Representing and Aggregating Conflicting Beliefs

Pedrito Maynard-Reid II
Department of Computer Science
Stanford University
Stanford, CA 94305, USA
pedmayn@cs.stanford.edu

Daniel Lehmann
School of Computer Science and Engineering
Hebrew University
Jerusalem 91904, Israel
lehmann@cs.huji.ac.il

Abstract

We consider the two-fold problem of representing collective beliefs and aggregating these beliefs. We propose modular, transitive relations for collective beliefs. They allow us to represent conflicting opinions and they have a clear semantics. We compare them with the quasi-transitive relations often used in Social Choice. Then, we describe a way to construct the belief state of an agent informed by a set of sources of varying degrees of reliability. This construction circumvents Arrow's Impossibility Theorem in a satisfactory manner. Finally, we give a simple set-theory-based operator for combining the information of multiple agents. We show that this operator satisfies the desirable invariants of idempotence, commutativity, and associativity, and, thus, is well-behaved when iterated, and we describe a computationally effective way of computing the resulting belief state.

Keywords: representation of beliefs, multi-agent systems

1 Introduction

We are interested in the multi-agent setting where agents are informed by sources of varying levels of reliability, and where agents can iteratively combine their belief states. This setting introduces three problems: (1) Finding an appropriate representation for collective beliefs; (2) Constructing an agent's belief state by aggregating the information from informant sources, accounting for the relative reliability of these sources; and, (3) Combining the information of multiple agents in a manner that is well-behaved under iteration.

The Social Choice community has dealt extensively with the first problem (although in the context of representing collective preferences rather than beliefs) (cf. (Sen 1986)). The classical approach has been to use quasi-transitive relations (of which total pre-orders are a special subclass) over the set of possible worlds. However, these relations do not distinguish between group indifference and group conflict, and this distinction can be crucial. Consider, for example, a situation in which all members of a group are indifferent between movie a and movie b. If some passerby expresses a preference for a, the group may very well choose to adopt this opinion for the group and borrow a. However, if the group was already divided over the relative merits of a and b, we would be wise to hesitate before choosing one over the other just because a new supporter of a appears on the scene. We propose a representation in which the distinction is explicit. We also argue that our representation solves some of the unpleasant semantical problems suffered by the earlier approach.

The second problem addresses how an agent should actually go about combining the information received from a set of sources to create a belief state. Such a mechanism should favor the opinions held by more reliable sources, yet allow less reliable sources to voice opinions when higher ranked sources have no opinion. True, under some circumstances it would not be advisable for an opinion from a less reliable source to override the agnosticism of a more reliable source, but often it is better to accept these opinions as default assumptions until better information is available. (Maynard-Reid II and Shoham 2000) provides a solution to this problem when belief states are represented as total pre-orders, but runs into Arrow's Impossibility Theorem (Arrow 1963) when there are sources of equal reliability. As we shall see, the generalized representation allows us to circumvent this limitation.

To motivate the third problem, consider the following dynamic scenario: A robot controlling a ship in space receives from a number of communication centers on Earth information about the status of its environment and tasks. Each center receives information from a group of sources of varying credibility or accuracy (e.g., nearby satellites and experts) and aggregates it. Timeliness of decision-making in space is often crucial, so we do not want the robot to have to wait while each center sends its information to some central location for it to be first combined before being forwarded to the robot. Instead, each center sends its aggregated information directly to the robot. Not only does this scheme reduce dead time, it also allows for "anytime" behavior on the robot's part: the robot incorporates new information as it arrives and makes the best decisions it can with whatever information it has at any given point. This distributed approach is also more robust since the degradation in performance is much more graceful should information from individual centers get lost or delayed.

In such a scenario, the robot needs a mechanism for combining or *fusing* the belief states of multiple agents potentially arriving at different times. Moreover, the belief state output by the mechanism should be invariant with respect to the order of agent arrivals. We will describe such a mechanism.

The paper is organized as follows: After some preliminary definitions and a discussion of the approach to aggregation taken in classical Social Choice, we introduce modular, transitive relations for representing generalized belief states. We then describe how to construct the belief state of an agent given the belief states of its informant sources when these sources are totally pre-ordered. Finally, we describe a simple set-theory-based operator for fusing agent belief states that satisfies the desirable invariants of idempotence, commutativity, and associativity, and we describe a computationally effective way of computing this belief state.

2 Preliminaries

We begin by defining various well-known properties of binary relations[1]; they will be useful to us throughout the paper.

Definition 1 *Suppose* $\leq$ *is a relation over a finite set* Ω*, i.e.,* $\leq \subseteq \Omega \times \Omega$*. We shall use* $x \leq y$ *to denote* $(x,y) \in \leq$ *and* $x \not\leq y$ *to denote* $(x,y) \notin \leq$*. The relation* $\leq$ *is:*

1. reflexive *iff* $x \leq x$ *for* $x \in \Omega$*. It is* irreflexive *iff* $x \not\leq x$ *for* $x \in \Omega$*.*
2. symmetric *iff* $x \leq y \Rightarrow y \leq x$ *for* $x, y \in \Omega$*. It is* asymmetric *iff* $x \leq y \Rightarrow y \not\leq x$ *for* $x, y \in \Omega$*. It is* anti-symmetric *iff* $x \leq y \wedge y \leq x \Rightarrow x = y$ *for* $x, y \in \Omega$*.*
3. *the* strict version *of a relation* $\leq'$ *over* Ω *iff* $x \leq y \Leftrightarrow x \leq' y \wedge y \not\leq' x$ *for* $x, y \in \Omega$*.*
4. total *iff* $x \leq y \vee y \leq x$ *for* $x, y \in \Omega$*.*
5. modular *iff* $x \leq y \Rightarrow x \leq z \vee z \leq y$ *for* $x, y, z \in \Omega$*.*
6. transitive *iff* $x \leq y \wedge y \leq z \Rightarrow x \leq z$ *for* $x, y, z \in \Omega$*.*
7. quasi-transitive *iff its strict version is transitive.*
8. *the* transitive closure *of a relation* $\leq'$ *over* Ω *iff* $x \leq y \Leftrightarrow \exists w_0, \ldots, w_n \in \Omega.\ x = w_0 \leq' \cdots \leq' w_n = y$ *for some integer* n*, for* $x, y \in \Omega$*.*
9. acyclic *iff* $\forall w_0, \ldots, w_n \in \Omega.\ w_0 < \cdots < w_n$ *implies* $w_n \not< w_0$ *for all integers* n*, where* $<$ *is the strict version of* $\leq$*.*
10. *a* total pre-order *iff it is total and transitive. It is a* total order *iff it is also anti-symmetric.*
11. *an* equivalence relation *iff it is reflexive, symmetric, and transitive.*

Proposition 1

1. *The transitive closure of a modular relation is modular.*
2. *Every transitive relation is quasi-transitive.*
3. *(Sen 1986) Every quasi-transitive relation is acyclic.*

Given a relation over a set of alternatives and a subset of these alternatives, we often want to pick the subset's "best" elements with respect to the relation. We define this set of "best" elements to be the subset's *choice set*:

Definition 2 *If* $\leq$ *is a relation over a finite set* Ω*,* $<$ *is its strict version, and* $X \subseteq \Omega$*, then the* choice set of X with respect to $\leq$ *is*

$$C(X, \leq) = \{x \in X : \nexists x' \in X.\ x' < x\}.$$

A *choice function* is one which assigns to every subset X a non-empty subset of X:

[1] We only use binary relations in this paper, so we will refer to them simply as relations.

Definition 3 *A* choice function *over a finite set* Ω *is a function* $f : 2^\Omega \setminus \emptyset \to 2^\Omega \setminus \emptyset$ *such that* $f(X) \subseteq X$ *for every* $X \subseteq \Omega$.

Now, every acyclic relation defines a choice function, one which assigns to each subset its choice set:

Proposition 2 *(Sen 1986) Given a relation* $\leq$ *over a finite set* Ω*, the choice set operation* C *defines a choice function iff* $\leq$ *is acyclic.*[2]

If a relation is not acyclic, elements involved in a cycle are said to be in a *conflict* because we cannot order them:

Definition 4 *Given a relation* $<$ *over a finite set* Ω*,* x *and* y *are in a* conflict wrt $<$ *iff there exist* $w_0, \ldots, w_n, z_0, \ldots, z_m \in \Omega$ *such that* $x = w_0 < \cdots < w_n = y = z_0 < \cdots < z_m = x$, *where* $x, y \in \Omega$.

3 Aggregation in Social Choice

We are interested in belief aggregation, but the community historically most interested in aggregation has been that of Social Choice theory. The aggregation is over preferences rather than beliefs, so the discussion in this subsection will focus on representing preferences; however, as we shall see, the results are equally relevant to representing beliefs. In the Social Choice community, the standard representation of an agent's preferences is a total pre-order. Each total pre-order $\preceq_i$ is interpreted as describing the weak preferences of an individual i, so that $x \preceq_i y$ means i considers alternative x to be at least as preferable as alternative y.[3] If $x \preceq_i y$ and $y \preceq_i x$, then i is indifferent between x and y.

Unfortunately, Arrow's Impossibility Theorem (Arrow 1963) showed that no aggregation operator over total pre-orders exists satisfying the following small set of desirable properties:

Definition 5 *Let* f *be an aggregation operator over the preferences* $\preceq_1, \ldots, \preceq_n$ *of* n *individuals, respectively, over a finite set of alternatives* Ω*, and let* $\preceq = f(\preceq_1, \ldots, \preceq_n)$.

- Restricted Range*: The range of* f *is the set of total pre-orders over* Ω.
- Unrestricted Domain*: The domain of* f *is the set of* n*-tuples of total pre-orders over* Ω.
- Pareto Principle*: If* $x \prec_i y$ *for all* i*, then* $x \prec y$.
- Independence of Irrelevant Alternatives (IIA)*: Suppose* $\preceq' = f(\preceq'_1, \ldots, \preceq'_n)$*. If, for* $x, y \in \Omega$*,* $x \preceq_i y$ *iff* $x \preceq'_i y$ *for all* i*, then* $x \preceq y$ *iff* $x \preceq' y$.
- Non-Dictatorship*: There is no individual* i *such that, for every tuple in the domain of* f *and every* $x, y \in \Omega$*,* $x \prec_i y$ *implies* $x \prec y$.

Proposition 3 *(Arrow 1963) There is no aggregation operator that satisfies restricted range, unrestricted domain, (weak) Pareto principle, independendence of irrelevant alternatives, and nondictatorship.*

This impossibility theorem led researchers to look for weakenings to Arrow's framework that would circumvent the result. One was to weaken the restricted range condition, requiring that the result of an aggregation only satisfy totality and quasi-transitivity rather than the full transitivity of a total pre-order. This weakening was sufficient to guarantee the existence of an aggregation function satisfying the other conditions, while still producing relations that defined choice functions (Sen 1986). However, this solution was not without its own problems.

First, total, quasi-transitive relations have unsatisfactory semantics. If $\preceq$ is total and quasi-transitive but not a total pre-order, its indifference relation is not transitive:

Proposition 4 *Let* $\preceq$ *be a relation over a finite set* Ω *and let* $\sim$ *be its symmetric restriction (i.e.,* $x \sim y$ *iff* $x \preceq y$ *and* $y \preceq x$*). If* $\preceq$ *is total and quasi-transitive but not transitive, then* $\sim$ *is not transitive.*

There has been much discussion as to whether or not indifference should be transitive; in many cases one feels indifference should be transitive. If Deb enjoys plums and mangoes equally and also enjoys mangoes and peaches equally, we would conclude that she also enjoys plums and peaches equally. It seems that total quasi-transitive relations that are not total pre-orders cannot be understood easily as preference or indifference.

