed*. This is operationalized as the requirement that if every individual’s opinion function is *consistent* with the assignment of probability 1 to some event (so that we cannot rule out the possibility of the individuals’ privately making that assignment), then the collective opinion function should also be consistent with it. Formally, we say that an opinion function P on X is *consistent with certainty of an event* A if there exists *some* extension \bar{P} of P such that $\bar{P}(A) = 1$.

Consensus compatibility. For all $A \in \sigma(X)$ and all $P_1, \dots, P_n \in \mathcal{P}_X$, if, for all i , P_i is consistent with certainty of A , then P_{P_1, \dots, P_n} is also consistent with certainty of A .

The rationale for this requirement is a precautionary one: if it is *possible* that

all individuals assign probability 1 to some event (though this may be unrevealed), the collective opinion function should not *rule out* certainty of A .

A third extension of consensus preservation concerns conditional beliefs. It looks more complicated than consensus compatibility, but it is less demanding. Its initial motivation is the idea that if all individuals are certain of some event in the agenda conditional on another event, then this conditional belief should be preserved collectively. For instance, if everyone is certain that there will be a famine, given a civil war, this belief should also be held collectively. Unfortunately, however, we cannot define individual i 's conditional probability of an event A , given another event B , simply as $P_i(A|B) = P_i(A \cap B)/P_i(B)$ (where $P_i(B) \neq 0$ and P_i is individual i 's opinion function). This is because, even when A and B are in X , the event $A \cap B$ may be outside X and thus outside the domain of P_i . So, we cannot know whether the individual is certain of A given B . But we can ask whether he or she *could* be certain of A given B , i.e., whether $\overline{P}_i(A|B) = 1$ for *some* extension \overline{P} of P .

This motivates the requirement that if each individual *could* be certain of A given B , then the collective opinion function should also be *consistent* with this 'conditional certainty'. Again, this can be interpreted as requiring the preservation of certain *unrevealed* beliefs. A unanimous assignment of conditional probability 1 to one event, given another, should not be overruled, *even if it is unrevealed*.

We capture this in the following way. Suppose there is a finite set of pairs of events in X – call them (A, B) , (A', B') , (A'', B'') , and so on – such that each individual could be simultaneously certain of A given B , of A' given B' , of A'' given B'' , and so on. Then the collective opinion function should also be consistent with conditional certainty of A given B , A' given B' , and so on. Formally, for any finite set S of pairs (A, B) of events in X , we say that an opinion function P on X is *consistent with conditional certainty of all (A, B) in S* if there exists *some* extension \overline{P} of P such that $\overline{P}(A|B) = 1$ for all (A, B) in S for which $P(B) \neq 0$.

Conditional consensus compatibility. For all finite sets S of pairs of events in X and all $P_1, \dots, P_n \in \mathcal{P}_X$, if, for all i , P_i is consistent with conditional certainty of all (A, B) in S , then P_{P_1, \dots, P_n} is also consistent with conditional certainty of all (A, B) in S .

The following proposition summarizes the logical relationships between the different consensus-preservation requirements; a proof is given in the Appendix.

- Proposition 1** (a) *Consensus preservation is implied by each of (i) implicit consensus preservation, (ii) consensus compatibility, and (iii) conditional consensus compatibility, and is equivalent to each of (i), (ii), and (iii) if the agenda X is a σ -algebra.*
- (b) *Consensus compatibility implies conditional consensus compatibility.*

Each of our characterization results below uses consensus preservation in either its original form or one of the strengthened forms. Implicit consensus preservation does not appear in any of our results; we have included it here for the sake of conceptual completeness.⁸

4 When is opinion pooling neutral?

We now show that, for many agendas, the neutral pooling functions are the only pooling functions satisfying independence and consensus preservation in either its original form or one of the strengthened forms. The stronger the consensus-preservation requirement, the larger the class of agendas for which our characterization of neutral pooling holds. For the moment, we set aside the question of whether independence and consensus preservation imply linearity as well as neutrality; we address this question in the next section.

4.1 Three theorems

We begin with the strongest of our consensus-preservation requirements, i.e., consensus compatibility. If we impose this requirement, our characterization of neutral pooling holds for a very large class of agendas: all *non-nested* agendas. We call an agenda X *nested* if it has the form $X = \{A, A^c : A \in X_+\}$ for some set X_+ ($\subseteq X$) that is linearly ordered by set-inclusion, and *non-nested* otherwise. For example, binary agendas of the form $X = \{A, A^c\}$ are nested: take $X_+ := \{A\}$, which is trivially linearly ordered by set-inclusion. Also, the agenda $X = \{(-\infty, t], (t, \infty) : t \in \mathbb{R}\}$ (where the set of possible worlds is $\Omega = \mathbb{R}$) is nested: take $X_+ := \{(-\infty, t] : t \in \mathbb{R}\}$, which is linearly ordered by set-inclusion.

By contrast, any agenda consisting of multiple logically independent pairs A, A^c is non-nested, i.e., X is non-nested if $X = \{A_k, A_k^c : k \in K\}$ with $|K| \geq 2$ such that every subset $S \subseteq X$ containing precisely one member of each pair $\{A_k, A_k^c\}$ (with $k \in K$) is consistent. As mentioned in the introduction, such agendas are of practical importance because many decision problems involve events that exhibit only probabilistic dependencies (correlations), but no logical ones. Another example of a non-nested agenda is the one in the expert-committee example above, containing $A, A \rightarrow B, B$, and their complements.

Theorem 1 (a) *For any non-nested agenda X , every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus compatibility is neutral.*

⁸An interesting fourth variant is the requirement obtained by combining the antecedent of implicit consensus preservation with the conclusion of consensus compatibility. This condition weakens both implicit consensus preservation and consensus compatibility, while still strengthening the initial consensus preservation requirement.

- (b) For any nested agenda X ($\neq \{\emptyset, \Omega\}$), there exists a non-neutral pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus compatibility.

Part (b) shows that the agenda condition used in part (a) – non-nestedness – is tight: whenever the agenda is nested, non-neutral pooling functions become possible. However, these pooling functions are non-neutral only in a limited sense: although the pooling criterion D_A need not be the same for all events $A \in X$, it must still be the same for all $A \in X_+$, and the same for all $A \in X \setminus X_+$ (with X_+ as defined above), so that pooling is ‘neutral within X_+ ’ and ‘neutral within $X \setminus X_+$ ’. This is clear from the proof.⁹

What happens if we weaken the requirement of consensus compatibility to conditional consensus compatibility? Both parts of Theorem 1 continue to hold, though part (a) becomes logically stronger, and part (b) logically weaker. Let us state the modified theorem explicitly:

Theorem 2 (a) For any non-nested agenda X , every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility is neutral.

- (b) For any nested agenda X ($\neq \{\emptyset, \Omega\}$), there exists a non-neutral pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility.

The situation changes once we weaken the consensus requirement further, namely to consensus preservation *simpliciter*. The class of agendas for which our characterization of neutrality holds shrinks significantly, namely to the class of *path-connected* agendas. Path-connectedness is an important condition in judgment-aggregation theory, where it was introduced by Nehring and Puppe (2010) (under the name ‘total blockedness’) and has been used, for example, to generalize Arrow’s theorem (Dietrich and List 2007a, Dokow and Holzman 2010).

To define path-connectedness, we require one preliminary definition. Given an agenda X , we say that an event $A \in X$ *conditionally entails* another event $B \in X$, written $A \vdash^* B$, if there exists 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. For instance, if $\emptyset \neq A \subseteq B \neq \Omega$, then $A \vdash^* B$ (take $Y = \emptyset$; in fact, this is even an *unconditional* entailment). Also, for the agenda of our expert committee, $X = \{A, A^c, A \rightarrow B, (A \rightarrow B)^c, B, B^c\}$, we have $A \vdash^* B$ (take $Y = \{A \rightarrow B\}$).

We call an agenda X *path-connected* if any two events $A, B \in X \setminus \{\emptyset, \Omega\}$ can be connected by a path of conditional entailments, i.e., there exist events $A_1, \dots, A_k \in X$ ($k \geq 1$) such that $A = A_1 \vdash^* A_2 \vdash^* \dots \vdash^* A_k = B$. An example of

⁹As a consequence, full neutrality follows *even for nested agendas* if independence is slightly strengthened by requiring that $D_A = D_{A^c}$ for some $A \in X \setminus \{\emptyset, \Omega\}$.

a path-connected agenda is $X := \{A, A^c : A \subseteq \mathbb{R} \text{ is a bounded interval}\}$, where the underlying set of worlds is $\Omega = \mathbb{R}$. For instance, there is a path of conditional entailments from $[0, 1] \in X$ to $[2, 3] \in X$ given by $[0, 1] \vdash^* [0, 3] \vdash^* [2, 3]$. To establish $[0, 1] \vdash^* [0, 3]$, it suffices to conditionalize on the empty set of events $Y = \emptyset$ (i.e., $[0, 1]$ even unconditionally entails $[0, 3]$). To establish $[0, 3] \vdash^* [2, 3]$, one may conditionalize on $Y = \{[2, 4]\}$.

Many agendas are not path-connected, including all nested agendas ($\neq \{\emptyset, \Omega\}$) and the agenda in our expert-committee example. The following result holds.

Theorem 3 (a) *For any path-connected agenda X , every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation is neutral.*
 (b) *For any non-path-connected agenda X (finite and distinct from $\{\emptyset, \Omega\}$), there exists a non-neutral pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation.*

4.2 Proof sketches

We now outline the proofs of Theorems 1 to 3. (Details are given in the Appendix.) We begin with part (a) of each theorem. Theorem 1(a) follows from Theorem 2(a), since both results apply to the same agendas but Theorem 1(a) uses a stronger consensus requirement.

To prove Theorem 2(a), we define a binary relation \sim on the set of all contingent events in the agenda. Recall that two events A and B are *exclusive* if $A \cap B = \emptyset$ and *exhaustive* if $A \cup B = \Omega$. For any $A, B \in X \setminus \{\emptyset, \Omega\}$, we define

there is a finite sequence $A_1, \dots, A_k \in X$ of length $k \geq 1$ with $A_1 = A$
 $A \sim B \Leftrightarrow$ and $A_k = B$ such that any adjacent A_j, A_{j+1} are neither exclusive nor
 exhaustive.

Theorem 2(a) then follows immediately from the following two lemmas (proved in the Appendix).

Lemma 1 *For any agenda X ($\neq \{\emptyset, \Omega\}$), the relation \sim is an equivalence relation on $X \setminus \{\emptyset, \Omega\}$, with exactly two equivalence classes if X is nested, and exactly one if X is non-nested.*

Lemma 2 *For any agenda X ($\neq \{\emptyset, \Omega\}$), a pooling function satisfying independence and conditional consensus compatibility is neutral on each equivalence class with respect to \sim (i.e., the local pooling criterion is the same for all events in the same equivalence class).*

The proof of Theorem 3(a) uses the following lemma (broadly analogous to a lemma in binary judgment-aggregation theory; e.g., Nehring and Puppe 2010 and Dietrich and List 2007a).

Lemma 3 *For any pooling function satisfying independence and consensus preservation, and all events A and B in the agenda X , if $A \vdash^* B$ then $D_A \leq D_B$, where D_A and D_B are the local pooling criteria for A and B , respectively. (Here $D_A \leq D_B$ means that, for all (p_1, \dots, p_n) , $D_A(p_1, \dots, p_n) \leq D_B(p_1, \dots, p_n)$.)*

To see why Theorem 3(a) follows, simply note that $D_A \leq D_B$ whenever there is a path of conditional entailments from $A \in X$ to $B \in X$ (by repeated application of the lemma); thus, $D_A = D_B$ whenever there are paths in both directions, as is guaranteed if the agenda is path-connected and $A, B \notin \{\emptyset, \Omega\}$.

Part (b) of each theorem can be proved by explicitly constructing a non-neutral pooling function – for an agenda of the relevant kind – which satisfies independence and the appropriate consensus-preservation requirement. In the case of Theorem 3(b), this pooling function is very complex, and hence we omit it in the main text. In the case of Theorems 1(a) and 1(b), the idea can be described informally. Recall that a nested agenda X can be partitioned into two subsets, X_+ and $X \setminus X_+ = \{A^c : A \in X_+\}$, each of which is linearly ordered by set-inclusion. The opinion pooling function constructed has the property that (i) all events A in X_+ have the same local pooling criterion $D = D_A$, which can be defined, for example, as the square of a linear pooling criterion (i.e., we first apply a linear pooling criterion and then take the square), and (ii) all events in $X \setminus X_+$ have the same ‘complementary’ pooling criterion D^* , defined as $D^*(x_1, \dots, x_n) = 1 - D(1 - x_1, \dots, 1 - x_n)$ for all $(x_1, \dots, x_n) \in [0, 1]^n$. Showing that the resulting pooling function is well-defined and satisfies all the relevant requirements involves some technicality, in part because we allow the agenda to have any cardinality.

5 When is opinion pooling linear?

As we have seen, for many agendas, only neutral pooling functions can satisfy our two requirements. But are these functions also linear? As we now show, the answer depends on the agenda. If we suitably restrict the class of agendas considered in part (a) of each of our previous theorems, we can derive linearity rather than just neutrality. Similarly, we can expand the class of agendas considered in part (b) of each theorem, and replace non-neutrality with non-linearity.

5.1 Three theorems

As in the previous section, we begin with the strongest consensus-preservation requirement, i.e., consensus compatibility. While this requirement leads to neutrality for all non-nested agendas (by Theorem 1), it leads to linearity for all non-nested agendas above a certain size.

