



Munich Personal RePEc Archive

Probabilistic opinion pooling generalized Part one: General agendas

Dietrich, Franz and List, Christian

Paris School of Economics, London School of Economics

May 2013

Online at <https://mpra.ub.uni-muenchen.de/57253/>
MPRA Paper No. 57253, posted 11 Jul 2014 14:56 UTC

Probabilistic opinion pooling generalized

Part one: General agendas*

Franz Dietrich
CNRS & UEA

Christian List
LSE

May 2013, revised July 2014

Abstract

How can different 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. Our unified theorems show why 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

*Although both authors are jointly responsible for this paper and project, Christian List wishes to note that Franz Dietrich should be considered the lead author, to whom the credit for the present mathematical proofs is due. This paper is the 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 Both papers build on our earlier, unpublished paper ‘Opinion pooling on general agendas’ (September 2007).

1 Introduction

This paper addresses the problem of *probabilistic opinion pooling*. Suppose several individuals submit probability assignments to some events. How can these individual probability assignments be aggregated into a corresponding collective probability assignment, while preserving probabilistic coherence? Although statisticians, economists, and philosophers have worked extensively on this problem, almost all contributions have made the assumption that 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 given σ -algebra. Instead, they focus on a limited set of relevant events, and this set does 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 simply not interested in the probability of such ‘artificial’ composite events. *Or* we – or the experts or decision-makers in question – 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 dependencies 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 without assuming that the agenda is a σ -algebra and consider instead general agendas: *any* set of events that is closed under negation (complementation) can qualify as an agenda. This 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, 2013a, Nehring and Puppe 2010, Dietrich and Mongin 2010, Dokow and Holzman 2010). We impose two axiomatic requirements on probabilistic opinion pooling:

- (i) the familiar ‘independence’ requirement, according to which the probability that a group assigns to an event should depend solely on the probabilities that the individuals assign to it;
- (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 peculiarity 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 many 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: the case of *probabilistic preference aggregation*. Here each individual assigns sub-

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 and 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 other extensions of (or alternatives to) the classic characterization of linear pooling, see also Wagner (1982/1985), Aczél, Ng, and Wagner (1984), Genest (1984), Mongin (1995), and Chambers (2007). All of these works, however, retain the assumption that the agenda is a σ -algebra. Genest and Zidek (1986) provide an excellent review 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 (and non-axiomatic) approach to the aggregation of *partial* probability assignments, where individuals do not assign probabilities to all events in the underlying σ -algebra, see Osherson and Vardi (2006).

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

jective 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 assignments can be interpreted, for instance, as beliefs about which preferences are the ‘correct’ ones (e.g., which correctly capture objective quality comparisons between the alternatives). The group must then arrive at collective probability assignments to the same events, which can be interpreted as the corresponding collective beliefs. (Alternatively, the probability assignments can be interpreted as vague or fuzzy preferences.)

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 fully 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 the present 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 suggests 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 set of events to which probabilities are to be assigned is called the *agenda*. Existing works on probability aggregation have focused on agendas that are σ -*algebras*, i.e., sets of events that are closed under complementation and countable union (and by implication also countable intersection). Here, we lift that restriction. As already noted, we may exclude some events from the agenda, *either* because they are of no interest, *or* because the individuals are unable to assign probabilities to them. For example, the agenda may contain (i) the event that global warming will continue, (ii) the event that the recession will continue, and (iii) the event that the UK will remain in the European Union, but not the disjunction of all three events, which

may be too difficult to assess. 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 the events A and B may or may not be logically related. As should be clear, X may contain A and B without containing $A \cup B$ or $A \cap B$.

An example of an agenda containing no conjunctive or disjunctive events. To give a simple example of an agenda that contains no conjunctions or disjunctions, suppose each possible world is specified by a vector of three binary characteristics. The first takes the value 1 if atmospheric CO₂ is above some critical 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 can 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 of course subsets of Ω .³) Although there are non-trivial overlaps and even logical connections between these events (in particular, A and $A \rightarrow B$ are inconsistent with B^c), the agenda contains no conjunctions or disjunctions. An expert committee may well be faced with an opinion pooling problem on this agenda.

Probabilistic opinions. Let us 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 of events 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

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).

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 each of the events A , $A \rightarrow B$, and B^c . This would be probabilistically incoherent since such an opinion function 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 simply the set of all probability functions on it.

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 $\frac{1}{n}P_1 + \dots + \frac{1}{n}P_n$. Here, we focus on linear and neutral pooling functions, as defined in the next paragraph.

Linearity and neutrality. We call a pooling function *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, we call a pooling function *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 (which explains the term ‘neutral’). 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 any linear average of opinion functions is an opinion function), it is far from clear whether a given non-linear function $D : [0, 1]^n \rightarrow [0, 1]$ defines a proper pooling function. In particular, formula (1) might not yield a well-defined opinion function, which has to be coherent. We will see that whether there can be neutral

⁴So, whenever 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. Sometimes $\sigma(X)$ is simply the set 2^Ω of *all* events; this happens when all events are constructible from events in X .

but non-linear pooling functions depends on the agenda in question. If the agenda is a σ -algebra, it is well known that the answer is negative (assuming $|X| > 4$). However, we will also identify many agendas for which the answer is positive.

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

3 Two kinds of applications

We can 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 rejected by the constitutional court’;
- (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 a particular farm.

Arguably, the study of 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 full σ -algebra. Here it would usually be implausible to replace the agenda X with the σ -algebra $\sigma(X)$, because many events in $\sigma(X)$ represent complex combinations of events, which we are neither interested in nor able to assess. Furthermore, even when $\sigma(X)$ is finite, it is often enormous: if X contains at least k logically independent events, then $\sigma(X)$ contains at least 2^{2^k} events, and so its size grows double-exponentially in k .⁵ This suggests that, unless k is small, $\sigma(X)$ is often far too large to be used as an agenda in practice.

By contrast, an application of type (b) plausibly gives rise to an agenda that is a σ -algebra, not only because the decision-makers may *need* a full probability distribution over the σ -algebra, but also because they may be *able* to specify such a distribution. For instance, a market analyst who has been asked to estimate next month’s distribution of Apple’s share price might decide to submit a log-normal distribution. This, in turn, requires the specification of ‘only’ two parameters:

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 $\{\bigcup_{C \in \mathcal{C}} C : \mathcal{C} \subseteq \mathcal{A}\}$ containing $2^{2^k} = 16$ events.

the mean and the variance of the exponential of the share price. Opinion pooling problems of type (b) are discussed in more detail in a companion paper (Dietrich and List 2013b), where they are one of our principal applications.

4 Axiomatic requirements on opinion pooling

We now introduce some requirements on opinion pooling functions for a general agenda X .

4.1 The independence requirement

Our first requirement, *independence*, says that the collective probability of each event in the agenda should depend only on the individual probabilities of *that* event. This requirement is familiar from the literature and 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 $P_1, \dots, P_n \in \mathcal{P}_X$,

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

The main justification for independence is the democratic idea that the collective view on any issue should depend only on individual views on that issue. This reflects a *local*, as opposed to *holistic*, understanding of democracy. (Under a *holistic* understanding, the collective view on an issue may also be influenced by individual views on other issues.) Independence, understood as a democratic requirement, 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 seem 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 disjunctive event ‘a blizzard *or* an interest-rate increase’. Here, however, we focus on general agendas, where the democratic 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, determining the collective view on any issue based on individual views on that 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

4.2 The consensus-preservation requirement

Our next requirement says that whenever 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 standard literature, where it is sometimes expressed as a zero-probability preservation requirement. In the case of general agendas, however, we can also formulate several strengthened variants of the requirement, which extend it to other forms of consensus. Although these strengthened variants are not as compelling as the original requirement, they are still defensible in some cases. Their close relationship with the original requirement is illustrated by the fact that when the agenda is a σ -algebra, they all collapse back into consensus preservation in its original form.

