

Franz Dietrich and [Christian List](#) Opinion pooling on general agendas

Working paper

Original citation:

Dietrich, Franz and List, Christian (2008) Opinion pooling on general agendas. Working Paper. London School of Economics, Government Department, London, UK.

This version available at: <http://eprints.lse.ac.uk/20127/>

Available in LSE Research Online: July 2008

© 2008 Franz Dietrich and Christian List.

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Opinion pooling on general agendas

Franz Dietrich and Christian List*

September 2007, revised May 2008

Abstract

How can different individuals' probability assignments to some events be aggregated into a collective probability assignment? Although there are several classic results on this problem, they all assume that the 'agenda' of relevant events forms a σ -algebra, an overly demanding assumption for many practical applications. We drop this assumption and explore probabilistic opinion pooling on general agendas. Our main theorems characterize linear pooling and neutral pooling for large classes of agendas, with standard results as special cases.

Keywords: opinion pooling, subjective probability, probability theory

1 Introduction

This paper addresses the classic problem of *probabilistic opinion pooling*: Different individuals' probability assignments to some events are to be aggregated into a collective probability assignment, while preserving probabilistic coherence. Although opinion pooling has been explored in some depth by statisticians, economists, and philosophers, almost all contributions so far assume that the set of relevant events – the *agenda* – forms a σ -algebra. Thus the agenda must be closed under both complement and countable union. This assumption is technically motivated by the fact that the standard coherence conditions on probability assignments, particularly σ -additivity, use a σ -algebra as their domain. But in realistic applications, many events in the σ -algebra, such as artificial conjunctions or disjunctions, may be irrelevant. Although it may be relevant whether it rains, and whether the interest rate goes up, the conjunction or disjunction of these events may be irrelevant. Irrelevance can mean one

*Although both authors are jointly responsible for this paper and project, Christian List wishes to note that Franz Dietrich should be considered the primary author, to whom the credit for the present mathematical proofs is due. Addresses: F. Dietrich, Department of Quantitative Economics, Maastricht University, and CPNSS, London School of Economics; C. List, Department of Government, London School of Economics.

of two things. On one interpretation, an event is irrelevant if the collective is not required to specify a probability for that event. Expert committees, for example, are hardly ever asked to provide a probability assignment to all events in a σ -algebra; rather, they are asked to specify the probabilities of a limited set of events of interest. On another, very different interpretation, an event is irrelevant if the conditions imposed on an opinion pooling function are not considered compelling when applied to that event. Strong conditions such as event-wise independence (‘the collective probability of an event must depend only on the individual probabilities of that event’) and zero-preservation (‘if all individuals assign a probability of zero to an event, the collective must also assign a probability of zero to it’) may be plausible when restricted to some events – perhaps those deemed ‘basic’ – but not when applied to others, such as artificial conjunctions or disjunctions.

We investigate opinion pooling without assuming that the agenda forms a σ -algebra, and consider instead general agendas. The notion of an agenda employed in this paper is borrowed from the theory of judgment aggregation (e.g., List and Pettit 2002, Dietrich 2006, Nehring and Puppe 2007, Dokow and Holzman forthcoming, Dietrich and List 2007). Our two main results are the following. For a large class of agendas – so-called *non-simple* ones, of which σ -algebras are only a very special case – any opinion pooling function satisfying two conditions must be *linear*, i.e., the collective probability of each event in the agenda must be a weighted linear average of the individuals’ probabilities of that event, where the weights are the same for all events. Second, for an even larger class of agendas – so-called *non-nested* ones, which include all non-simple agendas – the same two conditions lead to a *neutral* pooling function: here the collective probability of each event is some function of the individuals’ probabilities of that event, where the function is the same for all events. Each of these results uses a logically minimal agenda condition, i.e., for agendas violating non-simplicity, there are counterexamples to the first result and, for agendas violating non-nestedness, there are counterexamples to the second. Thus our results are the first agenda-characterization results in the theory of opinion pooling, in analogy with the recent agenda-characterization results in the theory of judgment aggregation (Nehring and Puppe 2007, Dokow and Holzman forthcoming, Dietrich and List 2007). Just as it has been an important step in the theory of judgment aggregation to characterize the class of agendas for which certain axiomatic characterizations of dictatorships hold, so the present results seek to address the parallel challenge in the theory of opinion pooling: namely to characterize the class of agendas for which certain axiomatic characterizations of salient opinion pooling functions hold – here, linear and neutral ones.

The two agenda conditions in our present results are surprisingly undemanding. As soon as the agenda contains two or more logically independent events, for example, it is non-nested; and if, in addition, it contains their conjunction or disjunction (or some other event logically connecting them), it is non-simple.

Logical independence means that any combination of truth-values across the events is consistent. In particular, agendas consisting only of logically independent events (and their complements) are already non-nested. Such agendas are common in many realistic decision problems, where there are often only probabilistic dependencies (i.e., correlations) between events but no logical ones, as in the case of rainfall and an interest-rate increase. Yet these agendas are fundamentally different from ones forming a σ -algebra. By focusing on σ -algebras alone, the standard treatments of opinion pooling have therefore been restricted in generality as well as excluded many realistic applications.

Our results not only go significantly beyond the standard results on opinion pooling in the literature and imply them as special cases, but we also present a new illustrative application of opinion pooling: the case of probabilistic preference aggregation. Here each individual submits a probability distribution over all possible linear orderings over some alternatives, interpretable, for example, as representing the individual's degrees of belief about which ordering is the 'correct' one (e.g., a quality ordering of candidates). The collective must then determine a single such probability distribution, interpretable as the corresponding collective degrees of belief. (Below we also suggest an alternative interpretation in terms of vague preferences.) We show that our characterization of linear pooling applies to this case when there are three or more alternatives. As explained below, this illustrative application of our result is different from existing results on probabilistic preference aggregation.

For the classical case in which the agenda is a σ -algebra, linearity and neutrality are among the most widely studied properties of opinion pooling functions (in the case of neutrality sometimes under names such as *strong label neutrality* or *strong setwise function property*). Linear pooling goes back to Stone (1961) or even Laplace, and neutral pooling to McConway (1981) and Wagner (1982). McConway's classic characterization of linear pooling is a special case of one of our results. As discussed in more detail below, the σ -algebra case has the interesting feature that every neutral pooling function is automatically linear, so that neutrality and linearity are equivalent here, if the σ -algebra contains more than four events (McConway 1981 and Wagner 1982; see also Mongin's 1995 linearity characterization). This peculiarity does not carry over to general agendas: some agendas allow for neutral yet non-linear opinion pooling. The reader is referred to Genest and Zidek's (1986) overview article for an excellent review of classical results on opinion pooling. For a recent 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).

To provide a pedagogical exposition of our proofs, we sketch the structure of two of our central proofs in the main text while presenting the underlying technical arguments in the appendix.

2 Model

Events and probabilities. Let Ω be a non-empty set of possible *worlds* (or *states*) and Σ a σ -algebra of *events*, i.e., subsets $A \subseteq \Omega$.¹ For example, Σ could be the power set of Ω . The *complement* of any event A is denoted $A^c := \Omega \setminus A$ and represents the negation of A . An event A is *contingent* if it is neither \emptyset (impossible) nor Ω (certain). The intersection $A \cap B$ and union $A \cup B$ of two events A and B represent their conjunction and disjunction, respectively. A set of events S is *consistent* if the intersection of the events in S is non-empty (i.e., $\bigcap_{A \in S} A \neq \emptyset$), and *inconsistent* otherwise; S *entails* another event B if that intersection is included in B (i.e., $\bigcap_{A \in S} A \subseteq B$). A *probability measure* over Σ is a function $P : \Sigma \rightarrow [0, 1]$, with the standard properties (i.e., σ -additivity and $P(\Omega) = 1$). Let \mathcal{P} be the set of all probability measures over Σ .