Since the existence of a choice function is generally sufficient for classical Social Choice problems, this issue was at least ignorable. However, in iterated aggregation, the result of the aggregation must not only be usable for making decisions, but must be interpretable as a new preference relation that may be involved in later

[2] Sen's uses a slightly stronger definition of choice sets, but the theorem still holds in our more general case.

[3] The direction of the relation symbol is unintuitive, but standard practice in the belief revision community.

aggregations; consequently, it must maintain clean semantics.

Secondly, the totality assumption is excessively restrictive for representing aggregate preferences. In general, a binary relation $\preceq$ can express four possible relationships between a pair of alternatives a and b: $a \preceq b$ and $b \not\preceq a$, $b \preceq a$ and $a \not\preceq b$, $a \preceq b$ and $b \preceq a$, and $a \not\preceq b$ and $b \not\preceq a$. Totality reduces this set to the first three which, under the interpretation of relations as representing weak preference, correspond to the two strict orderings of a and b, and indifference. However, consider the situation where a couple is trying to choose between an Italian and an Indian restaurant, but one strictly prefers Italian food to Indian food, whereas the second strictly prefers Indian to Italian. The couple's opinions are in conflict, a situation that does not fit into any of the three remaining categories. Thus, the totality assumption is essentially an assumption that conflicts do not exist. This, one may argue, is appropriate if we want to represent preferences of one agent (but see (Kahneman and Tversky 1979) for persuasive arguments that individuals *are* often ambivalent). However, the assumption is inappropriate if we want to represent aggregate preferences since individuals will almost certainly have differences of opinion.

4 Generalized Belief States

Let us turn to the domain of belief aggregation. A total pre-order over the set of possible worlds is a fairly well-accepted representation for a belief state in the belief revision community (Grove 1988; Katsuno and Mendelzon 1991; Lehmann and Magidor 1992; Gärdenfors and Makinson 1994). Instead of preference, relations represent relative likelihood, instead of indifference, equal likelihood. For the remainder of the paper, assume we are given some language $\mathcal{L}$ with a satisfaction relation $\models$ for $\mathcal{L}$. Let $\mathcal{W}$ be a finite, nonempty set of possible worlds (interpretations) over $\mathcal{L}$. Suppose $\preceq$ is a total pre-order on $\mathcal{W}$. The belief revision literature maintains that the conditional belief "if p then q" (where p and q are sentences in $\mathcal{L}$) holds if all the worlds in the choice set of those satisfying p also satisfy q; we write $Bel(p?q)$. The individual's unconditional beliefs are all those where p is the sentence *true*. If neither the belief $p?q$ nor its negation hold in the belief state, it is said to be *agnostic* with respect to $p?q$, written $Agn(p?q)$.

It should come as no surprise that belief aggregation is formally similar to preference aggregation and, as a result, is also susceptible to the problems described in the previous section. We propose a solution to these problems which generalizes the total pre-order representation so as to capture information about conflicts.

4.1 Modular, transitive states

We take strict likelihood as primitive. Since strict likelihood is not necessarily total, it is possible to represent agnosticism and conflicting opinions in the same structure. This choice deviates from that of most authors, but are similar to those of Kreps (Kreps 1990, p. 19) who is interested in representing both indifference and incomparability. Unlike Kreps, rather than use an asymmetric relation to represent strict likelihood (e.g., the strict version of a weak likelihood relation), we impose the less restrictive condition of modularity.

We formally define *generalized belief states*:

Definition 6 *A* generalized belief state $\prec$ *is a modular, transitive relation over* $\mathcal{W}$. *The set of possible generalized belief states over* $\mathcal{W}$ *is denoted* $\mathcal{B}$.

We interpret $a \prec b$ to mean "there is reason to consider a as strictly more likely than b." We represent equal likelihood, which we also refer to as "agnosticism," with the relationship $\sim$ defined such that $x \sim y$ if and only if $x \not\prec y$ and $y \not\prec x$. We define the conflict relation corresponding to $\prec$, denoted ∞, so that $x \infty y$ iff $x \prec y$ and $y \prec x$. It describes situations where there are reasons to consider either of a pair of worlds as strictly more likely than the other. In fact, one can easily check that ∞ precisely represents conflicts in a belief state in the sense of Definition 4.

For convenience, we will refer to generalized belief states simply as belief states for the remainder of the paper except when to do so would cause confusion.

4.2 Discussion

Let us consider why our choice of representation is justified. First, we agree with the Social Choice community that strict likelihood should be transitive.

As we discussed in the previous section, there is often no compelling reason why agnosticism/indifference should not be transitive; we also adopt this view. However, transitivity of strict likelihood by itself does not guarantee transitivity of agnosticism. A simple example is the following: $\prec = \{(a, c)\}$, so that $\sim = \{(a, b), (b, c)\}$. However, if we buy that strict likelihood should be transitive, then agnosticism is transitive identically when strict likelihood is also modular:

Proposition 5 *Suppose a relation* $\prec$ *is transitive and* $\sim$ *is the corresponding agnosticism relation. Then* $\sim$ *is transitive iff* $\prec$ *is modular.*

In summary, transitivity and modularity are necessary if strict likelihood and agnosticism are both required to be transitive.

We should point out that conflicts are also transitive in our framework. At first glance, this may appear undesirable: it is entirely possible for a group to disagree on the relative likelihood of worlds a and b, and b and c, yet agree that a is more likely than c. However, we note that this transitivity follows from the cycle-based definition of conflicts (Definition 4), not from our belief state representation. It highlights the fact that we are not only concerned with conflicts that arise from simple disagreements over pairs of alternatives, but those that can be inferred from a series of inconsistent opinions as well.

Now, to argue that modular, transitive relations are sufficient to capture relative likelihood, agnosticism, and conflicts among a group of information sources, we first point out that adding irreflexivity would give us the class of relations that are strict versions of total pre-orders, i.e., conflict-free. Let $\mathcal{T}$ be the set of total pre-orders over $\mathcal{W}$, $\mathcal{T}_<$, the set of their strict versions.

Proposition 6 *The set of irreflexive relations in $\mathcal{B}$ is isomorphic to $\mathcal{T}$ and, in fact, equals $\mathcal{T}_<$.*

Secondly, the following representation theorem shows that each belief state partitions the possible worlds into sets of worlds either all equally likely or all potentially involved in a conflict, and totally orders these sets; worlds in distinct sets have the same relation to each other as do the sets.

Proposition 7 $\prec \in \mathcal{B}$ *iff there is a partition* $\mathbf{W} = \langle W_0, \ldots, W_n \rangle$ *of* $\mathcal{W}$ *such that:*

1. *For every $x \in W_i$ and $y \in W_j$, $i \neq j$ implies $i < j$ iff $x \prec y$.*
2. *Every W_i is either fully connected ($w \prec w'$ for all $w, w' \in W_i$) or fully disconnected ($w \not\prec w'$ for all $w, w' \in W_i$).*

Figure 1 shows three examples of belief states: one which is a total pre-order, one which is the strict version of a total pre-order, and one which is neither.

Thus, generalized belief states are not a big change from the strict versions of total pre-orders. They merely generalize these by weakening the assumption that sets of worlds not strictly ordered are equally likely, allowing for the possibility of conflicts. Now we can distinguish between agnostic and conflicting conditional beliefs. A belief state $\prec$ is agnostic about

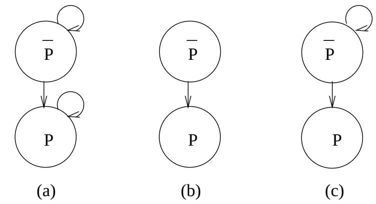

Figure 1: Three examples of generalized belief states: (a) a total pre-order, (b) the strict version of a total pre-order, (c) neither. (Each circle represents all the worlds in $\mathcal{W}$ which satisfy the sentence inside. An arc between circles indicates that $w \prec w'$ for every w in the head circle and w' in the tail circle; no arc indicates that $w \not\prec w'$ for each of these pairs. In particular, the set of worlds represented by a circle is fully connected if there is an arc from the circle to itself, fully disconnected otherwise.)

conditional belief $p?q$ (i.e., $Agn(p?q)$) if the choice set of worlds satisfying p contains both worlds which satisfy q and $\neg q$ and is fully disconnected. It is in conflict about this belief, written $Con(p?q)$, if the choice set is fully connected.

Finally, we compare the representational power of our definitions to those discussed in the previous section. First, $\mathcal{B}$ subsumes the class of total pre-orders:

Proposition 8 *$\mathcal{T} \subset \mathcal{B}$ and is the set of reflexive relations in $\mathcal{B}$.*

Secondly, $\mathcal{B}$ neither subsumes nor is subsumed by the set of total, quasi-transitive relations, and the intersection of the two classes is $\mathcal{T}$. Let $\mathcal{Q}$ be the set of total, quasi-transitive relations over $\mathcal{W}$, and $\mathcal{Q}_<$, the set of their strict versions.

Proposition 9

1. *$\mathcal{Q} \cap \mathcal{B} = \mathcal{T}$.*
2. *$\mathcal{B} \not\subseteq \mathcal{Q}$.*
3. *$\mathcal{Q} \not\subseteq \mathcal{B}$ if $\mathcal{W}$ has at least three elements.*
4. *$\mathcal{Q} \subset \mathcal{B}$ if $\mathcal{W}$ has one or two elements.*

Because modular, transitive relations represent strict preferences, it is probably fairer to compare them to the class of strict versions of total, quasi-transitive relations. Again, neither class subsumes the other, but this time the intersection is $\mathcal{T}_<$:

Proposition 10

1. *$\mathcal{Q}_< \cap \mathcal{B} = \mathcal{T}_<$.*

2. $\mathcal{B} \not\subseteq \mathcal{Q}_<$.

3. $\mathcal{Q}_< \not\subseteq \mathcal{B}$ *if* $\mathcal{W}$ *has at least three elements.*

4. $\mathcal{Q}_< \subset \mathcal{B}$ *if* $\mathcal{W}$ *has one or two elements.*

In the next section, we define a natural aggregation policy based on this new representation that admits clear semantics and obeys appropriately modified versions of Arrow's conditions.

5 Single-agent belief state construction

Suppose an agent is informed by a set of sources, each with its individual belief state. Suppose further that the agent has ranked the sources by level of credibility. We propose an operator for constructing the agent's belief state $\prec$ by aggregating the belief states of the sources in S while accounting for the credibility ranking of the sources.

Example 1 *We will use a running example from our space robot domain to help provide intuition for our definitions. The robot sends to earth a stream of telemetry data gathered by the spacecraft, as long as it receives positive feedback that the data is being received. At some point it loses contact with the automatic feedback system, so it sends a request for information to an agent on earth to find out if the failure was caused by a failure of the feedback system or by an overload of the data retrieval system. In the former case, it would continue to send data, in the latter, desist. As it so happens, there has been no overload, but the computer running the feedback system has hung. The agent consults the following three experts, aggregates their beliefs, and sends the results back to the robot:*

1. s_p*, the computer programmer that developed the feedback program, believes nothing could ever go wrong with her code, so there must have been an overload problem. However, she admits that if her program had crashed, the problem could ripple through to cause an overload.*

2. s_m*, the manager for the telemetry division, unfortunately has out-dated information that the feedback system is working. She was also told by the engineer who sold her the system that overloading could never happen. She has no idea what would happen if there was an overload or the feedback system crashed.*

3. s_t*, the technician working on the feedback system, knows that the feedback system crashed, but doesn't know whether there was a data-overload. Not being familiar with the retrieval system, she is also unable to speculate whether the data retrieval system would have overloaded if the feedback system had not failed.*

Let F *and* D *be propositional variables representing that the feedback and data retrieval systems, respectively, are okay. The belief states for the three sources are shown in Figure 2.*

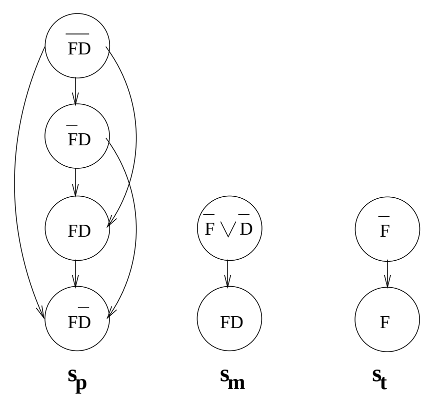

Figure 2: The belief states of s_p, s_m, and s_t in Example 1.

