Journal of Artificial Intelligence Research 19 (2003) 155-203 Submitted 1/03; published 9/03

Representing and Aggregating Conflicting Beliefs

Pedrito Maynard-Zhang MAYNARP@MUOHIO.EDU
Department of Computer Science and Systems Analysis
Miami University
Oxford, Ohio 45056, USA

Daniel Lehmann LEHMANN@CS.HUJI.AC.IL
School of Computer Science and Engineering
Hebrew University
Jerusalem 91904, Israel

Abstract

We consider the two-fold problem of representing collective beliefs and aggregating these beliefs. We propose a novel representation for collective beliefs that uses modular, transitive relations over possible worlds. They allow us to represent conflicting opinions and they have a clear semantics, thus improving upon the quasi-transitive relations often used in social choice. We then 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 by accounting for the explicitly encoded conflicts. 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. Finally, we extend our framework to incorporate voting.

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.

In addressing the first problem, we take as a starting point total preorders over possible worlds (i.e., interpretations of a specified language) used in the belief revision community to represent individuals' beliefs. The relations describe opinions on the relative likelihood of worlds and can be viewed as encoding all of an agent's conditional beliefs, i.e., not only what he believes now, but what he would believe under all other conditions. This representation is based on the semantical work (cf. Grove, 1988; Katsuno & Mendelzon, 1991) supporting the Alchourrón, Gärdenfors, and Makinson proposal (Alchourrón, Gärdenfors, & Makinson, 1985; Gärdenfors, 1988) (known as the *AGM theory*) for belief revision.

The social choice community has dealt extensively with the problem of representing collective preferences (cf. Sen, 1986). However, the problem is formally equivalent to that of representing collective beliefs, so the results are applicable. The classical approach has been to use quasi-transitive relations – relations whose asymmetric restrictions are transitive – over the set of objects. (Total preorders are a special subclass of these relations.) 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. Specifically, we propose modular, transitive relations and argue that they solve some of the unpleasant semantical problems suffered by the earlier approach. (We define modularity precisely later, but it can be viewed intuitively as a sufficient relaxation of the totality requirement on total preorders to make the distinction between indifference and conflict possible.)

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. We define a mechanism that does just this, relying on our generalized representation to circumvent Arrow's (1963) Impossibility Theorem when there are sources of equal reliability.

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 a simple set-theoretic mechanism that satisfies these requirements as well as a computationally effective way of computing the resulting belief state.

The aggregation and fusion mechanisms described so far take into account quality of support for opinions, but completely ignore *quantity* of support. However, the latter often provides sufficient information to resolve apparent conflicts. Take, for example, the situation where all the sources for the robot above have equal credibility and all except a small minority suggest the robot move the spaceship to avoid a potential collision with an oncoming asteroid. In such a situation, we often prefer to resolve the conflict by siding with the majority. To this end, we describe how to extend our framework to allow for voting, introducing a novel modular closure operation in the process.

After some preliminary definitions, we address each of these topics in turn.

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., $\leq \subseteq \Omega \times \Omega$.*[2] *We will 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 all $x \in \Omega$. It is* irreflexive *iff $x \not\leq x$ for all $x \in \Omega$.*
2. symmetric *iff $x \leq y \Rightarrow y \leq x$ for all $x, y \in \Omega$. It is* asymmetric *iff $x \leq y \Rightarrow y \not\leq x$ for all $x, y \in \Omega$. It is* anti-symmetric *iff $x \leq y \wedge y \leq x \Rightarrow x = y$ for all $x, y \in \Omega$.*
3. *the* asymmetric restriction *of a relation $\leq'$ over Ω iff $x \leq y \Leftrightarrow x \leq' y \wedge y \not\leq' x$ for all $x, y \in \Omega$. It is the* symmetric restriction *of $\leq'$ iff $x \leq y \Leftrightarrow x \leq' y \wedge y \leq' x$ for all $x, y \in \Omega$.*
4. total *iff $x \leq y \vee y \leq x$ for all $x, y \in \Omega$.*
5. modular *iff $x \leq y \Rightarrow x \leq z \vee z \leq y$ for all $x, y, z \in \Omega$.*
6. transitive *iff $x \leq y \wedge y \leq z \Rightarrow x \leq z$ for all $x, y, z \in \Omega$.*
7. quasi-transitive *iff its asymmetric restriction is transitive.*
8. *the* transitive closure *of a relation $\leq'$ over Ω iff, for some integer n,*