Opinion pooling functions. Consider a group of $n \geq 2$ individuals, labelled $i = 1, \dots, n$. An assignment of probability measures to the n individuals, (P_1, \dots, P_n) , is called a *profile*. An *opinion pooling function* is a function $F : \mathcal{P}^n \rightarrow \mathcal{P}$, which assigns to each profile (P_1, \dots, P_n) of individual probability measures a collective one $P = F(P_1, \dots, P_n)$, in short P_{P_1, \dots, P_n} . For example, P_{P_1, \dots, P_n} could be the arithmetic average $\frac{1}{n}P_1 + \dots + \frac{1}{n}P_n$, a case of *linear pooling*, as defined below. There are numerous other possible pooling functions, including geometric averages (of a weighted or non-weighted kind)² and expert rules (where $P_{P_1, \dots, P_n} := P_i$ with a fixed or profile-dependent ‘expert’ i).

Agenda and relevance. Unlike previous works on opinion pooling, we assume that only some events in Σ are *relevant*. As discussed in detail below, the relevant events can be interpreted in (at least) two ways: either as the events for which the group actually requires probabilities, or as those for which conditions imposed on the pooling function (independence, zero-preservation, etc., as stated below) are compelling. We call the set of relevant events the *agenda*, in analogy to the equally named notion in the theory of judgment aggregation. Formally, an *agenda* is a subset $X \subseteq \Sigma$ which is non-empty and closed under complement (i.e., if $A \in X$ then $A^c \in X$). Apart from non-emptiness and closure under complement, there are no other formal conditions on an agenda. Crucially, we allow an agenda to contain A and B without containing $A \cup B$ (or $A \cap B$): it may be relevant whether it rains, and whether the interest rate goes up, yet irrelevant whether it rains *or* the interest rate goes up. In the ‘classical’ case $X = \Sigma$, all events are treated as relevant and the conditions imposed on pooling functions below reduce to the standard conditions in the literature.

¹A σ -algebra over Ω is a non-empty set of subsets of Ω that is closed under both complement and countable union.

²Such pooling functions are defined on a suitable subdomain of \mathcal{P}^n .

An example of an agenda containing no conjunctions or disjunctions.

Suppose each possible world is specified by a vector of three binary characteristics. The first takes the value 1 if CO₂ emissions are above some critical threshold, and 0 otherwise. The second takes the value 1 if there is a mechanism whereby if CO₂ emissions are 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 imagine an expert committee faced with an opinion pooling problem on the 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,³ and B the event of a positive third characteristic. Although there are non-trivial overlaps and even logical connections between the events in this agenda (note that A and $A \rightarrow B$ are inconsistent with B^c), it contains no conjunctions or disjunctions and is therefore far removed from the σ -algebras to which the standard results in the literature apply. We come back to this agenda at various points below.

Two interpretations of the agenda. (1) We may interpret the agenda as the set of those events for which the group requires probabilities. As mentioned in the introduction, real-world expert committees, which give advice on complex factual questions, are hardly ever asked to assign explicit probabilities to all events in a σ -algebra; instead, they are usually required to come up with a partial probability assignment $P_X : X \rightarrow [0, 1]$ which is *coherent* in the sense of being extendible (though not necessarily uniquely) to a probability measure $P : \Sigma \rightarrow [0, 1]$. The expert committee of our global-warming example, for instance, may be asked to provide probabilities only for $A, A \rightarrow B$ and B (and their complements), but not for the entire underlying σ -algebra. The resulting opinion pooling problem is thus best captured by a *restricted opinion pooling function* $F_X : (\mathcal{P}_X)^n \rightarrow \mathcal{P}_X$, where \mathcal{P}_X is the set of all coherent partial probability assignments $P_X : X \rightarrow [0, 1]$, in analogy with a judgment aggregation function, which maps profiles of individual acceptance/rejection judgments on X into collective such judgments (List and Pettit 2002). Any opinion pooling function $F : \mathcal{P}^n \rightarrow \mathcal{P}$ satisfying independence for all events in X induces a well-defined restricted counterpart $F_X : (\mathcal{P}_X)^n \rightarrow \mathcal{P}_X$.

³Here the conditional \rightarrow in $A \rightarrow B$ is interpreted as a subjunctive conditional rather than a material one, since 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 interpret \rightarrow as a material conditional, we would have to rule out in addition the truth-value assignments $(0, 0, 1)$, $(1, 0, 1)$ and $(0, 0, 0)$. Under the material interpretation of \rightarrow , the event $A \rightarrow B$ would become $A^c \cup B$, and the agenda would no longer be free from conjunctions or disjunctions. However, the agenda would still not be a σ -algebra, and it would retain all the other properties discussed below (non-nestedness and non-simplicity).

(2) Under a second interpretation, the group is interested in probabilities of events both inside and outside the agenda (and such probabilities must therefore also be submitted by the individuals), but the conditions imposed on a pooling function – such as independence and zero-preservation – are deemed compelling only for events in the agenda. So here the agenda is defined by the scope we give to those conditions. For example, a committee may be interested in the probability of the conjunction $A \cap (A \rightarrow B)$ (e.g., for making some policy recommendation), but it may not wish to assign a probability to this proposition solely on the basis of the individual probability assignments to it. Rather, the committee may first wish to determine its probabilities for each of A and $A \rightarrow B$ and then derive the probability of the conjunction via some auxiliary assumptions (e.g., about the probabilistic dependence or independence of the conjuncts). Since conditions such as independence are meant to apply to A and $A \rightarrow B$ here but not their conjunction, the set X would be defined to include the former propositions (and their complements) but not the latter.

How realistic are the two interpretations? Typically, at least one of our two interpretations of the agenda applies. A paradigmatic class of opinion pooling problems in which the first interpretation applies (and arguably also the second) is the one in which events are specified by statements in natural language, such as ‘it will rain’ or ‘the new legislation X will be rejected by the constitutional court’. In this case, the number of events in Σ is often enormous but finite: k atomic statements lead to 2^k Boolean combinations (worlds in Ω), a number that grows exponentially in k . Unless k is small, individuals in the real world can hardly be asked to submit all 2^k probabilities, and even if they can, the group may not be interested in the probabilities of worlds that represent artificial conjunctions of unrelated statements. Thus the first interpretation is likely to apply here. The second interpretation may apply too. As argued in similar contexts in the theory of judgment aggregation, it is often implausible to form a collective belief on some conjunctive or disjunctive event solely on the basis of individual beliefs on that event, without giving attention to individual beliefs on, for example, the underlying conjuncts or disjuncts.

By contrast, if the problem is that of estimating the distribution of a real-valued or vector-valued random variable, such as rainfall or the number of insurance claims, the collective may well require a full probability measure, and it may also be realistic to assume that individuals can come up with a full probability measure themselves: they first choose some parametric class of probability measures (e.g., that of Gaussian or Poisson or binomial ones) and then estimate the relevant parameters (for example, there are only two such parameters in the case of a unidimensional Gaussian distribution). Here the first interpretation of the agenda clearly does not apply, but the second often does: it would seem highly artificial to apply a local pooling criterion on events in Σ such as the set of real numbers whose distance to the nearest prime exceeds 37, without giving

attention to other logically related events.⁴ The agenda would then include only those events on which the application of a local pooling criterion is meaningful. Events whose collective probability is determined in some derivative way would lie outside the agenda.

Moreover, if all individuals hold a probability measure from the same parametric class (say a Gaussian one), then an opinion pooling function satisfying independence on the entire σ -algebra, such as linear pooling, does not generally produce a collective probability measure from the same class. If we require the opinion pooling function to preserve the original parametric class, we therefore cannot apply conditions such as independence to the entire σ -algebra. Notice that such problems differ significantly from opinion pooling problems involving statements in natural language: Ω is often a (typically infinite) subset of \mathbf{R} or of a higher-dimensional Euclidean space \mathbf{R}^k , rather than a (finite) set of maximally consistent statements over which parametric families of probability measures are not readily available.

3 Two conditions on opinion pooling

Our characterization results use two conditions. The first, *independence*, requires that the collective probability of each relevant event depend only on the individual probabilities of that event. In the classical case in which every event is deemed relevant, this coincides with the equally named standard condition in the literature (sometimes also called *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 $P_1, \dots, P_n \in \mathcal{P}$, $P_{P_1, \dots, P_n}(A) = D_A(P_1(A), \dots, P_n(A))$.