Let us begin the formal development by defining sources:

Definition 7 $\mathcal{S}$ *is a finite set of* sources*. With each source* $s \in \mathcal{S}$ *is associated a belief state* $<^s \in \mathcal{B}$.

We denote the agnosticism and conflict relations of a source s by $\approx^s$ and $\bowtie^s$, respectively. It is possible to assume that the belief state of a source is conflict free, i.e., acyclic. However, this is not necessary if we allow sources to suffer from the human malady of "being torn between possibilities."

We assume that the agent's credibility ranking over the sources is a total pre-order:

Definition 8 $\mathcal{R}$ *is a totally ordered finite set of* ranks.

Definition 9 $rank : \mathcal{S} \rightarrow \mathcal{R}$ *assigns to each source a rank.*

Definition 10 $\sqsupseteq$ *is the total pre-order over* $\mathcal{S}$ *induced by the ordering over* $\mathcal{R}$*. That is,* $s \sqsupseteq s'$ *iff* $rank(s) \geq rank(s')$*; we say* s' *is* as credible as s. $\sqsupseteq_S$ *is the restriction of* $\sqsupseteq$ *to* $S \subseteq \mathcal{S}$.

We use $\sqsupset$ and $\equiv$ to denote the asymmetric and symmetric restrictions of $\sqsupseteq$, respectively.[4] The finiteness of $\mathcal{S}$ ($\mathcal{R}$) ensures that a maximal source (rank) always exists, which is necessary for some of our results. Weaker assumptions are possible, but at the price of unnecessarily complicating the discussion.

We are ready to consider the source aggregation problem. In the following, assume an agent is informed by a set of sources $S \subseteq \mathcal{S}$. We look at two special cases—equal-ranked and strictly-ranked source aggregation—before considering the general case.

5.1 Equal-ranked sources aggregation

Suppose all the sources have the same rank so that $\sqsupseteq_S$ is fully connected. Intuitively, we want take all offered opinions seriously, so we take the union of the relations:

Definition 11 *If $S \subseteq \mathcal{S}$, then $Un(S)$ is the relation $\bigcup_{s \in S} <^s$.*

By simply taking the union of the source belief states, we may lose transitivity. However, we do not lose modularity:

Proposition 11 *If $S \subseteq \mathcal{S}$, then $Un(S)$ is modular but not necessarily transitive.*

Thus, we know from Proposition 1 that we need only take the transitive closure of $Un(S)$ to get a belief state:

Definition 12 *If $S \subseteq \mathcal{S}$, then $AGRUn(S)$ is the relation $Un(S)^+$.*

Proposition 12 *If $S \subseteq \mathcal{S}$, then $AGRUn(S) \in \mathcal{B}$.*

Not surprisingly, by taking all opinions of all sources seriously, we may generate many conflicts, manifested as fully connected subsets of $\mathcal{W}$.

Example 2 *Suppose all three sources in the space robot scenario of Example 1 are considered equally credible, then the aggregate belief state will be the fully connected relation indicating that there are conflicts over every belief.*

5.2 Strictly-ranked sources aggregation

Next, consider the case where the sources are strictly ranked, i.e., $\sqsupseteq_S$ is a total order. We define an operator such that lower-ranked sources refine the belief states of higher ranked sources. That is, in determining the ordering of a pair of worlds, the opinions of higher-ranked sources generally override those of lower-ranked sources, and lower-ranked sources are consulted when higher-ranked sources are agnostic:

Definition 13 *If $S \subseteq \mathcal{S}$, then $AGRRf(S)$ is the relation*
$$\left\{(x, y) : \exists s \in S.\ x <^s y \wedge \left(\forall s' \sqsupset s \in S.\ x \approx^{s'} y\right)\right\}.$$

The definition of the $AGRRf$ operator does not rely on $\sqsupseteq_S$ being a total order, and we will use it in this more general setting in the following sub-section. However, in the case that $\sqsupseteq_S$ is a total order, the result of applying $AGRRf$ is guaranteed to be a belief state.

Proposition 13 *If $S \subseteq \mathcal{S}$ and $\sqsupseteq_S$ is a total order, then $AGRRf(S) \in \mathcal{B}$.*

Example 3 *Suppose, in the space robot scenario of Example 1, the technician is considered more credible than the manager who, in turn, is considered more credible than the programmer. The aggregate belief state, shown in Figure 3, informs the robot correctly that the feedback system has crashed, but that it shouldn't worry about an overload problem and should keep sending data.*

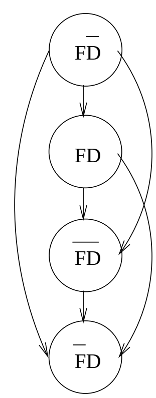

Figure 3: The belief state after aggregation in Example 3 when $s_t \sqsupset s_m \sqsupset s_p$.

Note that this case of strictly-ranked sources is almost exactly that considered in (Maynard-Reid II and Shoham 2000), except that the authors are not able to allow for conflicts in belief states. A surprising result they show is that standard AGM belief revision (Alchourrón *et al.* 1985) can be modeled as the aggregation of two sources, the informant and the informee, where the informant is considered more credible than the informee.

[4] Note that, unlike the relations representing belief states, $\geq$ and $\sqsupseteq$ are read in the intuitive way, that is, "greater" corresponds to "better."

5.3 General aggregation

In the general case, we may have several ranks represented and multiple sources of each rank. It will be instructive to first consider the following seemingly natural strawman operator, AGR^*: First combine equi-rank sources using $AGRUn$, then aggregate the strictly-ranked results using what is essentially $AGRRf$:

Definition 14 *Let $S \subseteq \mathcal{S}$. For any $r \in \mathcal{R}$, let $<_r = AGRUn(\{s \in S : rank(s) = r\})$ and $\approx_{r'}$, the corresponding agnosticism relation. Also, let $rank(S) = \{r \in \mathcal{R} : \exists s \in S.\ rank(s) = r\}$. $AGR^*(S)$ is the relation*

$$\left\{ (x,y) : \begin{array}{l} \exists r \in \mathcal{R}.\ x <_r y \wedge \\ (\forall r' > r \in ranks(S).\ x \approx_{r'} y) \end{array} \right\}$$

AGR^* indeed defines a legitimate belief state:

Proposition 14 *If $S \subseteq \mathcal{S}$, then $AGR^*(S) \in \mathcal{B}$.*

Unfortunately, a problem with this "divide-and-conquer" approach is it assumes the result of aggregation is independent of potential interactions between the individual sources of different ranks. Consequently, opinions that will eventually get overridden may still have an indirect effect on the final aggregation result by introducing superfluous opinions during the intermediate equi-rank aggregation step, as the following example shows:

Example 4 *Let $\mathcal{W} = \{a, b, c\}$. Suppose $S \subseteq \mathcal{S}$ such that $S = \{s_0, s_1, s_2\}$ with belief states $<^{s_0} = \{(b,a),(b,c)\}$ and $<^{s_1} = <^{s_2} = \{(a,b),(c,b)\}$, and where $s_2 \sqsupset s_1 \equiv s_0$. Then $AGR^*(S)$ is $\{(a,b),(c,b),(a,c),(c,a),(a,a),(b,b),(c,c)\}$. All sources are agnostic over a and c, yet (a,c) and (c,a) are in the result because of the transitive closure in the lower rank involving opinions ((b,c) and (b,a)) which actually get overridden in the final result.*

Because of these undesired effects, we propose another aggregation operator which circumvents this problem by applying refinement (as defined in Definition 13) to the set of source belief states before infering new opinions via closure:

Definition 15 *The* rank-based aggregation *of a set of sources $S \subseteq \mathcal{S}$ is $AGR(S) = AGRRf(S)^+$.*

Encouragingly, AGR outputs a valid belief state:

Proposition 15 *If $S \subseteq \mathcal{S}$, then $AGR(S) \in \mathcal{B}$.*

Example 5 *Suppose, in the space robot scenario of Example 1, the technician is still considered more credible than the manager and the programmer, but the latter two are considered equally credible. The aggregate belief state, shown in Figure 5, still gives the robot the correct information about the state of the system. The robot also learns for future reference that there is some disagreement over whether or not there would have been a data overload if the feedback system were working.*

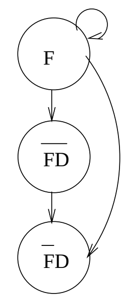

Figure 4: The belief state after aggregation in Example 5 when $s_t \sqsupset s_m \equiv s_p$.

We observe that AGR, when applied to the set of sources in Example 4, does indeed bypass the problem described above of extraneous opinion introduction:

Example 6 *Assume $\mathcal{W}$, S, and $\sqsupseteq$ are as in Example 4. $AGR(S) = \{(a,b),(c,b)\}$.*

We also observe that AGR behaves well in the special cases we've considered, reducing to $AGRUn$ when all sources have equal rank, and to $AGRRf$ when the sources are totally ranked:

Proposition 16 *Suppose $S \subseteq \mathcal{S}$.*

1. *If $\sqsupseteq_S$ is fully connected, $AGR(S) = AGRUn(S)$.*
2. *If $\sqsupseteq_S$ is a total order, $AGR(S) = AGRRf(S)$.*

5.4 Arrow, revisited

Finally, a strong argument in favor of AGR is that it satisfies appropriate modifications of Arrow's conditions. Let f be an operator which aggregates the belief states $<^{s_1}, \ldots, <^{s_n}$ over $\mathcal{W}$ of n sources $s_1, \ldots, s_n \in \mathcal{S}$, respectively, and let $\prec = f(<^{s_1}, \ldots, <^{s_n})$. We consider each condition separately.

Restricted range The output of the aggregation function will be a modular, transitive belief state rather than a total pre-order.

Definition 16 (modified) Restricted Range*: The range of f is $\mathcal{B}$.*

Unrestricted domain Similarly, the input to the aggregation function will be modular, transitive belief states of sources rather than total pre-orders.

Definition 17 (modified) Unrestricted Domain*: For each i, $<^{s_i}$ can be any member of $\mathcal{B}$.*

Pareto principle Generalized belief states already represent strict likelihood. Consequently, we use the actual input and output relations of the aggregation function in place of their strict versions to define the Pareto principle. Obviously, because we allow for the introduction of conflicts, AGR will not satisfy the original formal Pareto principle which essentially states that if all sources have an unconflicted belief that one world is strictly more likely than another, this must also be true of the aggregated belief state. Neither condition is necessarily stronger than the other.

Definition 18 (modified) Pareto Principle*: If $x <^{s_i} y$ for all i, then $x \prec y$.*

Independence of irrelevant alternatives Conflicts are defined in terms of cycles, not necessarily binary. By allowing the existence of conflicts, we effectively have made it possible for outside worlds to affect the relation between a pair of worlds, viz., by involving them in a cycle. As a result, we need to weaken IIA to say that the relation between worlds should be independent of other worlds **unless** these other worlds put them in conflict.