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 $B \supseteq A$.

⁶In the classical case in which 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 .

⁷In particular, when the agenda is a σ -algebra, independence is known to conflict with the preservation of unanimously held judgments of probabilistic independence, assuming non-dictatorial aggregation (see Genest and Wagner 1984; Bradley, Dietrich, and List forthcoming). 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.

- Individual i 's *unrevealed beliefs* are any 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 to $\sigma(X)$ which assign probability 1 to an event B but may also admit extensions which assign probability less than 1. In this case, individual i 's belief about B is unrevealed. From the perspective of a pooling function into which opinion functions like P_i are fed as input, such unrevealed beliefs are unobserved.

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, for every extension \bar{P} of P , we have $\bar{P}(A) = 1$.

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 (even 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 is more complicated to state than consensus compatibility, but 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 Appendix A8.

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

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.

5 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 invoked, 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.

5.1 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 great 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.*
 (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 a logically stronger claim, and part (b) a logically weaker claim. 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 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 the agenda of our expert committee and all nested agendas X ($\neq \{\emptyset, \Omega\}$). The following theorem holds.

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 $A \in X \setminus \{\emptyset, \Omega\}$.

- 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.*

5.2 Sketch proofs

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. First recall that two events A and B are *exclusive* if $A \cap B = \emptyset$ and *exhaustive* if $A \cup B = \Omega$. Now, for any events $A, B \in X \setminus \{\emptyset, \Omega\}$, we define

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

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).*

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 some linear pooling criterion, 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.

6 When is opinion pooling linear?

As we have just seen, for many agendas, only neutral pooling functions can satisfy our two requirements. But are these pooling functions also linear? As we now show, the answer depends on the agenda in question. 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.

6.1 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 In particular, to see how exactly non-simplicity strengthens non-nestedness, note the following fact (due to Dietrich 2013 in binary judgment-aggregation theory

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 the (frequent) case that the events in X represent 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 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 .

and proved here again in the present framework, for completeness):

- 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. 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.*

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 already 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 with another appropriate condition), 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 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.

6.2 Sketch proofs

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

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

¹²A generalized definition of partitionality is possible in Theorem 6: we could define an agenda 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 Ω ($= \bigcap_{B \in \emptyset} B$) 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.

¹³This assumes that the underlying set of worlds Ω satisfies $|\Omega| \geq 4$.

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.

7 Classic results on opinion pooling 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). The first thing to note is that, for a σ -algebra, all the agenda conditions we have used reduce to a simple condition on agenda *size*. It is easy to see that the following result holds:

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),*

- 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);*
- 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.*

8 An illustrative application: probabilistic preference aggregation

To illustrate the use of general agendas, we now apply our results to a rather different context, namely that of *probabilistic preference aggregation*, a probabilistic analogue of Arrovian 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.).

Now, 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 linear 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 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 not have

¹⁴To see for instance that X_K is non-simple if $|K| > 2$, choose three distinct alternatives $x, y, z \in K$ and note that the 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 these three events form a minimal inconsistent subset of X_K .

any negative flavour. Furthermore, 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 extendible 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 investigated. A 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). In such a model, 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 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.

9 A unified perspective

Finally, we wish to compare probabilistic opinion pooling, as discussed here, 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 1.¹⁵ 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*).¹⁶ 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 8, 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 very similar to those we have 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.

¹⁵Formally, 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).

¹⁶Specifically, 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.¹⁷

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 (2013a); 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.¹⁸

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*

¹⁷To 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.

¹⁸In the binary case, two of our unanimity-preservation requirements – namely implicit consensus preservation and consensus compatibility – become 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.¹⁹

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.²⁰

10 References

- Aczél, J. (1966) *Lectures on Functional Equations and their Applications*, Dover Publications Inc., United States of America
- Aczél, J., and C. Wagner (1980) “A characterization of weighted arithmetic means”, *SIAM Journal on Algebraic and Discrete Methods* 1: 259-260
- Aczél, J., C. T. Ng and C. Wagner (1984) “Aggregation Theorems for Allocation Problems”, *SIAM Journal on Algebraic and Discrete Methods* 5: 1-8
- Billot, A. (1991) “Aggregation of preferences: The fuzzy case”, *Theory and Decision* 30(1): 51-93
- Bradley, R., F. Dietrich, and C. List (forthcoming) “Aggregating Causal Judgments”, *Philosophy of Science*

¹⁹In 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).

²⁰Of 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
- 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
- Dietrich, F., and C. List (2013a) “Propositionwise judgment aggregation: the general case”, *Social Choice and Welfare* 40: 1067-1095
- Dietrich, F., and C. List (2013b) “Probabilistic opinion pooling generalized – Part 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: 153-63
- Genest, C., and J. V. Zidek (1986) “Combining Probability Distributions: A Critique and Annotated Bibliography”, *Statistical Science* 1(1): 113-135
- Genest, C., and C. Wagner (1984) “Further Evidence against Independence Preservation in Expert Judgement Synthesis”, *Technical Report* 84-10, Department of Statistics and Actuarial Science, University of Waterloo
- 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: 410-414
- Mongin, P. (1995) “Consistent Bayesian aggregation”, *Journal of Economic Theory* 66: 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
- Nehring, K., M. Pivato and C. Puppe (2013) “The Condorcet Set: Majority Voting over Interconnected Propositions”, working paper
- 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: 564-580
- Priest, G. (2001) *An Introduction to Non-classical Logic*, Cambridge University Press
- Sanver, M. R., and Ö. Selçuk (2009) “Sophisticated Preference Aggregation”, *Social Choice and Welfare* 33: 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: 207-220
- Wagner, C. (1985) “On the Formal Properties of Weighted Averaging as a Method of Aggregation”, *Synthese* 62: 97-108
- Weymark, J. (1997) “Aggregating Ordinal Probabilities on Finite Sets”, *Journal of Economic Theory* 75: 407-432

A Proofs

In this appendix, we prove all our results. Given the close mathematical connection between the present results and those in our companion paper on ‘premise-based’ opinion pooling for a σ -algebra agenda (Dietrich and List 2013b), there are in principle two plausible proof strategies: either we could prove our present results directly and those in the companion paper as corollaries, or *vice versa*. As it turns out, mixing these two strategies is the most natural approach. We will prove parts (a) of most of our present theorems directly (and use them in the companion paper to derive the results stated there), while we will derive parts (b) of most of our present theorems from the corresponding results in the companion paper.

Section A.1 provides some preparatory lemmas needed to establish a translation between the two papers. Sections A.2, A.3, ..., A.7 contain the proofs of Theorems 1, 2, ..., 6, together with some related results. Finally, Section A.8 contains the proof of Proposition 2.

A.1 The relationship with ‘premise-based’ opinion pooling for a σ -algebra agenda

We now relate opinion pooling on a general agenda X to opinion pooling on a σ -algebra agenda, as analysed in our companion paper. This technical connection will be exploited in our proofs.

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 set 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 is said to *induce* 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 $\sigma(X) = \Sigma$.²¹) We will call F^* the *inducing* pooling function, and F the *induced* pooling function.

Our axiomatic requirements on the induced pooling function F – i.e., independence and the various consensus requirements – can be related to the following axiomatic 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 $P_i(B) \neq 0$),

²¹In these cases, each opinion function in \mathcal{P}_X is extendable not just to a probability measure on $\sigma(X)$, but also to one on Σ . In general, extensions beyond $\sigma(X)$ may not always be possible, 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 extend to $\Sigma = 2^{\mathbb{R}}$ (a fact related to the Banach-Tarski paradox).

then $P_{P_1, \dots, P_n}(A|B) = 1$ (provided $P_{P_1, \dots, P_n}(B) \neq 0$).²²

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

Lemma 10 *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 .

In fact, this lemma is a corollary of a more general result on the correspondence between opinion pooling on a general agenda and opinion pooling on a larger σ -algebra agenda.²³

Lemma 11 *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).

