

Opinion pooling on general agendas

Franz Dietrich & Christian List

September 2007

1 Introduction

We consider the classical aggregation problem of *opinion pooling*: the probability assignments of different individuals are to be merged into collective probability assignments. While opinion pooling has been explored in some depth in the literature (by statisticians, economists, and philosophers), all contributions so far assume that the set of relevant events (the "agenda") forms a σ -algebra. This assumption has a technical motivation: the standard rationality conditions on probability assignments (in particular, σ -additivity) refer to a σ -algebra as their domain. On the other hand, the disjunction (union) or conjunction (intersection) of two relevant events need not be relevant in practice: it may well be relevant whether it rains, and relevant whether we are happy, but irrelevant whether the disjunction or conjunction of these two events holds. With irrelevance of events we either mean that the collective does not need a probability for them, or that conditions standardly imposed on opinion pooling (such as independence and zero-preservation) lose their normative appeal when applied to these (artificial) events.

In this paper we explore opinion pooling without assuming that the agenda of relevant events forms a σ -algebra. We show that, for a broad class of agendas (whose events may be much less interconnected than those of a σ -algebra), any opinion pooling operator with two properties must be *linear*, i.e. derive the collective probability of each relevant event as a weighted linear average of the individuals' probabilities of the event where the weights are event-independent. For an even broader class of agendas, we obtain a weaker conclusion: the pooling operator must be *neutral*, i.e. derive the collective probability of each relevant event as some (possibly non-linear) function of the individuals' probabilities of the event where the function is event-independent.

The latter neutrality result applies in particular to an agenda type that is frequent in practice: agendas consisting of *logically independent* events, i.e. decision problems of assigning probabilities to certain events, where any combination of truth values ('true' or 'false') of these events is consistent. This case of logically independent events is frequent because between events of interest (such as between rainfall and happiness) there often is no *logical* dependence but only a *probabilistic* dependence (correlation). Agendas of logically independent events can be viewed as diametrically opposed to agendas that form a σ -algebra. By focussing on σ -algebras, the literature has in effect excluded many realistic applications.

For the classical case that the agenda is a σ -algebra, linearity and neutrality are among the most studied properties of pooling operators (in the case of neutrality sometimes under other names such as *strong label neutrality* or *strong setwise function property*). Linear pooling goes back to Stone (1961) (or even to Laplace), and neutral pooling to McConway (1981) and Wagner (1982). The σ -algebra case has the interesting feature that every neutral pooling operator is automatically linear, so that neutrality is in fact equivalent to linearity, if the σ -algebra contains more than four events (McConway 1981 and Wagner 1982; see also Mongin's 1995 linearity characterisation). This peculiarity does not carry over to general agendas: some agendas allow for neutral yet non-linear opinion pooling, as seen below.

The reader is referred to Genest and Zidek's (1986) overview article for an excellent review of classical results on opinion pooling.

2 Model

Consider a group of $n \geq 2$ individuals, labelled $i = 1, \dots, n$. Let Ω be a non-empty set of *worlds* (or *states of affairs*) and Σ a σ -algebra of *events* $A \subseteq \Omega$. For instance, Σ could be the power set of Ω . We write $A^c := \Omega \setminus A$ for the complement (negation) of any $A \subseteq \Omega$. An event A is *contingent* if it is neither \emptyset nor Ω . An event A *entails* another one B if $A \subseteq B$. A set of events Y is *consistent* if $\bigcap_{A \in Y} A \neq \emptyset$; it is *inconsistent* if $\bigcap_{A \in Y} A = \emptyset$; and it *entails* an event B if $\bigcap_{A \in Y} A \subseteq B$.

The group has to find a "collective" (σ -additive) probability measure $P : \Sigma \rightarrow [0, 1]$, based on the "profile" (P_1, \dots, P_n) of individual probability measures $P_i : \Sigma \rightarrow [0, 1]$. Let \mathcal{P} denote the set of probability measures $P : \Sigma \rightarrow [0, 1]$. An (*opinion*) *pooling operator* is simply a mapping $F : \mathcal{P}^n \rightarrow \mathcal{P}$; it assigns to each profile $(P_1, \dots, P_n) \in \mathcal{P}^n$ of (individual) beliefs a collective belief $F(P_1, \dots, P_n)$, which we hereafter often denote P_{P_1, \dots, P_n} . For instance, P_{P_1, \dots, P_n} could be given by the arithmetic average $\frac{1}{n}P_1 + \dots + \frac{1}{n}P_n$, a case of *linear pooling* (as defined later). There are of course numerous other pooling operators, including geometric averages (of a weighted or non-weighted kind), expert rules (in which P_{P_1, \dots, P_n} is P_i with a fixed or profile-dependent "expert" i), median rules etc.