Definition 19 (modified) Independence of Irrelevant Alternatives (IIA)*: Suppose $s'_1, \ldots, s'_n \in \mathcal{S}$ such that $s_i \equiv s'_i$ for all i, and $\prec' = f(<^{s'_1}, \ldots, <^{s'_n})$. If, for $x, y \in \mathcal{W}$, $x <^{s_i} y$ iff $x <^{s'_i} y$ for all i, $x \not\propto y$, and $x \not\propto' y$, then $x \prec y$ iff $x \prec' y$.*

Non-dictatorship As with the Pareto principle definition, we use the actual input and output relations to define non-dictatorship since belief states represent strict likelihood. From this perspective, our setting requires that informant sources of the highest rank be "dictators" in the sense considered by Arrow. However, the setting originally considered by Arrow was one where all individuals are ranked equally. Thus, we make this explicit in our new definition of non-dictatorship by adding the pre-condition that all sources be of equal rank. Now, AGR treats a set of equi-rank sources equally by taking all their opinions seriously, at the price of introducing conflicts. So, intuitively, there are no dictators. However, because Arrow did not account for conflicts in his formulation, all the sources will be "dictators" by his definition. We need to modify the definition of non-dictatorship to say that no source can always push opinions through without them ever being contested.

Definition 20 (modified) Non-Dictatorship*: If $s_i \equiv s_j$ for all i, j, then there is no i such that, for every combination of source belief states and every $x, y \in \mathcal{W}$, $x <^{s_i} y$ and $y \not<^{s_i} x$ implies $x \prec y$ and $y \not\prec x$.*

We now show that AGR indeed satisfies these conditions:

Proposition 17 *Let $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ and $AGR_f(<^{s_1}, \ldots, <^{s_n}) = AGR(S)$. AGR_f satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

6 Multi-agent fusion

So far, we have only considered the case where a single agent must construct or update her belief state once informed by a set of sources. Multi-agent fusion is the process of aggregating the belief states of a set of agents, each with its respective set of informant sources. We proceed to formalize this setting.

An agent A is *informed by* a set of sources $S \subseteq \mathcal{S}$. Agent A's *induced belief state* is the belief state formed by aggregating the belief states of its informant sources, i.e., $AGR(S)$. Assume the set of agents to fuse agree upon *rank* (and, consequently, $\sqsupseteq$).[5] We define the fusion of this set to be an agent informed by the combination of informant sources:

Definition 21 *Let $\mathcal{A} = \{A_1, \ldots, A_n\}$ be a set of agents such that each agent A_i is informed by $S_i \subseteq \mathcal{S}$. The* fusion *of $\mathcal{A}$, written $\circledvee(\mathcal{A})$, is an agent informed by $S = \bigcup_{i=1}^{n} S_i$.*

Not surprisingly given its set-theoretic definition, fusion is idempotent, commutative, and associative. These properties guarantee the invariance required in multi-agent belief aggregation applications such as our space robot domain.

[5]We could easily extend the framework to allow for individual rankings, but we felt that the small gain in generality would not justify the additional complexity and loss of perspicuity. Similarly, we could consider each agent as having a credibility ordering only over its informant sources. However, it is unclear how, for example, crediblity orderings over disjoint sets of sources should be combined into a new credibility ordering since their union will not be total.

In the multi-agent space robot scenario described in Section 1, we only have a direct need for the belief states that result from fusion. We are only interested in the belief states of the original sources in as far as we want the fused belief state to reflect its informant history. An obvious question is whether it is possible to compute the belief state induced by the agents' fusion solely from their initial belief states, that is, without having to reference the belief states of their informant sources. This is highly desirable because of the expense of storing—or, as in the case of our space robot example, transmitting—all source belief states; we would like to represent each agent's knowledge as compactly as possible.

In fact, we can do this if all sources have equal rank. We simply take the transitive closure of the union of the agents' belief states:

Proposition 18 *Let $\mathcal{A}$ and S be as in Definition 21, $\prec^{A_i}$, agent A_i's induced belief state, and $\sqsupset_S$, fully connected. If $A = \circledvee(\mathcal{A})$, then $\left(\bigcup_{A_i \in \mathcal{A}} \prec^{A_i}\right)^+$ is A's induced belief state.*

Unfortunately, the equal rank case is special. If we have sources of different ranks, we generally cannot compute the induced belief state after fusion using only the agent belief states before fusion, as the following simple example demonstrates:

Example 7 *Let $\mathcal{W} = \{a, b\}$. Suppose two agents A_1 and A_2 are informed by sources s_1 with belief state $<^{s_1} = \{(a,b)\}$ and s_2 with belief state $<^{s_2} = \{(b,a)\}$, respectively. A_1's belief state is the same as s_1's and A_2's is the same as s_2's. If $s_1 \sqsupset s_2$, then the belief state induced by $\circledvee(A_1, A_2)$ is $<^{s_1}$, whereas if $s_2 \sqsupset s_1$, then it is $<^{s_2}$. Thus, just knowing the belief states of the fused agents is not sufficient for computing the induced belief state. We need more information about the original sources.*

However, if sources are totally pre-ordered by credibility, we can still do much better than storing all the original sources. It is enough to store for each opinion of $AGRRf(S)$ the rank of the highest-ranked source supporting it. We define *pedigreed belief states* which enrich belief states with this additional information:

Definition 22 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$. A's* pedigreed belief state *is a pair $(\prec, l)$ where $\prec = AGRRf(S)$ and $l : \prec \rightarrow \mathcal{R}$ such that $l((x,y)) = \max\{rank(s) : x <^s y, s \in S\}$. We use $\prec_r^A$ to denote the* restriction of A's pedigreed belief state to r*, that is, $\prec_r^A = \{(x,y) \in \prec : l((x,y)) = r\}$.*

We verify that a pair's label is, in fact, the rank of the source used to determine the pair's membership in $AGRRf(S)$, not that of some higher-ranked source:

Proposition 19 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$ and with pedigreed belief state $(\prec, l)$. Then*

$$x \prec_r^A y$$

$$\textit{iff}$$

$$\exists s \in S.\ x <^s y \wedge r = rank(s) \wedge \left(\forall s' \sqsupset s \in S.\ x \approx^{s'} y\right).$$

The belief state induced by a pedigreed belief state $(\prec, l)$ is, obviously, the transitive closure of $\prec$.

Now, given only the pedigreed belief states of a set of agents, we can compute the new pedigreed belief state after fusion. We simply combine the labeled opinions using our refinement techniques.

Proposition 20 *Let $\mathcal{A}$ and S be as in Definition 21, $\sqsupset_S$, a total pre-order, and $A = \circledvee(\mathcal{A})$. If*

1. *$\prec$ is the relation*
$$\left\{(x,y) : \begin{array}{l} \exists A_i \in \mathcal{A}, r \in \mathcal{R}.\ x \prec_r^{A_i} y \wedge \\ \left(\forall A_j \in \mathcal{A}, r' > r \in \mathcal{R}.\ x \sim_{r'}^{A_j} y\right) \end{array}\right\}$$
over $\mathcal{W}$,
2. *$l : \prec \rightarrow \mathcal{R}$ such that $l((x,y)) = \max\{r : x \prec_r^{A_i} y, A_i \in \mathcal{A}\}$, and*

then $(\prec, l)$ is A's pedigreed belief state.

From the perspective of the induced belief states, we are essentially discarding unlabeled opinions (i.e., those derived by the closure operation) before fusion. Intuitively, we are learning new information so we may need to retract some of our inferred opinions. After fusion, we re-apply closure to complete the new belief state. Interestingly, in the special case where the sources are strictly-ranked, the closure is unnecessary:

Proposition 21 *If $\mathcal{A}$ and S are as in Definition 21, $\sqsupset_S$ is a total order, and $(\prec, l)$ is the pedigreed belief state of $\circledvee(\mathcal{A})$, then $\prec^+ = \prec$.*

Example 8 *Let's look once more at the space robot scenario considered in Example 1. Suppose the arrogant programmer is not part of the telemetry team, but instead works for a company on the other side of the country. Then the robot has to request information* *from two separate agents, one to query the manager and technician and one to query the programmer. Assume that the agents and the robot all rank the sources the same, assigning the technician rank 2 and the other two agents rank 1, which induces the same credibility ordering used in Example 5. The agents' pedigreed belief states and the result of their fusion are shown in Figure 5.*

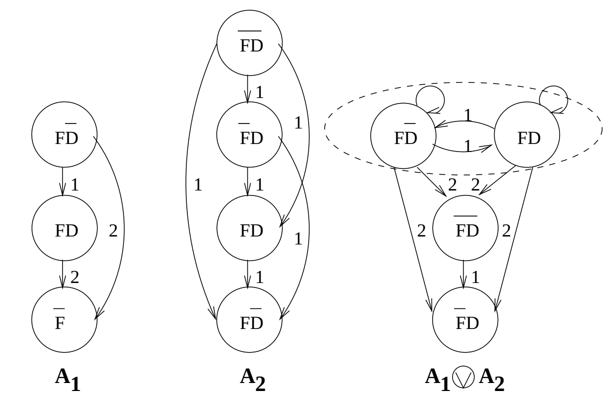

Figure 5: The pedigreed belief states of agent A_1 informed by s_m and s_t and of agent A_2 informed by s_p, and the result of their fusion in Example 8.

The first agent does not provide any information about overloading and the second agent provides incorrect information. However, we see that after fusing the two, the robot has a belief state that is identical to what it computed in Example 5 when there was only one agent informed by all three sources (we've only separated the top set of worlds so as to show the labeling). Consequently, it now knows the correct state of the system. And, satisfyingly, the final result does not depend on the order in which the robot receives the agents' reports.

The savings obtained in required storage space by this scheme can be substantial. Whereas explicitly storing all of an agent's informant sources S requires $\mathcal{O}(\|S\|2^{\mathcal{W}})$ amount of space in the worst case (when all the sources' belief states are fully connected relations), storing a pedigreed belief state only requires $\mathcal{O}(2^{\mathcal{W}})$ space in the worst case. Moreover, not only does the enriched representation allow us to conserve space, but it also provides for potential savings in the efficiency of computing fusion since, for each pair of worlds, we only need to consider the opinions of the *agents* rather than those of all the sources in the combined set of informants.

Incidentally, if we had used AGR^* as the basis for our general aggregation, simply storing the rank of the maximum supporting sources would not give us sufficient information to compute the induced belief state after fusion. To demonstrate this, we give an example where two pairs of sources induce the same annotated agent belief states, yet yield different belief states after fusion:

Example 9 *Let $\mathcal{W}$, $\mathcal{S}$, and $\sqsupseteq$ be as in Example 4. Suppose agents A_1, A_2, A_1', and A_2' are informed by sets of sources S_1, S_2, S_1', and S_2', respectively, where $S_1 = S_2 = \{s_2\}$, $S_1' = \{s_0, s_2\}$, and $S_2' = \{s_1, s_2\}$. AGR^* dictates that the pedigreed belief states of all four agents equal $<^{s_2}$ with all opinions annotated with $rank(s_2)$. In spite of this indistinguishability, if $A = \circledvee(\{A_1, A_2\})$ and $A' = \circledvee(\{A_1', A_2'\})$, then A's induced belief state equals $<^{s_2}$, i.e., $\{(a,b),(c,b)\}$, whereas A''s is $\{(a,b),(c,b),(a,c),(c,a),(a,a),(b,b),(c,c)\}$.*

7 Conclusion

We have described a semantically clean representation for aggregate beliefs which allows us to represent conflicting opinions without sacrificing the ability to make decisions. We have proposed an intuitive operator which takes advantage of this representation so that an agent can combine the belief states of a set of informant sources totally pre-ordered by credibility. Finally, we have described a mechanism for fusing the belief states of different agents which iterates well.