$$x \leq y \Leftrightarrow \exists w_0, \ldots, w_n \in \Omega.\ x = w_0 \leq' \cdots \leq' w_n = y$$

   for all $x, y \in \Omega$. (We generally use $\leq^+$ to denote the transitive closure of a relation $\leq$.)

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 asymmetric restriction of $\leq$.*
10. *a* total preorder *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.*
12. fully connected *iff $x \leq y$ for all $x, y \in \Omega$. It is* fully disconnected *iff $x \not\leq y$ for all $x, y \in \Omega$.*

Proposition 1

1. *The transitive closure of a modular relation is modular.*

1. We only use binary relations in this paper, so we will refer to them simply as relations.
2. For the reader's convenience, we have included in Appendix B a key to most of the notational symbols used throughout the paper.

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 asymmetric restriction, and $X \subseteq \Omega$, then the* choice set of X with respect to $\leq$ *is*

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

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

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

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* ch *defines a choice function iff $\leq$ is acyclic.*[3]

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 $x, y \in \Omega$.*

Finally, the cardinality of a set Ω will be denoted $\|\Omega\|$.

Assume we are given some language $\mathcal{L}$ with a satisfaction relation $\models$ for $\mathcal{L}$. Let $\mathcal{W}$ be a finite, non-empty set of possible worlds (interpretations) over $\mathcal{L}$. For a world $w \in \mathcal{W}$ and a sentence $p \in \mathcal{L}$, $w \models p$ iff p evaluates to true in w. Given a sentence p, $|p| = \{w \in \mathcal{W} \mid w \models p\}$.

3. Representing Collective Beliefs

Our representation of collective beliefs generalizes the representation developed in the belief revision community for the conditional beliefs of an individual, so we briefly review it. We then consider implications from social choice for representing collective beliefs. Finally, we describe our proposal and argue for its desirability.

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

3.1 Belief Revision Representation of Conditional Beliefs

Much of the belief revision field has built on the seminal work by Alchourrón, Gärdenfors, and Makinson (Alchourrón et al., 1985; Gärdenfors, 1988) refered to as the *AGM theory*. This work sought to formalize an "Occam's razor"-like principle of minimal change: the set of beliefs resulting from a revision should be one produced by modifying the original beliefs minimally to accomodate the new information. To capture this principle precisely, they proposed the famous AGM postulates which impose restrictions on belief change operators. Subsequent model-theoretic work (Grove, 1988; Katsuno & Mendelzon, 1991; Maynard-Reid II & Shoham, 2001) showed that accepting these postulates amounts to assuming that an individual's *belief state* is represented by a total preorder $\preceq$ over $\mathcal{W}$; revision of the individual's beliefs by a sentence $p \in \mathcal{L}$ then consists of computing $\text{ch}(|p|, \preceq)$.

Kraus, Lehmann, and Magidor (1990) and Lehmann and Magidor (1992) developed a similar central role for ordered structures in the semantics of nonmonotonic logics, and Gärdenfors and Makinson (1994) established the relation between the two topics. Semantically, $\preceq$ represents the weak relative likelihood of possible worlds: $x \preceq y$ means possible world x is considered to be at least as likely as possible world y.[4] If $x \preceq y$ and $y \preceq x$, then x and y are considered equally likely. We can also interpret $\preceq$ sententially using the famous *Ramsey Test* (Ramsey, 1931): it encodes a set of conditional beliefs, i.e., not only what is believed now (called the *belief set*), but all counterfactual beliefs as well (what would be believed if other conditions were the case). According to this criteria, the conditional belief "if p then q" holds if p and q are sentences in $\mathcal{L}$ and q is satisfied by all the worlds in $\text{ch}(|p|, \preceq)$; we write $\text{Bel}(p?q)$. If neither the belief $p?q$ nor the belief $p?\neg q$ hold in the belief state, it is said to be *agnostic* with respect to $p?q$, written $\text{Agn}(p?q)$. The belief set induced by the belief state consists of all those sentences q such that $\text{Bel}(true?q)$ holds.