Unlike in the literature, we assume that only certain events in Σ are *relevant*. As discussed below, the relevant events can be interpreted in (at least) two ways: either as the only events for which the group actually needs probabilities, or as the only events for which conditions placed on the pooling operator F (independence, zero-preservation, etc.) are normatively compelling. We call the set of relevant events the *agenda*, to stress the connection to agendas in social choice theory. Formally, an agenda is a non-empty set $X \subseteq \Sigma$ such that $A \in X$ implies $A^c \in X$. So non-emptiness and closure under taking complements are our only (plausible) constraints on the notion of relevance. Crucially, we allow X to contain A and B without containing $A \cup B$ (respectively $A \cap B$): it may be

relevant whether it rains, and also whether we are happy, yet irrelevant whether it rains *or* we are happy (respectively *and* we are happy).

We refer to the case $X = \Sigma$ as the "classical" case, as then all events are treated as relevant and the below conditions on aggregation reduce to standard conditions.

Two interpretations of relevance, i.e. of the agenda X .

1. We may interpret X as the set of events for which collective probabilities are ultimately needed; probabilities of events $A \notin \Sigma \setminus X$ are of no collective use. If only the restriction $P|_X$ of the collective probability measure $P : \Sigma \rightarrow [0, 1]$ is needed, then in the view of the independence condition below the individuals need only submit – indeed need only *hold* – probabilities of events in X . While the individuals need not make up their mind on other events, they must however ensure that the probabilities of events in X are consistent, i.e., can be extended to a full probability measure on Σ . In practice, a pooling operator is thus a mapping $\tilde{\mathcal{P}}^n \rightarrow \tilde{\mathcal{P}}$, where $\tilde{\mathcal{P}}$ is the set of functions $\tilde{P} : X \rightarrow [0, 1]$ that are extendible to a probability measure $P : \Sigma \rightarrow [0, 1]$.

2. Under a second interpretation, also probabilities of events outside X are of collective interest (and must be submitted by people), but only for events in X do conditions (such as independence) have normative force. So X is here defined by the scope we wish to give to aggregation conditions. Probabilities of artificial events like the conjunction $A \cap B$ of A : "Global warming kills animal species X " and B : "GDP growth will accelerate" might well be needed (say, for political decision-making); but one will not want to vote in isolation on $A \cap B$, i.e. to apply the independence condition to $A \cap B$. So this event $A \cap B$ will be in Σ but not in X .

3 Two conditions on opinion pooling and their interpretation

All our characterisation results are based on two conditions: *independence* and *implication-preservation*, as defined in a moment. In the classical case $X = \Sigma$, independence precisely matches the standard independence condition (also called *weak setwise function property*), and implication-preservation becomes equivalent to the standard condition of zero-preservation.

Independence. For every relevant event $A \in X$ there exists a function $D_A : [0, 1]^n \rightarrow [0, 1]$ (the "(local) decision rule" for A) such that $P_{P_1, \dots, P_n}(A) = D_A(P_1(A), \dots, P_n(A))$ for all $P_1, \dots, P_n \in \mathcal{P}$.

Implication-preservation. For all relevant 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 $P_{P_1, \dots, P_n}(A \setminus B) = 0$; i.e., if all individuals believe that A (probabilistically) implies B , so does the collective.

The main normative defence of independence is the democratic idea that the collective view on events/issues $A \in X$ should be determined by the individuals views on the issue.¹ Such a defence is compelling only if X does not contain "artificial" events (such as conjunctions of semantically unrelated events). So a democratic defence becomes difficult (perhaps impossible) in the classical case that $X = \Sigma$. This might explain why a democratic defence of independence has not (to our knowledge) been put forward in the literature (though similar arguments for independence conditions are common in other fields of aggregation theory). Aside from a democratic defence, two more pragmatic arguments for independence can also be given; these work also in the classical case $X = \Sigma$. First, voting on each issue/event in isolation is easier in practice. Second, independence prevents certain type of agenda manipulation.² An *objection* against independence in the classical case $X = \Sigma$ is its incompatibility with collectively preserving unanimous beliefs of probabilistic independence (see Genest and Wagner 1984).³ Whether the objection applies to our independence condition depends on the agenda at hand X . 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 the process of updating probabilities in the light of new information.

The idea underlying implication-preservation is intuitive: if all individuals believe that some relevant event implies another, e.g. that hail implies damage, or that political instability implies famine, then this belief is taken over collectively. In the classical case $X = \Sigma$, implication-preservation is equivalent to the following standard condition (take $B = \emptyset$):

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 the general case implication-preservation implies zero-preservation (take $B = A^c$) but not vice versa. Although restricted to implications between *relevant* events, implication-preservation in a sense reaches beyond X : it is (by $A \setminus B = A \cap B^c$ and as X is closed under taking complements) equivalent to a variant of zero-preservation extended to intersections of *two* events in X , hence

¹More precisely, independence reflects a *local* notion of democracy; under a more *global* notion of democracy, the collective view on an issue A may also be influenced by people's views on other issues whose semantic content is suitably related to A .