- Theorem 4** (a) *For any non-nested agenda X with $|X \setminus \{\Omega, \emptyset\}| > 4$, every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus compatibility is linear.*
- (b) *For any other agenda X ($\neq \{\emptyset, \Omega\}$), there exists a non-linear pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus compatibility.*

Next, let us weaken the requirement of consensus compatibility to conditional consensus compatibility. While this requirement leads to neutrality for all *non-nested* agendas (by Theorem 2), it leads to linearity only for *non-simple* agendas. Like path-connected agendas, non-simple agendas play an important role in binary judgment-aggregation theory, where they are the agendas susceptible to the analogues of Condorcet’s paradox: the possibility of inconsistent majority judgments (e.g., Dietrich and List 2007b, Nehring and Puppe 2007).

To define non-simplicity, we first require a preliminary definition. We call a set of events Y *minimal inconsistent* if it is inconsistent but every proper subset $Y' \subsetneq Y$ is consistent. Examples of minimal inconsistent sets are (i) $\{A, B, (A \cap B)^c\}$, where A and B are logically independent events, and (ii) $\{A, A \rightarrow B, B^c\}$, with A , B , and $A \rightarrow B$ as defined in the expert-committee example above. In each case, the three events are mutually inconsistent, but any two of them are mutually consistent. The notion of a minimal inconsistent set is useful for characterizing logical dependencies between the events in the agenda. Trivial examples of minimal inconsistent subsets of the agenda are those of the form $\{A, A^c\} \subseteq X$, where A is contingent. But many interesting agendas have more complex minimal inconsistent subsets. One may regard $\sup_{Y \subseteq X: Y \text{ is minimal inconsistent}} |Y|$ as a measure of the complexity of the logical dependencies in the agenda X . Given this idea, we call an agenda X *non-simple* if it has at least one minimal inconsistent subset $Y \subseteq X$ containing more than two (but not uncountably many¹⁰) events, and *simple* otherwise. For instance, the agenda consisting of A , $A \rightarrow B$, B and their complements in our expert-committee example is non-simple (take $Y = \{A, A \rightarrow B, B^c\}$).

Non-simplicity lies logically between non-nestedness and path-connectedness: it implies non-nestedness, and is implied by path-connectedness (if $X \neq \{\Omega, \emptyset\}$).¹¹

¹⁰This countability addition can often be dropped because all minimal inconsistent sets $Y \subseteq X$ are automatically finite or at least countable. This is so if X is finite or countably infinite, and also if the underlying set of worlds Ω is countable. It can further be dropped in case the events in X are represented by sentences in a language. Then, provided this language belongs to a compact logic, all minimal inconsistent sets $Y \subseteq X$ are finite (because any inconsistent set has a finite inconsistent subset). By contrast, if X is a σ -algebra and has infinite cardinality, then it usually contains events not representing sentences, because countably infinite disjunctions cannot be formed in a language. Such agendas often have uncountable minimal inconsistent subsets. For instance, if X is the σ -algebra of Borel-measurable subsets of \mathbb{R} , then its subset $Y = \{\mathbb{R} \setminus \{x\} : x \in \mathbb{R}\}$ is uncountable and minimal inconsistent. This agenda is nonetheless non-simple, since it also has many *finite* minimal inconsistent subsets Y with $|Y| \geq 3$ (e.g., $Y = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$).

¹¹To give an example of a non-nested but simple agenda X , let $X = \{A, A^c, B, B^c\}$, where the events A and B are logically independent, i.e., $A \cap B, A \cap B^c, A^c \cap B, A^c \cap B^c \neq \emptyset$. Clearly, this

To see how exactly non-simplicity strengthens non-nestedness, note the following fact (Dietrich 2013):

- Fact** (a) *An agenda X (with $|X \setminus \{\Omega, \emptyset\}| > 4$) is non-nested if and only if it has at least one subset Y with $|Y| \geq 3$ such that $(Y \setminus \{A\}) \cup \{A^c\}$ is consistent for each $A \in Y$.*
- (b) *An agenda X (with $|X \setminus \{\Omega, \emptyset\}| > 4$) is non-simple if and only if it has at least one inconsistent subset Y (of countable size) with $|Y| \geq 3$ such that $(Y \setminus \{A\}) \cup \{A^c\}$ is consistent for each $A \in Y$.*

Note that the characterizing condition in (b) can be obtained from the one in (a) simply by replacing ‘subset Y ’ with ‘inconsistent subset Y (of countable size)’.

- Theorem 5** (a) *For any non-simple agenda X with $|X \setminus \{\Omega, \emptyset\}| > 4$, every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility is linear.*
- (b) *For any simple agenda X (finite and distinct from $\{\emptyset, \Omega\}$), there exists a non-linear pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility.*

Finally, we turn to the least demanding consensus requirement, namely consensus preservation *simpliciter*. We have seen that this requirement leads to neutral pooling if the agenda is path-connected (by Theorem 3). To obtain a characterization of linear pooling, path-connectedness alone is not enough. In the following theorem, we impose an additional condition on the agenda. We call an agenda X *partitional* if it has a subset Y which partitions Ω into at least three non-empty events (where Y is finite or countably infinite), and *non-partitional* otherwise. (A subset Y of X *partitions* Ω if the elements of Y are individually non-empty, pairwise disjoint, and cover Ω .) For instance, X is partitional if it contains (non-empty) events A , $A^c \cap B$, and $A^c \cap B^c$; simply let $Y = \{A, A^c \cap B, A^c \cap B^c\}$.

- Theorem 6** (a) *For any path-connected and partitional agenda X , every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation is linear.*
- (b) *For any non-path-connected (finite) agenda X , there exists a non-linear pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation.*

agenda is non-nested. It is simple since its only minimal inconsistent subsets are $\{A, A^c\}$ and $\{B, B^c\}$. To give an example of a non-path-connected, but non-simple agenda, let X consist of $A, A \rightarrow B, B$ and their complements, as in our example above. We have already observed that it is non-simple. To see that it is not path-connected, note, for example, that there is no path of conditional entailments from B to B^c .

Part (b) shows that one of theorem’s agenda conditions, path-connectedness, is necessary for the characterization of linear pooling (which is unsurprising, as it is necessary for the characterization of neutral pooling). By contrast, the other agenda condition, partitionality, is not necessary: linearity also follows from independence and consensus preservation for some non-partitional but path-connected agendas. So, the agenda conditions of part (a) are non-minimal. We leave the task of finding minimal agenda conditions as a challenge for future research.¹²

Despite its non-minimality, the partiality condition in Theorem 6 is not redundant: if it were dropped (and not replaced by something else), part (a) would cease to hold. This follows from the following (non-trivial) proposition:

Proposition 2 *For some path-connected and non-partitional (finite) agenda X , there exists a non-linear pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence (even neutrality) and consensus preservation.*¹³

Readers familiar with binary judgment-aggregation theory will notice that the agenda which we construct to prove this proposition violates an important agenda condition from that area, namely *even-number negatability* (or *non-affineness*) (see Dietrich 2007, Dietrich and List 2007, Dokow and Holzman 2010). It would be intriguing if the same condition turned out to be the correct minimal substitute for partiality in Theorem 6.

5.2 Proof sketches

We now describe how Theorems 4 to 6 can be proved. (Again, details are given in the Appendix.) We begin with part (a) of each theorem. To prove Theorem 4(a), consider a non-nested agenda X with $|X \setminus \{\Omega, \emptyset\}| > 4$ and a pooling function F satisfying independence and consensus compatibility. We want to show that F is linear. Neutrality follows from Theorem 1(a). From neutrality, we can infer linearity by using two lemmas. The first contains the bulk of the work, and the second is an application of Cauchy’s functional equation (similar to its application in Aczél and Wagner 1980 and McConway 1981). Let us write $\mathbf{0}$ and $\mathbf{1}$ to denote the n -tuples $(0, \dots, 0)$ and $(1, \dots, 1)$, respectively.

Lemma 4 *If $D : [0, 1]^n \rightarrow [0, 1]$ is the local pooling criterion of a neutral and consensus-compatible pooling function for a non-nested agenda X with*

¹²A generalized definition of partitionality is possible in Theorem 6: we could define X to be *partitional* if there are finite or countably infinite subsets $Y, Z \subseteq X$ such that the set $\{A \cap C : A \in Y\}$, with $C = \bigcap_{B \in Z} B$, partitions C into at least three non-empty events. This definition generalizes the one in the main text, because if we take $Z = \emptyset$, then C becomes $\Omega (= \bigcap_{B \in \emptyset} B)$ and Y simply partitions Ω . But since we do not know whether this generalized definition renders partitionality logically minimal in Theorem 6, we use the simpler definition in the main text.

¹³In this proposition, we assume that the underlying set of worlds Ω satisfies $|\Omega| \geq 4$.

$|X \setminus \{\Omega, \emptyset\}| > 4$, then

$$D(x) + D(y) + D(z) = 1 \text{ for all } x, y, z \in [0, 1]^n \text{ with } x + y + z = \mathbf{1}. \quad (2)$$

Lemma 5 *If a function $D : [0, 1]^n \rightarrow [0, 1]$ with $D(\mathbf{0}) = 0$ satisfies (2), then it takes the linear form*

$$D(x_1, \dots, x_n) = \sum_{i=1}^n w_i x_i \text{ for all } x \in [0, 1]^n$$

for some non-negative weights w_1, \dots, w_n with sum 1.

The proof of Theorem 5(a) follows a similar strategy, but replaces Lemma 4 with the following lemma:

Lemma 6 *If $D : [0, 1]^n \rightarrow [0, 1]$ is the local pooling criterion of a neutral and conditional-consensus-compatible pooling function for a non-simple agenda X , then (2) holds.*

Finally, Theorem 6(a) can also be proved using a similar strategy, this time replacing Lemma 4 with the following lemma:

Lemma 7 *If $D : [0, 1]^n \rightarrow [0, 1]$ is the local pooling criterion of a neutral and consensus-preserving pooling function for a partitional agenda X , then (2) holds.*

Part (b) of each of Theorems 4 to 6 can be proved by constructing a suitable example of a non-linear pooling function. In the case of Theorem 4(b), we can re-use the non-neutral pooling function constructed to prove Theorem 1(b) as long as the agenda satisfies $|X \setminus \{\Omega, \emptyset\}| > 4$; for (small) agendas with $|X \setminus \{\Omega, \emptyset\}| \leq 4$, we construct a somewhat simplistic pooling function generating collective opinion functions that only assign probabilities of 0, $\frac{1}{2}$, or 1. The constructions for Theorems 5(b) and 6(b) are more difficult; the one for Theorem 5(b) also has the property that collective probabilities never take values other than 0, $\frac{1}{2}$, or 1.

6 Classic results as special cases

It is instructive to see how our present results generalize classic results in the literature, where the agenda is a σ -algebra (especially Aczél and Wagner 1980 and McConway 1981). Note that, for a σ -algebra, all the agenda conditions we have used reduce to a simple condition on agenda *size*:

Lemma 8 *For any agenda X ($\neq \{\Omega, \emptyset\}$) that is closed under pairwise union or intersection (i.e., any agenda that is an algebra), the conditions of non-nestedness, non-simplicity, path-connectedness, and partitionality are equivalent, and are each satisfied if and only if $|X| > 4$.*

Note, further, that when X is a σ -algebra, all of our consensus requirements become equivalent, as shown by Proposition 1(a). It follows that, in the special case of a σ -algebra, our six theorems reduce to two classical results:

- 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 $|X| < 4$ is uninteresting because it implies that $X = \{\emptyset, \Omega\}$, given that X is a σ -algebra. In fact, we can derive these classic theorems not only for σ -algebras, but also for algebras. This is because, given Lemma 8, Theorems 3 and 6 have the following implication:

Corollary 1 *For any agenda X that is closed under pairwise union or intersection (i.e., any agenda that is an algebra),*

- (a) *if $|X| > 4$, every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation is linear (and by implication neutral);*
- (b) *if $|X| = 4$, there exists a non-neutral (and by implication non-linear) pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and consensus preservation.*

7 Probabilistic preference aggregation

To illustrate the use of general agendas, we now present an application to *probabilistic preference aggregation*, a probabilistic analogue of Arrowian preference aggregation. A group seeks to rank a set K of at least two (mutually exclusive and exhaustive) alternatives in a linear order. Let Ω_K be the set of all strict orderings \succ over K (asymmetric, transitive, and connected binary relations). Informally, K can represent any set of distinct objects, e.g., policy options, candidates, social states, or distributions of goods, and an ordering \succ over K can have any interpretation consistent with a linear form (e.g., ‘better than’, ‘preferable to’, ‘higher than’, ‘more competent than’, ‘less unequal than’ etc.).

For any two distinct alternatives x and y in K , let $x \succ y$ denote the event that x is ranked above y ; i.e., $x \succ y$ denotes the subset of Ω_K consisting of all those orderings \succ in Ω_K such that $x \succ y$. We define the *preference agenda* as the set

$$X_K = \{x \succ y : x, y \in K \text{ with } x \neq y\},$$

which is non-empty and closed under complementation, as required for an agenda (this construction draws on Dietrich and List 2007a). In our opinion pooling problem, each individual i submits probability assignments for the events in X_K , and the group then determines corresponding collective probability assignments. An agent's opinion function $P : X_K \rightarrow [0, 1]$ can be interpreted as capturing the agent's *degrees of belief* about which of the various pairwise comparisons $x \succ y$ (in X_K) are 'correct'; call this the *belief interpretation*. Thus, for any two distinct alternatives x and y in K , $P(x \succ y)$ can be interpreted as the agent's degree of belief in the event $x \succ y$, i.e., the event that x is ranked above (preferable to, better than, higher than ...) y . (On a different interpretation, the *vague-preference interpretation*, $P(x \succ y)$ could represent the degree to which the agent prefers x to y , so that the present framework would capture vague preferences over alternatives as opposed to degrees of belief about how they are ranked in terms of the appropriate criterion.) A pooling function, as defined above, maps n individual such opinion functions to a single collective one.