The aggregation methods we have discussed here are just special cases of a more general framework based on voting. That is, we account not only for the ranking of the sources supporting or disagreeing with an opinion (i.e., the *quality* of support), but also the percentage of sources in each camp (the *quantity* of support). Such an extension allows for a much more refined approach to aggregation, one much closer to what humans often use in practice. Exploring this richer space is the subject of further research.

Another problem which deserves further study is developing a fuller understanding of the properties of the *Bel*, *Agn*, and *Con* operators and how they interrelate.

Acknowledgements

Pedrito Maynard-Reid II was partly supported by a National Physical Science Consortium Fellowship. The final version of this paper was written with the financial support of the Jean et Hélène Alfassa Fund for Research in Artificial Intelligence.

References

Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 2nd edition, 1963.

Peter Gärdenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(1):197–245, January 1994.

Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, March 1979.

Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.

Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, May 1992.

Pedrito Maynard-Reid II and Yoav Shoham. Belief fusion: Aggregating pedigreed belief states. *Journal of Logic, Language, and Information*, 2000. To appear.

Amartya Sen. Social choice theory. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics*, volume III, chapter 22, pages 1073–1181. Elevier Science Publishers, 1986.

A Proofs

Proposition 1

1. *The transitive closure of a modular relation is modular.*
2. *Every transitive relation is quasi-transitive.*
3. *(Sen 1986) Every quasi-transitive relation is acyclic.*

Proof:

1. Suppose a relation $\leq$ over finite set Ω is modular, and $\leq^+$ is the transitive closure of $\leq$. Suppose $x, y, z \in \Omega$ and $x \leq^+ y$. Then there exist $w_0, \ldots, w_n$ such that $x = w_0 \leq \cdots \leq w_n = y$. Since $\leq$ is modular and $w_0 \leq w_1$, either $w_0 \leq z$ or $z \leq w_1$. In the former case, $x = w_0 \leq z$, so $x \leq^+ z$. In the latter case, $z \leq w_1 \leq \cdots w_n = y$, so $z \leq^+ y$.
2. Suppose Ω is a finite set, $x, y, z \in \Omega$, $\leq$ is a transitive relation over Ω, and $<$ is its strict version. Suppose $x < y$ and $y < z$. Then $x \leq y$, $y \not\leq x$, $y \leq z$, and $z \not\leq y$. $x \leq y$ and $y \leq z$ imply $x \leq z$, and $y \leq z$ and $y \not\leq x$ imply $z \not\leq x$, both by transitivity. So $x < z$.

■

Proposition 2 *(Sen 1986) Given a relation $\leq$ over a finite set Ω, the choice set operation C defines a choice function iff $\leq$ is acyclic.*

Proof: See (Sen 1986). ■

Proposition 3 *(Arrow 1963) There is no aggregation operator that satisfies restricted range, unrestricted domain, (weak) Pareto principle, independendence of irrelevant alternatives, and nondictatorship.*

Proof: See (Arrow 1963). ■

Proposition 4 *Let $\preceq$ be a relation over a finite set Ω and let $\sim$ be its symmetric restriction (i.e., $x \sim y$ iff $x \preceq y$ and $y \preceq x$). If $\preceq$ is total and quasi-transitive but not transitive, then $\sim$ is not transitive.*

Proof: Let $\preceq$ be a total, quasi-transitive, nontransitive relation. First, such a relation exits: if $\Omega = \{a, b, c\}$, it is easily verified that the relation $\Omega \times \Omega \setminus \{(b, a)\}$ is total, quasi-transitive, but not transitive.

Suppose $x \preceq y$ and $y \preceq z$ but $x \not\preceq z$. By totality, $z \preceq x$, so $z \prec x$. If $x \prec y$, then $z \prec y$ by quasi-transitivity, a contradiction. Thus, $x \sim y$. Similarly, if $y \prec z$, then $y \prec x$, a contradiction, so $y \sim z$. But $z \prec x$, so $x \not\sim z$. Therefore, $\sim$ is not transitive. ■

Proposition 5 *Suppose a relation $\prec$ is transitive and $\sim$ is the corresponding agnosticism relation. Then $\sim$ is transitive iff $\prec$ is modular.*

Proof: Suppose $\sim$ is transitive and suppose $x \prec z$, $x, y, z \in \mathcal{W}$. We prove by contradiction: Suppose $x \not\prec y$ and $y \not\prec z$. By transitivity, $z \not\prec y$ and $y \not\prec x$, so $x \sim y$ and $y \sim z$. By assumption, $x \sim z$, so $x \not\prec z$, a contradiction.

Suppose, instead, $\prec$ is modular and suppose $x \sim y$ and $y \sim z$, $x, y, z \in \mathcal{W}$. Then $x \not\prec y$, $y \not\prec x$, $y \not\prec z$, and $z \not\prec y$. By modularity, $x \not\prec z$ and $z \not\prec x$, so $x \sim z$. ■

Proposition 6 *The set of irreflexive relations in $\mathcal{B}$ is isomorphic to $\mathcal{T}$ and, in fact, equals $\mathcal{T}_{<}$.*

Proof: Let $x, y, z \in \mathcal{W}$. Suppose $\prec \in \mathcal{B}$ is irreflexive. Let $\preceq$ be defined as $x \preceq y$ iff $y \not\prec x$. We first show that $\prec$ is the strict version of $\preceq$. Suppose $\prec'$ is the strict version of $\preceq$. If $x \prec' y$, then $x \preceq y$ and $y \not\preceq x$, so $x \prec y$. If, instead, $x \prec y$, then $y \not\preceq x$. By totality, $x \preceq y$, so $x \prec' y$.

We show that $\preceq \in \mathcal{T}$. If $x \not\prec y$ then $y \preceq x$. Otherwise, $x \prec y$. But since $\prec$ is irreflexive, $y \not\prec x$ (otherwise $x \prec x$ by transitivity), so $x \preceq y$ and $\preceq$ is total. Next, suppose $x \preceq y$ and $y \preceq z$. Then $y \not\prec x$ and $z \not\prec y$. By modularity, $z \not\prec x$, so $x \preceq z$, so $\preceq$ is transitive.

Now suppose $\preceq \in \mathcal{T}$ and $\prec$ is its strict version. First we show that $\prec$ is modular. Suppose $x \prec y$. Then $x \preceq y$ and $y \not\preceq x$. Since $\preceq$ is total, $x \preceq z$ or $z \preceq x$. Suppose $x \preceq z$. Whether $y \preceq z$ or $y \not\preceq z$, $z \not\preceq x$ by transitivity. Suppose, instead, $z \preceq x$. Then $z \preceq y$ and $y \not\preceq z$, both by transitivity. We conclude that $x \preceq z$ and $z \not\preceq x$, or $z \preceq y$ and $y \not\preceq z$, so $x \prec z$ or $z \prec y$. Second, transitivity of $\prec$ follows immediately from Proposition 1 and the transitivity of $\preceq$. Finally, $\prec$ is irreflexive since it is asymmetric. ■

Proposition 7 *$\prec \in \mathcal{B}$ iff there is a partition $\mathbf{W} = \langle W_0, \ldots, W_n \rangle$ of $\mathcal{W}$ such that:*

1. *For every $x \in W_i$ and $y \in W_j$, $i \neq j$ implies $i < j$ iff $x \prec y$.*
2. *Every W_i is either fully connected ($w \prec w'$ for all $w, w' \in W_i$) or fully disconnected ($w \not\prec w'$ for all $w, w' \in W_i$).*

Proof: We refer to the conditions in the proposition as conditions 1 and 2, respectively. We prove each direction of the proposition separately.

($\Longrightarrow$) Suppose $\prec \in \mathcal{B}$, that is, $\prec$ is a modular and transitive relation over $\mathcal{W}$. We use a series of definitions and lemmas to show that a partition of $\mathcal{W}$ exists satisfying conditions 1 and 2. We first define an equivalence relation by which we will partition $\mathcal{W}$. Two elements will be equivalent if they "look the same" from the perspective of every element of $\mathcal{W}$:

Definition 23 *$x \equiv y$ iff for every $z \in \mathcal{W}$, $x \prec z$ iff $y \prec z$ and $z \prec x$ iff $z \prec y$.*

Lemma 7.1 *$\equiv$ is an equivalence relation over $\mathcal{W}$.*

Proof: Suppose $x \in \mathcal{W}$. For every $z \in \mathcal{W}$, $x \prec z$ iff $x \prec z$ and $z \prec x$ iff $z \prec x$, so $x \equiv x$. Therefore, $\equiv$ is reflexive.

Suppose $x, y \in \mathcal{W}$ and $x \equiv y$. Then for every $z \in \mathcal{W}$, $x \prec z$ iff $y \prec z$ and $z \prec x$ iff $z \prec y$. But then for every $z \in \mathcal{W}$, $y \prec z$ iff $x \prec z$ and $z \prec y$ iff $z \prec x$. Therefore, $y \equiv x$, so $\equiv$ is symmetric.

Suppose $x, y, z \in \mathcal{W}$, $x \equiv y$, and $y \equiv z$. Suppose further that $w \in \mathcal{W}$. By definition of $\equiv$, $x \prec w$ iff $y \prec w$ and $w \prec x$ iff $w \prec y$, and $y \prec w$ iff $z \prec w$ and $w \prec y$ iff $w \prec z$. Therefore, $x \prec w$ iff $z \prec w$ and $w \prec x$ iff $w \prec z$. Since w is arbitrary, $x \equiv z$, so $\equiv$ is transitive. ■

$\equiv$ partitions $\mathcal{W}$ into its equivalence classes. We use $[w]$ to denote the equivalence class containing w, that is, the set $\{w' \in \mathcal{W} : w \equiv w'\}$. Observe that two worlds in conflict always appear in the same equivalence class:

Lemma 7.2 *If $x, y \in \mathcal{W}$ and $x \infty y$, then $[x] = [y]$.*

Proof: Suppose $x, y \in \mathcal{W}$ and $x \infty y$. Since $[x]$ is an equivalence class, it suffices to show that $y \in [x]$, that is, $x \equiv y$. Suppose $z \in \mathcal{W}$. By transitivity, if $x \prec z$, then $y \prec z$; if $y \prec z$, then $x \prec z$; if $z \prec x$, then $z \prec y$; and, if $z \prec y$ then $z \prec x$. Thus, $x \prec z$ iff $y \prec z$ and $z \prec x$ iff $z \prec y$, and since z is arbitrary, $x \equiv y$. ■

We now define a total order over these equivalence classes:

Definition 24 *For all $x, y \in \mathcal{W}$, $[x] \leq [y]$ iff $[x] = [y]$ or $x \prec y$.*

Lemma 7.3 *$\leq$ is well-defined, that is, if $x \equiv x'$ and $y \equiv y'$, then $x \prec y$ iff $x' \prec y'$, for all $x, x', y, y' \in \mathcal{W}$.*

Proof: Suppose $x \equiv x'$ and $y \equiv y'$, $x, x', y, y' \in \mathcal{W}$. By the definition of $\equiv$, for every $z \in \mathcal{W}$, $x \prec z$ iff

$x' \prec z$. In particular, $x \prec y$ iff $x' \prec y$. Also by the definition of $\equiv$, for every $z' \in \mathcal{W}$, $z' \prec y$ iff $z' \prec y'$. In particular, $x' \prec y$ iff $x' \prec y'$. Therefore, $x \prec y$ iff $x' \prec y'$. ■

Lemma 7.4 *$\leq$ is a total order over the equivalence classes of $\mathcal{W}$ defined by $\equiv$.*

Proof: Suppose $x, y, z \in \mathcal{W}$. We first show that $\leq$ is total. By definition of $\leq$, if $x \prec y$ or $y \prec x$, then $[x] \leq [y]$ or $[y] \leq [x]$, respectively. Suppose $x \not\prec y$ and $y \not\prec x$, and suppose $z \in \mathcal{W}$. By modularity of $\prec$, $x \prec z$ implies $y \prec z$, $y \prec z$ implies $x \prec z$, $z \prec x$ implies $z \prec y$, and $z \prec y$ implies $z \prec x$, so $x \equiv y$. Therefore, $[x] = [y]$, so $[x] \leq [y]$ by the definition of $\leq$.