²²When comparing this requirement to 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 when 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.

²³More precisely, Lemma 10 is a corollary of a slightly generalized statement of Lemma 11, in which Σ is *either* $\sigma(X)$ *or*, if X is finite, any σ -algebra which includes X . Our proof of Lemma 11 can be extended to this generalized statement (drawing on Lemma 12 and using an argument related to the ‘Claim’ in the proof of Theorem 1(b) of the companion paper).

The proof of Lemma 11 draws on the following measure-theoretic fact (in which the word ‘finite’ is essential):

Lemma 12 *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., of (\subseteq -)minimal events in $\Sigma' \setminus \{\emptyset\}$. Using the fact that Σ' is finite, it easily follows that \mathcal{A} partitions 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 in a point $\omega_A \in A$.) Then

$$P := \sum_{A \in \mathcal{A}} P'(A)Q_A$$

defines a probability function on Σ , because (by 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 all of Σ' . ■

Proof of Lemma 11. 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 any extension in \mathcal{P}_Σ , to be denoted \bar{P} . Consider the pooling function F^* for Σ defined by

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

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 therefore suffices to prove the ‘if’ claim with ‘some’, and the ‘only if’ claim with ‘every’.

First 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

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

Now suppose *some* inducing pooling function F^* is independent on X . Clearly, F inherits independence from F^* .

(c) As in part (b), replacing ‘every’ with ‘some’ strengthens the ‘if’ claim and weakens the ‘only if’ claim, so that it suffices to prove the ‘if’ claim with ‘some’, and the ‘only if’ claim with ‘every’. First, let F be consensus-preserving, and consider any inducing F^* . Then F^* is consensus-preserving on X because, for all $A \in X$ and all $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$ such that $P_1^*(A) = \dots = P_n^*(A) = 1$, we have

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

Now suppose *some* inducing F^* is consensus-preserving on X . It is clear that F inherits consensus preservation from F^* .

(d) First, let F be consensus-compatible, and let X be finite. We define F^* as follows. For any profile $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$, we first form the event A^* in Σ which is *smallest* (with respect to set-inclusion) subject to having probability one under each P_i^* ; to see why this event exists, note that the intersection of finitely many events of probability one has probability one, so that (since $\Sigma = \sigma(X)$ is finite) we can construct A^* as the intersection $\bigcap_{A \in \sigma(X): P_1^*(A) = \dots = P_n^*(A) = 1} A$. In other words, 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 $F(A_1^*|_X, \dots, A_n^*|_X)$ 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^* . The so-defined pooling function F^* clearly induces F . It is also consensus-preserving since, for any $P_1^*, \dots, P_n^* \in \mathcal{P}_\Sigma$ and any event $A \in \Sigma$, if $P_1^*(A) = \dots = P_n^*(A) = 1$, then $A \supseteq A^*$ (where A^* is the event constructed as above), and thus $F^*(P_1^*, \dots, P_n^*)(A) = 1$ since $F^*(P_1^*, \dots, P_n^*)(A^*) = 1$.

Conversely, assume *some* inducing F^* is consensus-preserving. To see why F is consensus-compatible, we consider any $P_1, \dots, P_n \in \mathcal{P}_X$ and any $A \in \Sigma$ such that each P_i has an extension $\overline{P}_i \in \mathcal{P}_\Sigma$ for which $\overline{P}_i(A) = 1$. We have to show that there is an extension $\overline{P} \in \mathcal{P}_\Sigma$ of $F(P_1, \dots, P_n)$ such that $\overline{P}(A) = 1$. Simply define \overline{P} as $F^*(\overline{P}_1, \dots, \overline{P}_n)$ and note that it is indeed the case that (i) \overline{P} extends $F(P_1, \dots, P_n)$, since F^* induces F , and (ii) $\overline{P}(A) = 1$ since F^* is consensus-preserving.

(e) First, let F be conditional-consensus-compatible, and let X be finite. We define F^* as follows. For any profile $(P_1^*, \dots, P_n^*) \in \mathcal{P}_\Sigma^n$, first form the (finite) set S of all 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). By the conditional consensus compatibility of F (and the fact that in the last sentence we could of course have replaced each ‘ P_i^* ’ by ‘ $P_i^*|_X$ ’), 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 $P^*(B) \neq 0$. We define $F^*(P_1^*, \dots, P_n^*)$ as 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 any $P_1, \dots, P_n \in \mathcal{P}_X$ and any finite set S of pairs (A, B) in X such that each P_i extends to a $\overline{P}_i \in \mathcal{P}_\Sigma$ with $\overline{P}_i(A|B) = 1$ (provided $\overline{P}_i(B) \neq 0$). We need to find 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 $P^*(B) \neq 0$. The function $P^* := F^*(\overline{P}_1, \dots, \overline{P}_n)$ is such an extension, because F^* induces F and is conditional-consensus-preserving on X . ■

Given a pooling function F^* for the σ -algebra agenda Σ , does it induce any pooling function for subagenda X ? The following result gives a sufficient condition.

Lemma 13 *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 .*

Proof of Lemma 13. Suppose pooling function F for σ -algebra agenda Σ is independent on the subagenda X , and that X is finite or $\sigma(X) = \Sigma$. Let $\Sigma' := \sigma(X)$. Notice that if X is finite, so is Σ' . Each $P \in \mathcal{P}_X$ by definition extends to a function in $\mathcal{P}_{\Sigma'}$, which, using Lemma 12 if Σ' is a (finite) σ -algebra distinct from Σ , extends to a function in \mathcal{P}_Σ . For any $Q \in \mathcal{P}_X$ let $\overline{Q} \in \mathcal{P}_\Sigma$ be an extension. Define a pooling function F' for X by

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

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

$$\begin{aligned} F'(P_1|_X, \dots, P_n|_X) &= F(\overline{P_1|_X}, \dots, \overline{P_n|_X})|_X \\ &= F(P_1, \dots, P_n)|_X \text{ (as } F \text{ is independent on } X\text{)}. \end{aligned}$$

Second, the identity $\mathcal{P}_X = \{P|_X : P \in \mathcal{P}_\Sigma\}$ holds: ‘ \supseteq ’ is trivial, and ‘ \subseteq ’ holds because each $P \in \mathcal{P}_X$ can be written as $\overline{P}|_X$. ■

A.2 Proof of Theorem 1

Proof of Theorem 1. (a) This part follows from Theorem 2(a), since consensus compatibility implies conditional consensus compatibility.

(b) We reduce this part to the companion paper’s Theorem 1(b). Consider a nested agenda $X \neq \{\emptyset, \Omega\}$. By the companion paper’s Theorem 1(b) (see also the footnote to that theorem), there is a pooling function F^* for the agenda $\Sigma := \sigma(X)$ which is independent on X , (globally) consensus preserving, but not neutral on X . By independence on X , this pooling function induces a pooling function for the (sub)agenda X (see Lemma 13). This pooling function is independent and consensus-compatible and not neutral (see Lemma 10). ■

A.3 Proof of Theorem 2

As mentioned, Theorem 2(a) is based on the above Lemmas 1 and 2. To prove these, we first show the following lemma.

Lemma 14 *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 of Lemma 14. (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 show 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 not exclusive and not exhaustive. It follows that 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 not exclusive (as $A_j^c \cap A_{j+1}^c = (A_j \cup A_{j+1})^c \neq \Omega^c = \emptyset$) and not 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 $A \sim B$ in virtue of 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 equivalent with respect to \sim . By $A \not\sim B$, either $A \cap B = \emptyset$ or $A \cup B = \Omega$. Without loss of generality, we may assume the former case, because in the latter case we may consider the complements A^c, B^c, C^c instead of A, B, C , using the fact that A^c, B^c, C^c are pairwise not equivalent with respect to \sim by (b) with $A^c \cap B^c = (A \cup B)^c = \Omega^c = \emptyset$. Now by $A \cap B = \emptyset$, 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. Suppose, for a contradiction, that Z is an equivalence class with respect to \sim containing the pair A, A^c . By assumption, Z is not the only equivalence class with respect to \sim , and 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 14(a), \sim is indeed an equivalence relation on $X \setminus \{\emptyset, \Omega\}$. By Lemma 14(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 that we may assume without loss of generality that $\emptyset, \Omega \notin X$.