²In the classical case $X = \Sigma$, McConway (1981) shows that independence (his "weak setwise function property") is equivalent to the "marginalization property" whereby (in short) aggregation should commute with the operation of reducing the σ -algebra to some sub- σ -algebra $\Sigma^* \subseteq \Sigma$. A similar result holds also for general agendas X . Thus independence prevents agenda setters from influencing the collective probability of some events by adding or removing other events in the agenda.

³Assuming the aggregation function is *non-dictatorial*, i.e. the collective does not always adopts the probability function of a fixed individual.

to certain irrelevant events as well (unless X is closed under such intersections). The following lemma summarises exactly how implication-preservation strengthens zero-preservation (the simple proof uses that $A \in X \Rightarrow A^c \in X$ and $A \setminus B = A \cap B^c = (A^c \cup B)^c$, for all $A, B \in \Sigma$).

Lemma 1 (a) *Zero-preservation is equivalent to the following condition:*

$$P_i(A) = 1 \ \forall i \Rightarrow P_{P_1, \dots, P_n}(A) = 1, \text{ for all } A \in X, P_1, \dots, P_n \in \mathcal{P}.$$

(b) *Implication-preservation is equivalent to each of the following conditions:*

$$P_i(A \cap B) = 0 \ \forall i \Rightarrow P_{P_1, \dots, P_n}(A \cap B) = 0, \text{ for all } A, B \in X, P_1, \dots, P_n \in \mathcal{P};$$

$$P_i(A \cup B) = 1 \ \forall i \Rightarrow P_{P_1, \dots, P_n}(A \cup B) = 1, \text{ for all } A, B \in X, P_1, \dots, P_n \in \mathcal{P}.$$

(c) *Implication-preservation implies zero-preservation, and is equivalent to it if the agenda X is closed under taking the union of two events.*

Note that an implication-preserving pooling operator 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.

4 Characterisation of neutral pooling

Let us first generalise the standard notions of linear and neutral pooling from the classical case $X = \Sigma$ to the case of a general agenda X .

A pooling operator F is *linear* if there are "weights" $w_1, \dots, w_n \geq 0$ with sum 1 such that

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

or in short $P_{P_1, \dots, P_n}|_X = \sum_{i=1}^n w_i P_i|_X$ for all $P_1, \dots, P_n \in \mathcal{P}$. In the classical case $X = \Sigma$, this reduces to $P_{P_1, \dots, P_n} = \sum_{i=1}^n w_i P_i$ for all $P_1, \dots, P_n \in \mathcal{P}$, i.e. to standard linearity.⁴

A pooling operator F is *neutral* if there is a "decision rule" $D : [0, 1]^n \rightarrow [0, 1]$ such that

$$P_{P_1, \dots, P_n}(A) = D(P_1(A), \dots, P_n(A)) \text{ for all relevant events } A \in X,$$

i.e. if F is independent with the *same* decision rule $D = D_A$ for all $A \in X$. So there is perfect symmetry in how different relevant events/issues are treated: the collective belief on whether it will rain is obtained via the same decision rule

⁴Also, 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 -relativised 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 $P_1, \dots, P_n \in \mathcal{P}$).

as the collective belief on whether we will be happy (assuming these events are relevant). Again, in the classical case $X = \Sigma$ our neutrality condition becomes the studied one in the literature.

Linearity is obviously a special case of neutrality, where the decision rule D moreover takes a linear form, i.e. is given by $D(t_1, \dots, t_n) = \sum_{i=1}^n w_i t_i$ for some non-negative weights w_i of sum 1.

In this section we ask for which agendas X pooling operators satisfying our two conditions (implication-preservation and independence) must be neutral. This question is interesting in its own right, but its answer will also help us later when characterising linear (rather than just neutral) pooling.

We call the agenda X *nested* if it takes 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 instance, binary agendas $X = \{A, A^c\}$ are of course nested: take $X_+ := \{A\}$. 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 indeed linearly ordered by set-inclusion.

An important type of non-nested agendas consists in the agendas X whose pairs $A, A^c \in X$ ("issues") are logically independent: X is non-nested if it takes the form $X = \{A_k, A_k^c : k \in K\}$ with $|K| \geq 2$ such that $\bigcap_{k \in K} A_k^* \neq \emptyset$ for every selection of events $A_k^* \in \{A_k, A_k^c\}$, $k \in K$. As mentioned in the introduction, such agendas are frequent in practice (and are somewhat opposed to σ -algebras with their highly interconnected events).

Nested agendas X are very special: all $A, B \in X$ are *logically dependent* (i.e. one of A, A^c entails one of B, B^c), a trivial case. Nested agendas might therefore also be called "pairwise connected" or "trivial". The following neutrality characterisation applies to all non-nested agendas, hence to all non-trivial opinion pooling situations.