Next, we show that $\leq$ is anti-symmetric. Suppose $[x] \leq [y]$ and $[y] \leq [x]$. Then $[x] = [y]$ or $x \prec y$ and $y \prec x$. In the former case we are done, in the latter, the result follows from Lemma 7.2.

Finally, we show that $\leq$ is transitive. Suppose $[x] \leq [y]$ and $[y] \leq [z]$. Obviously, if $[x] = [y]$ or $[y] = [x]$, then $[x] \leq [z]$. Suppose not. Then $x \prec y$ and $y \prec z$, so $x \prec z$ by the transitivity of $\prec$. Therefore, $[x] \leq [y]$ by the definition of $\leq$. ■

We name the members of the partition $W_0, \ldots, W_n$ such that $W_i \leq W_j$ iff $i \leq j$, where n is an integer. Such a naming exists since every finite, totally ordered set is isomorphic to some finite prefix of the integers.

We now check that this partition satisfies the two conditions. For the first condition, suppose $x \in W_i$, $y \in W_j$, and $i \neq j$. We want to show that $i < j$ iff $x \prec y$. Since $i \neq j$, $[x] \neq [y]$. Suppose $i < j$. Then $i \leq j$, so $[x] \leq [y]$. Since $[x] \neq [y]$, $x \prec y$ by the definition of $\leq$. Now suppose, instead, that $x \prec y$. Then $[x] \leq [y]$ by the definition of $\leq$, so $i \leq j$. Since $[x] \neq [y]$, $y \not\prec x$ by Lemma 7.2. Since $[x] \neq [y]$ and $y \not\equiv x$, $[x] \not\leq [y]$ by the definition of $\leq$, so $j \not\leq i$. Thus, $i < j$.

Finally, we show that each W_i is either fully connected or fully disconnected. Suppose $x, y, z \in W_i$ so that $x \equiv y \equiv z$. It suffices to show that $x \prec x$ iff $y \prec z$. By the definition of $\equiv$, $x \prec x$ iff $y \prec x$, and $x \prec x$ iff $x \prec z$. Suppose $x \prec x$. Then, $y \prec x$ and $x \prec z$, so $y \prec z$ by transitivity of $\prec$. Suppose now, $x \not\prec x$. Then, $y \not\prec x$ and $x \not\prec z$, so $y \not\prec z$ by modularity of $\prec$.

($\Longleftarrow$) Suppose $\mathbf{W} = \langle W_0, \ldots, W_n \rangle$ is a partition of $\mathcal{W}$ and $\prec$ is a relation over $\mathcal{W}$ satisfying the given conditions. We want to show that $\prec$ is modular and transitive. We first give the following lemma:

Lemma 7.5 *Suppose $\mathbf{W}$ is a partition of $\mathcal{W}$ and $\prec$ is a relation over $\mathcal{W}$ satisfying condition 1. If $W_i, W_j \in \mathbf{W}$, $x \in W_i$, $y \in W_j$, and $x \prec y$, then $i \leq j$.*

Proof: If $i = j$, we're done. Suppose $i \neq j$. Then, since $x \prec y$, $i < j$ by condition 1. ■

We now show $\prec$ is modular. Suppose $x \in W_i$, $y \in W_j$, and $x \prec y$. Then $i \leq j$ by Lemma 7.5. Suppose $z \in W_k$. Then $i \leq k$ or $k \leq j$ by the modularity of $\leq$. Suppose $i < k$ or $k < j$. Then $x \prec z$ or $z \prec y$ by condition 1. Otherwise $i = k = j$, so $x, y, z \in W_i$. Since $x \prec y$, W_i is fully connected by condition 2, so $x \prec z$ (and $z \prec y$).

Finally, we show that $\prec$ is transitive. Suppose $x \in W_i$, $y \in W_j$, $z \in W_k$, $x \prec y$, and $y \prec z$. By Lemma 7.5, $i \leq j$ and $j \leq k$, so $i \leq k$ by the transitivity of $\leq$. Suppose $i < k$. Then $x \prec z$ by condition 1. Otherwise $i = k = j$, so $x, y, z \in W_i$. Since $x \prec y$, W_i is fully connected by condition 2, so $x \prec z$. ■

(END OF PROPOSITION 7 PROOF)

Proposition 8 *$\mathcal{T} \subset \mathcal{B}$ and is the set of reflexive relations in $\mathcal{B}$.*

Proof: We first show that $\mathcal{T} \subset \mathcal{B}$. Let $\preceq \in \mathcal{T}$ and $x, y, z \in \mathcal{W}$. By definition, $\preceq$ is transitive. Suppose $x \preceq y$. Since $\preceq$ is total, $x \preceq z$ or $z \preceq x$. If $z \preceq x$, then $z \preceq y$ by transitivity, so $\preceq$ is modular. On the other hand, the empty relation over $\mathcal{W}$ is modular and transitive, but not total and, consequently, not in $\mathcal{T}$.

Now we show that $\prec \in \mathcal{B}$ is in $\mathcal{T}$ if and only if it is reflexive. If $\prec \in \mathcal{T}$, it is total, so it is reflexive. If, instead, $\prec$ is reflexive, then $x \prec x$ so, by modularity, $x \prec y$ or $y \prec x$. Thus, $\prec$ is total. And, since $\prec \in \mathcal{B}$, it is transitive. ■

Proposition 9

1. *$\mathcal{Q} \cap \mathcal{B} = \mathcal{T}$.*
2. *$\mathcal{B} \not\subseteq \mathcal{Q}$.*
3. *$\mathcal{Q} \not\subseteq \mathcal{B}$ if $\mathcal{W}$ has at least three elements.*
4. *$\mathcal{Q} \subset \mathcal{B}$ if $\mathcal{W}$ has one or two elements.*

Proof:

1. Suppose $\preceq \in \mathcal{Q} \cap \mathcal{B}$. Then $\preceq$ is total and transitive and, hence, in $\mathcal{T}$. Suppose $\preceq \in \mathcal{T}$. By definition, $\preceq$ is total. Also by definition, it is transitive, so by Proposition 1, it is quasi-transitive and, thus, in $\mathcal{Q}$. By Proposition 8, $\preceq \in \mathcal{B}$ and, so, in $\mathcal{Q} \cap \mathcal{B}$.

2. The empty relation is modular and transitive, but not total and, so, not in $\mathcal{Q}$.

3. Suppose a and b are distinct elements of $\mathcal{W}$. The relation $\mathcal{W} \times \mathcal{W} \setminus \{(b,a)\}$ is total, and, since the strict version is $\{(a,b)\}$ which is transitive, it is also quasi-transitive. However, if there are at least three elements in $\mathcal{W}$, it is not transitive and, so, not in $\mathcal{B}$.

4. Suppose $\mathcal{W}$ has one element. Then $\mathcal{B}$ contains both possible relations over $\mathcal{W}$, whereas $\mathcal{Q}$ contains only the fully connected relation over $\mathcal{W}$.

   Suppose $\mathcal{W}$ has two elements a and b. Then $\mathcal{B}$ contains empty relation, the fully connected relation, and all the remaining eight relations which contain either (a,b) or (b,a), but not both. $\mathcal{Q}$, on the other hand, only contains the three reflexive relations containing either (a,b) or (b,a).

■

Proposition 10

1. $\mathcal{Q}_< \cap \mathcal{B} = \mathcal{T}_<$.

2. $\mathcal{B} \not\subseteq \mathcal{Q}_<$.

3. $\mathcal{Q}_< \not\subseteq \mathcal{B}$ *if* $\mathcal{W}$ *has at least three elements.*

4. $\mathcal{Q}_< \subset \mathcal{B}$ *if* $\mathcal{W}$ *has one or two elements.*

Proof:

1. Suppose $\prec \in \mathcal{Q}_< \cap \mathcal{B}$. Since $\prec \in \mathcal{Q}_<$, it is irreflexive, so since it is in $\mathcal{B}$, it is in $\mathcal{T}_<$ by Proposition 6. Suppose, instead, $\prec \in \mathcal{T}_<$. By Proposition 6, $\prec \in \mathcal{B}$. Let $\preceq \in \mathcal{T}$ be a relation such that $\prec$ is its strict version. (Obviously such a relation must exist.) From Proposition 9, $\preceq \in \mathcal{Q}$, so $\prec \in \mathcal{Q}_<$. Thus, $\prec \in \mathcal{Q}_< \cap \mathcal{B}$.

2. The fully connected relation over $\mathcal{W}$ is in $\mathcal{B}$, but not asymmetric and, so, not in $\mathcal{Q}_<$.

3. Suppose a and b are distinct elements of $\mathcal{W}$. If $\mathcal{W}$ has at least three elements, the relation $\{(a,b)\}$ is not modular and, thus, not in $\mathcal{B}$, yet it is the strict version of the relation $\mathcal{W} \times \mathcal{W} \setminus \{(b,a)\}$ which is total and quasi-transitive (since $\{(a,b)\}$ is transitive).

4. Suppose $\mathcal{W}$ has one element. Then $\mathcal{B}$ contains both possible relations over $\mathcal{W}$, whereas $\mathcal{Q}_<$ contains only the empty relation over $\mathcal{W}$.

   Suppose $\mathcal{W}$ has two elements a and b. Then $\mathcal{B}$ contains empty relation, the fully connected relation, and all eight of the remaining relations which contain either (a,b) or (b,a), but not both. $\mathcal{Q}_<$, on the other hand, only contains the three irreflexive relations.

■

Proposition 11 *If* $S \subseteq \mathcal{S}$, *then* $Un(S)$ *is modular but not necessarily transitive.*

Proof: Let $\prec = Un(S)$. Suppose $x, y, z \in \mathcal{W}$ and $x \prec y$. Then there is some $s \in S$ such that $x <^s y$. By assumption, $<^s$ is modular, so $x <^s z$ or $z <^s y$. By the definition of $Un(S)$, $x \prec z$ or $z \prec y$, so $\prec$ is modular.

Suppose $a, b, c \in \mathcal{W}$ and $S = \{s_1, s_2\}$ such that $<^{s_1} = \{(a,b),(a,c)\}$ and $<^{s_2} = \{(b,a),(c,a)\}$. $Un(S)$ is not transitive. ■

Proposition 12 *If* $S \subseteq \mathcal{S}$, *then* $AGRUn(S) \in \mathcal{B}$.

Proof: The transitive closure of any relation is transitive. Since $Un(S)$ is modular, the transitive closure of $Un(S)$ is also modular by Proposition 1. ■

Proposition 13 *If* $S \subseteq \mathcal{S}$ *and* $\sqsupseteq_S$ *is a total order, then* $AGRRf(S) \in \mathcal{B}$.

Proof: We first prove that $AGRRf(S)$ is modular. Note that the proof does not depend on $\sqsupseteq_S$ being a total order.