3.2 Social Choice Implications

Our first inclination, then, would be to use total preorders to represent collective beliefs since they work so well for individuals' beliefs. Unfortunately, such an approach is inherently problematic as was discovered early on in the social choice community. That community's interest lies in representing collective preferences rather than collective beliefs; however, the results are equally relevant since the classical representation of an individual's preferences is also a total preorder. Instead of relative likelihood, relations represent relative preference; instead of equal likelihood, indifference.

Arrow's (1963) celebrated Impossibility Theorem showed that no aggregation operator over total preorders exists satisfying the following small set of desirable properties:

Definition 5 *Let f be an aggregation operator over the relations $\preceq_1, \ldots, \preceq_n$ 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 preorders over Ω.*
- Unrestricted Domain*: The domain of f is the set of n-tuples of total preorders over Ω.*

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

- Pareto Principle*: If $x \prec_i y$ for all i, then $x \prec y$.*[5]
- 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, 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 preorder. 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, and perhaps most obviously, the domain and the range of the aggregation operator are different, violating what is known in the belief revision literature as the *principle of categorical matching* (cf. Gardenfors and Rott's 1995 survey). This problem is closely related to the second which is that total, quasi-transitive relations have unsatisfactory semantics. If $\preceq$ is total and quasi-transitive but not a total preorder, 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. 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 is indifferent between plums and mangoes and also indifferent between mangoes and peaches, we would be greatly surprised were she to profess a strong preference for plums over peaches.[6] Thus, it seems that total quasi-transitive relations that are not total preorders cannot be understood easily as preference or indifference. Since the existence of a choice function is generally sufficient for classical social choice problems, these issues were 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 aggregations and, consequently, must maintain clean semantics.

Third, 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

5. Technically, this is known as the *weak* Pareto principle. The strong Pareto principle states that $x \prec y$ if there exists i such that $x \prec_i y$ and $x \preceq_i y$ for all i. Obviously, the strong version implies the weak version, so Arrow's theorem applies to it as well.
6. However, see Luce's (1956) work on *semiorders* for some of the opposing arguments in the transitivity of indifference debate.

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 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's (1979) 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.

3.3 Generalized Belief States

Because belief aggregation is formally similar to preference aggregation, it is also susceptible to the problems faced by the social choice community. We take the view that much of the difficulty encountered in previous attempts to define acceptable aggregation policies has been the lack of explicit representations of conflicts among the individuals. We generalize the total preorder representation so as to capture information about conflicts. This generalization opens the way for semantically clear aggregation policies, with the added benefit of focusing attention on the culprit sets of worlds.

3.3.1 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 is similar to that of 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 asymmetric restriction 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 $\bowtie$, so that $x \bowtie y$ 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 $\bowtie$ 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 except when to do so would cause confusion.

3.3.2 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 above, 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: $\mathcal{W} = \{a, b, c\}$ and $\prec = \{(a, c)\}$, 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 asymmetric restrictions of total preorders, i.e., conflict-free. Let $\mathcal{T}$ be the set of total preorders over $\mathcal{W}$, $\mathcal{T}_<$, the set of their asymmetric restrictions.

Proposition 6 *$\mathcal{T}_< \subset \mathcal{B}$ and is the set of irreflexive relations in $\mathcal{B}$.*

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 preorder, one which is the asymmetric restrictions of a total preorder, and one which is 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.)