The second condition, *conditional zero-preservation*, requires that if all individuals assign a conditional probability of zero to some relevant event given another, then a zero conditional probability should also be assigned collectively.⁵ In the classical case – more generally, whenever the agenda exhibits certain closure properties – this becomes equivalent to the standard condition of *zero-preservation*, as discussed below.

Conditional zero-preservation. For all events $A, B \in X$ and $P_1, \dots, P_n \in \mathcal{P}$, if $P_i(A|B) = 0$ for all individuals i then $P_{P_1, \dots, P_n}(A|B) = 0$, with the stipulation that, for any $P \in \mathcal{P}$, $P(A|B) := 0$ when $P(B) = 0$.⁶

⁴Rather, one might want to determine some salient parameters of the collective probability measure P_{P_1, \dots, P_n} (e.g., its expectation $E(P_{P_1, \dots, P_n}) \in \mathbf{R}_+ \cup \{\infty\}$ if $\Omega = \mathbf{R}_+$) as a function of the same parameters of the individual measures P_1, \dots, P_n .

⁵We are indebted to Richard Bradley for suggesting this formulation of the condition.

⁶This standard convention implies that, when $P(B) = 0$, $P(\bullet|B)$ does not define a proba-

How can these two conditions be motivated? The main normative defence of independence is the democratic idea that the collective view on any issue should be determined by individual views on that issue. This reflects a *local*, as opposed to *holistic*, notion of democracy; under a *holistic* notion, the collective view on an issue may also be influenced by individual views on other related issues. Such a defence of independence is compelling only if the agenda does not contain ‘artificial’ events, such as conjunctions of intuitively unrelated events. For this reason, a democratic defence of independence is difficult in the classical case where the agenda is the entire σ -algebra. This might explain why a democratic defence of independence has not, to our knowledge, been put forward in the standard literature on opinion pooling (although similar arguments for independence are common in other fields of aggregation theory).

Apart from a democratic defence, two pragmatic arguments for independence can be given, which apply regardless of how large or small the agenda is. First, determining the collective view on any issue solely on the basis of the individual views on that issue is informationally less demanding than a holistic approach and thus easier in practice. Second, independence prevents certain types of agenda manipulation.⁷

However, an objection against independence in the classical case is its incompatibility with collectively preserving unanimous beliefs of probabilistic independence (see Genest and Wagner 1984; Bradley, Dietrich and List 2006).⁸ Whether the objection applies to our independence condition depends on the precise nature of the agenda. Finally, some authors reject independence – in the classical case $\Sigma = X$ and presumably also in our general case – as they prefer to require *external Bayesianity*, whereby aggregation should commute with Bayesian updating of probabilities in the light of new information.

The idea underlying conditional zero-preservation is intuitive: if all individuals believe that the probability of some relevant event given another is zero, e.g., that the probability of no economic impact given war is zero, or that the probability of famine given democracy is zero (on a strong version of Amartya Sen’s thesis), then this belief should be preserved collectively. Conditional zero-preservation implies the following standard condition, its unconditional counterpart:⁹

bility measure. The convention can be avoided by using an alternative formulation of *conditional zero-preservation*: For all events $A, B \in X$ and $P_1, \dots, P_n \in \mathcal{P}$, if every individual i has $P_i(A|B) = 0$ provided $P_i(B) \neq 0$, then also $P_{P_1, \dots, P_n}(A|B) = 0$ provided $P_{P_1, \dots, P_n}(B) \neq 0$.

⁷In the classical case $X = \Sigma$, 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 $\Sigma^* \subseteq \Sigma$. A similar result holds for general agendas X . Thus independence prevents agenda setters from influencing the collective probability assignment to some events by adding or removing other events to or from the agenda.

⁸Assuming the aggregation function is *non-dictatorial*, i.e., the collective does not always adopt the probability function of a fixed individual.

⁹In the special case in which $\Omega \in X$, the implication is immediately obvious, since we can

Zero-preservation. For all $A \in X$ and all $P_1, \dots, P_n \in \mathcal{P}$, if $P_i(A) = 0$ for all individuals i then $P_{P_1, \dots, P_n}(A) = 0$.

In fact, as shown in the appendix, if the agenda is closed under pairwise intersection or union – e.g., if it is the entire σ -algebra – conditional zero-preservation is equivalent to zero-preservation; in the general case, conditional zero-preservation is stronger.

4 Characterization of neutral pooling

What is the class of opinion pooling functions satisfying the two conditions we have introduced? The answer to this question depends on the nature of the agenda. In this section, we show that, if (and only if) the agenda is what we call *non-nested*, our two conditions characterize the class of *neutral* opinion pooling functions. In the next section, we show that, if (and only if) the agenda is what we call *non-simple*, a condition that strengthens non-nestedness, our two conditions characterize the class of *linear* opinion pooling functions. All formal definitions are given in what follows.

Let us begin with our first condition on the agenda. 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 \subseteq , and *non-nested* otherwise. For example, binary agendas $X = \{A, A^c\}$ are nested: take $X_+ := \{A\}$, which is trivially linearly ordered by set-inclusion. Also, the agenda $X = \{(-\infty, t], (t, \infty) : t \in \mathbf{R}\}$ (a subset of the Borel- σ -algebra Σ over the real line $\Omega = \mathbf{R}$) is nested: take $X^+ := \{(-\infty, t] : t \in \mathbf{R}\}$, 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 (i.e., 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.

To see that non-nestedness is a very weak condition, notice that nested agendas have the very special property that all $A, B \in X$ are *logically dependent* (i.e., one of A, A^c entails one of B, B^c). They may thus also be described as ‘pairwise connected’ or ‘trivial’, implying that any non-trivial agenda is non-nested.

We are now in a position to state our first main result. Call a pooling function F *neutral* if there exists a single function $D : [0, 1]^n \rightarrow [0, 1]$ such that,

conditionalize on $B = \Omega$. The general argument is given in the appendix.

for every event $A \in X$ and every profile $(P_1, \dots, P_n) \in \mathcal{P}^n$,

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

i.e., F is independent with the *same* local pooling criterion $D = D_A$ for all events $A \in X$, thereby treating all relevant events perfectly symmetrically. In the classical case where the agenda is the entire σ -algebra, this condition becomes the one studied in the literature.

Theorem 1 (a) *For a non-nested agenda X , a conditional zero-preserving pooling function is independent if and only if it is neutral.*
 (b) *For a nested agenda X , finite and not $\{\emptyset, \Omega\}$, there exists a conditional zero-preserving pooling function that is independent but not neutral.*

Part (b) shows that the agenda condition of the characterization result in part (a) is essentially tight; an aggregation function with the required properties is explicitly constructed in the appendix. However, although nested agendas X thus permit non-neutral pooling functions, only a limited kind of non-neutrality is possible: As the proof below indicates, the pooling criterion D_A 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 full neutrality follows even in the nested case if independence is slightly strengthened by requiring that $D_A = D_{A^c}$ for some $A \in X \setminus \{\emptyset, \Omega\}$.

Let us outline the proof of part (a); the full technical details are provided in the appendix. We begin by defining a binary relation \sim on the set of all contingent events in the agenda. As a preliminary definition, call two events $A, B \in X$ *exclusive* if $A \cap B = \emptyset$ and *exhaustive* if $A \cup B = \Omega$. Now, for any $A, B \in X \setminus \{\emptyset, \Omega\}$, we define

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

Our proof proceeds via three lemmas.

Lemma 1 *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 $A, A^c \in X \setminus \{\emptyset, \Omega\}$.

Lemma 2 *An agenda $X \neq \{\emptyset, \Omega\}$ is nested if and only if there are exactly two equivalence classes with respect to \sim , and non-nested if and only if there is exactly one.*

Call a pooling function *neutral on a set* $Z \subseteq X$ if it is independent with the same pooling criterion for all events in Z .