What are the structural properties of this preference agenda?

Lemma 9 *For a preference agenda X_K , the conditions of non-nestedness, non-simplicity, and path-connectedness are equivalent, and are each satisfied if and only if $|K| > 2$; the condition of partitionality is violated for any K .*

The proof that the preference agenda is non-nested if and only if $|K| > 2$ is trivial. The analogous claims for non-simplicity and path-connectedness are well-established in binary judgment-aggregation theory, to which we refer the reader.¹⁴ Finally, it is easy to show that any preference agenda violates partitionality.

Since the preference agenda is non-nested, non-simple, and path-connected when $|K| > 2$, Theorems 1(a), 2(a), 3(a), 4(a), and 5(a) apply; but Theorem 6(a) does not, because partitionality is violated. Let us here focus on Theorem 5. This theorem has the following corollary for the preference agenda:

Corollary 2 *For a preference agenda X_K ,*

- (a) *if $|K| > 2$, every pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility is linear;*
- (b) *if $|K| = 2$, there exists a non-linear pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ satisfying independence and conditional consensus compatibility.*

It is interesting to compare this result with Arrow's classic theorem. While Arrow's theorem yields a negative conclusion if $|K| > 2$ (showing that only dictatorial aggregation functions satisfy its requirements), our linearity result does

¹⁴To see that X_K is non-simple if $|K| > 2$, choose three distinct alternatives $x, y, z \in K$ and note that the three events $x \succ y$, $y \succ z$, and $z \succ x$ in X_K are mutually inconsistent, but any pair of them is consistent, so that they form a minimal inconsistent subset of X_K .

not have any negative flavour. We obtain this positive result despite the fact that our axiomatic requirements are comparable to Arrow's. Independence, in our framework, is the probabilistic analogue of Arrow's independence of irrelevant alternatives: for any pair of distinct alternatives x, y in K , the collective probability for $x \succ y$ should depend only on individual probabilities for $x \succ y$. Conditional consensus compatibility is a strengthened analogue of Arrow's weak Pareto principle (an exact analogue would be consensus preservation): it requires that, for any two pairs of distinct alternatives, $x, y \in K$ and $v, w \in K$, if all individuals are certain that $x \succ y$ given that $v \succ w$, then this agreement should be preserved at the collective level. The analogues of Arrow's universal domain and collective rationality are built into our definition of a pooling function, whose domain and co-domain are defined as the set of all (by definition coherent) opinion functions over X_K .

Thus our result points towards an alternative escape-route from Arrow's impossibility theorem (though it may be practically applicable only in special contexts): *if* we enrich Arrow's informational framework by allowing *degrees of belief* over different possible linear orderings as input and output of the aggregation (or alternatively, *vague preferences*, understood probabilistically), *then* we can avoid Arrow's dictatorship conclusion. Instead, we obtain a positive characterization of linear pooling, despite imposing requirements on the pooling function that are stronger than Arrow's classic requirements (in so far as conditional consensus compatibility is stronger than the analogue of the weak Pareto principle).

On the belief interpretation, the present informational framework is meaningful so long as there exists a fact of the matter about which of the orderings \succ in Ω_K is the 'correct' one (e.g., an objective quality ordering), so that it makes sense to form beliefs about this fact. On the vague-preference interpretation, our framework requires that vague preferences over pairs of alternatives are extendable to a coherent probability distribution over the set of 'crisp' orderings \succ in Ω_K .

There are, of course, substantial bodies of literature on avoiding Arrow's dictatorship conclusion in richer informational frameworks and on probabilistic or vague preference aggregation. It is well known, for example, that the introduction of interpersonally comparable preferences (of an ordinal or cardinal type) is sufficient for avoiding Arrow's negative conclusion (e.g., Sen 1970/1979). Also, different models of probabilistic or vague preference aggregation have been proposed.¹⁵ A typical assumption is that, for any pair of alternatives $x, y \in K$, each individual prefers x to y to a certain degree between 0 and 1. However, the standard constraints on vague or fuzzy preferences do not require individuals to hold

¹⁵A model in which individuals and the collective specify probabilities of selecting each of the alternatives in K (as opposed to probability assignments over events of the form ' x is ranked above y ') has been studied, for instance, by Intriligator (1973), who has characterized a version of linear averaging in it. Similarly, a model in which individuals have vague or fuzzy preferences has been studied, for instance, by Billot (1991) and more recently by Piggins and Perote-Peña (2007) (see also Sanver and Selçuk 2009).

probabilistically coherent opinion functions in our sense; hence the literature has tended to generate Arrow-style impossibility results. By contrast, it is illuminating to see that a *possibility* result on probabilistic preference aggregation can be derived as a corollary of one of our new results on probabilistic opinion pooling.

8 A unified perspective

Finally, we wish to compare probabilistic opinion pooling with binary judgment aggregation and Arrovian preference aggregation in its original form. Thanks to the notion of a general agenda, we can represent each of these other aggregation problems within the present framework.

- To represent *binary judgment aggregation*, we simply need to restrict attention to binary opinion functions, i.e., opinion functions that take only the values 0 and 1.¹⁶ Binary opinion functions correspond to consistent and complete judgment sets in judgment-aggregation theory, i.e., sets of the form $J \subseteq X$ which satisfy $\bigcap_{A \in J} A \neq \emptyset$ (*consistency*) and contain a member of each pair $A, A^c \in X$ (*completeness*).¹⁷ A *binary opinion pooling function* assigns to each profile of binary opinion functions a collective binary opinion function. Thus, binary opinion pooling functions correspond to standard judgment aggregation functions (with universal domain and consistent and complete outputs).
- To represent *preference aggregation*, we need to restrict attention both to the preference agenda, as introduced in Section 7, and to binary opinion functions, as just defined. Binary opinion functions for the preference agenda correspond to linear preference orders, as familiar from preference aggregation theory in the tradition of Arrow. Here, binary opinion pooling functions correspond to Arrovian social welfare functions.

The literature on binary judgment aggregation contains several theorems that use axiomatic requirements similar to those used here. In the binary case, however, these requirements lead to dictatorial, rather than linear, aggregation, as in Arrow’s original impossibility theorem in preference-aggregation theory. In fact, Arrow-like theorems are immediate corollaries of the results on judgment aggregation, when applied to the preference agenda (e.g., Dietrich and List 2007a, List and Pettit 2004). In particular, the independence requirement reduces to Arrow’s independence of irrelevant alternatives, and the unanimity-preservation requirements reduce to variants of the Pareto principle.

¹⁶Formally, a binary opinion function is a function $f : X \rightarrow \{0, 1\}$ that is extendible to a probability function on $\sigma(X)$, or equivalently, to a truth-function on $\sigma(X)$ (i.e., a $\{0, 1\}$ -valued function on $\sigma(X)$ that is logically consistent).

¹⁷Specifically, a binary opinion function $f : X \rightarrow \{0, 1\}$ corresponds to the consistent and complete judgment set $\{A \in X : f(A) = 1\}$.

How can the same axiomatic requirements lead to a positive conclusion – linearity – in the probabilistic framework and to a negative one – dictatorship – in the binary case? The reason is that, in the binary case, linearity collapses into dictatorship because the only well-defined linear pooling functions are dictatorial here. Let us explain this point. Linearity of a binary opinion pooling function F is defined just as in the probabilistic framework: there exist real-valued weights $w_1, \dots, w_n \geq 0$ with $w_1 + \dots + w_n = 1$ such that, for every profile (P_1, \dots, P_n) of binary opinion functions, the collective truth-value of any given event A in the agenda X is the weighted arithmetic average $w_1 P_1(A) + \dots + w_n P_n(A)$. Yet, for this to define a proper binary opinion pooling function, some individual i must get a weight of 1 and all others must get a weight of 0, since otherwise the average $w_1 P_1(A) + \dots + w_n P_n(A)$ could fall strictly between 0 and 1, violating the binary restriction. In other words, linearity is equivalent to dictatorship here.¹⁸

We can obtain a unified perspective on several distinct aggregation problems by combining this paper’s linearity results with the corresponding dictatorship results from the existing literature (adopting the unification strategy proposed in Dietrich and List 2010). This yields several unified characterization theorems applicable to probability aggregation, judgment aggregation, and preference aggregation. Let us state these results. The first combines Theorem 4 with a result due to Dietrich (2013); the second combines Theorem 5 with a result due to Dietrich and List (2013); and the third combines Theorem 6 with the analogue of Arrow’s theorem in judgment aggregation (Dietrich and List 2007a and Dokow and Holzman 2010). In the binary case, the independence requirement and our various unanimity requirements are defined as in the probabilistic framework, but with a restriction to binary opinion functions.¹⁹

Theorem 4⁺ (a) *For any non-nested agenda X with $|X \setminus \{\Omega, \emptyset\}| > 4$, every **binary or probabilistic** opinion pooling function satisfying independence and consensus compatibility is linear (where linearity reduces to dictatorship in the binary case).*

(b) *For any other agenda X ($\neq \{\emptyset, \Omega\}$), there exists a non-linear **binary or probabilistic** opinion pooling function satisfying independence and consensus compatibility.*

Theorem 5⁺ (a) *For any non-simple agenda X with $|X \setminus \{\Omega, \emptyset\}| > 4$, every*

¹⁸To be precise, for (trivial) agendas with $X \setminus \{\Omega, \emptyset\} = \emptyset$, the weights w_i may differ from 1 and 0. But it still follows that every linear binary opinion pooling function (in fact, every binary opinion pooling function) is dictatorial here, for the trivial reason that there is only one binary opinion function and thus only one (dictatorial) binary opinion pooling function.

¹⁹In the binary case, two of our unanimity-preservation requirements (implicit consensus preservation and consensus compatibility) are equivalent, because every binary opinion function is *uniquely* extendible to $\sigma(X)$. Also, conditional consensus compatibility can be stated more easily in the binary case, namely in terms of a single conditional judgment rather than a finite set of conditional judgments.

binary or probabilistic opinion pooling function satisfying independence and conditional consensus compatibility is linear (where linearity reduces to dictatorship in the binary case).

- (b) For any simple agenda X (finite and distinct from $\{\emptyset, \Omega\}$), there exists a non-linear **binary or probabilistic** opinion pooling function satisfying independence and conditional consensus compatibility.

Theorem 6⁺ (a) For any path-connected and partitional agenda X , every **binary or probabilistic** opinion pooling function satisfying independence and consensus preservation is linear (where linearity reduces to dictatorship in the binary case).

- (b) For any non-path-connected (finite) agenda X , there exists a non-linear **binary or probabilistic** opinion pooling function satisfying independence and consensus preservation.²⁰

By Lemma 9, Theorems 4⁺, 5⁺, and 6⁺ are relevant to preference aggregation insofar as the preference agenda X_K satisfies each of non-nestedness, non-simplicity, and path-connectedness if and only if $|K| > 2$, where K is the set of alternatives. Recall, however, that the preference agenda is never partitional, so that part (a) of Theorem 6⁺ never applies. By contrast, the binary result on which part (a) is based applies to the preference agenda, as it uses the weaker condition of even-number-negatability (or non-affineness) instead of partitionality (and that weaker condition is satisfied by X_K if $|K| > 2$). As noted above, it remains an open question how far partitionality can be weakened in the probabilistic case.²¹

9 References

- Aczél, J. (1966) *Lectures on Functional Equations and their Applications*, New York and London: Academic Press
- Aczél, J., C. T. Ng, and C. Wagner (1984) “Aggregation Theorems for Allocation Problems”, *SIAM Journal on Algebraic and Discrete Methods* 5(1): 1-8
- Aczél, J., and C. Wagner (1980) “A characterization of weighted arithmetic means”, *SIAM Journal on Algebraic and Discrete Methods* 1(3): 259-260
- Billot, A. (1991) “Aggregation of preferences: The fuzzy case”, *Theory and Decision* 30(1): 51-93
- Bradley, R., F. Dietrich, and C. List (2014) “Aggregating Causal Judgments”, *Philosophy of Science* 81(4): 491-515.

²⁰In the binary case in part (a), partiality can be weakened to *even-number negatability* or *non-affineness*. See Dietrich and List (2007a) and Dokow and Holzman (2010).

²¹Of course, one could also state unified versions of Theorems 1 to 3 on neutral opinion pooling, by combining these theorems with existing results on binary judgment aggregation. We would simply need to replace the probabilistic opinion pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ with a *binary or probabilistic* such function.