Theorem 1 *For a non-nested agenda X , an implication-preserving pooling operator is independent if and only if it is neutral.*

This characterisation of neutral pooling assumes a non-nested agenda. Is this assumption tight or just "needed for our proof"? It is tight, at least in the finite case:

Theorem 2 *For a nested agenda X , finite and not $\{\emptyset, \Omega\}$, there exists an implication-preserving pooling operator that is independent but not neutral.*

Although nested agendas X allow for non-neutral pooling, only a limited kind of non-neutrality is possible: as will be clear from the proof, the decision rule 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 once independence is strengthened by requiring that $D_A = D_{A^c}$ for all $A \in X$ (or at least for *some* $A \in X \setminus \{\emptyset, \Omega\}$).

5 Proof of the neutrality characterisation

We now prove Theorems 1 and 2. The following binary relation on the agenda is used later to characterise nested agendas.

Definition 1 For any relevant events $A, B \in X$, write $A \sim B$ if there exists a finite sequence of relevant events $A_1, \dots, A_k \in X$ with $A_1 = A$ and $A_k = B$ such that any neighbours A_l, A_{l+1} are neither exclusive nor exhaustive (i.e. $A_l \cap A_{l+1} \neq \emptyset$ and $A_l \cup A_{l+1} \neq \Omega$).

Lemma 2 Consider any agenda X .

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

Proof. (a) Reflexivity, symmetry and transitivity on $X \setminus \{\emptyset, \Omega\}$ are all obvious (where we 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 $A, B \in X$ 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_t, A_{t+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_t^c, A_{t+1}^c are not exclusive (as $A_t^c \cap A_{t+1}^c = (A_t \cup A_{t+1})^c \neq \Omega^c = \emptyset$) and not exhaustive (as $A_t^c \cup A_{t+1}^c = (A_t \cap A_{t+1})^c \neq \emptyset^c = \Omega$).

(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 \sim -equivalence classes is not one. As $X \setminus \{\emptyset, \Omega\} \neq \emptyset$ this number is not zero. So it is at least two. We show two claims:

Claim 1. There are exactly two \sim -equivalence classes.

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

Proof of Claim 1. For a contradiction, let $A, B, C \in X \setminus \{\emptyset, \Omega\}$ be pairwise not \sim -equivalent. By $A \not\sim B$, either $A \cap B = \emptyset$ or $A \cup B = \Omega$. 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 that A^c, B^c, C^c are pairwise not \sim -equivalent 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 a \sim -equivalence class containing the pair A, A^c . By assumption Z is not the only \sim -equivalence class, 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. ■

The nested agendas are precisely the agendas with two \sim -equivalence classes:

Lemma 3 *An agenda $X \neq \{\emptyset, \Omega\}$ is nested if and only if it has two \sim -equivalence classes, and non-nested if and only if it has a single one.*

Proof. Consider an agenda $X \neq \{\emptyset, \Omega\}$. By Lemma 2(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, and X has two \sim -equivalence classes if and only if $X \setminus \{\emptyset, \Omega\}$ does. So we may assume without loss of generality that $\emptyset, \Omega \notin X$.

First suppose there are two \sim -equivalence classes. Let X_+ be one of them. By Lemma 2(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 2(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 2(d) exactly two \sim -equivalence classes, 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_t, A_{t+1} , $A_t \cap A_{t+1} \neq \emptyset$ and $A_t \cup A_{t+1} \neq \Omega$. 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_t, A_{t+1} such that $A_t, A_{t+1}^c \in X^+$. So, as X^+ is linearly ordered by \subseteq , either $A_t \subseteq A_{t+1}^c$ or $A_{t+1}^c \subseteq A_t$. In the first case, $A_t \cap A_{t+1} = \emptyset$, a contradiction. In the second case, $A_t \cup A_{t+1} = \Omega$, also a contradiction. ■

The above characterisation of nestedness in terms of \sim -equivalence classes is important largely for the following reason.

Lemma 4 *An independent implication-preserving pooling operator is neutral on each \sim -equivalence class.*

With neutrality *on a set* Z we of course mean the analogue of neutrality with the agenda X replaced by the set Z .

Proof. Let F be independent and implication-preserving. Let $D_A, A \in X$, be the decision rules 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 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. ■

With these lemmas in place, we now turn to our neutrality characterisation.