Thus, generalized belief states are not a big change from the asymmetric restrictions of total preorders. 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

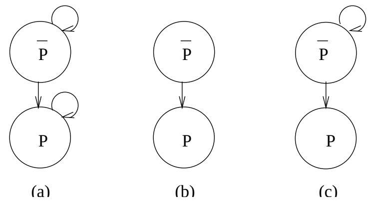

Figure 1: Three examples of generalized belief states: (a) a total preorder, (b) the asymmetric restriction of a total preorder, (c) neither.

agnostic about conditional belief $p?q$ (i.e., $\text{Agn}(p?q)$) 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 $\text{Con}(p?q)$, if the choice set is fully connected.

Finally, we compare the representational power of our definitions to those discussed in the previous section. First, as a companion result to Proposition 6, it is obvious that $\mathcal{B}$ subsumes the class of total preorders $\mathcal{T}$ and, in fact, $\mathcal{T}$ is the set of reflexive relations in $\mathcal{B}$.

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 asymmetric restrictions.

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 asymmetric restrictions 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.*

Note that generalized belief states as described are extremely rich and would require optimization in practice to avoid high maintenance cost. Although this issue is somewhat outside the scope of this paper, we do address (in the respective sections) ways to minimize the further explosion of this complexity when the complications of fusion and voting are introduced.

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.

4. 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 by aggregating the belief states of the sources 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.*

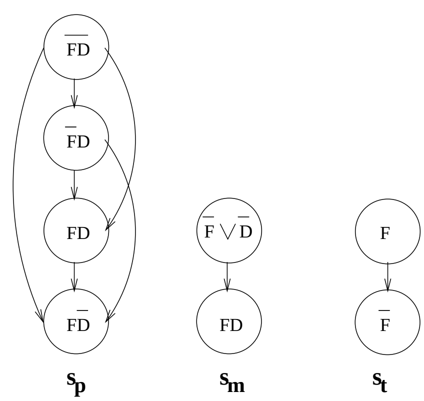

Figure 2: The belief states of s_p, s_m, and s_t in Example 1.

4.1 Sources

Let us begin the formal development by defining sources and their belief states:

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 preorder built on a totally ordered set of ranks (e.g., integers).

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

Definition 9 $\mathrm{rank} : \mathcal{S} \to \mathcal{R}$ *assigns to each source a rank. Also, for $S \subseteq \mathcal{S}$,* $\mathrm{ranks}(S)$ *denotes the set $\{r \in \mathcal{R} : \exists s \in S.\ \mathrm{rank}(s) = r\}$.*

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

We use $\sqsupset$ and $\equiv$ to denote the asymmetric and symmetric restrictions of $\sqsupseteq$, respectively.[7] 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. Also observe that $\mathcal{R}$ can be any arbitrary totally ordered set. Thus, not only does it allow for numeric ranking systems (such as the integers), but non-numeric systems as well (e.g., military ranks). Furthermore, this generality allows our proposal to easily accommodate applications where new ranks need to be dynamically added and it is inconvenient or impossible to change the rank labels of existing sources (e.g., a large workers' union where members are ranked by relative level and quality of experience).

We are now 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—aggregation of equally ranked and strictly ranked sources—before considering the general case.

4.2 Aggregating Equally Ranked Sources

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

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

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

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

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

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

Intuitively, we are simply inferring opinions implied by the conflicts introduced by the aggregation. We will show this formally when we consider the more general aggregation operator below.

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.*

4.3 Aggregating Strictly Ranked Sources

Next, consider the case where the sources are strictly ranked, i.e., $\sqsupseteq_S$ is a total order. We define a *lexicographic* 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* $\mathrm{AGRRf}(S)$ *is the relation*

$$\left\{(x,y) : \exists s \in S.\ x <^s y \wedge \left(\forall s' \in S.\ s' \sqsupset_S s \Rightarrow x \approx^{s'} y\right)\right\}.$$

As with $\mathrm{AGRUn}(S)$, $\mathrm{AGRRf}(S)$ is not guaranteed to be transitive, but it is always modular:

Proposition 13 *If $S \subseteq \mathcal{S}$, then* $\mathrm{AGRRf}(S)$ *is modular but not necessarily transitive.*

However, in the case that $\sqsupseteq_S$ is a total order, the result of applying AGRRf is guaranteed to be a belief state.