- Chambers, C. (2007) “An ordinal characterization of the linear opinion pool”, *Economic Theory* 33(3): 457-474
- Clemen, R. T., and R. L. Winkler (1999) “Combining Probability Distributions from Experts in Risk Analysis”, *Risk Analysis* 19(2): 187-203
- Dietrich, F. (2006) “Judgment Aggregation: (Im)Possibility Theorems”, *Journal of Economic Theory* 126(1): 286-298
- Dietrich, F. (2007) “A generalised model of judgment aggregation”, *Social Choice and Welfare* 28(4): 529-565
- Dietrich, F. (2010) “Bayesian group belief”, *Social Choice and Welfare* 35(4): 595-626
- Dietrich, F. (2013) “Judgment aggregation and agenda manipulation”, working paper
- Dietrich, F., and C. List (2007a) “Arrow’s theorem in judgment aggregation”, *Social Choice and Welfare* 29(1): 19-33
- Dietrich, F., and C. List (2007b) “Judgment aggregation by quota rules: majority voting generalized”, *Journal of Theoretical Politics* 19(4): 391-424
- Dietrich, F., and C. List (2010) “The aggregation of propositional attitudes: towards a general theory”, *Oxford Studies in Epistemology* 3: 215-234 [with “Corrigendum” on the authors’ webpages]
- Dietrich, F., and C. List (2013) “Propositionwise judgment aggregation: the general case”, *Social Choice and Welfare* 40(4): 1067-1095
- Dietrich, F., and C. List (2014) “Probabilistic opinion pooling generalized – Part two: The premise-based approach”, working paper
- Dietrich, F., and P. Mongin (2010) “The premise-based approach to judgment aggregation”, *Journal of Economic Theory* 145(2): 562-582
- Dokow, E., and R. Holzman (2010) “Aggregation of binary evaluations”, *Journal of Economic Theory* 145(2): 495-511
- Genest, C. (1984) “Pooling operators with the marginalization property”, *Canadian Journal of Statistics* 12(2): 153-163
- Genest, C., and C. Wagner (1987) “Further Evidence against Independence Preservation in Expert Judgement Synthesis”, *Aequationes Mathematicae* 32(1): 74-86
- Genest, C., and J. V. Zidek (1986) “Combining Probability Distributions: A Critique and Annotated Bibliography”, *Statistical Science* 1(1): 114-135
- Intriligator, M. D. (1973) “A Probabilistic Model of Social Choice”, *Review of Economic Studies* 40(4): 553-560
- 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., and P. Pettit (2004) “Aggregating sets of judgments: Two impossibility results compared”, *Synthese* 140(1-2): 207-235
- McConway, K. J. (1981) “Marginalization and Linear Opinion Pools”, *Journal of the American Statistical Association* 76(374): 410-414

- Mongin, P. (1995) “Consistent Bayesian aggregation”, *Journal of Economic Theory* 66(2): 313-351
- Nehring, K., and C. Puppe (2007) “The structure of strategy-proof social choice – Part I: General characterization and possibility results on median spaces”, *Journal of Economic Theory* 135(1): 269-305
- Nehring, K., and C. Puppe (2010) “Abstract Arrovian Aggregation”, *Journal of Economic Theory* 145(2): 467-494
- Osherson, D., and M. Y. Vardi (2006) “Aggregating disparate estimates of chance”, *Games and Economic Behavior* 56(1): 148-173
- Piggins, A., and J. Perote-Peña (2007) “Strategy-proof Fuzzy Aggregation Rules”, *Journal of Mathematical Economics* 43(5): 564-580
- Priest, G. (2001) *An Introduction to Non-classical Logic*, Cambridge: Cambridge University Press
- Sanver, M. R., and Ö. Selçuk (2009) “Sophisticated Preference Aggregation”, *Social Choice and Welfare* 33(1): 73-86
- Sen, A. (1970 / republished 1979) *Collective Choice and Social Welfare*, San Francisco: Holden-Day, and Amsterdam: North Holland
- Wagner, C. (1982) “Allocation, Lehrer Models, and the Consensus of Probabilities”, *Theory and Decision* 14(2): 207-220
- Wagner, C. (1985) “On the Formal Properties of Weighted Averaging as a Method of Aggregation”, *Synthese* 62(1): 97-108
- Weymark, J. (1997) “Aggregating Ordinal Probabilities on Finite Sets”, *Journal of Economic Theory* 75(2): 407-432

A Proofs

We now prove all our results. In light of the mathematical connection between the present results and those in our companion paper on ‘premise-based’ opinion pooling for σ -algebra agendas (Dietrich and List 2014), one might imagine two possible proof strategies: either one could prove our present results directly and those in the companion paper as corollaries, or *vice versa*. In fact, we will mix those two strategies. We will prove parts (a) of all present theorems directly (and use them in the companion paper to derive the corresponding results), while we will prove parts (b) directly in some cases and as corollaries of corresponding results from the companion paper in others.

This Appendix is organised as follows. In Sections A.1 to A.5, we prove parts (a) of Theorems 2 to 6, along with related results. Theorem 1(a) requires no independent proof, as it follows from Theorem 2(a). In Section A.6, we clarify the connection between the two papers, and then prove parts (b) of all present theorems. Finally, in Section A.7, we prove Propositions 1 and 2.

A.1 Proof of Theorem 2(a)

As explained in the main text, Theorem 2(a) follows from Lemmas 1 and 2. We now prove these lemmas. To do so, we will also prove some preliminary results.

Lemma 10 *Consider any agenda X .*

- (a) \sim defines an equivalence relation on $X \setminus \{\emptyset, \Omega\}$.
- (b) $A \sim B \Leftrightarrow A^c \sim B^c$ for all events $A, B \in X \setminus \{\emptyset, \Omega\}$.
- (c) $A \subseteq B \Rightarrow A \sim B$ for all events $A, B \in X \setminus \{\emptyset, \Omega\}$.
- (d) If $X \neq \{\emptyset, \Omega\}$, the relation \sim has
 - either a single equivalence class, namely $X \setminus \{\emptyset, \Omega\}$,
 - or exactly two equivalence classes, each one containing exactly one member of each pair $A, A^c \in X \setminus \{\emptyset, \Omega\}$.

Proof. (a) Reflexivity, symmetry, and transitivity on $X \setminus \{\emptyset, \Omega\}$ are all obvious (we have excluded \emptyset and Ω to ensure reflexivity).

(b) It suffices to prove one direction of implication (as $(A^c)^c = A$ for all $A \in X$). Let $A, B \in X \setminus \{\emptyset, \Omega\}$ with $A \sim B$. Then there is a path $A_1, \dots, A_k \in X$ from A to B such that any neighbours A_j, A_{j+1} are non-exclusive and non-exhaustive. So A_1^c, \dots, A_k^c is a path from A^c to B^c , where any neighbours A_j^c, A_{j+1}^c are non-exclusive (as $A_j^c \cap A_{j+1}^c = (A_j \cup A_{j+1})^c \neq \Omega^c = \emptyset$) and non-exhaustive (as $A_j^c \cup A_{j+1}^c = (A_j \cap A_{j+1})^c \neq \emptyset^c = \Omega$). So, $A^c \sim B^c$.

(c) Let $A, B \in X \setminus \{\emptyset, \Omega\}$. If $A \subseteq B$, then $A \sim B$ due to a direct connection, because A, B are neither exclusive (as $A \cap B = A \neq \emptyset$) nor exhaustive (as $A \cup B = B \neq \Omega$).

(d) Let $X \neq \{\emptyset, \Omega\}$. Suppose the number of equivalence classes with respect to \sim is not one. As $X \setminus \{\emptyset, \Omega\} \neq \emptyset$, it is not zero. So it is at least two. We show two claims:

Claim 1. There are exactly two equivalence classes with respect to \sim .

Claim 2. Each class contains exactly one member of any pair $A, A^c \in X \setminus \{\emptyset, \Omega\}$.

Proof of Claim 1. For a contradiction, let $A, B, C \in X \setminus \{\emptyset, \Omega\}$ be pairwise not (\sim -)equivalent. By $A \not\sim B$, either $A \cap B = \emptyset$ or $A \cup B = \Omega$. We may assume the former case, because in the latter case we may consider A^c, B^c, C^c instead of A, B, C . (Note that A^c, B^c, C^c are again pairwise non-equivalent by (b) and $A^c \cap B^c = (A \cup B)^c = \Omega^c = \emptyset$.) Now, since $A \cap B = \emptyset$, we have $B \subseteq A^c$, whence $A^c \sim B$ by (c). By $A \not\sim C$, there are two cases:

- either $A \cap C = \emptyset$, which implies $C \subseteq A^c$, whence $C \sim A^c$ by (c), so that $C \sim B$ (as $A^c \sim B$ and \sim is transitive by (a)), a contradiction;
- or $A \cup C = \Omega$, which implies $A^c \subseteq C$, whence $A^c \sim C$ by (c), so that again we derive the contradiction $C \sim B$, which completes the proof of Claim 1.

Proof of Claim 2. For a contradiction, let Z be an (\sim -)equivalence class containing the pair A, A^c . By assumption, Z is not the only equivalence class, so there is a $B \in X \setminus \{\emptyset, \Omega\}$ with $B \not\sim A$ (hence $B \not\sim A^c$). Then either $A \cap B = \emptyset$ or $A \cup B = \Omega$. In the first case, $B \subseteq A^c$, so that $B \sim A^c$ by (c), a contradiction. In the second case, $A^c \subseteq B$, so that $A^c \sim B$ by (c), a contradiction. ■

Proof of Lemma 1. Consider an agenda $X \neq \{\emptyset, \Omega\}$. By Lemma 10(a), \sim is indeed an equivalence relation on $X \setminus \{\emptyset, \Omega\}$. By Lemma 10(d), it remains to prove that X is nested if and only if there are exactly two equivalence classes. Note that X is nested if and only if $X \setminus \{\emptyset, \Omega\}$ is nested. So we may assume without loss of generality that $\emptyset, \Omega \notin X$.

First, suppose there are two equivalence classes. Let X_+ be one of them. By Lemma 10(d), $X = \{A, A^c : A \in X_+\}$. To complete the proof that X is nested, we show that X_+ is linearly ordered by set-inclusion \subseteq . Clearly, \subseteq is reflexive, transitive, and anti-symmetric. We must show that it is connected. So, let $A, B \in X_+$; we prove that $A \subseteq B$ or $B \subseteq A$. Since $A \not\sim B^c$ (by Lemma 10(d)), either $A \cap B^c = \emptyset$ or $A \cup B^c = \Omega$. So, either $A \subseteq B$ or $B \subseteq A$.

Conversely, let X be nested. So $X = \{A, A^c : A \in X_+\}$ for some set $X_+ \subseteq \Sigma$ that is linearly ordered by set inclusion. Let $A \in X_+$. We show that $A \not\sim A^c$, implying that X has at least – so by Lemma 10(d) *exactly* – two equivalence classes. For a contradiction, suppose $A \sim A^c$. Then there is a path $A_1, \dots, A_k \in X$ from $A = A_1$ to $A^c = A_k$ such that, for all neighbours A_j, A_{j+1} , $A_j \cap A_{j+1} \neq \emptyset$ and $A_j \cup A_{j+1} \neq \Omega$. Since each event $C \in X$ either is in X^+ or has its complement in X^+ , and since $A_1 = A \in X^+$ and $A_k^c = A \in X^+$, there are neighbours A_j, A_{j+1} such that $A_j, A_{j+1}^c \in X^+$. So, as X^+ is linearly ordered by \subseteq , either $A_j \subseteq A_{j+1}^c$ or $A_{j+1}^c \subseteq A_j$, i.e., either $A_j \cap A_{j+1} = \emptyset$ or $A_j \cup A_{j+1} = \Omega$, a contradiction. ■

We now give a useful re-formulation of the requirement of conditional consensus compatibility for opinion pooling on a general agenda X . Note first that an opinion function is consistent with certainty of A ($\in X$) given B ($\in X$) if and only if it is consistent with certainty of the event ‘ B implies A ’ (i.e., with *zero* probability of the event $B \setminus A$ or ‘ B but not A ’). This observation yields the following re-formulation of conditional consensus compatibility (in which the roles of A and B have been interchanged):

Implication preservation. For all $P_1, \dots, P_n \in \mathcal{P}_X$, and all finite sets S of pairs (A, B) of events in X , if every opinion function P_i is consistent with certainty that A implies B for all (A, B) in S (i.e., some extension $\bar{P}_i \in \mathcal{P}_{\sigma(X)}$ of P_i satisfies $\bar{P}_i(A \setminus B) = 0$ for all pairs $(A, B) \in S$), then so is the collective opinion function P_{P_1, \dots, P_n} .

Proposition 3 *For any agenda X , a pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ is conditional consensus compatible if and only if it is implication preserving.*

Proof of Lemma 2. Let F be an independent and conditional-consensus-compatible pooling function for agenda X . For all $A \in X$, let D_A be the pooling criterion given by independence. We show that $D_A = D_B$ for all $A, B \in X$ with $A \cap B \neq \emptyset$ and $A \cup B \neq \Omega$. This will imply that $D_A = D_B$ whenever $A \sim B$ (by induction on the length of a path from A to B), which completes the proof.

So, let $A, B \in X$ with $A \cap B \neq \emptyset$ and $A \cup B \neq \Omega$. Notice that $A \cap B$, $A \cup B$, and $A \setminus B$ need not belong to X . Let $x \in [0, 1]^n$; we show that $D_A(x) = D_B(x)$. As $A \cap B \neq \emptyset$ and $A^c \cap B^c = (A \cup B)^c \neq \emptyset$, there are $\bar{P}_1, \dots, \bar{P}_n \in \mathcal{P}_{\sigma(X)}$ such that

$$\bar{P}_i(A \cap B) = x_i \text{ and } \bar{P}_i(A^c \cap B^c) = 1 - x_i \text{ for all } i = 1, \dots, n.$$

Now consider the opinion functions $P_1, \dots, P_n \in \mathcal{P}_X$ given by $P_i := \bar{P}_i|_X$. Since $\bar{P}_i(A \setminus B) = 0$ and $\bar{P}_i(B \setminus A) = 0$ for all i , the collective opinion function P_{P_1, \dots, P_n} has an extension $\bar{P}_{P_1, \dots, P_n} \in \mathcal{P}_{\sigma(X)}$ such that $\bar{P}_{P_1, \dots, P_n}(A \setminus B) = \bar{P}_{P_1, \dots, P_n}(B \setminus A) = 0$, by implication preservation (which is equivalent to conditional consensus compatibility by Proposition 3). So $\bar{P}_{P_1, \dots, P_n}(A) = \bar{P}_{P_1, \dots, P_n}(A \cap B) = \bar{P}_{P_1, \dots, P_n}(B)$, and hence, $P_{P_1, \dots, P_n}(A) = P_{P_1, \dots, P_n}(B)$. So, using the fact that $P_{P_1, \dots, P_n}(A) = D_A(x)$ (as $P_i(A) = x_i$ for all i) and $P_{P_1, \dots, P_n}(B) = D_B(x)$ (as $P_i(B) = x_i$ for all i), we have $D_A(x) = D_B(x)$. ■

A.2 Proof of Theorem 3(a)

As explained in the main text, Theorem 3(a) follows from Lemma 3, which we now prove.