Proof of Theorem 1. The claim follows by combining Lemmas 3 and 4. ■

To prove Theorem 2, we first recall a simple fact of probability theory (in which the word "finite" is of course essential):

Lemma 5 *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 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 $B = \Sigma^*$ 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 . ■

Proof of Theorem 2. Consider a finite nested agenda $X \neq \{\emptyset, \Omega\}$. We construct a pooling operator $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ 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 Σ^* ($\subseteq \Sigma$) be the σ -algebra generated by X , and \mathcal{P}^* the set of probability measures on Σ^* . By assumption, there exists a pooling operator $F^* : (\mathcal{P}^*)^n \rightarrow \mathcal{P}^*$, $(P_1^*, \dots, P_n^*) \mapsto P_{P_1^*, \dots, P_n^*}^*$ 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 5 be extended to one on Σ ; call it $P_{P_1^*, \dots, P_n^*}^*|_{\Sigma}$. Now define a pooling operator $F : \mathcal{P}^n \rightarrow \mathcal{P}$, $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ 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 operator (of course there is one, for instance dictatorship by individual 1, given by $(P_1, \dots, P_n) \mapsto P_1$), and call its decision rule $D : [0, 1]^n \rightarrow [0, 1]$. 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

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

We now transform this neutral pooling operator into a non-neutral one (that is still independent and implication-preserving). To do so, we consider a func-

tion $T : [0, 1] \rightarrow [0, 1]$ 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 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 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 just defined pooling operator $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ is independent, implication-preserving, but not neutral.

Independence. Applied to any event of type $A_j \in X$, independence holds with decision rule D_{A_j} defined as $T \circ D$, by (2). Applied to any event of type $A_j^c \in X$, independence holds with decision rule method $D_{A_j^c}$ 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 decision rule $D_A : [0, 1]^n \rightarrow [0, 1]$. We show that the rules 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 decision rule of 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. ■

6 Characterisation of linear pooling

While for non-nested agendas a pooling operator with our two properties must be neutral, it need not be linear. We now identify exactly for which non-nested agendas linearity (rather than just neutrality) follows.

A set of events $Y \subseteq \Sigma$ is *minimal inconsistent* if $\bigcap_{A \in Y} A = \emptyset$ but $\bigcap_{A \in Y'} A \neq \emptyset$ for every proper subset $Y' \subsetneq Y$. For instance, the set of events {"it rains", "if it rains we get wet", "we do not get wet"} is minimal inconsistent: these three events are mutually inconsistent, but any two of them are mutually consistent. The logical interrelations within the agenda X are perhaps best understood in terms of the minimal inconsistent sets $Y \subseteq X$. Trivial examples of minimal inconsistent sets are those of type $\{A, A^c\} \subseteq X$ (with $A \neq \emptyset, \Omega$). Most interesting agendas Y contain other minimal inconsistent sets Y with $|Y| \geq 3$. One might regard $\sup_{Y \subseteq X: Y \text{ is minimal inconsistent}} |Y|$ as a measure of the complexity of the (interconnections within the) agenda X at hand.

We call the agenda X *non-simple* if there is a minimal inconsistent set $Y \subseteq X$ containing more than two (but not uncountably many⁵) events, and *simple* otherwise.

Theorem 3 *For a non-simple agenda, an implication-preserving pooling operator is independent if and only if it is linear.*

As in our earlier neutrality characterisation, the agenda assumption is tight in the finite case:

⁵The countability condition can often be dropped because all minimal inconsistent sets $Y \subseteq X$ are automatically countable or even finite. This is so not only if X is finite or countably infinite, but 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.

Theorem 4 For a simple agenda, finite and not $\{\emptyset, \Omega\}$, there exists an implication-preserving pooling operator that is independent but not linear.

7 Proof of the linearity characterisation

We now prove Theorems 3 and 4, using above theorems and lemmas.

Proof of Theorem 3. Let X be non-simple, and F implication-preserving. We write $\mathbf{0}$ and $\mathbf{1}$ for the n -tuples $(0, \dots, 0)$ and $(1, \dots, 1)$, respectively.

Obviously, if F is linear then F is independent.

Now suppose F is independent; we show linearity. By Theorem 1 and since non-simple agendas are non-nested, F is neutral, say with decision rule $D : [0, 1]^n \rightarrow [0, 1]$ for all events $A \in X$.

1. In this part of the proof we derive some properties of D .

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

To show this, 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.

Claim 2. $D(\mathbf{0}) = 0$ and $D(\mathbf{1}) = 1$.

Indeed, since the pooling operator is implication-preserving and hence zero-preserving, $D(\mathbf{0}) = 0$, so that by Claim 1 $D(\mathbf{1}) = 1 - D(\mathbf{0}) = 1$.

Claim 3. If

$$D(x) + D(y) + D(z) = 1 \text{ for all } x, y, z \in [0, 1]^n \text{ with } x + y + z = \mathbf{1}, \quad (3)$$

then F is linear.