Lemma 13.1 *If* $S \subseteq \mathcal{S}$, *then* $AGRRf(S)$ *is modular.*

Proof: Let $\prec = AGRRf(S)$. Suppose $x, y, z \in \mathcal{W}$ and $x \prec y$. Then there exists $s \in S$ such that $x <^s y$ and, for every $s' \in S$, $s' \sqsupset s$ implies $x \not<^{s'} y$ and $y \not<^{s'} x$. Consider a source s' such that $x <^{s'} z$, $z <^{s'} x$, $y <^{s'} z$, or $z <^{s'} y$ and, for every $s'' \in S$, $s'' \sqsupset s'$ implies $x \not<^{s''} z$, $z \not<^{s''} x$, $y \not<^{s''} z$, and $z \not<^{s''} y$. We know such a source exists since, by modularity, $x <^s z$ or $z <^s y$. Furthermore, since s' is a maximal rank such source, $s' \sqsupseteq s$. We consider the four cases and show that, in each, either $x \prec z$ or $z \prec y$:

Case 1: $x <^{s'} z$. Since, for every $s'' \in S$, $s'' \sqsupset s'$ implies $x \not<^{s''} z$ and $z \not<^{s''} x$, $x \prec z$.

Case 2: $z <^{s'} x$. By modularity, $z <^{s'} y$ or $y <^{s'} x$.

Suppose $z <^{s'} y$. Since, for every $s'' \in S$, $s'' \sqsupset s'$ implies $z \not<^{s''} y$ and $y \not<^{s''} z$, $z \prec y$.

Suppose $y <^{s'} x$. Then $s \sqsupseteq s'$, so $s \equiv s'$. Since $x <^s y$, $x <^s z$ or $z <^s y$ by modularity. Thus, substituting s for s' above, either $x <^s z$ and, for every $s'' \in S$, $s'' \sqsupset s$ implies $x \not<^{s''} z$ and $z \not<^{s''} x$ so that $x \prec z$, or

$z <^s y$ and, for every $s'' \in S$, $s'' \sqsupset s$ implies $z \not<^{s''} y$ and $y \not<^{s''} z$ so that $z \prec y$.

Case 3: $y <^{s'} z$. By modularity, $y <^{s'} x$ or $x <^{s'} z$ which we have already considered in cases 2 and 1, respectively.

Case 4: $z <^{s'} y$. We have already considered this in case 2. ■

It remains to show that $\prec$ is transitive. Suppose $x \prec y$ and $y \prec z$. Then there exists $s_1 \in S$ such that $x <^{s_1} y$ and, for every $s_1' \in S$, $s_1' \sqsupset s$ implies $x \not<^{s_1'} y$ and $y \not<^{s_1'} x$, and there exists $s_2 \in S$ such that $y <^{s_2} z$ and, for every $s_2' \in S$, $s_2' \sqsupset s$ implies $y \not<^{s_2'} z$ and $z \not<^{s_2'} y$. Suppose $s_1 \sqsupset s_2$ (the case $s_2 \sqsupset s_1$ is similar). Then $y \not<^{s_1} z$ and $z \not<^{s_1} y$. By modularity, since $x <^{s_1} y$ and $z \not<^{s_1} y$, $x <^{s_1} z$. Let $s' \in S$ and $s' \sqsupset s_1$. Then $x \not<^{s'} y$ and $y \not<^{s'} x$. And, since $s_1 \sqsupset s_2$, $s' \sqsupset s_2$, so $y \not<^{s'} z$ and $z \not<^{s'} y$. By modularity, $x \not<^{s'} z$ and $z \not<^{s'} x$. Therefore, $x \prec z$. ■

Proposition 14 *If $S \subseteq \mathcal{S}$, then $AGR^*(S) \in \mathcal{B}$.*

Proof: By Proposition 12, $\prec_r \in \mathcal{B}$ for every $r \in ranks(S)$. For convenience, we assume the existence of a "virtual" source s_r corresponding to each $\prec_r$. Precisely, for each $r \in ranks(S)$, assume there exists a source $s_r \in \mathcal{S}$ such that $<^{s_r} = \prec_r$ and $rank(s_r) = r$, and let S' be the set of these sources. Then,

$$\begin{aligned} & AGR^*(S) \\ = & \left\{ (x,y) : \begin{array}{l} \exists r \in \mathcal{R}.\ x <_r y \wedge \\ (\forall r' > r \in ranks(S).\ x \approx_{r'} y) \end{array} \right\} \\ = & \left\{ (x,y) : \begin{array}{l} \exists s \in S'.\ x <^s y \wedge \\ \left(\forall s' \sqsupset s \in S'.\ x \approx^{s'} y\right) \end{array} \right\} \\ = & AGRRf(S'). \end{aligned}$$

Since there is one source in S' per rank r, and since $>$ is a total order over $\mathcal{R}$, $\sqsupset_{S'}$ is a total order. The result follows from Proposition 13. ■

Proposition 15 *If $S \subseteq \mathcal{S}$, then $AGR(S) \in \mathcal{B}$.*

Proof: By Lemma 13.1, $AGRRf(S)$ is modular. $AGRRf(S)^+$ is obviously transitive, and, by Proposition 1, it is modular as well. ■

Proposition 16 *Suppose $S \subseteq \mathcal{S}$.*

1. *If $\sqsupset_S$ is fully connected, $AGR(S) = AGRUn(S)$.*
2. *If $\sqsupset_S$ is a total order, $AGR(S) = AGRRf(S)$.*

Proof:

1. Suppose $\sqsupset_S$ is fully connected. Then the second half of the definition of $AGRRf$ is vacuously true so that $AGRRf(S)$ simplifies to $\{(x,y) : \exists s \in S.\ x <^s y\}$. But this is exactly $\bigcup_{s \in S} <^s$, i.e., $Un(S)$, so $AGR(S) = AGRRf(S)^+ = Un(S)^+ = AGRUn(S)$.
2. Suppose $\sqsupset_S$ is an total order. By Proposition 13, $AGRRf(S)$ is transitive, so $AGR(S) = AGRRf(S)^+ = AGRRf(S)$.

■

Proposition 17 *Let $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ and $AGR_f(<^{s_1}, \ldots, <^{s_n}) = AGR(S)$. AGR_f satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

Proof: Let $\prec = AGR_f(<^{s_1}, \ldots, <^{s_n})$. Then $\prec = AGR(S)$.

Restricted range: AGR_f satisfies restricted range by Proposition 15.

Unrestricted domain: AGR_f satisfies unrestricted domain by Definition 7.

Pareto principle: Suppose $x <^{s_i} y$ for all s_i. In particular, $x <^s y$ where s is a maximal rank source of S. Since s is maximal, it is vacuously true that for every $s' \sqsupset s \in S$, $x \not<^s y$ and $y \not<^{s'} x$. Therefore, $x \prec y$, so AGR_f satisfies the Pareto principle.

IIA: Let $S' = \{s_1', \ldots, s_n'\}$. First note that $AGRRf$ satisfies IIA:

Lemma 17.1 *Suppose $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$, $S' = \{s_1', \ldots, s_n'\} \subseteq \mathcal{S}$, $s_i \equiv s_i'$ for all i, $\prec_* = AGRRf(S)$, and $\prec_*' = AGRRf(S')$. If, for $x, y \in \mathcal{W}$, $x <^{s_i} y$ iff $x <^{s_i'} y$ for all i, then $x \prec_* y$ iff $x \prec_*' y$.*

Proof: Suppose $s_i \equiv s_i'$, and $x <^{s_i} y$ iff $x <^{s_i'} y$, for all i. Then $x \prec_* y$ iff $x \prec_*' y$ since Definition 13 only relies on the relative ranking of the sources and the relations between x and y in their belief states to determine the relation between x and y in the aggregated state. ■

Thus, IIA can only be disobeyed when the closure step of AGR introduces new opinions. (Note that IIA is satisfied when there are no sources of equal rank since, by Proposition 16, the closure step introduces no new opinions under these conditions.) However, these new opinions are only added between worlds already involved in a conflict, as the following two lemmas show:

Lemma 17.2 *Suppose $S \subseteq \mathcal{S}$ and $\prec_* = AGRRf(S)$. For every integer $n \geq 2$, if $x, y \in \mathcal{W}$, $x \not\prec_* y$, there exist $x_0, \ldots, x_n \in \mathcal{W}$ such that $x = x_0 \prec_* \cdots \prec_* x_n = y$, and n is the smallest integer such that this is true, then $x_n \prec_* \cdots \prec_* x_0$.*

Proof: Suppose $x, y \in \mathcal{W}$, $x \not\prec_* y$, and there exist $x_0, \ldots, x_n \in \mathcal{W}$ such that $x = x_0 \prec_* \cdots \prec_* x_n = y$, and n is the smallest integer such that this is true. Consider any triple x_{i-1}, x_i, x_{i+1}, where $1 \leq i \leq n-1$. First, $x_{i-1} \not\prec_* x_{i+1}$, otherwise there would be a chain of shorter length than n between x and y. Now, since $x_{i-1} \prec_* x_i$, there exists $s_1 \in S$ such that $x_{i-1} <^{s_1} x_i$ and, for all $s' \sqsupset s_1 \in S$, $x_{i-1} \approx^{s'} x_i$. Similarly, there exists $s_2 \in S$ such that $x_i <^{s_2} x_{i+1}$ and, for all $s' \sqsupset s_2 \in S$, $x_i \approx^{s'} x_{i+1}$. Thus, all sources with higher rank than $\max(s_1, s_2)$ are agnostic with respect to x_{i-1} and x_{i+1}.

Suppose $s_1 \sqsupset s_2$. Then $x_i \approx^{s_1} x_{i+1}$ so, by transitivity, $x_{i-1} <^{s_1} x_{i+1}$. But then $x_{i-1} \prec_* x_{i+1}$, a contradiction. Similarly, we derive a contradiction if $s_2 \sqsubset s_1$. Thus, $s_1 \equiv s_2$.

Now, since $x_{i-1} \not\prec_* x_{i+1}$ and all sources with rank higher than s_1 and s_2 are agnostic with respect to x_{i-1} and x_{i+1}, $x_{i-1} \not<^{s_1} x_{i+1}$. By modularity, $x_{i+1} <^{s_1} x_i$. Since $s_1 \equiv s_2$, and all the sources with higher rank than s_2 are agnostic with respect to x_i and x_{i+1}, $x_{i+1} \prec_* x_i$. Similarly, $x_i <^{s_2} x_{i-1}$, so $x_i \prec_* x_{i-1}$. Since i was chosen arbitrarily between 1 and $n-1$, $x_n \prec_* \cdots \prec_* x_0$. And, in fact, all the opinions between these worlds originate from sources of the same rank. ■

Lemma 17.3 *Suppose $S \subseteq \mathcal{S}$, $\prec_* = AGRRf(S)$, $\prec = AGR(S)$, and $x \not\prec_* y$ for $x, y \in \mathcal{W}$. If $x \prec y$, then $x \infty y$.*

Proof: Suppose $x \not\prec_* y$. If $x \prec y$, then there exist $x_0, \ldots, x_n$ such that $x = x_0 \prec_* \cdots \prec_* x_n = y$ and n is the smallest positive integer such that this is true. Then, by Lemma 17.2, $y = x_n \prec_* \cdots \prec_* x_0 = x$, so $y \prec x$ and $x \infty y$. ■

Now, suppose $x, y \in \mathcal{W}$, $x <^{s_i} y$ iff $x <^{s'_i} y$ for all i, $x \not\infty y$, and $x \not\infty' y$. We show that $x \prec y$ implies $x \prec' y$ (the other direction is identical). Suppose $x \prec y$. Let $\prec_* = AGRRf(S)$ and $\prec'_* = AGRRf(S')$. Since $x \not\infty y$, $x \prec_* y$ by Lemma 17.3. But then $x \prec'_* y$ by Lemma 17.1, so $x \prec' y$.
(END OF IIA SUB-PROOF)

Non-dictatorship: Suppose $\sqsupseteq_S$ is fully connected and suppose $x <^{s_i} y$ and $y \not<^{s_i} x$. Let s_j be such that $y <^{s_j} x$. Then $x \prec y$ and $y \prec x$, so s_i is not a dictator. ■

(END OF PROPOSITION 17 PROOF)

Proposition 18 *Let $\mathcal{A}$ and S be as in Definition 21, $\prec^{A_i}$, agent A_i's induced belief state, and $\sqsupseteq_S$, fully connected. If $A = \circledvee(\mathcal{A})$, then $\left(\bigcup_{A_i \in \mathcal{A}} \prec^{A_i}\right)^+$ is A's induced belief state.*

Proof: We will use the following lemma:

Lemma 18.1 *If Π is a set of relations over an arbitrary finite set Ω, then*

$$\left(\bigcup_{\leq \in \Pi} \leq^+\right)^+ = \left(\bigcup_{\leq \in \Pi} \leq\right)^+$$

where $\leq^+$ is the transitive closure of $\leq$.