First suppose there are two equivalence classes with respect to \sim . Let X_+ be one of them. By Lemma 14(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 . As \subseteq is reflexive, transitive, and anti-symmetric, what we have to show is connectedness. So, suppose $A, B \in X_+$, and let us show that $A \subseteq B$ or $B \subseteq A$. Since $A \not\sim B^c$ (by Lemma 14(d)), either $A \cap B^c = \emptyset$ or $A \cup B^c = \Omega$. In the first case, $A \subseteq B$. In the second case, $B \subseteq A$.

Conversely, let X be nested, i.e., of the form $X = \{A, A^c : A \in X_+\}$ for some set $X_+ \subseteq \Sigma$ that is linearly ordered by set-inclusion \subseteq . Consider any $A \in X_+$. We show that $A \not\sim A^c$, which shows that X has more than one – hence by Lemma 14(d) exactly two – equivalence classes with respect to \sim , as desired. 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 $A_j \cup A_{j+1} \neq \Omega$. As each event $C \in X$ is either in X^+ or has complement in X^+ , and as $A_1 = A \in X^+$ 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 $A_{j+1}^c \subseteq A_j$. In the first case, $A_j \cap A_{j+1} = \emptyset$, a contradiction. In the second case, $A_j \cup A_{j+1} = \Omega$, also a contradiction. ■

Before proving Lemma 2, we give a useful re-formulation of the condition 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 $B \setminus A$, the event that B holds without A . So, conditional consensus compatibility can be re-formulated as the following condition (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} .

The following proposition states what has just been (informally) shown:

Proposition 3 *For any agenda X , conditional consensus compatibility is equivalent to implication preservation.*

We are ready to prove Lemma 2.

Proof of Lemma 2. Let F be an independent and conditional-consensus-compatible pooling function for agenda X . For each $A \in X$, let D_A be the pooling criterion as 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 implies immediately that $D_A = D_B$ whenever $A \sim B$ (by induction on the length k of a path from A to B), completing the proof.

So, suppose $A, B \in X$ with $A \cap B \neq \emptyset$ and $A \cup B \neq \Omega$. Notice that the events $A \cap B$, $A \cup B$ and $A \setminus B$ need not belong to X . Consider any $x \in [0, 1]^n$, and let us 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 exist probability functions $P_1^*, \dots, P_n^* \in \mathcal{P}_{\sigma(X)}$ such that

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

Now consider the opinion functions $P_1, \dots, P_n \in \mathcal{P}_X$ given by $P_i := P_i^*|_X$ for each individual i . Since $P_i^*(A \setminus B) = 0$ and $P_i^*(B \setminus A) = 0$ for all i , the collective opinion function P_{P_1, \dots, P_n} has an extension $P_{P_1, \dots, P_n}^* \in \mathcal{P}_{\sigma(X)}$ such that

$$P_{P_1, \dots, P_n}^*(A \setminus B) = P_{P_1, \dots, P_n}^*(B \setminus A) = 0$$

by implication preservation (which is equivalent to conditional consensus compatibility by Proposition 3). It follows that

$$P_{P_1, \dots, P_n}^*(A) = P_{P_1, \dots, P_n}^*(A \cap B) = P_{P_1, \dots, P_n}^*(B).$$

Hence,

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

So, using the fact that $P_{P_1, \dots, P_n}(A) = D_A(x)$ (because $P_i(A) = x_i$ for all i) and $P_{P_1, \dots, P_n}(B) = D_B(x)$ (because $P_i(B) = x_i$ for all i), it follows that $D_A(x) = D_B(x)$, as desired. ■

Proof of Theorem 2. (a) This part follows immediately from Lemmas 1 and 2.

(b) This part follows from Theorem 1(b) because consensus compatibility implies conditional consensus compatibility (by Proposition 1). ■

A.4 Proof of Theorem 3

Proof of Lemma 3. Let $F : \mathcal{P}_X^n \rightarrow \mathcal{P}_X$ be independent and consensus-preserving, and consider $A, B \in X$. Suppose $A \vdash^* B$, say in virtue of (countable) set $Y \subseteq X$. Write D_A and D_B for the pooling criteria 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 $\bigcap_{C \in \{A\} \cup Y} C$ is non-empty but has empty intersection with B^c (by the conditional entailment), it equals its intersection with B , and so $\bigcap_{C \in \{A, B\} \cup Y} C \neq \emptyset$. Similarly, as $\bigcap_{C \in \{B^c\} \cup Y} C$ is non-empty but has empty intersection with A , it equals its intersection with A^c , so that $\bigcap_{C \in \{A^c, B^c\} \cup Y} C \neq \emptyset$. Hence there are worlds $\omega \in \bigcap_{C \in \{A, B\} \cup Y} C$ and $\omega' \in \bigcap_{C \in \{A^c, B^c\} \cup Y} C$. For each individual i , consider the probability function $P_i^* : \sigma(X) \rightarrow [0, 1]$ defined by

$$P_i^* := x_i \delta_\omega + (1 - x_i) \delta_{\omega'},$$

where $\delta_\omega, \delta_{\omega'} : \sigma(X) \rightarrow [0, 1]$ denote the Dirac-measures in ω and ω' , respectively; and consider the opinion function $P_i := P_i^*|_X$. As each P_i satisfies $P_i(A) = P_i(B) = x_i$, we have

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

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), and hence $P_{P_1, \dots, P_n}(\bigcap_{C \in Y} C) = 1$ since the intersection of *countably* many events of probability one has again probability one. So

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

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

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

where the equality holds by an earlier argument. ■

Proof of Theorem 3. (a) Let X be path-connected, and let F be independent and consensus-preserving. We show neutrality. Without loss of generality, let $X \neq \{\emptyset, \Omega\}$. We write D_A for the local pooling criterion of any contingent event $A \in X \setminus \{\emptyset, \Omega\}$. As X is path-connected, repeated application of Lemma 3 yields $D_A \leq D_B$ for all $A, B \in X \setminus \{\emptyset, \Omega\}$, and hence $D_A = D_B$ for all $A, B \in X \setminus \{\emptyset, \Omega\}$. Define D as the common pooling criterion D_A of all $A \in X \setminus \{\emptyset, \Omega\}$. We complete the neutrality proof by showing that D also works as a pooling criterion for Ω and \emptyset . This follows from the fact that, as a consequence of consensus preservation, $D(\mathbf{1}) = 1$ and $D(\mathbf{0}) = 0$.

(b) This part is reducible to the companion paper's Theorem 3(b). Consider a non-path-connected finite agenda X . By the companion paper's Theorem 3(b), there is a pooling function F^* for agenda $\Sigma := \sigma(X)$ which, on X , is independent and consensus-preserving but not neutral. By independence on X , this pooling function induces a pooling function for (sub)agenda X (see Lemma 13), which is independent and consensus-preserving but not neutral (see Lemma 10). ■

A.5 Proof of Theorem 4

The proof of Theorem 4(a) (specifically, of Lemma 4) draws on the following agenda characterization result.

Proposition 4 *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 $A \in Y$.