To show this, suppose (3). 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 (3) and the second from Claim 1. So

$$D(x + y) = D(x) + D(y) \text{ for all } x, y \in [0, 1]^n \text{ with } x + y \in [0, 1]^n. \quad (4)$$

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 the " t " occurs at position i in " $(0, \dots, 0, t, 0, \dots, 0)$ ". By (4), 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; whence 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 (4)), 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 Claim 2. So F is a linear pooling operator, as desired.

2. In this second (longer) part of the proof we ultimately show that (3) holds, which by Claim 3 completes the proof. So 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,

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

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 4-8), which then allow us to show that $D(x) + D(y) + D(z) = 1$ (Claim 9), as desired.

Claim 4. $P(\bigcap_{D \in Y \setminus \{A, B, C\}} D) = 1$.

For all $D \in Y \setminus \{A, B, C\}$ we have $D \supseteq A^* \cup B^* \cup C^*$, so that by (5) we have $P_1(D) = \dots = P_n(D) = 1$, and hence $P(D) = 1$ by Lemma 1. This implies Claim 4 because the intersection of countably many events of probability one has probability one.

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

As $A \cap B \cap C$ is disjoint from the event $\bigcap_{D \in Y \setminus \{A, B, C\}} D$, which by Claim 4 has P -probability one, we have $P(A \cap B \cap C) = 0$. This implies Claim 5 because $A^c \cup B^c \cup C^c$ is the complement of $A \cap B \cap C$.

Claim 6. $P((A^c \cap B^c) \cup (A^c \cap C^c) \cup (B^c \cap C^c)) = 0$.

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 (5). So $P_i(A^c \cap B^c) = 0$ for all i . Hence $P(A^c \cap B^c) = 0$ by Lemma 1(b). For analogous reasons, $P(A^c \cap C^c) = 0$ and $P(B^c \cap C^c) = 0$. Now Claim 6 follows since the

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

Claim 7. $P((A^c \cap B \cap C) \cup (A \cap B^c \cap C) \cup (A \cap B \cap C^c)) = 1$

By Claims 5 and 6, 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 7 states.

Claim 8. $P(A^*) + P(B^*) + P(C^*) = P(A^* \cup B^* \cup C^*) = 1.$

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 4 and 7. So $P(A^* \cup B^* \cup C^*) = 1$, as desired.

Claim 9. $D(x) + D(y) + D(z) = 1$ (which completes the proof by Claim 3).

As $P(A^* \cup B^* \cup C^*) = 1$ by Claim 8, and as the intersection of A^c with $A^* \cup B^* \cup C^*$ is A^* , we have

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

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 (6) imply that $P(A^*) = D(x)$. By similar arguments, $P(B^*) = D(y)$ and $P(C^*) = D(z)$. So Claim 9 follows from Claim 8. ■

Proof of Theorem 4. Let the agenda $X (\subseteq \Sigma)$ be simple, finite, and not $\{\emptyset, \Omega\}$. We construct a non-linear pooling operator 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 2 (proved using Lemma 5) holds analogously here as well.

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

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

and we 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 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 $X_{= 1/2}^1 = \emptyset$, 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 $\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$ and $\cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$ are non-empty by Claim 1, there exist worlds $\omega^1 \in \cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$ and $\omega^2 \in \cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$, where we assume that $\omega^1 = \omega^2$ if $X_{=1/2} = \emptyset$, i.e. if $\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C = \cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C = \cap_{C \in X_{>1/2}} C$. (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 operator $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ has the desired properties, as shown in the next two claims.

Independence. We show that the pooling operator is neutral (hence independent) with the decision rule 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 $\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$, as this set contains ω^1 and is (by Claim 2) an atom of Σ . So $A^c \cap [\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C] = \emptyset$. 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 operator 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 Lemma 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 $A \supseteq \bigcap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$ and $B \supseteq \bigcap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$, 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. ■

8 References

Aczél, J. (1966) *Lectures on Functional Equations and their Applications*, Dover Publications Inc., United States of America

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

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

McConway (1981) “Marginalization and Linear Opinion Pools”, *Journal of 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 (2002) Strategy-proof social choice on single-peaked domains: possibility, impossibility and the space between, working paper, Univ California at Davies

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 Results without assuming implication-preservation

How far do we get if we replace implication-preservation by the weaker condition of zero-preservation? For particular agendas, we still obtain neutrality of the pooling operator, but not linearity; this is now shown. Our two theorems suggest that without invoking implication-preservation it is hard to defend linearity but

still possible to defend neutrality – provided the relevant events are (at least) indirectly connected, in a sense to be defined now.