Lemma 3 *An independent conditional zero-preserving pooling function is neutral on each equivalence class with respect to \sim .*

Part (a) of theorem 1 now follows immediately. As already mentioned, the proof of part (b) consists in explicitly constructing a non-neutral but independent and conditional zero-preserving pooling function on any nested agenda. To convey the idea informally, recall that a nested agenda can be partitioned into two non-empty sequences of events that are nested by inclusion. The opinion pooling function we construct has the property that (i) all events in one of the two nested sequences in the agenda have the same pooling criterion $D : [0, 1]^n \rightarrow [0, 1]$, defined, for example, as the square of a linear pooling criterion, and (ii) all events in the complementary nested sequence have the same pooling criterion, defined as $D^*(x_1, \dots, x_n) = 1 - D(1 - x_1, \dots, 1 - x_n)$ for all $(x_1, \dots, x_n) \in [0, 1]^n$. Since the construction and verification of the relevant properties involve some technicality, we give the full proof in the appendix.

5 Characterization of linear pooling

Let us now introduce a stronger, yet still surprisingly undemanding condition on the agenda. One preliminary definition is needed. Call a set of events $Y \subseteq \Sigma$ *minimal inconsistent* if it is inconsistent but every proper subset $Y' \subsetneq Y$ is consistent. Examples of minimal inconsistent sets are $\{A, B, (A \cap B)^c\}$, where A and B are logically independent events, and $\{A, A \rightarrow B, B^c\}$, with $A, B, A \rightarrow B$ 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, 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,

¹⁰The countability condition can often be dropped because all minimal inconsistent sets $Y \subseteq X$ are automatically countable or even finite. This is so if X is finite or countably infinite, and also in the (frequent) case that the events in X represent sentences in a language: then, provided the 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, the σ -algebra Σ often contains events not representing a sentence, so that the (unnatural) agenda $X = \Sigma$ often has infinite minimal inconsistent subsets.

and *simple* otherwise.¹¹ It follows immediately that the agenda consisting of $A, A \rightarrow B, B$ and complements in our expert committee example is non-simple. Moreover, non-simplicity implies non-nestedness, but not vice-versa.¹²

We can now state our second main result. Call a pooling function *linear* if there exist ‘weights’ $w_1, \dots, w_n \geq 0$ with sum 1 such that, for every event $A \in X$ and every profile $(P_1, \dots, P_n) \in \mathcal{P}^n$,

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

or in short, $P_{P_1, \dots, P_n}|_X = \sum_{i=1}^n w_i P_i|_X$. As before, this reduces to the standard definition in the classical case.¹³ Linearity is obviously a special case of neutrality, where the pooling criterion D takes a linear form.

Theorem 2 (a) *For a non-simple agenda X , a conditional zero-preserving pooling function is independent if and only if it is linear.*
 (b) *For a simple agenda X , finite and not $\{\emptyset, \Omega\}$, there exists a conditional zero-preserving pooling function that is independent but not linear.*

As before, part (b) shows that the agenda condition of the characterization result in part (a) is essentially tight. For the proof of part (b), which is of significant difficulty, the reader is referred to the appendix. Again, let us outline the proof of part (a), with details in the appendix.

Consider a non-simple agenda X and a conditional zero-preserving pooling function F . Let us write $\mathbf{0}$ and $\mathbf{1}$ to denote the n -tuples $(0, \dots, 0)$ and $(1, \dots, 1)$, respectively. Obviously, if F is linear then F is independent. So let us suppose F is independent and show linearity. By theorem 1(a) and since non-simple agendas are non-nested, F is neutral, say with pooling criterion $D : [0, 1]^n \rightarrow [0, 1]$ for all events $A \in X$. The proof now consists of three lemmas. The first establishes some simple properties of D ; the second contains the bulk of the work; and the third is an application of Cauchy’s functional equation (similar to one made by McConway 1981).

Lemma 4 (a) $D(x) + D(\mathbf{1} - x) = 1$ for all $x \in [0, 1]^n$.

¹¹Non-simplicity is a variant of a standard condition in the theory of judgment aggregation (Nehring and Puppe 2007, Dietrich and List 2007).

¹²To give an example of a non-nested but simple agenda, let $\Omega = \{x, y, z, w\}$ and $\Sigma =$ power set of Ω . Now define $X = \{\{x, y\}, \{z, w\}, \{w, x\}, \{y, z\}\}$. Clearly, this agenda cannot be expressed as $X = \{A, A^c : A \in X_+\}$ for some set X_+ linearly ordered by set-inclusion \subseteq , but its largest minimal inconsistent subsets, i.e., $\{\{x, y\}, \{z, w\}\}$ and $\{\{w, x\}, \{y, z\}\}$, each contain only two events.

¹³Also, if the agenda X is such that every probability measure $P \in \mathcal{P}$ is uniquely determined by the probabilities of relevant events, our X -relativized linearity notion is equivalent to the standard global linearity notion (because then $P_{P_1, \dots, P_n}|_X = \sum_{i=1}^n w_i P_i|_X$ implies $P_{P_1, \dots, P_n} = \sum_{i=1}^n w_i P_i$, for all $(P_1, \dots, P_n) \in \mathcal{P}^n$).

(b) $D(\mathbf{0}) = 0$ and $D(\mathbf{1}) = 1$.

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

Lemma 6 If $D(x) + D(y) + D(z) = 1$ for all $x, y, z \in [0, 1]^n$ with $x + y + z = \mathbf{1}$, then there exist non-negative weights w_i with sum 1 such that, for all $x \in [0, 1]^n$, $D(x_1, \dots, x_n) = \sum_{i=1}^n w_i x_i$.

Part (a) of theorem 2 now follows immediately. In the proof of part (b), we explicitly construct a non-linear but independent and conditional zero-preserving pooling function on any simple agenda. The constructed pooling function has the property that its pooling criterion D takes only three values, namely 0, 1/2 and 1.

It is instructive to see how our present results generalize the standard results in the literature, where the agenda is a σ -algebra (in particular, McConway 1981). Notice the following simple fact:

Lemma 7 When X is closed under pairwise intersection or union (e.g., $X = \Sigma$), non-nestedness and non-simplicity are equivalent, and X satisfies both if and only if $|X| > 4$.

Given the equivalence of conditional zero-preservation and zero-preservation on such an agenda, theorems 1(a) and 2(a) thus yield the following corollary (compare the results cited in the introduction):

Corollary 1 Let X be closed under pairwise intersection or union (e.g., $X = \Sigma$).

- (a) If $|X| > 4$, a zero-preserving pooling function is independent if and only if it is linear (and if and only if it is neutral).
- (b) If $|X| \leq 4$, but $X \neq \{\emptyset, \Omega\}$, there exists a zero-preserving pooling function that is independent but neither linear nor neutral.

6 An illustrative application: probabilistic preference aggregation

To illustrate the generality of our results, let us finally show how they apply in a rather different context, namely that of *probabilistic preference aggregation*. A group faces a (finite non-empty) set K of (mutually exclusive and exhaustive) alternatives, and seeks to form a collective view on how to rank them in a linear order. Let Ω_K be the set of all strict orderings \succ over K (asymmetric, transitive and connected binary relations), and let Σ_K be the power set of Ω_K . Informally, K can represent any set of distinct objects, e.g., policy options,

candidates, social states, distributions of goods, or artifacts, 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.).

A probability measure $P : \Sigma_K \rightarrow [0, 1]$ can now be interpreted as capturing an agent’s *degrees of belief* about which of the orderings \succ in Ω_K is the ‘correct’ one; call this the *belief interpretation*. In particular, for any pair of alternatives $x, y \in K$, $P(x \succ y)$ can be interpreted as the agent’s degree of belief in the event $x \succ y$, defined as the subset of Ω_K consisting of all those orderings \succ in which x is ranked above y . (On an entirely 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 probability measures to a single collective one.