Proof. Let X be an agenda. The equivalence between (b) and (c) is obvious (to see why (b) implies (c), simply replace the set Y in (b) by a three-member subset of it). It is also relatively easy to see why (c) implies (a). Indeed, whenever (a) is violated, so is (c), by the following argument. First, if $|X \setminus \{\Omega, \emptyset\}| \leq 4$, then (c) is violated since every three-element set $Y \subseteq X$ either contains Ω or \emptyset (so that $Y \setminus \{A\} \cup \{A^c\}$ is inconsistent for some $A \in Y$) or takes the form $Y = \{B, B^c, A\}$ (so that $Y \setminus \{A\} \cup \{A^c\}$ is inconsistent). Second, if the agenda X is nested, say $X = \{C, C^c : C \in Z\}$ for some subset $Z \subseteq X$ that is linearly ordered by set-inclusion, condition (c) is violated since any three-element set $Y \subseteq X$ has elements $A \neq B$ which both belong to Z or both belong to $\{A^c : A \in Z\}$, so that (by the linearity of Z and of $\{C^c : C \in Z\}$) $A \subseteq B$ or $B \subseteq A$, and hence, $(Y \setminus \{B\}) \cup \{B^c\}$ is inconsistent or $(Y \setminus \{A\}) \cup \{A^c\}$ is inconsistent.

It remains to show that (a) implies (c). Let X be non-nested with $|X \setminus \{\Omega, \emptyset\}| > 4$; we show (c). Without loss of generality, we may assume that $\Omega, \emptyset \notin X$ (since otherwise it suffices to do the proof for $X \setminus \{\Omega, \emptyset\}$ instead of X , drawing on the fact that each of the conditions (a) and (c) holds for X if and only if it holds for $X \setminus \{\Omega, \emptyset\}$). We distinguish between two cases.

Case 1: Some $A, B \in X$ are logically independent, i.e., all of $A \cap B$, $A \cap B^c$, $A^c \cap B$ and $A^c \cap B^c$ are non-empty. Consider such $A, B \in X$. Since $|X| > 4$ there is a $C \in X \setminus \{A, A^c, B, B^c\}$. As C is non-empty, it intersects with at least one of $A \cap B$, $A \cap B^c$, $A^c \cap B$ and $A^c \cap B^c$. We may assume without loss of generality that $A \cap B^c \cap C \neq \emptyset$ (otherwise, simply interchange A with A^c and/or B with B^c). Our argument distinguishes between two subcases.

Subcase 1.1: $A^c \cap B^c \cap C^c, A \cap B \cap C^c \neq \emptyset$. In this case, condition (c) holds for $Y := \{A, B^c, C^c\}$, since $A^c \cap B^c \cap C^c \neq \emptyset$, $A \cap B \cap C^c \neq \emptyset$ and $A \cap B^c \cap C \neq \emptyset$.

Subcase 1.2: $A^c \cap B^c \cap C^c = \emptyset$ or $A \cap B \cap C^c = \emptyset$ (perhaps both). We assume without loss of generality that $A \cap B \cap C^c = \emptyset$, i.e., that $A \cap B \subseteq C$. (The proof is analogous in the other case.) There are three subsubcases.

Subsubcase 1.2.1: $A^c \cap B \cap C^c, A \cap B^c \cap C^c \neq \emptyset$. Here, condition (c) holds for $Y := \{A, B, C^c\}$, since $A^c \cap B \cap C^c \neq \emptyset$, $A \cap B^c \cap C^c \neq \emptyset$ and $A \cap B \cap C = A \cap B \neq \emptyset$.

Subsubcase 1.2.2: $A^c \cap B \cap C^c = \emptyset$. So $A^c \cap B \subseteq C$. As also $A \cap B \subseteq C$, we have $B \subseteq C$. We distinguish between cases:

- First assume $A^c \cap B^c \cap C^c \neq \emptyset$. Then condition (c) holds with $Y = \{A^c, B^c, C\}$, because $A \cap B^c \cap C \neq \emptyset$, $A^c \cap B \cap C \neq \emptyset$ (as $A^c \cap B \neq \emptyset$ and $B \subseteq C$) and $A^c \cap B^c \cap C^c \neq \emptyset$.
- Second assume $A^c \cap B^c \cap C^c = \emptyset$, i.e., $A^c \cap B^c \subseteq C$. Since also $B \subseteq C$, we have $C^c \subseteq A \cap B^c$. Condition (c) holds with $Y = \{A, B^c, C\}$, because $A^c \cap B^c \cap C \neq \emptyset$ (as $\emptyset \neq A^c \cap B^c \subseteq C$), $A \cap B \cap C \neq \emptyset$ (as $\emptyset \neq A \cap B \subseteq C$) and $A \cap B^c \cap C^c \neq \emptyset$ (as $\emptyset \neq C^c \subseteq A \cap B^c$).

Subsubcase 1.2.3: $A \cap B^c \cap C^c = \emptyset$. (If in the following proof for the current subcase we interchange A and B , then we obtain an alternative, but longer, proof for Subsubcase 1.2.2.) Since $A \cap B^c \cap C^c = \emptyset$ we have $A \cap B^c \subseteq C$. As also $A \cap B \subseteq C$, it follows that $A \subseteq C$. So, since $C \neq A, \Omega$, we have $A \subsetneq C \subsetneq \Omega$. We now show that

$$(*) A^c \cap B \cap C, A^c \cap B^c \cap C^c \neq \emptyset \text{ or } (**) A^c \cap B^c \cap C, A^c \cap B \cap C^c \neq \emptyset. \quad (3)$$

To show this, we assume that (*) is violated and show that (**) holds, by distinguishing between two cases:

- First, assume that $A^c \cap B \cap C = \emptyset$. It follows, on the one hand, that $A^c \cap B \cap C^c \neq \emptyset$ (as $A^c \cap B \neq \emptyset$), and, on the other hand, that $A^c \cap B^c \cap C \neq \emptyset$ (as otherwise $A^c \cap C = \emptyset$, i.e., $C \subseteq A$, a contradiction since $A \subsetneq C$). This proves (**).
- Second, assume that $A^c \cap B \cap C \neq \emptyset$. Then $A^c \cap B^c \cap C^c = \emptyset$ as (*) is violated. It follows, on the one hand, that $A^c \cap B^c \cap C \neq \emptyset$ (as $A^c \cap B^c \neq \emptyset$), and, on the other hand, that $A^c \cap B \cap C^c \neq \emptyset$ (as otherwise $A^c \cap C^c = \emptyset$, i.e., $A^c \subseteq C$, a contradiction since $A \subseteq C \subsetneq \Omega$). This proves (**).

We can now prove condition (c). In the case of (*), (c) holds with $Y = \{A^c, B^c, C\}$, since $A \cap B^c \cap C \neq \emptyset$ (as $A \cap B^c \neq \emptyset$ and $A \subseteq C$), $A^c \cap B \cap C \neq \emptyset$ (by (*)) and $A^c \cap B^c \cap C^c \neq \emptyset$ (by (**)). In the case of (**), (c) holds with $Y = \{A^c, B, C\}$, since $A \cap B \cap C \neq \emptyset$ (as $A \cap B \neq \emptyset$ and $A \subseteq C$), $A^c \cap B^c \cap C \neq \emptyset$ (by (**)) and $A^c \cap B \cap C^c \neq \emptyset$ (by (**)).

Case 2: No $A, B \in X$ are logically independent.

Claim 2.1. There exists a (with respect to set-inclusion) maximal nested (sub-)agenda $X^* \subseteq X$.

This follows from Zorn's Lemma using the fact that the set \mathcal{V} of nested sub-agendas $V \subseteq X$ is non-empty (because $X \neq \emptyset$ and for any $A \in X$ we have $\{A, A^c\} \in \mathcal{V}$) and the fact that every subset $\emptyset \neq \mathcal{W} \subseteq \mathcal{V}$ that is linearly ordered with respect to set-inclusion (a *chain*) has an upper bound in \mathcal{V} (namely the agenda $\cup_{V \in \mathcal{W}} V$, which is indeed nested, as is easily shown). This proves Claim 2.1.

Since X^* is nested, we may express it as $X^* = \{A, A^c : A \in X_+^*\}$ where X_+^* is a subset of X^* which contains exactly one member of each pair $A, A^c \in X^*$ and is linearly ordered with respect to set-inclusion.

Claim 2.2. There exists $D \in X \setminus X^*$ such that $D \cap A \neq \emptyset$ for all $A \in X_+^*$.