Proposition 14 *If $S \subseteq \mathcal{S}$ and $\sqsupseteq_S$ is a total order, then* $\mathrm{AGRRf}(S) \in \mathcal{B}$.

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.*

4.4 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 equally ranked sources using AGRUn, then aggregate the strictly ranked results using what is essentially AGRRf.

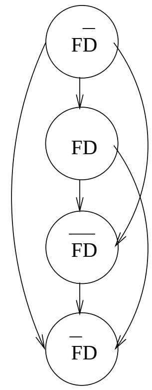

Figure 3: The belief state after aggregation in Example 3 when $s_t \sqsupset s_m \sqsupset s_p$.

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

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

AGR* indeed defines a legitimate belief state:

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

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 equal-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* $\mathrm{AGR}^*(S)$ *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 inferring new opinions via closure:

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

Encouragingly, AGR outputs a valid belief state:

Proposition 16 *If $S \subseteq \mathcal{S}$, then* $\mathrm{AGR}(S) \in \mathcal{B}$.

The output for our running space robot example is also reasonable:

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*

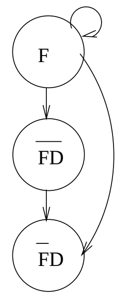

Figure 4: The belief state after aggregation in Example 5 when $s_t \sqsupset s_m \equiv s_p$.

equally credible. The aggregate belief state, shown in Figure 4, 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.

Furthermore, 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;* $\mathrm{AGR}(S) = \{(a, b), (c, b)\}$ *as desired. The concerned reader may note that s_2 is a "dictator" in the sense that s_2's opinions override all opposing opinions. However, this is reasonable in the context because all other sources have strictly lower rank.*

We 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 17 *Suppose $S \subseteq \mathcal{S}$.*

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

Another property of AGR is that its transitive closure part minimally extends the result of AGRRf to make it complete (i.e., all conflicts represented explicitly) in the sense that new opinions are only added between worlds already involved in a conflict:

Proposition 18 *Suppose $S \subseteq \mathcal{S}$, $\prec_* = \mathrm{AGRRf}(S)$, $\prec = \mathrm{AGR}(S)$, and $x \not\prec_* y$ for $x, y \in \mathcal{W}$. If $x \prec y$, then $x \bowtie y$.*

One small observation: $\mathrm{AGR}(\emptyset) = \emptyset$ is a property of our definition, reflecting the fact that we should not generate opinions out of nothing.

4.5 Arrow, Revisited

Finally, a strong argument in favor of AGR is that it satisfies Arrow's conditions. Technically, our setting is slightly different from that of Arrow's, so we need to modify each condition so that it is appropriate for our setting, yet retains the intended spirit of the original condition. Let f be an operator which aggregates the belief states $<^{s_1}, \ldots, <^{s_n}$ over $\mathcal{W}$ of n sources $s_1, \ldots, s_n \in S \subseteq \mathcal{S}$, respectively, let $\prec = f(<^{s_1}, \ldots, <^{s_n})$, and let $\sqsupseteq_S$ be a total preorder over S. We consider each condition separately.

Restricted range For our setting, the output of the aggregation function will be a modular, transitive belief state rather than a total preorder considered by Arrow.

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 preorders.

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

Pareto principle In Arrow's setting, the relations represented non-strict relative likelihood (preference, actually) so that the asymmetric restrictions of the relations were used to define the Pareto principle. However, in our setting, generalized belief states already represent strict likelihood. Consequently, we use the actual input and output relations of the aggregation function in place of their asymmetric restrictions to define the Pareto principle. Obviously, because of AGR's ability to introduce conflicts, it will not satisfy the original formal Pareto principle which would essentially require that if all sources have an unconflicted belief of one world being strictly more likely than another, this must also be true in the aggregate 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. This makes intuitive sense: if two worlds are put into conflict after aggregation due to a cycle involving other worlds, we may need to access these other worlds to be able to detect the conflict.