Proof: Let $\prec = \left(\bigcup_{\leq \in \Pi} \leq^+\right)^+$, $\prec' = \left(\bigcup_{\leq \in \Pi} \leq\right)^+$, and $a, b \in \Omega$. Suppose $a \prec b$. Then there exist $\leq_0, \ldots, \leq_{n-1} \in \Pi$ and $w_0, \ldots, w_n \in \Omega$ such that

$$a = w_0 \leq_0^+ \cdots \leq_{n-1}^+ w_n = b$$

Thus, there exist $x_{00}, \ldots, x_{0m_0}, \ldots, x_{(n-1)0}, \ldots, x_{(n-1)m_{n-1}}$ in Ω such that

$$\begin{aligned} a &= w_0 = x_{00} \leq_0 \cdots \leq_0 x_{0m_0} = w_1 = x_{10} \cdots \\ &= w_{n-1} = x_{(n-1)0} \leq_{n-1} \cdots \leq_{n-1} x_{(n-1)m_{n-1}} \\ &= w_n = b \end{aligned}$$

so $a \leq' b$.

Now suppose $a \prec' b$. Then there exist $\leq_0, \ldots, \leq_{n-1} \in \Pi$ and $w_0, \ldots, w_n \in \Omega$ such that

$$a = w_0 \leq_0 \cdots \leq_{n-1} w_n = b$$

Obviously, this implies that

$$a = w_0 \leq_0^+ \cdots \leq_{n-1}^+ w_n = b$$

which implies that

$$a = w_0 \leq_*^+ \cdots \leq_*^+ w_n = b$$

where $\leq_* = \left(\bigcup_{\leq \in \Pi} \leq\right)$, so $a \leq b$. ■

Now, let $\prec$ be the belief state induced by $\circledvee(\mathcal{A})$. Then $\prec = AGR(S)$. By Proposition 16, $\prec = AGRUn(S)$, so

$$\begin{aligned} \prec &= Un(S)^+ \\ &= \left(\bigcup_{s \in S} <^s\right)^+ \\ &= \left(\bigcup_{s \in \bigcup_{i=1}^n S_i} <^s\right)^+ \\ &= \left(\bigcup_{A_i \in \mathcal{A}} \bigcup_{s \in S_i} <^s\right)^+ \end{aligned}$$

By the lemma,

$$\begin{aligned}\prec &= \left(\bigcup_{A_i\in\mathcal{A}}\left(\bigcup_{s\in S_i} <^s\right)^+\right)^+ \\ &= \left(\bigcup_{A_i\in\mathcal{A}} AGRUn(S_i)\right)^+ \\ &= \left(\bigcup_{A_i\in\mathcal{A}} \prec^{A_i}\right)^+\end{aligned}$$

■

Proposition 19 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$ and with pedigreed belief state $(\prec, l)$. Then*

$$x \prec_r^A y$$
$$\textit{iff}$$
$$\exists s \in S.\ x <^s y \wedge r = rank(s) \wedge \left(\forall s' \sqsupset s \in S.\ x \approx^{s'} y\right).$$

Proof: Suppose $x \prec_r^A y$. Then $x \prec y$ and $l((x,y)) = r$. By Definitions 13 and 22, there exists $s \in S$ such that $x <^s y$ and for every $s' \sqsupset s \in S$, $x \approx^{s'} y$. In particular, if $x <^{s'} y$ for some $s' \in S$, then $s \sqsupset s'$, so $rank(s) \geq rank(s')$. Thus,

$$\begin{aligned} r &= l((x,y)) \\ &= \max\{rank(s') : x <^{s'} y, s' \in S\} \\ &= rank(s).\end{aligned}$$

Now suppose there exists $s \in S$ such that $x <^s y$, $r = rank(s)$, and, for every $s' \sqsupset s \in S$, $x \approx^{s'} y$. Then $x \prec y$. Moreover, since for every $s' \in S$, $x <^{s'} y$ implies $s \sqsupseteq s'$ which implies $rank(s) \geq rank(s')$,

$$\begin{aligned} l((x,y)) &= \max\{rank(s') : x <^{s'} y, s' \in S\} \\ &= rank(s) \\ &= r.\end{aligned}$$

Therefore, $x \prec_r^A y$. ■

Proposition 20 *Let $\mathcal{A}$ and S be as in Definition 21, $\sqsupseteq_S$, a total pre-order, and $A = \circledvee(\mathcal{A})$. If*

1. *$\prec$ is the relation*

$$\left\{(x,y) : \begin{array}{l}\exists A_i \in \mathcal{A}, r \in \mathcal{R}.\ x \prec_r^{A_i} y \wedge \\ \left(\forall A_j \in \mathcal{A}, r' > r \in \mathcal{R}.\ x \sim_{r'}^{A_j} y\right)\end{array}\right\}$$

over $\mathcal{W}$,

2. *$l : \prec \rightarrow \mathcal{R}$ such that $l((x,y)) = \max\{r : x \prec_r^{A_i} y, A_i \in \mathcal{A}\}$, and*

then $(\prec, l)$ is A's pedigreed belief state.

Proof: Let $\prec' = AGRRf(S)$ and $l' : \prec' \rightarrow \mathcal{R}$ such that $l'((x,y)) = \max\{rank(s) : x <^s y, s \in S\}$. It suffices to show that $\prec = \prec'$ and $l = l'$.

Suppose $x \prec y$. We show that $x \prec' y$, i.e., there exists $s \in S$ such that $x <^s y$ and, for every $s' \sqsupset s \in S$, $x \not<^{s'} y$ and $y \not<^{s'} x$, and that $l'((x,y)) = l((x,y))$. Since $x \prec y$, there exists A_i and r such that $x \prec_r^{A_i} y$ and, for every $A_j \in \mathcal{A}$ and $r' > r \in \mathcal{R}$, $x \not\prec_{r'}^{A_j} y$ and $y \not\prec_{r'}^{A_j} x$. Since $x \prec_r^{A_i} y$, there exists $s \in S_i$ such that $x <^s y$, $rank(s) = r$, and, for every $s_1 \sqsupset s \in S_i$, $x \not<^{s_1} y$ and $y \not<^{s_1} x$. $S_i \subseteq S$, so there exists $s \in S$ such that $x <^s y$. Now suppose s' is a maximal rank source of S with $x <^{s'} y$ or $y <^{s'} x$. Such an s' exists since $x <^s y$. Since $\sqsupseteq$ is a total pre-order, it suffices to show that $s \sqsupseteq s'$. Suppose $s' \in S_j$. Since $S_j \subseteq S$, s' is also a maximal rank source of S_j with $x <^{s'} y$ or $y <^{s'} x$, so $x \prec_{rank(s')}^{A_j} y$ or $y \prec_{rank(s')}^{A_j} x$. But since $x \prec_r^{A_i} y$, $r = rank(s) \geq rank(s')$, so $s \sqsupseteq s'$. Furthermore, $l'((x,y)) = rank(s) = r = l((x,y))$.

Now suppose $x \prec' y$. We show that $x \prec y$, i.e., there exists A_i and r such that $x \prec_r^{A_i} y$ and, for every $A_j \in \mathcal{A}$ and $r' > r \in \mathcal{R}$, $x \not\prec_{r'}^{A_j} y$ and $y \not\prec_{r'}^{A_j} x$, and that $l((x,y)) = l'((x,y))$. Since $x \prec' y$, there exists $s \in S$ such that $x <^s y$ and, for every $s' \sqsupset s \in S$, $x \not<^{s'} y$ and $y \not<^{s'} x$. Suppose $s \in S_i$. Since $S_i \subseteq S$, it is also the case that for every $s' \sqsupset s \in S_i$, $x \not<^{s'} y$ and $y \not<^{s'} x$, so $x \prec_{rank(s)}^{A_i} y$. Now, let A_j and r' be such that $x \prec_{r'}^{A_j} y$ or $y \prec_{r'}^{A_j} x$. It suffices to show that $rank(s) \geq r'$. By Proposition 19, there exists $s' \in S_j$ such that $x <^{s'} y$ or $y <^{s'} x$ and $rank(s') = r'$. But then $s \sqsupseteq s'$, so $rank(s) \geq rank(s') = r'$. Furthermore, $l((x,y)) = rank(s) = l'((x,y))$. ■

Proposition 21 *If $\mathcal{A}$ and S are as in Definition 21, $\sqsupseteq_S$ is a total order, and $(\prec, l)$ is the pedigreed belief state of $\circledvee(\mathcal{A})$, then $\prec^+ = \prec$.*

Proof: Since $\sqsupseteq_S$ is a total order, $AGR(S) = AGRRf(S)$ by Proposition 16. Thus, $\prec = AGRRf(S) = AGR(S) = AGRRf(S)^+ = \prec^+$. ■

B Notation key

Ω: arbitrary finite set
$a, b, c, \ldots$: specific elements of a set
$x, y, z, \ldots$: arbitrary elements of a set
$A, B, C, \ldots$: specific subsets of a set
$X, Y, Z, \ldots$: arbitrary subsets of a set
Π: arbitrary set of relations
$\leq$: arbitrary relation
$C(X, \leq)$: choice set of W wrt $\leq$

$\mathcal{W}$: finite set of possible worlds/alternatives
w, W: element, subset of $\mathcal{W}$, respectively

$\mathcal{B}$: set of generalized belief states (modular, transitive relations)
$\prec$: element of $\mathcal{B}$, strict likelihood/preference
$\preceq$: weak likelihood/preference
$\sim$: equal likelihood, agnosticism/indifference
∞: conflict

$\mathcal{T}$: set of total pre-orders
$\mathcal{T}_{<}$: strict versions of total pre-orders
$\mathcal{Q}$: set of total, quasi-transitive relations
$\mathcal{Q}_{<}$: strict versions of total, quasi-transitive relations

$\mathcal{S}$: set of sources
s, S: element, subset of $\mathcal{S}$, respectively
$<^s$: belief state of source s
$\approx$: source agnosticism
$\bowtie$: source conflict
$\mathcal{R}$: set of ranks
r: element of $\mathcal{R}$
$\sqsupseteq$, $\sqsupseteq_S$: credibility ordering over $\mathcal{S}$, $S \subseteq \mathcal{S}$, respectively

$\mathcal{A}$: set of agents
A: element of $\mathcal{A}$
$\prec^A$: A's induced belief state
$(\prec, l)$: pedigreed belief state
l: labeling function of a pedigreed belief state
$\prec_r^A$: restriction of A's pedigreed belief state to rank r