Since X^* is nested but X is not, we have $X^* \subsetneq X$, and thus there are $C, C^c \in X \setminus X^*$. It suffices to show that at least one of the sets C and C^c intersects all $A \in X_+^*$. This is true because otherwise there would exist $A, A' \in X_+^*$ such that $C \cap A = C^c \cap A' = \emptyset$, which (recalling that $A \subseteq A'$ or $A' \subseteq A$, and writing A'' for the smaller one of A and A') implies that $C \cap A'' = C^c \cap A'' = \emptyset$, hence, that $A'' = \emptyset$, a contradiction since $\emptyset \notin X$. This proves Claim 2.2.

Let

$$\begin{aligned} Y_1 & : = \{A \in X_+^* : A \subseteq D\}, \\ Y_2 & : = \{A \in X_+^* : A^c \subseteq D\}. \end{aligned}$$

Claim 2.3. $Y_1 \cap Y_2 = \emptyset$, and $Y_1 \cup Y_2 = X_+^*$.

First, $Y_1 \cap Y_2 = \emptyset$, because otherwise there would be an $A \in X_+^*$ such that $A \subseteq D$ and $A^c \subseteq D$, a contradiction as $D \neq \Omega$. Second, suppose for a contradiction that $A \in X_+^* \setminus (Y_1 \cup Y_2)$. Since A and D are not logically independent (by assumption of Case 2), and since $A \cap D \neq \emptyset$ (by Claim 2.2), $A \cap D^c \neq \emptyset$ (as $A \notin Y_1$) and $A^c \cap D^c \neq \emptyset$ (as $A \notin Y_2$), it follows that $A^c \cap D = \emptyset$, i.e., $D \subseteq A$. We next show that D is included not just in A , but also in all other events in $X_+^* \setminus Y_1$:

$$D \subseteq A' \text{ for all } A' \in X_+^* \setminus Y_1. \quad (4)$$

To show this, let $A' \in X_+^* \setminus Y_1$, and note first that $A'^c \cap D^c \supseteq A'^c \cap A^c \cap D^c$, which (by the fact that $D \subseteq A$, i.e., $A^c \cap D^c = A^c$) reduces to $A'^c \cap A^c$, which in turn is non-empty since either $A^c \supseteq A'^c \neq \emptyset$ or $A'^c \supseteq A^c \neq \emptyset$. So, $A'^c \cap D^c \neq \emptyset$. Since A' and D are not logically independent (by assumption of Case 2), and since $A' \cap D \neq \emptyset$ (by Claim 2.2), $A' \cap D^c \neq \emptyset$ (as $A' \notin Y_1$) and $A'^c \cap D^c \neq \emptyset$ (as just shown), it follows that $A'^c \cap D = \emptyset$, so that $D \subseteq A'$. This proves (4).

Note that, for every event A' in X_+^* , either $A' \subseteq D$ (if $A' \in Y_1$) or $D \subseteq A'$ (if $A' \notin Y_1$, by (4)). So the augmented (sub)agenda $X^* \cup \{D, D^c\}$ is nested. This is a contradiction as X^* is a maximal nested subagenda of X . Claim 2.3 is thus established.

Claim 2.4. $Y_1, Y_2 \neq \emptyset$.

By Claim 2.3 we may equivalently show that $Y_1, Y_2 \neq X_+^*$. Suppose for a contradiction that $Y_1 = X_+^*$ or $Y_2 = X_+^*$. Then $X^* \cup \{D, D^c\}$ is a nested agenda which is larger than X^* (since $D \notin X^*$), a contradiction since X^* is a maximal nested subagenda of X . This proves Claim 2.4.

The proof of condition (c) is completed by combining Claim 2.4 with the following observation:

Claim 2.5. For all $B \in Y_1$ and $C \in Y_2$, the set $Y := \{B^c, C, D\}$ satisfies the requirements of condition (c), i.e., $|Y| = 3$ and $(Y \setminus \{A\}) \cup \{A^c\}$ is consistent for each $A \in Y$.

Consider any $B \in Y_1$ and $C \in Y_2$ and let $Y := \{B^c, C, D\}$. To see why $|Y| = 3$, note that $B^c \neq C$ since $C \in X_+^*$ while $B^c \notin X_+^*$, and that $D \neq B^c, C$ since $B^c, C \in X^*$ while $D \notin X^*$. Further:

- $\{B, C, D\}$ is consistent by the follow argument. First, $B \subseteq C$ as B and C belong to the linearly ordered set X_+^* and as $C \not\subseteq B$ by the fact that $B \in Y_1$ and $C \notin Y_1$. So $B \cap C \cap D = B \cap D$. The latter set is indeed non-empty by Claim 2.2.
- $\{B^c, C^c, D\}$ is consistent by the following argument. First, since $B \subseteq C$ (as just shown), $C^c \subseteq B^c$. Also, $C^c \subseteq D$ since $C \in Y_2$. So $B^c \cap C^c \cap D = C^c$, where this set is of course non-empty as $\emptyset \notin X$.
- $\{B^c, C, D^c\}$ is consistent for the following reason. First, $D^c \subseteq B^c$ since $B \subseteq D$ (as $B \in Y_1$). Second, $D^c \subseteq C$ since $C^c \subseteq D$ (as $C \in Y_2$). So, $B^c \cap C \cap D^c = D^c$, which is non-empty as $\emptyset \notin X$. ■

As mentioned in the main text, Theorem 4(a) is based on Lemmas 4 and 5. To prove these, we first show a simple lemma:

Lemma 15 *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 of Lemma 15. (a) Note that, as $X \neq \{\Omega, \emptyset\}$, X contains an event $A \neq \emptyset, \Omega$. For each $x \in [0, 1]^n$ there are (by $A \neq \emptyset, \Omega$) opinion functions $P_1, \dots, P_n \in \mathcal{P}_X$ such that $(P_1(A), \dots, P_n(A)) = x$, and hence $(P_1(A^c), \dots, P_n(A^c)) = \mathbf{1} - x$, which implies that

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

as desired.

(b) Since the pooling function is consensus-preserving, then $D(\mathbf{1}) = 1$, whence by part (a) $D(\mathbf{0}) = 1 - D(\mathbf{1}) = 0$. ■

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

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

By construction,

$$P_i^*(A^* \cup B^* \cup C^*) = x_i + y_i + z_i = 1 \text{ for all } i. \quad (5)$$

Now consider the restriction $P_i := P_i^*|_X$ for each individual i . For the so-defined profile $(P_1, \dots, P_n) \in \mathcal{P}_X^n$, we consider the collective opinion function P_{P_1, \dots, P_n} . The proof is completed by showing two claims.

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

The first equality holds for *all* extensions $P^* \in \mathcal{P}_{\sigma(X)}$ of P , by the pairwise disjointness of the events A^*, B^*, C^* . Regarding the second equality, note that each individual i 's opinion function 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^* ,

$$P^*(A^c) = P^*(A^*). \quad (6)$$

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 (6) imply that $P^*(A^*) = D(x)$. By similar arguments, $P^*(B^*) = D(y)$ 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

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

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

$$D(x) + D(1 - x) = 1 \text{ for all } x \in [0, 1] \quad (8)$$

(since $D(x) + D(\mathbf{1} - x) + D(\mathbf{0}) = 1$ where $D(\mathbf{0}) = 0$). Using (7) and then (8), 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,

$$D(x + y) = D(x) + D(y) \text{ for all } x, y \in [0, 1]^n \text{ with } x + y \in [0, 1]^n. \quad (9)$$

For any $i \in \{1, \dots, n\}$, consider the function $D_i : [0, 1] \rightarrow [0, 1]$ defined 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 (9), D_i satisfies $D_i(s + t) = D_i(s) + D_i(t)$ 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. Hence there exists a $w_i \geq 0$ such that $\bar{D}_i(t) = w_i t$ for all $t \geq 0$ by a well-known theorem (see Aczél 1966, Theorem 1). Now for all $x \in [0, 1]^n$, we have $D(x) = \sum_{i=1}^n D_i(x_i)$ (by repeated application of (9)), and so (by $D_i(x_i) = \bar{D}_i(x_i) = w_i x_i$) $D(x) = \sum_{i=1}^n w_i x_i$. Applying the latter with $x = \mathbf{1}$ yields $D(\mathbf{1}) = \sum_{i=1}^n w_i$, hence $\sum_{i=1}^n w_i = 1$. ■

Proof of Theorem 4. (a) This part is proved by first using Theorem 1(a) to obtain neutrality, and then inferring linearity from Lemmas 4 and 5 (and from Lemma 15(b), which of course applies as consensus compatibility implies consensus preservation).