Definition 19 (modified) Independence of Irrelevant Alternatives (IIA)*: Suppose $s'_1, \ldots, s'_n \in \mathcal{S}$ such that $s_i \equiv_S s'_i$ for all i, and $\prec' = f(<^{s'_1}, \ldots, <^{s'_n})$. Further suppose $x <^{s_i} y$ iff $x <^{s'_i} y$ for all i, $x \not\bowtie y$, and $x \not\bowtie' y$. 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 equally ranked 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 should always push opinions through without them ever being contested.

Definition 20 (modified) Non-Dictatorship*: If $s_i \equiv_S s_j$ 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 19 *Let $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ and $AGR_f(<^{s_1}, \ldots, <^{s_n}) = \mathrm{AGR}(S)$. AGR_f satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

5. 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.

5.1 Formalization

An agent A is *informed by* a set of sources $S \subseteq \mathcal{S}$.[8] Agent A's *induced belief state* is the belief state formed by aggregating the belief states of its informant sources, i.e., $\mathrm{AGR}(S)$. We will use $A_\emptyset$ and $A_\mathcal{S}$ to denote special agents informed by $\emptyset$ and $\mathcal{S}$, respectively.

Assume the set of agents to fuse agree upon rank (and, consequently, $\sqsupseteq$).[9] 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 $\oplus(\mathcal{A})$, 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.

5.2 Computing Fusion Efficiently

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 so 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 20 *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 $A = \oplus(\mathcal{A})$, then $\left(\bigcup_{A_i \in \mathcal{A}} \prec^{A_i}\right)^+$ is A's induced belief state.*

8. Each source can be thought of as a primitive agent with fixed belief state.
9. 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.

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 $\oplus(A_1, A_2)$ 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 to maintain more information about each agent's informants. The question is whether we can do better than storing all the original sources.

We might wonder whether it is possible to somehow compute a credibility rank for each agent based on the credibility of her informant sources, then simply apply AGR to the agents' induced belief states. This works fine if, for every pair of agents, all the informants of one are more credible than those of the other. However, this does not work in general if each agent can have informants both more and less credible than those of another agent as the following example demonstrates:

Example 8 *Let $\mathcal{W} = \{a, b, c\}$. Suppose agent A_1 is informed by source s_1 with belief state $<^{s_1} = \{(a, b), (b, c), (a, c)\}$, and suppose agent A_2 is informed by sources s_0 and s_2 with belief states $<^{s_0} = \{(c, b), (b, a), (c, a)\}$ and $<^{s_2} = \{(b, a), (c, a)\}$, respectively. Further suppose that $s_2 \sqsupset s_1 \sqsupset s_0$. Then A_1's induced belief state is $<^{s_1}$ and A_2's is $<^{s_0}$. The belief state induced by $\oplus(A_1, A_2)$ is $\{(b, c), (c, a), (b, a)\}$. On the otherhand, if we rank A_1 over A_2 and apply* AGR *to their induced belief states, we get $<^{s_1}$; if we rank A_2 over A_1, we get $<^{s_0}$; and, if we rank them equally, we get the fully connected belief state. All of these are obviously incorrect.*

Hence, we need to store more information about the source of each opinion. However, we can still do better than keeping the sources around if sources are totally preordered by credibility. It is enough to store for each opinion of $\text{AGRRf}(S)$ 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 $\prec = \text{AGRRf}(S)$ and $l : \prec \rightarrow \mathcal{R}$ such that $l((x, y)) = \max(\{\text{rank}(s) : x <^s y, s \in S\})$. We use $\prec^A_r$ to denote the* restriction of A's pedigreed belief state to r*, that is, $\prec^A_r = \{(x, y) \in \prec : l((x, y)) = r\}$.*

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

Proposition 21 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$ and with pedigreed belief state $(\prec, l)$. Then $\prec^A_r$ is the relation*

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

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. We call this operation *pedigreed fusion*:

Definition 23 *Let S and $\mathcal{A}$ be as in Definition 21, $\sqsupseteq_S$, a total preorder, and $\mathcal{P}_{\mathcal{A}}$, the set of pedigreed belief states of the agents in $\mathcal{A}$. The* pedigreed fusion *of $\mathcal{P}_{\mathcal{A}}$, written $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}})$, is $(\prec, l)$ where*