Proof of Lemma 3. Let $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ be independent and consensus-preserving. Let $A, B \in X$ such that $A \vdash^* B$, say in virtue of (countable) set $Y \subseteq X$. Write D_A and D_B for the pooling criterion for A and B , respectively. Let $x = (x_1, \dots, x_n) \in [0, 1]^n$. We show that $D_A(x) \leq D_B(x)$. As $\cap_{C \in \{A\} \cup Y} C$ is non-empty but has empty intersection with B^c (by the conditional entailment), it equals its intersection with B , so $\cap_{C \in \{A, B\} \cup Y} C \neq \emptyset$. Similarly, as $\cap_{C \in \{B^c\} \cup Y} C$ is non-empty but has empty intersection with A , it equals its intersection with A^c , so $\cap_{C \in \{A^c, B^c\} \cup Y} C \neq \emptyset$. Hence there exist $\omega \in \cap_{C \in \{A, B\} \cup Y} C$ and $\omega' \in \cap_{C \in \{A^c, B^c\} \cup Y} C$. For each individual i , we define a probability function $P_i^* : \sigma(X) \rightarrow [0, 1]$ by $P_i^* := x_i \delta_\omega + (1 - x_i) \delta_{\omega'}$ (where $\delta_\omega, \delta_{\omega'} : \sigma(X) \rightarrow [0, 1]$ are the Dirac-measures at ω and ω' , respectively), and we then let $P_i := P_i^*|_X$. As each P_i satisfies $P_i(A) = P_i(B) = x_i$,

$$\begin{aligned} P_{P_1, \dots, P_n}(A) &= D_A(P_1(A), \dots, P_n(A)) = D_A(x), \\ P_{P_1, \dots, P_n}(B) &= D_B(P_1(B), \dots, P_n(B)) = D_B(x). \end{aligned}$$

Further, for each P_i and each $C \in Y$, we have $P_i(C) = 1$, so that $P_{P_1, \dots, P_n}(C) = 1$ (by consensus preservation). Hence $P_{P_1, \dots, P_n}(\cap_{C \in Y} C) = 1$, since ‘countable inter-

sections preserve probability one'. So,

$$\begin{aligned} P_{P_1, \dots, P_n}(\cap_{C \in \{A\} \cup Y} C) &= P_{P_1, \dots, P_n}(A) = D_A(x), \\ P_{P_1, \dots, P_n}(\cap_{C \in \{B\} \cup Y} C) &= P_{P_1, \dots, P_n}(B) = D_B(x). \end{aligned}$$

To prove that $D_A(x) \leq D_B(x)$, it suffices to show that $P_{P_1, \dots, P_n}(\cap_{C \in \{A\} \cup Y} C) \leq P_{P_1, \dots, P_n}(\cap_{C \in \{B\} \cup Y} C)$. This is true because

$$\cap_{C \in \{A\} \cup Y} C = \cap_{C \in \{A, B\} \cup Y} C \subseteq \cap_{C \in \{B\} \cup Y} C,$$

where the identity holds by an earlier argument. ■

A.3 Proof of Theorem 4(a)

As explained in the main text, Theorem 4(a) follows from Theorem 1(a) via Lemmas 4 and 5.²² It remains to prove both lemmas. We draw on a known agenda characterization result and a technical lemma.

Proposition 4 (Dietrich 2013) *For any agenda X , the following are equivalent:*

- (a) X is non-nested with $|X \setminus \{\Omega, \emptyset\}| > 4$;
- (b) X has a (consistent or inconsistent) subset Y with $|Y| \geq 3$ such that $(Y \setminus \{A\}) \cup \{A^c\}$ is consistent for each $A \in Y$;
- (c) X has a (consistent or inconsistent) subset Y with $|Y| = 3$ such that $(Y \setminus \{A\}) \cup \{A^c\}$ is consistent for each $A \in Y$.

Lemma 11 *If $D : [0, 1]^n \rightarrow [0, 1]$ is the local pooling criterion of a neutral pooling function for an agenda X ($\neq \{\Omega, \emptyset\}$), then*

- (a) $D(x) + D(\mathbf{1} - x) = 1$ for all $x \in [0, 1]^n$,
- (b) $D(\mathbf{0}) = 0$ and $D(\mathbf{1}) = 1$, provided the pooling function is consensus preserving.

Proof. (a) As $X \neq \{\Omega, \emptyset\}$, we may pick some $A \in X \setminus \{\Omega, \emptyset\}$. For each $x \in [0, 1]^n$, there exist (by $A \neq \emptyset, \Omega$) opinion functions $P_1, \dots, P_n \in \mathcal{P}_X$ such that $(P_1(A), \dots, P_n(A)) = x$, which implies that $(P_1(A^c), \dots, P_n(A^c)) = \mathbf{1} - x$ and

$$D(x) + D(\mathbf{1} - x) = P_{P_1, \dots, P_n}(A) + P_{P_1, \dots, P_n}(A^c) = 1.$$

(b) Given consensus-preservation $D(\mathbf{1}) = 1$. By part (a), $D(\mathbf{0}) = 1 - D(\mathbf{1})$. So $D(\mathbf{0}) = 0$. ■

²²This uses Lemma 11(b) below, where consensus preservation holds by consensus compatibility.

Proof of Lemma 4. Let D be the local pooling criterion of such a pooling function for such an agenda X . Consider any $x, y, z \in [0, 1]^n$ with sum $\mathbf{1}$. By Proposition 4, there exist $A, B, C \in X$ such that each of the sets

$$A^* := A^c \cap B \cap C, B^* := A \cap B^c \cap C, C^* := A \cap B \cap C^c$$

is non-empty. For all individuals i , since $x_i + y_i + z_i = 1$ and since A^*, B^*, C^* are pairwise disjoint non-empty members of $\sigma(X)$, there exists a $P_i^* \in \mathcal{P}_{\sigma(X)}$ such that $P_i^*(A^*) = x_i$, $P_i^*(B^*) = y_i$ and $P_i^*(C^*) = z_i$. By construction,

$$P_i^*(A^* \cup B^* \cup C^*) = x_i + y_i + z_i = 1 \text{ for all } i. \quad (3)$$

Let $P_i := P_i^*|_X$ for each individual i . For the profile $(P_1, \dots, P_n) \in \mathcal{P}_X^n$ thus defined, we consider the collective opinion function P_{P_1, \dots, P_n} . We complete the proof by proving two claims.

Claim 1. $P^*(A^*) + P^*(B^*) + P^*(C^*) = P^*(A^* \cup B^* \cup C^*) = 1$ for some $P^* \in \mathcal{P}_{\sigma(X)}$ extending P_{P_1, \dots, P_n} .

The first identity holds for *all* extensions $P^* \in \mathcal{P}_{\sigma(X)}$ of P , by pairwise disjointness of A^*, B^*, C^* . For the second identity, note that each P_i has an extension $P_i^* \in \mathcal{P}_{\sigma(X)}$ for which $P_i^*(A^* \cup B^* \cup C^*) = 1$, so that by consensus compatibility P_{P_1, \dots, P_n} also has such an extension.

Claim 2. $D(x) + D(y) + D(z) = 1$.

Consider an extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P_{P_1, \dots, P_n} of the kind in Claim 1. As $P^*(A^* \cup B^* \cup C^*) = 1$, and as the intersection of A^c with $A^* \cup B^* \cup C^*$ is A^* ,

$$P^*(A^c) = P^*(A^*). \quad (4)$$

Since $A^c \in X$, we further have $P^*(A^c) = P_{P_1, \dots, P_n}(A^c) = D(P_1(A^c), \dots, P_n(A^c))$, where $P_i(A^c) = P_i^*(A^c) = x_i$ for each individual i . So, $P^*(A^c) = D(x)$. This and (4) imply that $P^*(A^*) = D(x)$. Analogously, $P^*(B^*) = D(y)$ and $P^*(C^*) = D(z)$. So, Claim 2 follows from Claim 1. ■

Proof of Lemma 5. Consider any $D : [0, 1]^n \rightarrow [0, 1]$ such that $D(\mathbf{0}) = 0$ and

$$D(x) + D(y) + D(z) = 1 \text{ for all } x, y, z \in [0, 1]^n \text{ with } x + y + z = \mathbf{1}. \quad (5)$$

We have $D(\mathbf{1}) = 1$ (since $D(\mathbf{1}) + D(\mathbf{0}) + D(\mathbf{0}) = 1$ where $D(\mathbf{0}) = 0$) and

$$D(x) + D(1 - x) = 1 \text{ for all } x \in [0, 1] \quad (6)$$

(since $D(x) + D(1 - x) + D(\mathbf{0}) = 1$ where $D(\mathbf{0}) = 0$). Using (5) and then (6), for all $x, y \in [0, 1]^n$ with $x + y \in [0, 1]^n$,

$$1 = D(x) + D(y) + D(\mathbf{1} - x - y) = D(x) + D(y) + 1 - D(x + y).$$

So,

$$D(x + y) = D(x) + D(y) \text{ for all } x, y \in [0, 1]^n \text{ with } x + y \in [0, 1]^n. \quad (7)$$

Consider any individual i . We define $D_i : [0, 1] \rightarrow [0, 1]$ by $D_i(t) = D(0, \dots, 0, t, 0, \dots, 0)$, where t occurs at position i in $(0, \dots, 0, t, 0, \dots, 0)$. By (7), $D_i(s+t) = D_i(s) + D_i(t)$ for all $s, t \geq 0$ with $s + t \leq 1$. As one can easily check, D_i can be extended to a function $\bar{D}_i : [0, \infty) \rightarrow [0, \infty)$ such that $\bar{D}_i(s+t) = \bar{D}_i(s) + \bar{D}_i(t)$ for all $s, t \geq 0$, i.e., such that \bar{D}_i satisfies the non-negative version of Cauchy's functional equation. So, there is some $w_i \geq 0$ such that $\bar{D}_i(t) = w_i t$ for all $t \geq 0$ (by Theorem 1 in Aczél 1966). Now, for all $x \in [0, 1]^n$, $D(x) = \sum_{i=1}^n D_i(x_i)$ (by repeated application of (7)), and so (as $D_i(x_i) = \bar{D}_i(x_i) = w_i x_i$) $D(x) = \sum_{i=1}^n w_i x_i$. Applying the latter with $x = \mathbf{1}$ yields $D(\mathbf{1}) = \sum_{i=1}^n w_i$, hence $\sum_{i=1}^n w_i = 1$. ■

A.4 Proof of Theorem 5(a)

As explained in the main text, Theorem 5(a) follows from Theorem 2(a) via Lemmas 6 and 5.²³ It remains to prove Lemma 6.

Proof of Lemma 6. Let D be the local pooling criterion of a neutral and conditional-consensus-compatible pooling function for a non-simple agenda X . Consider any $x, y, z \in [0, 1]^n$ with sum $\mathbf{1}$. As X is non-simple, there is a (countable) minimal inconsistent set $Y \subseteq X$ with $|Y| \geq 3$. Pick pairwise distinct $A, B, C \in Y$. Let

$$A^* := \bigcap_{E \in Y \setminus \{A\}} E, \quad B^* := \bigcap_{E \in Y \setminus \{B\}} E, \quad C^* := \bigcap_{E \in Y \setminus \{C\}} E.$$

As $\sigma(X)$ is closed under countable intersections, $A^*, B^*, C^* \in \sigma(X)$. For each i , as $x_i + y_i + z_i = 1$ and as A^*, B^*, C^* are (by Y 's minimal inconsistency) pairwise disjoint non-empty members of $\sigma(X)$, there exists a $P_i^* \in \mathcal{P}_{\sigma(X)}$ such that

$$P_i^*(A^*) = x_i, \quad P_i^*(B^*) = y_i, \quad P_i^*(C^*) = z_i.$$

By construction,

$$P_i^*(A^* \cup B^* \cup C^*) = x_i + y_i + z_i = 1 \text{ for all } i. \quad (8)$$

Now let $P_i := P_i^*|_X$ for each individual i , and let $P := P_{P_1, \dots, P_n}$. We derive four properties of P (Claims 1-4), which then allow us to show that $D(x) + D(y) + D(z) = 1$ (Claim 5).

Claim 1. $P^*(\bigcap_{E \in Y \setminus \{A, B, C\}} E) = 1$ for all extensions $P^* \in \mathcal{P}_{\sigma(X)}$ of P .

²³This uses Lemma 11(b), where consensus preservation holds by conditional consensus compatibility.