To specify the problem further, define the *preference agenda* to be $X_K = \{x \succ y : x, y \in K \text{ with } x \neq y\}$, which is non-empty and closed under complement, as required for an agenda.¹⁴ We obtain a probabilistic analogue of Arrow’s preference aggregation problem if each individual i submits probabilities on the events in X_K , and the group determines corresponding collective probabilities. Our independence condition then requires that, for any pair of distinct alternatives $x, y \in K$, the collective probability for $x \succ y$ depend only on individual probabilities for $x \succ y$: the probabilistic analogue of Arrow’s independence of irrelevant alternatives. Conditional zero-preservation requires that, for any two pairs of distinct alternatives, $x, y \in K$ and $v, w \in K$, if all individuals agree that the probability of $x \succ y$ given $v \succ w$ is zero, then this agreement be preserved at the collective level: a strengthened probabilistic analogue of the weak Pareto principle (the direct analogue would be zero-preservation). The analogues of Arrow’s universal domain and collective rationality conditions are built into our definition of a pooling function, whose domain and co-domain are defined by the set of all coherent probability measures over Σ_K . To apply our earlier results, note the following easy lemma.

Lemma 8 *If $|K| > 2$, then the preference agenda X_K is non-simple.*

Proof. Consider X_K with $|K| > 2$. Since K has at least three distinct elements x, y, z , the events $x \succ y$, $y \succ z$, and $z \succ x$ are all contained in X_K . But as they are mutually inconsistent but any pair of them is consistent, they form a minimal inconsistent subset of X_K of size greater than two, as required. ■

Now theorem 2(b) yields an immediate corollary:

¹⁴This construction draws on Dietrich and List (2007).

Corollary 2 *For the preference agenda X_K with $|K| > 2$, a conditional zero-preserving pooling function is independent if and only if it is linear.*

This result not only shows that a probabilistic preference aggregation problem can be solved by linear pooling but it also points towards an alternative logically possible escape-route from Arrow’s impossibility theorem (though it may be practically applicable only in special contexts): *If Arrow’s informational framework is enriched so as to allow *degrees of belief* over different possible preference orderings as input and output of the aggregation (or alternatively, *vague preferences*, understood probabilistically), then Arrow’s dictatorship conclusion can be avoided.* Instead, we obtain a positive characterization of linear pooling, despite imposing conditions on the pooling function that are somewhat stronger than Arrow’s classic conditions (in so far as conditional zero-preservation is stronger than the analogue of the weak Pareto principle).

On the belief interpretation, the informational framework employed here 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 dictatorial conclusion in richer informational frameworks and on probabilistic or vague preference aggregation. It is well known that the introduction of interpersonally comparable preferences (of an ordinal or cardinal type) is sufficient for avoiding Arrow’s 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 each specify probabilities of selecting each of the alternatives in K (as opposed to a probability distribution over all possible orderings over K) has been investigated, for example, by Intriligator (1973), who has characterized a version of linear averaging in it. A model in which individuals have vague or fuzzy preferences in the sense that, for any pair of alternatives $x, y \in K$, each individual prefers x to y to a certain degree between 0 and 1 has been investigated, for example, by Billot (1991) and more recently by Piggins and Perote-Peña (2007) (see also Sanver and Selçuk 2007), but the standard constraints on vague or fuzzy preferences, unlike the constraints on probabilities in our present model, do not require individuals to hold a coherent probability distribution over the set of possible crisp orderings; hence the literature has tended to generate Arrow-style impossibility results.

For the purposes of our present paper, 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 opinion pooling.

7 References

- Aczél, J. (1966) *Lectures on Functional Equations and their Applications*, Dover Publications Inc., United States of America
- Billot, A. (1991) “Aggregation of preferences: The fuzzy case”, *Theory and Decision* 30(1): 51-93
- Bradley, R., F. Dietrich and C. List (2006) “Aggregating Causal Judgments”, paper presented at the 2006 Conference of the Philosophy of Science Association
- Dietrich, F. (2006) “Judgment Aggregation: (Im)Possibility Theorems”, *Journal of Economic Theory* 126(1): 286-298
- Dietrich, F. and C. List (2007) “Arrow’s theorem in judgment aggregation”, *Social Choice and Welfare* 29(1): 19-33
- Dokow, E. and R. Holzman (forthcoming) “The aggregation of binary evaluations”, *Journal of Economic Theory*
- Genest, C. and J. V. Zidek (1986) “Combining Probability Distributions: A Critique and Annotated Bibliography”, *Statistical Science* 1 (1): 113-135
- Genest, C. and K. Wagner (1984) “Further Evidence against Independence Preservation in Expert Judgement Synthesis”, *Technical Report* 84-10, Dept. 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
- McConway (1981) “Marginalization and Linear Opinion Pools”, *Journal of the American Statistical Association* 76: 410-14
- Mongin, P. (1995) “Consistent Bayesian aggregation”, *Journal of Economic Theory* 66, p. 313-351
- Nehring, K. and C. Puppe (2007) “Abstract Arrovian Aggregation”, working paper, University of Karlsruhe
- 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
- Sanver, M. R. and Ö. Selçuk (2007) “Sophisticated Preference Aggregation”, working paper, Istanbul Bilgi University
- 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

A Appendix: proofs

A.1 Properties of conditional zero-preservation

We have already pointed out the relationship between conditional zero-preservation and the classic condition of zero-preservation; it is re-stated formally below. For our proofs, it is also useful to introduce a condition that is logically equivalent to conditional zero-preservation. It requires that if all individuals agree that some relevant event probabilistically implies another (i.e., the probability of the one occurring without the other is zero), then that agreement be preserved at the collective level:

Implication-preservation. For all events $A, B \in X$ and all $P_1, \dots, P_n \in \mathcal{P}$, if $P_i(A \setminus B) = 0$ for all individuals i then $P_{P_1, \dots, P_n}(A \setminus B) = 0$.

To see the equivalence between conditional zero-preservation and implication-preservation, note that, for any $P \in \mathcal{P}$ and any $A, B \in X$,

$$\begin{aligned} P(A|B) = 0 &\Leftrightarrow \left(\frac{P(A \cap B)}{P(B)} = 0 \text{ and } P(B) \neq 0 \right) \text{ or } P(B) = 0 \\ &\quad \text{(since, by stipulation, } P(A|B) = 0 \text{ when } P(B) = 0) \\ &\Leftrightarrow P(A \cap B) = 0 \text{ or } P(B) = 0 \\ &\Leftrightarrow P(A \cap B) = 0 \\ &\Leftrightarrow P(A \setminus B^c) = 0. \end{aligned}$$

Since X is complement-closed, we can substitute B^C for B in the antecedent and consequent of implication-preservation, and the equivalence follows.

This equivalence also illuminates the relationship between conditional zero-preservation and its unconditional counterpart, because the relationship between implication-preservation and zero-preservation is very transparent. To see that implication-preservation implies zero-preservation, simply observe that if we take $B = A^c$ then $P(A \setminus B) = P(A \setminus A^c) = P(A)$. To see the converse when X is closed under pairwise intersection or union (and recalling that it is closed under taking complements), note that, for any $A, B \in X$, $A \setminus B = A \cap B^c \in X$. More generally, implication-preservation is equivalent to a variant of zero-preservation extended to the intersections of any two events in X . Note, however, that an implication-preserving pooling function need not preserve a unanimously held zero probability of a *union* of two relevant events, or of an intersection or union of *more than two* relevant events.

The following proposition summarizes the logical relationships:

- Proposition 1** (a) *Conditional zero-preservation is equivalent to implication-preservation.*
- (b) *Zero-preservation is equivalent to the following condition:*
 $[\forall i P_i(A) = 1] \Rightarrow P_{P_1, \dots, P_n}(A) = 1$, for all $A \in X$, $P_1, \dots, P_n \in \mathcal{P}$.
- (c) *Implication-preservation (and thereby conditional zero-preservation) is equivalent to each of the following conditions:*
 $[\forall i P_i(A \cap B) = 0] \Rightarrow P_{P_1, \dots, P_n}(A \cap B) = 0$, for all $A, B \in X$, $P_1, \dots, P_n \in \mathcal{P}$;
 $[\forall i P_i(A \cup B) = 1] \Rightarrow P_{P_1, \dots, P_n}(A \cup B) = 1$, for all $A, B \in X$, $P_1, \dots, P_n \in \mathcal{P}$.
- (d) *Conditional zero-preservation implies zero-preservation, and is equivalent to it if the agenda X is closed under taking the union of any two events.*

A.2 Proof of theorem 1(a)

Proof of lemma 1. (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 $A_j^c \cup A_{j+1}^c = (A_j \cap A_{j+1})^c \neq \emptyset^c = \Omega$).