1. *$\prec$ is the relation*

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

   over $\mathcal{W}$, and

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

Proposition 22 *Let $\mathcal{A}$, $\mathcal{P}_{\mathcal{A}}$, S, and $\sqsupseteq_S$ be as in Definition 23. Then $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}})$ is the pedigreed belief state of $\oplus(\mathcal{A})$.*

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 23 *If $\mathcal{A}$, $\mathcal{P}_{\mathcal{A}}$, and S are as in Definition 23, $\sqsupseteq_S$ is a total order, and $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}}) = (\prec, l)$, then $\prec^+ = \prec$.*

Let us return once more to the space robot scenario considered in Example 1 to illustrate pedigreed fusion.

Example 9 *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.*

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. Suppose S is the set of an agent's informant sources, $n = \|\mathcal{W}\|$, and $m = \|S\|$. Explicitly storing S (along with the rank of each source) requires $O(n^2 m)$ amount of space; this worst case bound is reached when all the sources' belief states are fully connected relations. On the

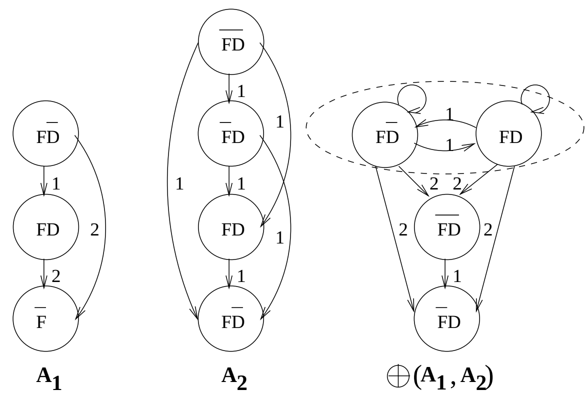

Figure 5: The pedigreed belief states of agent A_1 informed by s_m and s_t and of agent A_2 informed by s_p, and the result of their fusion in Example 9.

other hand, storing a pedigreed belief state only requires $O(n^2)$ space.[10] Moreover, not only does the enriched representation allow us to conserve space, 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 the strawman 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 10 *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 $\mathrm{rank}(s_2)$. In spite of this indistinguishability, if $A = \oplus(\{A_1, A_2\})$ and $A' = \oplus(\{A_1', A_2'\})$, then A's induced belief state equals $<^{s_2}$, i.e., $\{(a,b),(c,b)\}$, whereas A''s is $\{(a,b),(c,b),(a,c),(c,a),(a,a),(b,b),(c,c)\}$.*

Also notice that Maynard-Reid II and Shoham (2001) consider essentially the special case of fusing two agents informed by strictly-ranked sources. They show the surprising result that standard AGM belief revision can be modeled as the fusion of two agents, the informant and the informee, where the informant's sources are all strictly more credible than the informee's. Furthermore, they show that, because of its clean set-theoretic semantics, fusion provides a very attractive, semantically well-behaved solution to the difficult problem of iterated belief revision. Our general fusion definition satisfies all the examples of iterated fusion they describe.

10. These bounds assume that the amount of space needed to store each rank is bounded by some small constant.

6. Incorporating Voting

A potential drawback of the framework we have described is it does not account for "strength" of support. For example, we cannot differentiate between the situation where one thousand sources of the highest rank support $a < b$ and only one source of that rank supports $b < a$, and the situation where the one source supports $a < b$ and the thousand other sources support $b < a$. In both cases our framework yields a simple conflict between a and b rather than acknowledging the overwhelming support one way or the other. This additional information about strength of support is often sufficient to resolve what would otherwise have appeared to be a conflict.

To address this problem, we generalize our framework to incorporate voting. We first describe a family of aggregation operators based on voting of which AGR is a special case. In the process, we introduce a novel modular closure operator. We discuss properties of special members of this family including indiscriminate aggregation, simple majority, and unanimity, as well as attractive properties of the family as a whole. We then describe an extension of our setting to accommodate ranked individuals so that individuals of higher rank are given precedence during aggregation. Finally, we consider fusion.