For all $E \in Y \setminus \{A, B, C\}$, we have $E \supseteq A^* \cup B^* \cup C^*$, so that by (8) $P_1(E) = \dots = P_n(E) = 1$, and hence $P(E) = 1$ (by consensus preservation, which follows from conditional consensus compatibility by Proposition 1(a)). So, for any extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P , we have $P^*(E) = 1$ for all $E \in Y \setminus \{A, B, C\}$. Thus $P^*(\bigcap_{E \in Y \setminus \{A, B, C\}} E) = 1$, as ‘countable intersections preserve probability one’.

Claim 2. $P^*(A^c \cup B^c \cup C^c) = 1$ for all extensions $P^* \in \mathcal{P}_{\sigma(X)}$ of P .

Let $P^* \in \mathcal{P}_{\sigma(X)}$ be an extension of P . Since $A \cap B \cap C$ is disjoint from $\bigcap_{E \in Y \setminus \{A, B, C\}} E$, which has P^* -probability one by Claim 1, $P^*(A \cap B \cap C) = 0$. This implies Claim 2, since

$$A^c \cup B^c \cup C^c = (A \cap B \cap C)^c.$$

Claim 3. $P^*((A^c \cap B \cap C) \cup (A \cap B^c \cap C) \cup (A \cap B \cap C^c)) = 1$ for *some* extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P .

As $A^c \cap B^c$ is disjoint with each of A^*, B^*, C^* , it is disjoint with $A^* \cup B^* \cup C^*$, which has P_i^* -probability of one for all individuals i by (8). So, $P_i^*(A^c \cap B^c) = 0$, i.e., $P_i^*(A^c \setminus B) = 0$, for all i . Analogously, $P_i^*(A^c \setminus C) = 0$ and $P_i^*(B^c \setminus C) = 0$ for all i . Since, as just shown, each P_i has an extension P_i^* which assigns zero probability to $A^c \setminus B$, $A^c \setminus C$ and $B^c \setminus C$, by conditional consensus compatibility (and Proposition 3) the collective opinion function P also has an extension $P^* \in \mathcal{P}_{\sigma(X)}$ assigning zero probability to these three events, and hence, to their union $(A^c \setminus B) \cup (A^c \setminus C) \cup (B^c \setminus C) = (A^c \cap B^c) \cup (A^c \cap C^c) \cup (B^c \cap C^c)$. In other words, with P^* -probability of zero *at least two* of A^c, B^c, C^c hold. Further, with P^* -probability of one *at least one* of A^c, B^c, C^c holds (by Claim 2). So, with P^* -probability of one *exactly one* of A^c, B^c, C^c holds. This is precisely what had to be shown.

Claim 4. $P^*(A^*) + P^*(B^*) + P^*(C^*) = P^*(A^* \cup B^* \cup C^*) = 1$ for *some* extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P .

Consider an extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P of the kind in Claim 3. The first identity follows from the pairwise disjointness of A^*, B^*, C^* . Regarding the second identity, note that $A^* \cup B^* \cup C^*$ is the intersection of the events $\bigcap_{E \in Y \setminus \{A, B, C\}} E$ and $(A^c \cap B \cap C) \cup (A \cap B^c \cap C) \cup (A \cap B \cap C^c)$, each of which has P^* -probability of one by Claims 1 and 3. So $P^*(A^* \cup B^* \cup C^*) = 1$.

Claim 5. $D(x) + D(y) + D(z) = 1$.

Consider an extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P of the kind in Claim 4. As $P^*(A^* \cup B^* \cup C^*) = 1$ by Claim 4, and as the intersection of A^c with $A^* \cup B^* \cup C^*$ is A^* ,

$$P^*(A^c) = P^*(A^*). \quad (9)$$

Since $A^c \in X$, we also have

$$P^*(A^c) = P_{P_1, \dots, P_n}(A^c) = D(P_1(A^c), \dots, P_n(A^c)),$$

where $P_i(A^c) = P_i^*(A^c) = x_i$ for all individuals i . So $P^*(A^c) = D(x)$. This and (9) imply that $P^*(A^*) = D(x)$. Similarly, $P^*(B^*) = D(y)$ and $P^*(C^*) = D(z)$. So Claim 5 follows from Claim 4. ■

A.5 Proof of Theorem 6(a)

As explained in the main text, Theorem 6(a) follows from Theorem 3(a) via Lemmas 7 and 5 (while applying Lemma 11(b)). It remains to prove Lemma 7.

Proof of Lemma 7. Let D be the local pooling criterion for such a pooling function for a partitional agenda X . Consider any $x, y, z \in [0, 1]^n$ with sum $\mathbf{1}$. Since X is partitional, some countable $Y \subseteq X$ partitions Ω into at least three non-empty events. Choose distinct $A, B, C \in Y$. For each individual i , since $x_i + y_i + z_i = 1$ and since A, B and C are pairwise disjoint and non-empty, there is some $P_i \in \mathcal{P}_X$ such that

$$P_i(A) = x_i, P_i(B) = y_i, P_i(C) = z_i.$$

Let P be the collective opinion function for this profile. Since Y is a countable partition of Ω and P can be extended to a (σ -additive) probability function, $\sum_{E \in Y} P(E) = 1$. Now, for each $E \in Y \setminus \{A, B, C\}$, we have $P(E) = 0$ by consensus preservation (as $P_i(E) = 0$ for all i). So $P(A) + P(B) + P(C) = 1$. Hence $D(x) + D(y) + D(z) = 1$ because

$$\begin{aligned} P(A) &= D(P_1(A), \dots, P_n(A)) = D(x), \\ P(B) &= D(P_1(B), \dots, P_n(B)) = D(y), \\ P(C) &= D(P_1(C), \dots, P_n(C)) = D(z). \blacksquare \end{aligned}$$

A.6 Proof of parts (b) of all theorems

Parts (b) of three of the six theorems will be proved by reduction to results in the companion paper. To prepare this reduction, we first relate opinion pooling on a general agenda X to premise-based opinion pooling on a σ -algebra agenda, as analysed in the companion paper. Consider any agenda X and any σ -algebra agenda Σ of which X is a subagenda. (A *subagenda* of an agenda is a subset which is itself an agenda, i.e., a non-empty subset closed under complementation.) For instance, Σ could be $\sigma(X)$. We can think of the pooling function F for X as being induced by a pooling function F^* for the larger agenda Σ . Formally, a pooling function $F^* : \mathcal{P}_\Sigma^n \rightarrow \mathcal{P}_\Sigma$ for agenda Σ *induces* the pooling function $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ for (sub)agenda X if F^* and F generate the same collective opinions within X , i.e.,

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

(Strictly speaking, we further require that $\mathcal{P}_X = \{P|_X : P \in \mathcal{P}_\Sigma\}$, but this requirement holds automatically in standard cases, e.g., if X is finite or $\sigma(X) = \Sigma$.²⁴) We call F^* the *inducing* pooling function, and F the *induced* one. Our

²⁴In these cases, each opinion function in \mathcal{P}_X is extendable not just to a probability function on $\sigma(X)$, but also to one on Σ . In general, extensions beyond $\sigma(X)$ may not always be possible,

axiomatic requirements on the induced pooling function F – i.e., independence and the various consensus requirements – can be related to the following requirements on the inducing pooling function F^* for the agenda Σ (introduced and discussed in the companion paper):

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

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

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

Conditional consensus preservation on X . For all A, B in subagenda X and all $P_1, \dots, P_n \in \mathcal{P}_\Sigma$, if, for each individual i , $P_i(A|B) = 1$ (provided $P_i(B) \neq 0$), then $P_{P_1, \dots, P_n}(A|B) = 1$ (provided $P_{P_1, \dots, P_n}(B) \neq 0$).²⁵

The following lemma establishes some key relationships between the properties of the induced and the inducing pooling functions:

Lemma 12 *Suppose a pooling function F^* for a σ -algebra agenda Σ induces a pooling function F for a subagenda 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 on the correspondence between opinion pooling on general agendas and on σ -algebra agendas.²⁶

as is well-known from measure theory. For instance, if $\Omega = \mathbb{R}$, X consists of all intervals or complements thereof, and $\Sigma = 2^{\mathbb{R}}$, then $\sigma(X)$ contains the Borel-measurable subsets of \mathbb{R} , and it is well-known that measures on $\sigma(X)$ may not be extendable to $\Sigma = 2^{\mathbb{R}}$ (a fact related to the Banach-Tarski paradox).

²⁵If one compares this requirement with that of conditional consensus compatibility for a general agenda X , one might wonder why the new requirement involves only a single conditional certainty (i.e., that of A given B), whereas the earlier requirement involves an entire set of conditional certainties (which must be respected simultaneously). The key point is that if each P_i is a probability function on Σ , then the simplified requirement as stated here *implies* the more complicated requirement from the main text.

²⁶More precisely, Lemma 12 is a corollary of a slightly generalized statement of Lemma 13, in

Lemma 13 Consider an agenda X and the corresponding σ -algebra agenda $\Sigma = \sigma(X)$. 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 further 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 further be replaced by ‘some’;
- (d) consensus-compatible if and only if some inducing pooling function for agenda Σ is consensus-preserving;
- (e) conditional-consensus-compatible if and only if some inducing pooling function for agenda Σ is conditional-consensus-preserving on X

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

Proof of Lemma 13. Consider an agenda X , the generated σ -algebra $\Sigma = \sigma(X)$, and a pooling function F for X .

(a) For each $P \in \mathcal{P}_X$, fix an extension in \mathcal{P}_Σ denoted \bar{P} . Consider the pooling function F^* for Σ defined by $F^*(P_1^*, \dots, P_n^*) = \overline{F(P_1^*|_X, \dots, P_n^*|_X)}$ for all $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$. Clearly, F^* induces F (regardless of how the extensions \bar{P} of $P \in \mathcal{P}_X$ were chosen).

(b) We give a proof for the ‘independence’ case; the proofs for the ‘neutrality’ and ‘linearity’ cases are analogous. Note (using part (a)) that replacing ‘every’ by ‘some’ strengthens the ‘if’ claim and weakens the ‘only if’ claim. It thus suffices to prove the ‘if’ claim with ‘some’, and the ‘only if’ claim with ‘every’. Clearly, if *some* inducing F^* is independent on X , then F inherits independence. Now let F be independent with pooling criteria D_A , $A \in X$. Consider any $F^* : \mathcal{P}_\Sigma^n \rightarrow \mathcal{P}_\Sigma^n$ inducing F . Then F^* is independent on X with the same pooling criteria as for F because for all $A \in X$ and all $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$ we have

$$\begin{aligned} F^*(P_1^*, \dots, P_n^*)(A) &= F(P_1^*|_X, \dots, P_n^*|_X)(A) \text{ as } F^* \text{ induces } F \\ &= D_A(P_1^*|_X(A), \dots, P_n^*|_X(A)) \text{ by } F\text{'s independence} \\ &= D_A(P_1^*(A), \dots, P_n^*(A)). \end{aligned}$$

(c) As in part (b), it suffices to prove the ‘if’ claim with ‘some’, and the ‘only if’ claim with ‘every’. Clearly, if *some* inducing F^* is consensus-preserving on X , F inherits consensus preservation. Now let F be consensus-preserving and induced by F^* . Then F^* is consensus-preserving on X because, for all $A \in X$ and

which Σ is *either* $\sigma(X)$ *or*, if X is finite, any σ -algebra which includes X . Our proof of Lemma 13 can be extended to this generalized statement (drawing on Lemma 15 and using an argument related to the ‘Claim’ in the proof of Theorem 1(b) of the companion paper).

$P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$ such that $P_1^*(A) = \dots = P_n^*(A) = 1$, we have

$$\begin{aligned} F^*(P_1^*, \dots, P_n^*)(A) &= F(P_1^*|_X, \dots, P_n^*|_X)(A) \text{ as } F^* \text{ induces } F \\ &= 1 \text{ as } F \text{ is consensus preserving.} \end{aligned}$$

(d) First, let F be consensus-compatible and X finite. We define F^* as follows. For any $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$, consider the event A^* in Σ which is *smallest* subject to having probability one under each P_i^* . This event exists and is constructible as $A^* = \bigcap_{A \in \sigma(X): P_1^*(A) = \dots = P_n^*(A) = 1} A$, drawing on finiteness of $\Sigma = \sigma(X)$ and the fact that intersections of finitely many events of probability one have probability one. Clearly, A^* is the union of the supports of the functions P_i^* . We define $F^*(P_1^*, \dots, P_n^*)$ as any extension in \mathcal{P}_Σ of $F(P_1^*|_X, \dots, P_n^*|_X)$ assigning probability one to A^* . Such an extension exists because F is consensus-compatible and each $P_i^*|_X$ is extendable to a probability function (namely P_i^*) assigning probability one to A^* . Clearly, F^* induces F . It also is consensus-preserving: for all $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$ and $A \in \Sigma$, if $P_1^*(A) = \dots = P_n^*(A) = 1$, then A includes the above-constructed event A^* , whence $F^*(P_1^*, \dots, P_n^*)(A) = 1$ as $F^*(P_1^*, \dots, P_n^*)(A^*) = 1$.

Conversely, let *some* inducing pooling function F^* be consensus-preserving. To see why F is consensus-compatible, consider $P_1, \dots, P_n \in \mathcal{P}_X$ and $A \in \Sigma$ such that each P_i has an extension $\bar{P}_i \in \mathcal{P}_\Sigma$ for which $\bar{P}_i(A) = 1$. We show that some extension $\bar{P} \in \mathcal{P}_\Sigma$ of $F(P_1, \dots, P_n)$ satisfies $\bar{P}(A) = 1$. Simply let \bar{P} be $F^*(\bar{P}_1, \dots, \bar{P}_n)$ and note that \bar{P} is indeed an extension of $F(P_1, \dots, P_n)$ (as F^* induces F) and $\bar{P}(A) = 1$ (as F^* is consensus-preserving).