(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 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 2. Consider an agenda $X \neq \{\emptyset, \Omega\}$. By lemma 1(d), the two claims are equivalent. So it suffices to prove the first one. Note that X is nested if and only if $X \setminus \{\emptyset, \Omega\}$ is nested. So we may assume without loss of generality that $\emptyset, \Omega \notin X$.

First suppose there are two equivalence classes with respect to \sim . Let X_+ be one of them. By lemma 1(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 of course 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 1(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 1(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 to A^c 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 $A_1 \in X^+$ and $A_k^c \in X^+$, 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. ■

Proof of lemma 3. Let F be independent and implication-preserving. Let $D_A, A \in X$ be the pooling criteria 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 $A \cup B \neq \Omega$. 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 measures $P_1, \dots, P_n \in \mathcal{P}$ such that

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

We have $P_i(A \setminus B) = 0$ for all i , so that by implication-preservation $P_{P_1, \dots, P_n}(A \setminus B) = 0$; and we have $P_i(B \setminus A) = 0$ for all i , so that by implication-preservation $P_{P_1, \dots, P_n}(B \setminus A) = 0$. So

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

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

A.3 Proof of theorem 1(b)

We first recall a simple fact of probability theory (in which the word ‘finite’ is of course essential).

Lemma 9 *Every probability measure on a finite sub- σ -algebra of Σ can be extended to a probability measure on Σ .*

Proof. Let $\Sigma^* \subseteq \Sigma$ be a finite sub- σ -algebra of Σ , and $P^* : \Sigma^* \rightarrow [0, 1]$ a probability measure. 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} is a partition of Ω , and so that $\sum_{A \in \mathcal{A}} P^*(A) = 1$. For each atom $A \in \mathcal{A}$, consider a world $\omega_A \in A$, and the associated Dirac measure $\delta_{\omega_A} : \Sigma \rightarrow [0, 1]$ (defined, for all $B \in \Sigma$, by $\delta_{\omega_A}(B) = 1$ if $\omega_A \in B$ and $\delta_{\omega_A}(B) = 0$ if $\omega_A \notin B$). Then

$$P := \sum_{A \in \mathcal{A}} P^*(A) \delta_{\omega_A}$$

defines a probability measure on Σ , as it is by $\sum_{A \in \mathcal{A}} P^*(A) = 1$ a convex combination of the probability measures $\delta_{\omega_1}, \dots, \delta_{\omega_k}$. Further, P extends P^* because for all $B \in \Sigma^*$ we have

$$P(B) = \sum_{A \in \mathcal{A}: \omega_A \in B} P^*(A) = \sum_{A \in \mathcal{A}: A \subseteq B} P^*(A) = P^*(B),$$

where the first equality holds by definition of P , and the last equality by additivity of P^* and the fact that $\{A \in \mathcal{A} : A \subseteq B\}$ forms a partition of B . ■

To prove part (b) of theorem 1, consider a finite nested agenda $X \neq \{\emptyset, \Omega\}$. We construct a pooling function $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ with the relevant properties. Without loss of generality, we suppose that $\emptyset, \Omega \in X$, and that the σ -algebra generated by X is Σ , drawing on the following fact:

Claim. If the theorem holds when Σ is generated by X , it holds in general.

Indeed, suppose the theorem holds in the special case. Let $\Sigma^* (\subseteq \Sigma)$ be the σ -algebra generated by X , and \mathcal{P}^* the set of probability measures on Σ^* . By assumption, there exists a pooling function $F^* : (\mathcal{P}^*)^n \rightarrow \mathcal{P}^*$, $(P_1^*, \dots, P_n^*) \mapsto P_{P_1^*, \dots, P_n^*}^*$ with the relevant properties. For all $P_1^*, \dots, P_n^* \in \mathcal{P}^*$, the collective probability measure $P_{P_1^*, \dots, P_n^*}^* : \Sigma^* \rightarrow [0, 1]$ can by lemma 9 be extended to one on

Σ ; call it $P_{P_1^*, \dots, P_n^*}^\Sigma$. Now define a pooling function $F : \mathcal{P}^n \rightarrow \mathcal{P}$, $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ by

$$P_{P_1, \dots, P_n} := P_{P_1|_{\Sigma^*}, \dots, P_n|_{\Sigma^*}}^\Sigma$$

(i.e., the P_i 's are first restricted to Σ^* , then pooled using F^* into a probability measure on Σ^* , which is then extended to Σ). F inherits from F^* all relevant properties (independence, non-neutrality, and implication-preservation), essentially because these properties refer only to probabilities of events that are in Σ^* (more precisely, that are in X or – in the case of implication-preservation – that are differences of events in X). This proves the claim.

As X is nested and finite, we may write it as $X = \{A_0, \dots, A_k, A_1^c, \dots, A_k^c\}$ with events $\emptyset = A_0 \subsetneq A_1 \subsetneq \dots \subsetneq A_k = \Omega$.

Consider any neutral implication-preserving pooling function whose pooling criterion $D : [0, 1]^n \rightarrow [0, 1]$ is at least weakly increasing in each argument (e.g., dictatorship by individual 1, given by $(P_1, \dots, P_n) \mapsto P_1$). As $X \neq \{\emptyset, \Omega\}$, there is a contingent event $A \in X$. As A is contingent, there are $P_1, \dots, P_n \in \mathcal{P}$ that all assign probability $1/2$ to A (hence to A^c), so that the collective probabilities of A and of A^c are each given by $D(1/2, \dots, 1/2)$. As these probabilities sum to 1, it follows that

$$D(1/2, 1/2, \dots, 1/2) = 1/2. \quad (1)$$

We now transform this neutral pooling function into a non-neutral one (that is still independent and implication-preserving). To do so, we consider a function $T : [0, 1] \rightarrow [0, 1]$ such that (i) $T(1/2) \neq 1/2$, (ii) $T(0) = 0$ and $T(1) = 1$, and (iii) T is strictly increasing (e.g., $T(x) = x^2$ for all $x \in [0, 1]$).

Now consider any $P_1, \dots, P_n \in \mathcal{P}$. We have to define the collective probability measure $P_{P_1, \dots, P_n} : \Sigma \rightarrow [0, 1]$. As the σ -algebra Σ is generated by X , hence by $\{A_j : j = 0, \dots, k\}$, the atoms of Σ (i.e., the \subseteq -minimal elements of $\Sigma \setminus \{\emptyset\}$) are the differences $A_j \setminus A_{j-1}$, $j = 1, \dots, k$. We define the measure $P_{P_1, \dots, P_n} : \Sigma \rightarrow [0, 1]$ by specifying its value on the atoms as follows:

$$P_{P_1, \dots, P_n}(A_j \setminus A_{j-1}) := T \circ D(P_1(A_j), \dots, P_n(A_j)) - T \circ D(P_1(A_{j-1}), \dots, P_n(A_{j-1}))$$

for all $j \in \{1, \dots, k\}$. As each A_j ($j \in \{0, \dots, k\}$) is partitioned into the sets $A_l \setminus A_{l-1}$, $l = 1, \dots, j$, its measure is given by

$$P_{P_1, \dots, P_n}(A_j) = \sum_{l=1}^j [T \circ D(P_1(A_l), \dots, P_n(A_l)) - T \circ D(P_1(A_{l-1}), \dots, P_n(A_{l-1}))],$$

which (by cancelling out and using the fact that $A_0 = \emptyset$, that $D(0, \dots, 0) = 0$, and that $T(0) = 0$) reduces to

$$P_{P_1, \dots, P_n}(A_j) = T \circ D(P_1(A_j), \dots, P_n(A_j)) \text{ for all } j = 0, \dots, k. \quad (2)$$

To see why P_{P_1, \dots, P_n} is indeed a *probability* measure, note that each atom has non-negative measure (using the fact that T and D are increasing functions),

and that $P_{P_1, \dots, P_n}(\Omega) = P_{P_1, \dots, P_n}(A_k) = 1$ (by (2) and since $D(1, \dots, 1) = 1$ and $T(1) = 1$).