A relevant event $A \in X$ *conditionally entails* another one $B \in X$ (written $A \vdash^* B$) if $\{A\} \cup Y$ entails B (i.e. $\bigcap_{C \in \{A\} \cup Y} C \subseteq B$) for some countable set $Y \subseteq X$ that is consistent with A (i.e. $\bigcap_{C \in \{A\} \cup Y} C \neq \emptyset$) and with B^c (i.e. $\bigcap_{C \in \{B^c\} \cup Y} C \neq \emptyset$). The agenda X is *pathconnected* if for any two events $A, B \in X \setminus \{\emptyset, \Omega\}$ there exist events $A_1, \dots, A_k \in X$ ($k \geq 1$) such that $A = A_1 \vdash^* A_2 \vdash^* \dots \vdash^* A_k = B$. In other words, any two contingent events in the agenda can be connected by a path of conditional entailments.⁶ For instance, $X := \{A, A^c : A \subseteq \mathbf{R} \text{ is a bounded interval}\}$ is a pathconnected agenda (a subset of the Borel- σ -algebra Σ over $\Omega = \mathbf{R}$).⁷ One easily shows that pathconnected agendas are non-simple; but many non-simple agendas are not pathconnected.

We now give a characterisation of neutral pooling based on requiring zero-preservation rather than implication-prerservation.

Theorem 5 *For a pathconnected agenda X , a zero-preserving pooling operator is independent if and only if it is neutral.*

In this theorem, independence leads to neutrality. Does it even lead to linearity? The answer is negative, as the next theorem shows.

Theorem 6 *For some pathconnected agenda X (in some σ -algebra Σ over some set of worlds Ω) there exists a zero-preserving pooling operator that is neutral but not linear.*

The following lemma is central for proving Theorem 5.

Lemma 6 *For any independent and zero-preserving pooling operator, $A \vdash^* B$ implies $D_A \leq D_B$ for all relevant events $A, B \in X$ (where D_A and D_B are decision rules for A and B , respectively).*

Proof. Let F, A, B, D_A, D_B be as specified, and assume $A \vdash^* B$, say in virtue of the set $Y \subseteq X$. Let $x = (x_1, \dots, x_n) \in [0, 1]^n$. We show that $D_A(x) \leq D_B(x)$. As $\bigcap_{C \in \{A\} \cup Y} C$ has empty intersection with B^c (by the conditional entailment), it equals its intersection with B ; in particular, $\bigcap_{C \in \{A, B\} \cup Y} C \neq \emptyset$. Similarly, as $\bigcap_{C \in \{B^c\} \cup Y} C$ has empty intersection with A , it equals its intersection with A^c ; in particular, $\bigcap_{C \in \{A^c, B^c\} \cup Y} C \neq \emptyset$. Hence there are worlds $\omega \in \bigcap_{C \in \{A, B\} \cup Y} C$ and

⁶Pathconnectedness is closely related (not identical) to the *total blockedness* condition used in a different context by Nehring and Puppe (2002).

⁷For example, a path of conditional entailments between the intervals $[0, 1]$ and $[2, 3]$ can be constructed as follows: $[0, 1] \vdash^* [0, 3]$ (one may conditionalise on the empty set of events $Y = \emptyset$, i.e. the entailment is *unconditional*), and $[0, 3] \vdash^* [2, 3]$ (one may conditionalise on $Y = \{[2, 4]\}$).

$\omega' \in \cap_{C \in \{A^c, B^c\} \cup Y} C$. For each individual i , consider the probability measure $P_i : \Sigma \rightarrow [0, 1]$ defined by

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

where $\delta_\omega, \delta_{\omega'} : \Sigma \rightarrow [0, 1]$ denote the Dirac-measures in ω and ω' , respectively. 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 zero-preservation; see Lemma 1), and hence $P_{P_1, \dots, P_n}(\cap_{C \in Y} C) = 1$ since the intersection of *countably* many events of probability one has again probability one. So

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

Now $P_{P_1, \dots, P_n}(\cap_{C \in \{A\} \cup Y} C) \leq P_{P_1, \dots, P_n}(\cap_{C \in \{B\} \cup Y} C)$ since $\cap_{C \in \{A\} \cup Y} C = \cap_{C \in \{A, B\} \cup Y} C \subseteq \cap_{C \in \{B\} \cup Y} C$ (for the equality, see an earlier argument). So $D_A(x) \leq D_B(x)$, as desired. ■

Proof of Theorem 5. Let X be pathconnected and F zero-preserving. Obviously, if F is neutral then it is independent. Now let F be independent. If $X = \{\emptyset, \Omega\}$, F is obviously neutral, as desired. Now let $X \neq \{\emptyset, \Omega\}$ and write D_A for the decision rule of any contingent event $A \in X \setminus \{\emptyset, \Omega\}$. As X is pathconnected, repeated application of Lemma 6 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 decision rule D_A of all $A \in X \setminus \{\emptyset, \Omega\}$. We complete the neutrality proof by showing that D also works as a decision rule for \emptyset and Ω . Consider any $P_1, \dots, P_n \in \mathcal{P}$. By definition of probability measures,

$$\begin{aligned} P_1(\emptyset) &= \dots = P_n(\emptyset) = P_{P_1, \dots, P_n}(\emptyset) = 0, \\ P_1(\Omega) &= \dots = P_n(\Omega) = P_{P_1, \dots, P_n}(\Omega) = 1. \end{aligned}$$