(e) First, let F be conditional-consensus-compatible, and let X be finite. We define F^* as follows. For a profile $(P_1^*, \dots, P_n^*) \in \mathcal{P}_\Sigma^n$, consider the (finite) set S of pairs (A, B) in X such that $P_i^*(A|B) = 1$ for each i with $P_i^*(B) \neq 0$ (equivalently, such that $P_i^*(B \setminus A) = 0$ for each i). Since F is conditional-consensus-compatible (and since in the last sentence we can replace each ' P_i^* ' with ' $P_i^*|_X$ '), there is an extension $P^* \in \mathcal{P}_\Sigma$ of $F(P_1^*|_X, \dots, P_n^*|_X)$ such that $P^*(A|B) = 1$ for all $(A, B) \in S$ for which $P^*(B) \neq 0$. Let $F^*(P_1^*, \dots, P_n^*) := P^*$. Clearly, F^* induces F and is conditional-consensus-preserving on X .

Conversely, let *some* inducing F^* be conditional-consensus-preserving on X . To check that F is conditional-consensus-compatible, consider $P_1, \dots, P_n \in \mathcal{P}_X$ and a finite set S of pairs (A, B) in X such that each P_i can be extended to $\bar{P}_i \in \mathcal{P}_\Sigma$ with $\bar{P}_i(A|B) = 1$ (provided $\bar{P}_i(B) \neq 0$). We require an extension $P^* \in \mathcal{P}_\Sigma$ of $F(P_1, \dots, P_n)$ such that $P^*(A|B) = 1$ for all $(A, B) \in S$ for which $P^*(B) \neq 0$. Now $P^* := F^*(\bar{P}_1, \dots, \bar{P}_n)$ is such an extension, since F^* induces F and is conditional-consensus-preserving on X . ■

Which pooling functions for Σ induce ones for X ? Here is a sufficient condition:

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

The proof draws on a measure-theoretic fact in which the word ‘finite’ is essential:

Lemma 15 *Every probability function on a finite sub- σ -algebra of σ -algebra Σ can be extended to a probability function on Σ .*

Proof. Let $\Sigma' \subseteq \Sigma$ be a finite sub- σ -algebra of σ -algebra Σ , and consider any $P' \in \mathcal{P}_{\Sigma'}$. Let \mathcal{A} be the set of atoms of Σ' , i.e., (\subseteq -)minimal events in $\Sigma' \setminus \{\emptyset\}$. As Σ' is finite, \mathcal{A} must partition Ω . So, $\sum_{A \in \mathcal{A}} P'(A) = 1$. For each $A \in \mathcal{A}$, let Q_A be a probability function on Σ such that $Q_A(A) = 1$. (Such functions exist, since each Q_A could for instance be the Dirac measure at some $\omega_A \in A$.) Then $P := \sum_{A \in \mathcal{A}} P'(A)Q_A$ defines a probability function on Σ , because (given the identity $\sum_{A \in \mathcal{A}: P'(A) \neq 0} P'(A) = 1$) it is a convex combination of probability functions on Σ . Further, P extends P' , because it agrees with P' on \mathcal{A} , hence on Σ' . ■

Proof of Lemma 14. Suppose the pooling function F for σ -algebra agenda Σ is independent on subagenda X , and that X is finite or $\sigma(X) = \Sigma$. Let $\Sigma' := \sigma(X)$. If X is finite, so is Σ' . Each $P \in \mathcal{P}_X$ can by definition be extended to a function in $\mathcal{P}_{\Sigma'}$, which (by Lemma 15 in case Σ' is a finite σ -algebra distinct from Σ) can be extended to a function in \mathcal{P}_{Σ} . For any $Q \in \mathcal{P}_X$, pick an extension $\bar{Q} \in \mathcal{P}_{\Sigma}$. Define a pooling function F' for X by

$$F'(Q_1, \dots, Q_n) := F(\bar{Q}_1, \dots, \bar{Q}_n)|_X \text{ for all } Q_1, \dots, Q_n \in \mathcal{P}_X.$$

Now F induces F' for two reasons. First, for all $P_1, \dots, P_n \in \mathcal{P}_{\Sigma}$,

$$F'(P_1|_X, \dots, P_n|_X) = F(\overline{P_1|_X}, \dots, \overline{P_n|_X})|_X = F(P_1, \dots, P_n)|_X,$$

where the second ‘=’ holds as F is independent on X . Second, $\mathcal{P}_X = \{P|_X : P \in \mathcal{P}_{\Sigma}\}$, where ‘ \supseteq ’ is trivial and ‘ \subseteq ’ holds because each $P \in \mathcal{P}_X$ equals $\bar{P}|_X$. ■

Proof of parts (b) of Theorems 1-6. First, Theorems 2(b) and 6(b) follow directly from Theorems 1(b) and 3(b), respectively, since consensus compatibility implies conditional consensus compatibility (by Proposition 1) and as non-neutrality implies non-linearity.

Second, we derive Theorems 1(b), 3(b) and 5(b) from the corresponding results in the companion paper, namely Theorems 1(b), 3(b), and 5(b), respectively. The

derivations are similar for the three results; we thus spell out the derivation only for Theorem 1(b). Consider a nested agenda $X \neq \{\Omega, \emptyset\}$. By the companion paper's Theorem 1(b) (see also the footnote to it), some pooling function F^* for agenda $\Sigma := \sigma(X)$ is independent on X , (globally) consensus preserving and non-neutral on X . By Lemma 14, F^* induces a pooling function for (sub)agenda X , which by Lemma 12 is independent, consensus-compatible, and non-neutral.

Finally, we prove Theorem 4(b) directly rather than by reduction. Consider an agenda $X \neq \{\emptyset, \Omega\}$ which is nested or satisfies $|X \setminus \{\emptyset, \Omega\}| \leq 4$. If X is nested, the claim follows from Theorem 1(b), since non-neutrality implies non-linearity. Now let X be non-nested and $|X \setminus \{\emptyset, \Omega\}| \leq 4$. We may assume without loss of generality that $\emptyset, \Omega \notin X$ (as any independent, consensus-compatible, and non-neutral pooling function for agenda $X' = X \setminus \{\emptyset, \Omega\}$ induces one for agenda X). Since $|X| \leq 4$, and since $|X| > 2$ (as X is non-nested), we have $|X| = 4$, say $X = \{A, A^c, B, B^c\}$. By non-nestedness, A and B are logically independent, i.e., the events $A \cap B, A \cap B^c, A^c \cap B, A^c \cap B^c$ are all non-empty. On \mathcal{P}_X^n , consider the function $F : (P_1, \dots, P_n) \mapsto T \circ P_1$, where $T(p)$ is 1 if $p = 1$, 0 if $p = 0$, and $\frac{1}{2}$ if $p \in (0, 1)$. We complete the proof by establishing that (i) F maps into \mathcal{P}_X , i.e., is a proper pooling function, (ii) F is consensus-compatible, (iii) F is independent, and (iv) F is non-linear. Claims (iii) and (iv) hold trivially.

Proof of (i): Let $P_1, \dots, P_n \in \mathcal{P}_X$ and $P := F(P_1, \dots, P_n) = T \circ P_1$. We need to extend P to a probability function on $\sigma(X)$. For each atom C of $\sigma(X)$ (i.e., each $C \in \{A \cap B, A \cap B^c, A^c \cap B, A^c \cap B^c\}$), let P_C be the unique probability function on $\sigma(X)$ assigning probability one to C . We distinguish between three (exhaustive) cases.

Case 1: $P_1(E) = 1$ for two events E in X . Without loss of generality, let $P_1(A) = P_1(B) = 1$, and hence, $P_1(A^c) = P_1(B^c) = 0$. It follows that $P(A) = P(B) = 1$ and $P(A^c) = P(B^c) = 0$. So P extends (in fact, uniquely) to a probability function on $\sigma(X)$, namely to $P_{A \cap B}$.

Case 2: $P_1(E) = 1$ for exactly one event E in X . Without loss of generality, assume $P_1(A) = 1$ (hence, $P_1(A^c) = 0$) and $P_1(B), P_1(B^c) \in (0, 1)$. Hence, $P(A) = 1, P(A^c) = 0$ and $P(B) = P(B^c) = \frac{1}{2}$. So P extends (again uniquely) to a probability function on $\sigma(X)$, namely to $\frac{1}{2}P_{A \cap B} + \frac{1}{2}P_{A \cap B^c}$.

Case 3: $P_1(E) = 1$ for no event E in X . Then $P_1(A), P_1(A^c), P_1(B), P_1(B^c) \in (0, 1)$, and so $P(A) = P(A^c) = P(B) = P(B^c) = \frac{1}{2}$. Hence, P extends (non-uniquely) to a probability function on $\sigma(X)$, e.g., to $\frac{1}{2}P_{A \cap B} + \frac{1}{2}P_{A^c \cap B^c}$ or $\frac{1}{4}P_{A \cap B} + \frac{1}{4}P_{A^c \cap B} + \frac{1}{4}P_{A \cap B^c} + \frac{1}{4}P_{A^c \cap B^c}$.

Proof of (ii): Let $P_1, \dots, P_n \in \mathcal{P}_X$ and consider any $C \in \sigma(X)$ such that each P_i extends to some $P_i^* \in \mathcal{P}_{\sigma(X)}$ such that $P_i^*(C) = 1$. (It only matters that P_1 has such an extension, given the definition of F .) We have to show that $P := F(P_1, \dots, P_n) = T \circ P_1$ is extendable to a $P^* \in \mathcal{P}_{\sigma(X)}$ such that $P^*(C) = 1$. We verify the claim in each of the three cases considered in the proof of (i). In

Cases 1 and 2, the claim holds because the (unique) extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P has the same support as P_1^* . (In fact, in Case 1 $P^* = P_1^*$.) In Case 3, C must intersect with each event in X (otherwise some event in X would have zero probability under P_1 , in contradiction with Case 3) and include more than one of the atoms $A \cap B$, $A \cap B^c$, $A^c \cap B$, and $A^c \cap B^c$ (again by Case 3). As is easily checked, $C \supseteq (A \cap B) \cup (A^c \cap B^c)$ or $C \supseteq (A \cap B^c) \cup (A^c \cap B)$. So, to ensure that the extension P^* or P satisfies $P^*(C) = 1$, it suffices to specify P^* as $\frac{1}{2}P_{A \cap B} + \frac{1}{2}P_{A^c \cap B^c}$ in the first case, and as $\frac{1}{2}P_{A \cap B^c} + \frac{1}{2}P_{A^c \cap B}$ in the second case. ■

A.7 Proof of Propositions 1 and 2

Proof of Proposition 1. Consider an opinion pooling function for an agenda X . We first prove part (b), by showing that conditional consensus compatibility is equivalent to the restriction of consensus compatibility to events A expressible as $(\cup_{(C,D) \in S} (C \setminus D))^c$ for finite $S \subseteq X \times X$. This fact follows from the equivalence of conditional consensus compatibility and implication preservation (Proposition 3) and the observation that, for any such set S , an opinion function is consistent with zero probability of all $C \setminus D$ with $(C, D) \in S$ if and only if it is consistent with zero probability of $\cup_{(C,D) \in S} (C \setminus D)$, i.e., probability one of $(\cup_{(C,D) \in S} (C \setminus D))^c$.

We now prove part (a) The claims made about implicit consensus preservation and consensus compatibility have already been proved (informally) in the main text. It remains to show that conditional consensus compatibility implies consensus preservation and is equivalent to it if $X = \sigma(X)$. As just shown, conditional consensus compatibility is equivalent to the restriction of consensus compatibility to events A of the form $(\cup_{(C,D) \in S} (C \setminus D))^c$ for some finite set $S \subseteq X \times X$. Note that, for any $A \in X$, we may define S as $\{(A^c, A)\}$, so that $(\cup_{(C,D) \in S} (C \setminus D))^c = (A^c \setminus A)^c = A$. So, conditional consensus compatibility implies consensus preservation and is equivalent to it if $X = \sigma(X)$. ■

Proof of Proposition 2. Assume $|\Omega| \geq 4$. We can thus partition Ω into four non-empty events and let X consist of any union of *two* of these four events. The set X is indeed an agenda since $A \in X \Leftrightarrow A^c \in X$. Since nothing depends on the sizes of the four events, we assume without loss of generality that they are singleton, i.e., that $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and $X = \{A \subseteq \Omega : |A| = 2\}$.

Step 1. We here show that X is path-connected and non-partitional. Non-partitionality is trivial. To establish path-connectedness, we consider events $A, B \in X$ and must construct a path of conditional entailments from A to B . This is done by distinguishing between three cases.

Case 1: $A = B$. Then the path is trivial, since $A \vdash^* A$ (take $Y = \emptyset$).

Case 2: A and B have exactly one world in common. Call it ω , and let ω' be the unique world in $\Omega \setminus (A \cup B)$. Then $A \vdash^* B$ in virtue of $Y = \{\{\omega, \omega'\}\}$.

Case 3: A and B have no world in common. We may then write $A = \{\omega_A, \omega'_A\}$ and $B = \{\omega_B, \omega'_B\}$ with $\omega_A, \omega'_A, \omega_B, \omega'_B$ pairwise distinct. Now $\{\omega_A, \omega'_A\} \vdash^* \{\omega_A, \omega_B\}$ (take $Y = \{\{\omega_A, \omega'_B\}\}$) and $\{\omega_A, \omega_B\} \vdash^* \{\omega_B, \omega'_B\}$ (take $Y = \{\{\omega_B, \omega'_A\}\}$).