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

Independence. Applied to any event of type $A_j \in X$, independence holds with pooling criterion D_{A_j} defined as $T \circ D$, by (2). Applied to any event of type $A_j^c \in X$, independence holds with pooling criterion $D_{A_j^c}$ defined by $(t_1, \dots, t_n) \mapsto 1 - T \circ D(1 - t_1, \dots, 1 - t_n)$, because for all $P_1, \dots, P_n \in \mathcal{P}$ we have

$$\begin{aligned} P_{P_1, \dots, P_n}(A_j^c) &= 1 - P_{P_1, \dots, P_n}(A_j) = 1 - T \circ D(P_1(A_j), \dots, P_n(A_j)) \\ &= 1 - T \circ D(1 - P_1(A_j^c), \dots, 1 - P_n(A_j^c)). \end{aligned}$$

Non-neutrality. By independence, the decision on any $A \in X$ is made via a pooling criterion $D_A : [0, 1]^n \rightarrow [0, 1]$. We show that the pooling criteria D_A are not all identical – or, more precisely, cannot be chosen to be all identical. As $X \neq \{\emptyset, \Omega\}$, there is a pair $A_j, A_j^c \in X$ of contingent events. As is easily checked, the pooling criterion for any contingent event is unique; so D_{A_j} and $D_{A_j^c}$ must be defined as in our independence proof above. We show that $D_{A_j} \neq D_{A_j^c}$. Using (1), we have

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

So, as $T(1/2) \neq 1/2$ (by assumption on T), $D_{A_j}(1/2, \dots, 1/2) \neq D_{A_j^c}(1/2, \dots, 1/2)$, and hence $D_{A_j} \neq D_{A_j^c}$, as desired.

Implication-preservation. Consider any $A, B \in X$ and $P_1, \dots, P_n \in \mathcal{P}$ such that $P_i(A \setminus B) = 0$ for all individuals i . As one easily checks, $A \setminus B$ takes the form $A_m \setminus A_l$ for some $m, l \in \{0, \dots, k\}$ with $m \geq l$. Hence

$$\begin{aligned} P_{P_1, \dots, P_n}(A \setminus B) &= P_{P_1, \dots, P_n}(A_m) - P_{P_1, \dots, P_n}(A_l) \text{ (by } A_l \subseteq A_m) \\ &= T \circ D(P_1(A_m), \dots, P_n(A_m)) - T \circ D(P_1(A_l), \dots, P_n(A_l)). \end{aligned}$$

In the last expression, each individual i has $P_i(A_m) = P_i(A_l)$, as $A_l \subseteq A_m$ with $P_i(A_m \setminus A_l) = P(A \setminus B) = 0$. So the expression equals zero, i.e., $P_{P_1, \dots, P_n}(A \setminus B) = 0$, as desired. ■

A.4 Proof of theorem 2(a)

Proof of lemma 4.

(a) Note that as X is non-simple it contains an event A for which $A \neq \emptyset, \Omega$. For each $x \in [0, 1]^n$ there are (by $A \neq \emptyset, \Omega$) probability functions $P_1, \dots, P_n \in \mathcal{P}$

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 implication-preserving and hence zero-preserving, $D(\mathbf{0}) = 0$, so that by part (a) $D(\mathbf{1}) = 1 - D(\mathbf{0}) = 1$. ■

Proof of lemma 5.

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

$$A^* := A^c \cap \left(\bigcap_{D \in Y \setminus \{A\}} D \right), \quad B^* := B^c \cap \left(\bigcap_{D \in Y \setminus \{B\}} D \right), \quad C^* := C^c \cap \left(\bigcap_{D \in Y \setminus \{C\}} D \right).$$

As Σ is closed under countable intersections, $A^*, B^*, C^* \in \Sigma$. For all i , as $x_i + y_i + z_i = 1$ and as A^*, B^*, C^* are pairwise disjoint non-empty members of Σ , there exists a $P_i \in \mathcal{P}$ with

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

By construction,

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

For the so-defined profile (P_1, \dots, P_n) , we consider the collective probability function $P := P_{P_1, \dots, P_n}$. We now derive five properties of P (claims 1-5), which then allow us to show that $D(x) + D(y) + D(z) = 1$ (claim 6), as desired.

Claim 1. $P(\bigcap_{D \in Y \setminus \{A, B, C\}} D) = 1$.

For all $D \in Y \setminus \{A, B, C\}$ we have $D \supseteq A^* \cup B^* \cup C^*$, so that by (3) we have $P_1(D) = \dots = P_n(D) = 1$, and hence $P(D) = 1$ by proposition 1. This implies claim 1 because the intersection of countably many events of probability one has probability one.

Claim 2. $P(A^c \cup B^c \cup C^c) = 1$.

As $A \cap B \cap C$ is disjoint from the event $\bigcap_{D \in Y \setminus \{A, B, C\}} D$, which by claim 1 has P -probability one, we have $P(A \cap B \cap C) = 0$. This implies claim 2 because $A^c \cup B^c \cup C^c$ is the complement of $A \cap B \cap C$.

Claim 3. $P((A^c \cap B^c) \cup (A^c \cap C^c) \cup (B^c \cap C^c)) = 0$.

As $A^c \cap B^c$ is disjoint with each of A^*, B^*, C^* , it is disjoint with the event $A^* \cup B^* \cup C^*$ to which each individual i assigns probability one by (3). So $P_i(A^c \cap B^c) = 0$ for all i . Hence $P(A^c \cap B^c) = 0$ by proposition 1(c). For analogous reasons, $P(A^c \cap C^c) = 0$ and $P(B^c \cap C^c) = 0$. Now claim 3 follows

since the union of finitely (or countably) many events of probability zero has probability zero.

Claim 4. $P((A^c \cap B \cap C) \cup (A \cap B^c \cap C) \cup (A \cap B \cap C^c)) = 1$

By claims 2 and 3, there is a P -probability of one that *at least one* of A^c, B^c, C^c holds, but a P -probability of zero that *at least two* of A^c, B^c, C^c hold. So with P -probability of one *exactly one* of A^c, B^c, C^c holds. This is precisely what claim 4 states.

Claim 5. $P(A^*) + P(B^*) + P(C^*) = P(A^* \cup B^* \cup C^*) = 1.$

The first equality follows from the pairwise disjointness of the events A^*, B^*, C^* and the additivity of P . Regarding the second equality, note that $A^* \cup B^* \cup C^*$ is the intersection of the events $\bigcap_{D \in Y \setminus \{A, B, C\}} D$ 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 4. So $P(A^* \cup B^* \cup C^*) = 1$, as desired.

Claim 6. $D(x) + D(y) + D(z) = 1.$

As $P(A^* \cup B^* \cup C^*) = 1$ by claim 5, and as the intersection of A^c with $A^* \cup B^* \cup C^*$ is A^* , we have

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

By $A^c \in X$ we moreover have

$$P(A^c) = D(P_1(A^c), \dots, P_n(A^c)) = D(P_1(A^*), \dots, P_n(A^*)) = D(x).$$

This and (4) imply that $P(A^*) = D(x)$. By similar arguments, $P(B^*) = D(y)$ and $P(C^*) = D(z)$. So claim 6 follows from claim 5. ■

Proof of lemma 6. Suppose

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

Then, for all $x, y \in [0, 1]^n$ with $x + y \in [0, 1]^n$,

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

where the first equality follows from (5) and the second from lemma 4(a). So

$$D(x + y) = D(x) + D(y) \text{ for all } x, y \in [0, 1]^n \text{ with } x + y \in [0, 1]^n. \quad (6)$$

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 (6), 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 easily checks, 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 (6)), 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 to $x = \mathbf{1}$ yields $D(\mathbf{1}) = \sum_{i=1}^n w_i$, hence $\sum_{i=1}^n w_i = 1$ by lemma 4(b). So F is a linear pooling function, as desired. ■

A.5 Proof of theorem 2(b)

Let the agenda $X (\subseteq \Sigma)$ be simple, finite, and not $\{\emptyset, \Omega\}$. We construct a non-linear pooling function that is independent (in fact, neutral) and implication-preserving. We may assume without loss of generality that the σ -algebra generated by X is Σ , because the ‘claim’ in the proof of theorem 1(b) (proved using lemma 9) holds analogously here as well.