(b) Consider any 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 assume X is non-nested and $|X \setminus \{\emptyset, \Omega\}| \leq 4$. We may assume without loss of generality that $\emptyset, \Omega \notin X$ (since any independent, consensus-compatible and non-neutral pooling function for agenda $X' = X \setminus \{\emptyset, \Omega\}$ immediately extends to such a pooling function 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 $A \cap B$, $A \cap B^c$, $A^c \cap B$ and $A^c \cap B^c$ are all non-empty. On \mathcal{P}_X^n , consider the function $F : (P_1, \dots, P_n) \mapsto T \circ P_1$, where

$$T(p) := \begin{cases} 1 & \text{if } p = 1 \\ 0 & \text{if } p = 0 \\ \frac{1}{2} & \text{if } p \in (0, 1). \end{cases}$$

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 let $P := F(P_1, \dots, P_n) = T \circ P_1$. 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 $\sigma(X)$ which assigns 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, assume $P_1(A) = P_1(B) = 1$, 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 that P extends (in fact, uniquely) to a probability function on $\sigma(X)$, namely the function $P_{A \cap B}$.

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 $P_1(B), P_1(B^c) \in (0, 1)$. It follows that $P(A) = 1$, $P(A^c) = 0$ and $P(B) = P(B^c) = \frac{1}{2}$, so that P extends (again, uniquely) to a probability function on $\sigma(X)$, namely the function $\frac{1}{2}P_{A \cap B} + \frac{1}{2}P_{A \cap B^c}$.

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 (this time non-uniquely) to a probability function on $\sigma(X)$, for instance to the function $\frac{1}{2}P_{A \cap B} + \frac{1}{2}P_{A^c \cap B^c}$ or the function $\frac{1}{4}P_{A \cap B} + \frac{1}{4}P_{A^c \cap B} + \frac{1}{4}P_{A \cap B^c} + \frac{1}{4}P_{A^c \cap B^c}$.

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 $P_i^*(C) = 1$. (What only matters for us is 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 one can easily check, it follows that $C \supseteq (A \cap B) \cup (A^c \cap B^c)$ 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.6 Proof of Theorem 5

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 $|Y| \geq 3$. So, there are pairwise distinct $A, B, C \in Y$. Define

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

As $\sigma(X)$ is closed under countable intersections, $A^*, B^*, C^* \in \sigma(X)$. For all 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

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

By construction,

$$P_i^*(A^* \cup B^* \cup C^*) = x_i + y_i + z_i = 1 \text{ for all } i. \quad (10)$$

Now consider the restriction $P_i := P_i^*|_X$ for each individual i . For the so-defined profile $(P_1, \dots, P_n) \in \mathcal{P}_X^n$, we consider the collective opinion function $P := P_{P_1, \dots, P_n}$. We now derive four properties of P (Claims 1-4), which then allow us to show that $D(x) + D(y) + D(z) = 1$ (Claim 5), as desired.

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

For all $E \in Y \setminus \{A, B, C\}$, we have $E \supseteq A^* \cup B^* \cup C^*$, so that by (10) we have $P_1(E) = \dots = P_n(E) = 1$, 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 $E \in Y \setminus \{A, B, C\}$, and thus $P^*(\cap_{E \in Y \setminus \{A, B, C\}} E) = 1$, since the intersection of countably many events of probability one has 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 .

Consider any extension $P^* \in \mathcal{P}_{\sigma(X)}$ of P . As $A \cap B \cap C$ is disjoint from the event $\cap_{E \in Y \setminus \{A, B, C\}} E$, which by Claim 1 has P^* -probability one, $P^*(A \cap B \cap C) = 0$. This implies Claim 2 because $A^c \cup B^c \cup C^c = (A \cap B \cap C)^c$.

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 the event $A^* \cup B^* \cup C^*$ which has P_i^* -probability of one for each individual i by (10). So, $P_i^*(A^c \cap B^c) = 0$, i.e., $P_i^*(A^c \setminus B) = 0$, for all i . For analogous reasons, $P_i^*(A^c \setminus C) = 0$ and $P_i^*(B^c \setminus C) = 0$ for all i . Since, as just shown, each individual i 's opinion function P_i has an extension P_i^* which assigns zero probability to the events $A^c \setminus B$, $A^c \setminus C$ 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, there is a P^* -probability of zero that *at least two* of A^c, B^c, C^c hold. Further, there is a P^* -probability of one that *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 equality follows from the pairwise disjointness of the events A^*, B^*, C^* . Regarding the second equality, note that $A^* \cup B^* \cup C^*$ is the intersection of the events $\cap_{E \in Y \setminus \{A, B, C\}} E$ 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 $P^*(A^* \cup B^* \cup C^*) = 1$, as claimed.

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^* , we have

$$P^*(A^c) = P^*(A^*). \quad (11)$$

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 each individual i . So, $P^*(A^c) = D(x)$. This and (11) imply that $P^*(A^*) = D(x)$. By similar arguments, $P^*(B^*) = D(y)$ and $P^*(C^*) = D(z)$. So, Claim 5 follows from Claim 4. ■

Proof of Theorem 5. (a) This part is shown by first deducing neutrality from Theorem 2(a) (and the fact that non-simple agendas are non-nested), and then inferring linearity from Lemmas 6 and 5 (and from Lemma 15(b), which applies because conditional consensus compatibility implies consensus preservation).

(b) This part is reducible to the companion paper's Theorem 5(b). Consider a simple agenda, finite and not $\{\emptyset, \Omega\}$. By the companion paper's Theorem 5(b), there is a pooling function F^* for agenda $\Sigma := \sigma(X)$ which, on X , is independent and conditional-consensus-preserving but not neutral. By its independence on X , it induces a pooling function for (sub)agenda X (see Lemma 13), which is independent and conditional-consensus-compatible but not neutral (see Lemma 10). ■

A.7 Proof of Theorem 6

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 members $A, B, C \in Y$. For all individuals i , since $x_i + y_i + z_i = 1$ and since A, B and C are pairwise disjoint and non-empty, there exists a $P_i \in \mathcal{P}_X$ such that

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

We write P for the collective opinion function under this profile. Since Y is a countable partition of Ω , and since P extends to a (σ -additive) probability function, we have $\sum_{E \in Y} P(E) = 1$. Note that, 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,

$$P(A) + P(B) + P(C) = 1.$$

So, since

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

we have $D(x) + D(y) + D(z) = 1$. ■

Proof of Theorem 6. (a) This part follows by first deducing neutrality from Theorem 3(a), and then inferring linearity from Lemmas 7 and 5 (and Lemma 15(b)).

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

A.8 Proof of Propositions 1 and 2

Proof of Proposition 1. Consider an opinion pooling function for an agenda X . We first show part (b) and then part (a).

(b) We show this claim by proving 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 some finite set S of pairs (C, D) of events in X . This fact follows from the equivalence of conditional consensus compatibility to implication preservation (see Proposition 3) and the observation that, for any such set S , an opinion function is consistent with zero probability of each event $C \setminus D$ (with $(C, D) \in S$) *if and only if* it is consistent with zero probability of the event $\cup_{(C,D) \in S} (C \setminus D)$, i.e., with certainty of the event $(\cup_{(C,D) \in S} (C \setminus D))^c$.