Step 2. We now construct a pooling function $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ that is independent (in fact, neutral), consensus-preserving, and non-linear. As an ingredient of the construction, consider first a linear pooling function $L : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ (for instance the dictatorial one given by $(P_1, \dots, P_n) \mapsto P_1$). We shall transform L into a non-linear pooling function that is still neutral and consensus-preserving. First, fix a transformation $T : [0, 1] \rightarrow [0, 1]$ such that:

- (i) $T(1 - x) = 1 - T(x)$ for all $x \in [0, 1]$ (hence $T(1/2) = 1/2$),
- (ii) $T(0) = 0$ (hence by (i) $T(1) = 1$),
- (iii) T is strictly concave on $[0, 1/2]$ (hence by (i) strictly convex on $[1/2, 1]$).

(Such a T exists; e.g. $T(x) = 4(x - 1/2)^3 + 1/2$ for all $x \in [0, 1]$.) Now, for any $P_1, \dots, P_n \in \mathcal{P}_X$ and $A \in X$, let $P_{P_1, \dots, P_n}(A) := T(L(P_1, \dots, P_n)(A))$. We must prove that, for any $P_1, \dots, P_n \in \mathcal{P}_X$, the function P_{P_1, \dots, P_n} , as just defined, can indeed be extended to a probability function on $\sigma(X) = 2^\Omega$. This completes the proof, as it establishes that we have defined a proper pooling function and this pooling function is neutral (since L is neutral), consensus-preserving (since L is consensus-preserving and $T(1) = 1$), and non-linear (since L is linear and T a non-linear transformation).

To show that P_{P_1, \dots, P_n} can be extended to a probability function on $\sigma(X) = 2^\Omega$, we consider any probability function Q on 2^Ω and show that $T \circ Q|_X$ extends to a probability function on 2^Ω (which completes our task, since $Q|_X$ could be $L(P_1, \dots, P_n)$ for $P_1, \dots, P_n \in \mathcal{P}_X$). It suffices to prove that there exist real numbers $p_k = p_k^Q$, $k = 1, 2, 3, 4$, such that the function on 2^Ω assigning p_k to each $\{\omega_k\}$ is a *probability* function and extends $T \circ Q|_X$, i.e., such that

- (a) $p_1, p_2, p_3, p_4 \geq 0$ and $p_1 + p_2 + p_3 + p_4 = 1$,
- (b) for all $A \in X$, $T(Q(A)) = \sum_{k: \omega_k \in A} p_k$.

For all $k \in \{1, 2, 3, 4\}$, let $q^k := Q(\{\omega_k\})$; and for all $k, l \in \{1, 2, 3, 4\}$ with $k < l$, let $q_{kl} := Q(\{\omega_k, \omega_l\})$. In order for p_1, \dots, p_4 to satisfy (b), they must satisfy the system

$$p_k + p_l = T(q_{kl}) \text{ for all } k, l \in \{1, 2, 3, 4\} \text{ with } k < l.$$

Given $p_1 + p_2 + p_3 + p_4 = 1$, three of these six equations are redundant. Indeed, consider $k, l \in \{1, 2, 3, 4\}$, $k < l$, and define $k', l' \in \{1, 2, 3, 4\}$, $k' < l'$, by $\{k', l'\} = \{1, 2, 3, 4\} \setminus \{k, l\}$. As $p_k + p_l = 1 - p_{k'} - p_{l'}$ and $T(q_{kl}) = T(1 - q_{k'l'}) = 1 - T(q_{k'l'})$, the equation $p_k + p_l = T(q_{kl})$ is equivalent to $p_{k'} + p_{l'} = T(q_{k'l'})$. So (b) reduces (given $p_1 + p_2 + p_3 + p_4 = 1$) to the system

$$p_1 + p_2 = T(q_{12}), p_1 + p_3 = T(q_{13}), p_2 + p_3 = T(q_{23}).$$

This is a system of three linear equations in three variables $p_1, p_2, p_3 \in \mathbb{R}$. To solve it, let $t_{kl} := T(q_{kl})$ for all $k, l \in \{1, 2, 3, 4\}$, $k < l$. We first bring the coefficient

matrix of our three-equation system into triangular form:

$$\begin{aligned} \begin{pmatrix} 1 & 1 & t_{12} \\ 1 & & 1 & t_{13} \\ & 1 & 1 & t_{23} \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 1 & t_{12} \\ & -1 & 1 & t_{13} - t_{12} \\ & & 2 & t_{23} + t_{13} - t_{12} \end{pmatrix} \\ &\rightarrow \begin{pmatrix} 1 & 1 & t_{12} \\ & 1 & -1 & t_{12} - t_{13} \\ & & 1 & \frac{t_{23} + t_{13} - t_{12}}{2} \end{pmatrix}. \end{aligned}$$

The system therefore has the following solution:

$$p_3 = \frac{t_{23} + t_{13} - t_{12}}{2} \quad (10)$$

$$p_2 = t_{12} - t_{13} + \frac{t_{23} + t_{13} - t_{12}}{2} = \frac{t_{12} + t_{23} - t_{13}}{2} \quad (11)$$

$$p_1 = t_{12} - \frac{t_{12} + t_{23} - t_{13}}{2} = \frac{t_{12} + t_{13} - t_{23}}{2}$$

Recalling that $p_4 = 1 - (p_1 + p_2 + p_3)$, we also have

$$p_4 = 1 - \frac{t_{12} + t_{13} + t_{23}}{2}. \quad (12)$$

By their construction, the numbers p_1, \dots, p_4 given by (10)-(12) satisfy condition (b) and equation $p_1 + \dots + p_4 = 1$. To complete the proof of conditions (a)-(b), it remains to show that $p_1, \dots, p_4 \geq 0$. We do this by proving two claims.

Claim 1. $p_4 \geq 0$, i.e., $\frac{t_{12} + t_{13} + t_{23}}{2} \leq 1$.

We have to prove that $T(q_{12}) + T(q_{13}) + T(q_{23}) \leq 2$. Note that

$$q_{12} + q_{13} + q_{23} = q^1 + q^2 + q^1 + q^3 + q^2 + q^3 = 2(q^1 + q^2 + q^3) \leq 2.$$

We distinguish three cases.

Case 1: All of q_{12}, q_{13}, q_{23} are all at least $1/2$. Then, by (i)-(iii), $T(q_{12}) + T(q_{13}) + T(q_{23}) \leq q_{12} + q_{13} + q_{23} \leq 2$, as desired.

Case 2: At least two of q_{12}, q_{13}, q_{23} are below $1/2$. Then, again using (i)-(iii), $T(q_{12}) + T(q_{13}) + T(q_{23}) < 1/2 + 1/2 + 1 = 2$, as desired.

Case 3: Exactly one of q_{12}, q_{13}, q_{23} is below $1/2$. Suppose $q_{12} < 1/2 \leq q_{13} \leq q_{23}$ (otherwise just switch the roles of q_{12}, q_{13}, q_{23}). For all $\delta \geq 0$ such that $q_{23} + \delta \leq 1$, the properties (i)-(iii) of T imply that

$$T(q_{13}) + T(q_{23}) \leq T(q_{13} - \delta) + T(q_{23} + \delta). \quad (13)$$

Since the graphical intuition for (13) is clear, let us only give an informal proof, stressing visualisation. Dividing by 2, we have to show that the average value $a_1 := \frac{1}{2}T(q_{13}) + \frac{1}{2}T(q_{23})$ is at most the average value $a_2 := \frac{1}{2}T(q_{13} - \delta) + \frac{1}{2}T(q_{23} + \delta)$.

Let SL be the straight line segment in \mathbb{R}^2 joining the points $(q_{13} - \delta, T(q_{13} - \delta))$ and $(q_{23} + \delta, T(q_{23} + \delta))$, and let SL^* be the straight line segment joining the points $(q_{13}, T(q_{13}))$ and $(q_{23}, T(q_{23}))$. Since a_1 and a_2 are, respectively, the second coordinates of the points on SL^* and SL with the first coordinate $\frac{1}{2}q_{13} + \frac{1}{2}q_{23}$, it suffices to show that SL^* is ‘below’ SL . This follows once we prove that T ’s graph is ‘below’ SL (as T is convex on $[1/2, 1]$ and SL^* joins two points on T ’s graph on $[1/2, 1]$). If $q_{13} - \delta \geq 1/2$, this is trivial by T ’s convexity on $[1/2, 1]$. Now let $q_{13} - \delta < 1/2$. Let SL' be the straight line segments joining the points $(q_{13} - \delta, T(q_{13} - \delta))$ and $(1 - (q_{13} - \delta), T(1 - (q_{13} - \delta)))$, and let SL'' be the straight line segment joining the points $(1 - (q_{13} - \delta), T(1 - (q_{13} - \delta)))$ and $(q_{23} + \delta, T(q_{23} + \delta))$. Check using T ’s properties that LS' passes through the point $(1/2, 1/2)$. This implies that (*) T ’s graph is ‘below’ SL' on $[1/2, 1]$, and that (**) SL'' is steeper than SL' (by T ’s convexity on $[1/2, 1]$). Also, (***) T ’s graph is ‘below’ SL'' (again by T ’s convexity on $[1/2, 1]$). In sum, on $[1/2, 1]$, T ’s graph is (by (*) and (***)) ‘below’ both SL' and SL'' which are both ‘below’ SL by (**). So, still on $[1/2, 1]$, T ’s graph is ‘below’ SL . This proves (13). Applying (13) with $\delta = 1 - q_{23}$, we obtain

$$T(q_{13}) + T(q_{23}) \leq T(q_{13} - (1 + q_{23})) + T(1).$$

On the right side, $T(1) = 1$ and (as $q_{13} - (1 + q_{23}) \leq 1 - q_{12}$ and as T is increasing) $T(q_{13} - (1 + q_{23})) \leq T(1 - q_{12}) = 1 - T(q_{12})$. So $T(q_{13}) + T(q_{23}) \leq 1 + 1 - T(q_{12})$, i.e., $T(q_{12}) + T(q_{13}) + T(q_{23}) \leq 2$, as claimed.

Claim 2. $p_k \geq 0$ for all $k = 1, 2, 3$.

We only show that $p_1 \geq 0$, as the proofs for p_2 and p_3 are analogous. We have to prove that $t_{13} + t_{23} - t_{12} \geq 0$, i.e., that $T(q_{13}) + T(q_{23}) \geq T(q_{12})$, or equivalently, that $T(q^1 + q^3) + T(q^2 + q^3) \geq T(q^1 + q^2)$. As T is increasing, it suffices to establish that $T(q^1) + T(q^2) \geq T(q^1 + q^2)$. We again consider three cases.

Case 1: $q^1 + q^2 \leq 1/2$. Suppose $q^1 \leq q^2$ (otherwise swap the roles of q^1 and q^2). For all $\delta \geq 0$ such that $q^1 - \delta \geq 0$, we have

$$T(q^1) + T(q^2) \geq T(q^1 - \delta) + T(q^2 + \delta),$$

as T is concave on $[0, 1/2]$ and $0 \leq q^1 - \delta \leq q^1 \leq q^2 \leq q^2 + \delta \leq 1/2$. So, for $\delta = q^1$,

$$T(q^1) + T(q^2) \geq T(0) + T(q^2 + q^1) = T(q^1 + q^2).$$

Case 2: $q^1 + q^2 > 1/2$ but $q^1, q^2 \leq 1/2$. By (i)-(iii),

$$T(q^1) + T(q^2) \geq q^1 + q^2 \geq T(q^1 + q^2).$$

Case 3: $q^1 > 1/2$ or $q^2 > 1/2$. Suppose $q^2 > 1/2$ (otherwise swap q^1 and q^2 in the proof). Then $q^1 < 1/2$, since otherwise $q^1 + q^2 > 1$. Let $y := 1 - q^1 - q^2$. As $y < 1/2$, an argument analogous to that in Case 1 yields $T(q^1) + T(y) \geq T(q^1 + y)$, i.e., $T(q^1) + T(1 - q^1 - q^2) \geq T(1 - q^2)$. So, by (i), $T(q^1) + 1 - T(q^1 + q^2) \geq 1 - T(q^2)$, i.e., $T(q^1) + T(q^2) \geq T(q^1 + q^2)$. ■

One might wonder why the pooling function constructed in this proof violates conditional consensus compatibility. (It must do so, because otherwise pooling would be linear – hence neutral – by Theorem 5(a).) Let Ω and X be as in the proof, and consider a profile with complete unanimity: all individuals i assign probability 0 to ω_1 , 1/4 to ω_2 , 1/4 to ω_3 , and 1/2 to ω_4 . As $\{\omega_1\}$ is the difference of two events in X (e.g. $\{\omega_1, \omega_2\} \setminus \{\omega_2, \omega_3\}$), implication preservation (which is equivalent to conditional consensus compatibility) would require ω_1 's collective probability to be 0 as well. But ω_1 's collective probability is (in the notation of the proof) given by

$$p_1 = \frac{t_{12} + t_{13} - t_{23}}{2} = \frac{T(q_{12}) + T(q_{13}) - T(q_{23})}{2}.$$

Here, q_{kl} is the collective probability of $\{\omega_k, \omega_l\}$ under a linear pooling function, so that q_{kl} is the probability which each individual assigns to $\{\omega_k, \omega_l\}$. So

$$p_1 = \frac{T(1/4) + T(1/4) - T(1/2)}{2} = T(1/4) - \frac{T(1/2)}{2},$$

which is strictly positive as T is strictly concave on $[0, 1/2]$ with $T(0) = 0$.