So it suffices to show that $D(0, \dots, 0) = 0$ and $D(1, \dots, 1) = 1$, which follows from zero-preservation (see also Lemma 1). ■

Proof of Theorem 6. Our counterexample uses the state space $\Omega := \{\omega_1, \omega_2, \omega_3, \omega_4\}$ (with pairwise distinct ω_k 's), the σ -algebra $\Sigma := \{A : A \subseteq \Omega\}$ (the power set of Ω), and the agenda $X := \{A \subseteq \Omega : |A| = 2\}$ (the set of binary events). As X is negation-closed and non-empty, it is indeed an agenda.

1. In this part of the proof, we show that X is pathconnected. Consider any events $A, B \in X$. We construct a path from A to B , by distinguishing 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 $Y = \{\{\omega_B, \omega'_A\}\}$).

2. In this part, we construct a pooling operator $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ that is zero-preserving, neutral, but not linear. As an ingredient to the construction, consider first a linear pooling operator $L : \mathcal{P}^n \rightarrow \mathcal{P}$. We show that L can be transformed into a non-linear pooling operator that is still neutral and zero-preserving. We use an (arbitrary) 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 $x \in [0, 1]$.)

We prove that for every probability measure $Q \in \mathcal{P}$ (thought of as the outcome of applying the linear pooling operator L) there exist real numbers $p_k = p_k^Q$, $k = 1, 2, 3, 4$ (thought of as the new probabilities of the states ω_k , $k = 1, 2, 3, 4$, after transforming Q) 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 $A \in X$, $\sum_{k: \omega_k \in A} p_k = T(Q(A))$.

This completes the proof, because by (a) a pooling operator $F : \mathcal{P}^n \rightarrow \mathcal{P}$, $(P_1, \dots, P_n) \mapsto P_{P_1, \dots, P_n}$ can be defined by letting

$$P_{P_1, \dots, P_n}(A) := \sum_{k: \omega_k \in A} p_k^{L(P_1, \dots, P_n)} \text{ for all } A \in \Sigma,$$

which by (b) satisfies

$$P_{P_1, \dots, P_n}(A) = T(L(P_1, \dots, P_n)(A)) \text{ for all } A \in X,$$

implying that F is neutral (as L is neutral), zero-preserving (as L is zero-preserving and $T(0) = 0$), and non-linear (as L is linear and T a non-linear transformation).

Let $Q \in \mathcal{P}^n$. For any $k \in \{1, 2, 3, 4\}$, put $q^k := Q(\{\omega_k\})$; and for any $k, l \in \{1, 2, 3, 4\}$, $k < l$, put $q_{kl} = Q(\{\omega_k, \omega_l\})$.

In order for 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, suppose that $k, l \in \{1, 2, 3, 4\}$, $k < l$, and define $k', l' \in \{1, 2, 3, 4\}$, $k' < l'$, by $\{k', l'\} = \{1, 2, 3, 4\} \setminus \{k, l\}$. By $p_k + p_l = 1 - p_{k'} - p_{l'}$ 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 $(p_1, p_2, p_3) \in \mathbf{R}^3$. Write $t_{kl} := T(q_{kl})$ for all $k, l \in \{1, 2, 3, 4\}$, $k < l$.

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

So we have

$$\begin{aligned} p_3 &= \frac{t_{23} + t_{13} - t_{12}}{2}, \\ p_2 &= t_{12} - t_{13} + \frac{t_{23} + t_{13} - t_{12}}{2} = \frac{t_{12} + t_{23} - t_{13}}{2}, \\ p_1 &= t_{12} - \frac{t_{12} + t_{23} - t_{13}}{2} = \frac{t_{12} + t_{13} - t_{23}}{2}, \\ p_4 &= 1 - (p_1 + p_2 + p_3) = 1 - \frac{t_{12} + t_{13} + t_{23}}{2}. \end{aligned}$$

We have to show that the numbers p_1, \dots, p_4 so-defined satisfy not only (b) and $p_1 + \dots + p_4 = 1$ but also the remaining condition in (a), i.e. non-negativity. 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 $\geq 1/2$. 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 $< 1/2$. 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 $< 1/2$. 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 gives us

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

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 $\delta = q^1$ 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 $y := 1 - q^1 - q^2$. As also $y < 1/2$, 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 operator constructed in the proof of Theorem 6 violates implication-preservation – which it must do since Theorem 3 tells us that implication-preserving independent pooling operators must be linear (for non-simple, hence in particular for pathconnected agendas). Let Ω, Σ, 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 would require the collective probability of ω_1 to be 0 too. 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 operator, 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 $T(0) = 0$.