As an ingredient to the construction, we use an arbitrary linear implication-preserving pooling function $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}^{\text{lin}}$ (e.g., that defined by $(P_1, \dots, P_n) \mapsto P_1$), and denote by D^{lin} its pooling criterion for all events $A \in X$. To anticipate, the pooling function $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ to be constructed will have the pooling criterion $D : [0, 1]^n \rightarrow [0, 1]$ (for every $A \in X$) given by

$$D(t_1, \dots, t_n) := \begin{cases} 0 & \text{if } D^{\text{lin}}(t_1, \dots, t_n) < 1/2 \\ 1/2 & \text{if } D^{\text{lin}}(t_1, \dots, t_n) = 1/2 \\ 1 & \text{if } D^{\text{lin}}(t_1, \dots, t_n) > 1/2. \end{cases} \quad (7)$$

Consider any $P_1, \dots, P_n \in \mathcal{P}$. We have to define P_{P_1, \dots, P_n} . We write collective probabilities under the linear function simply as

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

and define

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

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

To define P_{P_1, \dots, P_n} , we first need to prove two claims (using the fact that X is simple).

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

To show this, note first that $X_{= 1/2}$ has of course a subset Y such that $p(A \cap B) > 0$ for all $A, B \in Y$ (e.g., $Y = \emptyset$). Among all such subsets $Y \subseteq$

$X_{=1/2}$, let $X_{=1/2}^1$ a *maximal* one (with respect to set-inclusion), and let $X_{=1/2}^2 := X_{=1/2} \setminus X_{=1/2}^1$. By definition, $X_{=1/2}^1$ and $X_{=1/2}^2$ form a partition of $X_{=1/2}$. We show that (i) and (ii) hold.

(i) Property (i) holds by definition for $X_{=1/2}^1$, and holds for $X_{=1/2}^2$ too by the following argument. Let $A, B \in X_{=1/2}^2$ and suppose for a contradiction that $p(A \cap B) = 0$. By definition of $X_{=1/2}^2$, there are $A', B' \in X_{=1/2}^1$ such that $p(A \cap A') = 0$ and $p(B \cap B') = 0$. In particular, $p(A \cap C) = p(B \cap C) = 0$ for $C := A' \cap B'$. Since the intersection of any two of the sets A, B, C has zero p -probability, we have

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

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

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

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

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

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

To show this, first write X as $X = \{C_j^0, C_j^1 : j = 1, \dots, J\}$, where $J = |X|/2$ and where each pair C_j^0, C_j^1 consists of an event and its complement. We may write Σ as the set of all unions of intersections of the form $C_1^{k_1} \cap \dots \cap C_J^{k_J}$, i.e., as

$$\Sigma = \{\cup_{(k_1, \dots, k_J) \in K} (C_1^{k_1} \cap \dots \cap C_J^{k_J}) : K \subseteq \{0, 1\}^J\}. \quad (8)$$

Recalling that Σ is the σ -algebra generated by X , the inclusion ' \supseteq ' in (8) is obvious, and the inclusion ' \subseteq ' holds because the right hand side of (8) includes X (as any $C_j^k \in X$ can be written as the union of all intersections $C_1^{k_1} \cap \dots \cap C_J^{k_J}$ for which $k_j = k$) and is a σ -algebra (as it is closed under taking unions and complements: just take the unions respectively complements of the corresponding sets $K \subseteq \{0, 1\}^J$).

From (8) and the pairwise disjointness of the intersections of the form $C_1^{k_1} \cap \dots \cap C_J^{k_J}$, it is clear that every consistent such intersection is an atom of Σ . Now $\cap_{C \in X_{=1/2}^j \cup X_{>1/2}} C$ is (for $j \in \{0, 1\}$) precisely such a consistent intersections. Indeed, $\cap_{C \in X_{=1/2}^j \cup X_{>1/2}} C$ is consistent by claim 1, and contains a member of each pair A, A^c in X , if $p(A) = p(A^c) = 1/2$ by claim 1 and if $p(A) \neq p(A^c)$ since there

then is a $B \in \{A, A^c\}$ with $p(B) > 1/2$, i.e., with $B \in X_{>1/2} \subseteq X_{=1/2}^j \cup X_{>1/2}$. This proves claim 2.

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

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

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

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

Independence. We show that the pooling function is neutral (hence independent) with the pooling criterion D given in (7). To do so, consider any $P_1, \dots, P_n \in \mathcal{P}$ and any $A \in X$, and write $(t_1, \dots, t_n) := (P_1(A), \dots, P_n(A))$. We have to show that $P_{P_1, \dots, P_n}(A) = D(t_1, \dots, t_n)$. To do this, we consider three cases, and use $p, X_{>1/2}, X_{=1/2}^1, X_{=1/2}^2, \omega^1, \omega^2$ as defined above.

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

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

Case 3. $p(A) = D^{\text{lin}}(t_1, \dots, t_n) = 1/2$. Then $D(t_1, \dots, t_n) = 1/2$. So we must prove that $P_{P_1, \dots, P_n}(A) = 1/2$, i.e., that A contains exactly one of ω^1 and ω^2 . As $p(A) = 1/2$, exactly one of $X_{=1/2}^1$ and $X_{=1/2}^2$ contains A and the other one

contains A^c , by claim 1. Say $A \in X_{=1/2}^1$ and $A^c \in X_{=1/2}^2$ (the proof is analogous if instead $A \in X_{=1/2}^2$ and $A^c \in X_{=1/2}^1$). So $A \supseteq \cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$, and hence $\omega^1 \in A$. On the other hand, $\omega^2 \notin A$, because A is disjoint with A^c , hence with its subset $\cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$, which contains ω^2 .

Non-linearity. As $X \neq \{\emptyset, \Omega\}$, there is a contingent event $A \in X$, hence a probability measure $P \in \mathcal{P}$ with $t := P(A) \notin \{0, 1/2, 1\}$. Now assume all individuals submit this P . If the pooling function were linear, the collective probability of A would again be t ($\notin \{0, 1/2, 1\}$), a contradiction since the collective probability is given by $D(t, \dots, t)$ ($\in \{0, 1/2, 1\}$), as just shown.

Implication-preservation. We assume that $A, B \in X$ and $P_1, \dots, P_n \in \mathcal{P}$ such that $P_i(A \cup B) = 1$ for all i , and show that $P_{P_1, \dots, P_n}(A \cup B) = 1$, which by proposition 1 establishes implication-preservation. For all i we have $P_i(A) + P_i(B) \geq P_i(A \cup B) = 1$, and hence $P(A_i) \geq 1 - P_i(B) = P_i(B^c)$. So, as $D^{\text{lin}} : [0, 1]^n \rightarrow [0, 1]$ takes a linear form with non-negative coefficients and hence is weakly increasing in every component,

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

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

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

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

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

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

Analogously, it cannot be that A and B are both in $X_{=1/2}^2$. So one of A and B is in $X_{=1/2}^1$ and the other one in $X_{=1/2}^2$; say $A \in X_{=1/2}^1$ and $B \in X_{=1/2}^2$ (the proof is analogous otherwise). So $A \supseteq \cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$ and $B \supseteq \cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$, and hence $\omega^1 \in A$ and $\omega^2 \in B$. So $A \cup B$ contains both ω^1 and ω^2 , whence $P_{P_1, \dots, P_n}(A \cup B) = 1$, as desired. ■