(a) The claims made in this part about implicit consensus preservation and consensus compatibility have already been shown (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 expressible as $(\cup_{(C,D) \in S} (C \setminus D))^c$ for some finite set $S \subseteq X \times X$. Note that, for any $A \in X$, we may let S take the form $S = \{(A^c, A)\}$, in which case $(\cup_{(C,D) \in S} (C \setminus D))^c = (A^c \setminus A)^c = A$. So, conditional consensus compatibility implies consensus preservation and is equivalent to it if $X = \sigma(X)$. ■

Proof of Proposition 2. We assume that $|\Omega| \geq 4$. We can therefore consider a partition of Ω into four non-empty events and define X to be the agenda consisting of any union of *two* of these four events. (Note that $A \in X \Leftrightarrow A^c \in X$, so that X is indeed an agenda.) Nothing hinges on the size of the four events, so that we may assume without loss of generality that each of them is singleton, i.e., that $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and $X = \{A \subseteq \Omega : |A| = 2\}$.

Part 1. In this part of the proof, we show that X is path-connected and non-partitional. Non-partitionality is trivial. To show path-connectedness, we consider any events $A, B \in X$ and have to 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. We may then write $A = \{\omega_A, \omega\}$ and $B = \{\omega_B, \omega\}$ with $\omega_A, \omega_B, \omega$ pairwise distinct. We have $\{\omega_A, \omega\} \vdash^* \{\omega\}$ (take $Y = \{\{\omega, \omega'\}\}$, where ω' is the element of $\Omega \setminus \{\omega_A, \omega_B, \omega\}$) and $\{\omega\} \vdash^* \{\omega_B, \omega\}$ (take $Y = \emptyset$).

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. We have $\{\omega_A, \omega'_A\} \vdash^* \{\omega_A, \omega_B\}$ (take $Y = \{\{\omega_A, \omega'_B\}\}$) and $\{\omega_A, \omega_B\} \vdash^* \{\omega_B, \omega'_B\}$ (take $Y = \{\{\omega_B, \omega'_A\}\}$).

Part 2. In this part, we construct a pooling function $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ that is independent (in fact, neutral) and consensus-preserving, but not linear. As an ingredient to 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. For this purpose, we use a fixed 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 indeed exists; e.g. $T(x) = 4(x - 1/2)^3 + 1/2$ for all $x \in [0, 1]$.)

For any $P_1, \dots, P_n \in \mathcal{P}_X$, we define the collective probability of any agenda event $A \in X$ to be $P_{P_1, \dots, P_n}(A) := T(L(P_1, \dots, P_n)(A))$. We now prove that, for any $P_1, \dots, P_n \in \mathcal{P}_X$, the just-defined function P_{P_1, \dots, P_n} on X does indeed extend to a probability function on $\sigma(X) = 2^\Omega$. This completes the proof, since it shows that we have defined a proper pooling function, which is of course neutral (as L is neutral), consensus-preserving (as L is consensus-preserving and $T(1) = 1$), and non-linear (as L is linear and T a non-linear transformation).

Consider any probability function Q on 2^Ω (think of $Q|_X$ as $L(P_1, \dots, P_n)$ for some given $P_1, \dots, P_n \in \mathcal{P}_X$, and of Q as an extension to a probability function on 2^Ω). We show that the transformed function $T \circ Q|_X$ extends to a probability function on 2^Ω . To do so, it suffices to show 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 any $k \in \{1, 2, 3, 4\}$, we put $q^k := Q(\{\omega_k\})$; and for any $k, l \in \{1, 2, 3, 4\}$ with $k < l$, we put $q_{kl} := Q(\{\omega_k, \omega_l\})$. In order for the numbers 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 equalities 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 $\{k', l'\} = \{1, 2, 3, 4\} \setminus \{k, l\}$. Since $p_k + p_l = 1 - p_{k'} - p_{l'}$ and $T(q_{kl}) = T(1 - q_{k'l'}) = 1 - T(q_{k'l'})$, the equality $p_k + p_l = T(q_{kl})$ 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

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

We now solve this system of three linear equations in $(p_1, p_2, p_3) \in \mathbb{R}^3$. Write $t_{kl} := T(q_{kl})$ 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:

$$p_3 = \frac{t_{23} + t_{13} - t_{12}}{2} \tag{12}$$

$$p_2 = t_{12} - t_{13} + \frac{t_{23} + t_{13} - t_{12}}{2} = \frac{t_{12} + t_{23} - t_{13}}{2} \tag{13}$$

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

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

$$p_4 = 1 - \frac{t_{12} + t_{13} + t_{23}}{2}. \tag{14}$$

By their construction, the numbers p_1, \dots, p_4 given by (12)-(14) 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 $q_{13} - \delta, q_{23} + \delta \in [1/2, 1]$, the convexity of T on $[1/2, 1]$ implies that

$$\begin{aligned} T(q_{13}) &\leq \frac{1}{2} [T(q_{13} - \delta) + T(q_{23} + \delta)] \\ \text{and } T(q_{23}) &\leq \frac{1}{2} [T(q_{13} - \delta) + T(q_{23} + \delta)], \end{aligned}$$

so that (by adding these two inequalities)

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

This inequality may be applied to $\delta = 1 - q_{23}$, since

$$q_{13} - (1 - q_{23}) = (q_{13} + q_{23} + q_{12}) - q_{12} - 1 \leq 2 - q_{12} - 1 = 1 - q_{12} \in [1/2, 1];$$

which implies that

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

On the right hand side of this inequality, we have $T(1) = 1$ and, by $q_{13} - (1 + q_{23}) \leq 1 - q_{12}$ and T 's increasingness, $T(q_{13} - (1 + q_{23})) \leq T(1 - q_{12}) = 1 - T(q_{12})$. So, we obtain $T(q_{13}) + T(q_{23}) \leq 1 + 1 - T(q_{12})$, i.e., $T(q_{12}) + T(q_{13}) + T(q_{23}) \leq 2$, as desired.

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 an increasing function, it suffices to establish $T(q^1) + T(q^2) \geq T(q^1 + q^2)$. Again, we consider three cases.

Case 1: $q^1 + q^2 \leq 1/2$. Suppose $q^1 \leq q^2$ (otherwise the roles of q^1 and q^2 get swapped). For all $\delta \geq 0$ such that $q^1 - \delta, q^2 + \delta \in [0, 1/2]$, the concavity of T on $[0, 1/2]$ implies that

$$\begin{aligned} T(q^1) &\geq \frac{1}{2} [T(q^1 - \delta) + T(q^2 + \delta)] \\ \text{and } T(q^2) &\geq \frac{1}{2} [T(q^1 - \delta) + T(q^2 + \delta)], \end{aligned}$$

so that (by adding these inequalities)

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

Applying this to the case $\delta = q^1$ yields $T(q^1) + T(q^2) \geq T(0) + T(q^2 + q^1) = T(q^1 + q^2)$, as desired.

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

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

as desired.

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 $q^1 < 1/2$, as otherwise $q^1 + q^2 > 1$. Define $y := 1 - q^1 - q^2$. Since we also have $y < 1/2$, 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 – which it must do since Theorem 5 tells us that independent conditional-consensus-compatible pooling functions must be linear (for non-simple, hence in particular for path-connected agendas). Let Ω and X be as in the proof, and consider a profile with complete unanimity: all individuals i give ω_1 probability 0, each of ω_2, ω_3 probability $1/4$, and hence ω_4 probability $1/2$. 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 the collective probability of ω_1 to be 0 as well. But the collective probability of ω_1 is (in the notation of the proof) given by

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

where q_{kl} is the collective probability of $\{\omega_k, \omega_l\}$ under a linear pooling function, so that q_{kl} equals the unanimous individual probability of $\{\omega_k, \omega_l\}$. 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 $T(0) = 0$.