6.1 Voting Functions

We will use a pairwise voting strategy similar to the well-known Condorcet's method. (For more on the Condorcet method and the other methods and results from standard voting theory we cite, see Black's (1958) classical reference on voting theory). Condorcet's method considers each pair of worlds separately, ranking world x over world y in the aggregate if and only if there are more votes for that ranking than there are for y over x. If one world beats or ties all other worlds, it is known as the *Condorcet winner*. We deviate from this method in that we use a fixed threshold proportion of support to decide on the acceptance of an opinion in the aggregate rather than the size of its support relative to that of the opposite opinion. Let $\mathrm{count}_S(x, y) = \|\{s \in S : x <^s y\}\|$ for any $S \subseteq \mathcal{S}$.

Definition 24 *Let $S \subseteq \mathcal{S}$. For $p \in [0, 1]$, the* voting function for p, *written* vt_p, *maps S to the relation*

$$\{(x, y) : \mathrm{count}_S(x, y) > 0 \textit{ and } \mathrm{count}_S(x, y)/\|S\| \geq p\}.$$

This definition falls under the class of voting systems Black (1958) calls *absolute majority* systems. It is motivated by the observation that relative support is many times less relevant than strength of support. The support for the two possible rankings of two worlds may be so low that neither can justifiably be considered part of the aggregate belief state. Similarly, the support for both alternative rankings may be so high that it may be more reasonable to introduce both and create a conflict rather than choose one with slightly higher support. Our strategy will not be appropriate for all applications, of course, but there are many instances when it is most appropriate. Also, our method satisfies a generalization of the *Condorcet criterion*, a widely accepted criteria of "good" voting systems that requires the Condorcet winner, if it exists, be the most likely world in the aggregate belief state. Our method never produces a strict ranking of two worlds opposite to that of Condorcet's method, although there will be cases where Condorcet's method ranks one world strictly more likely than another and our method produces agnosticism or a conflict. As a result, the Condorcet winner will always be among the most likely worlds. At any rate, our aggregation results do not depend significantly on the choice of voting strategy; one could easily replace it with a different strategy if desired.

That said, we make a few observations about our voting functions. First, the voting function definition requires that we accept an opinion if at least p proportion of the sources support it. However, we often want to specify that opinions only be accepted if *strictly more* than some cutoff proportion of sources support it. For example, the best-known voting function is the majority function where we only accept an opinion if it gets more than 50% of the vote. We can easily specify majority vote with the function $\mathrm{vt}_{0.5+\epsilon}(S)$ where $0 < \epsilon < 0.5/\|S\|$ (e.g., $\epsilon = 0.25/\|S\|$) so that tied opinions are rejected. In general, to only accept opinions garnering more than p proportion of the vote (for $0 \le p < 1$), it suffices to use the function $\mathrm{vt}_{p+\epsilon}(S)$ where

$$0 < \epsilon < \frac{1 - p\|S\| + \lfloor p\|S\| \rfloor}{\|S\|}$$

e.g., $\epsilon = (1 - p\|S\| + \lfloor p\|S\|\rfloor)/(2\|S\|)$.[11]

Second, it is immediately obvious that the aggregate relation may contain conflicts if $p \le 0.5$, even if the original source belief states are conflict-free. In fact, it is possible to get conflicts in the aggregate of conflict-free belief states even for larger p, as the following famous example demonstrates:

Example 11 *Let $\mathcal{W} = \{a, b, c\}$ and S be such that $1/3$ of the sources have belief state $\{(a, b), (b, c), (a, c)\}$, $1/3$ have $\{(b, c), (c, a), (b, a)\}$, and $1/3$ have $\{(c, a), (a, b), (c, b)\}$. Then $\mathrm{vt}_p(S) = \{(a, b), (b, c), (c, a)\}$, a cycle, for $1/3 < p \le 2/3$. This is known as the* Condorcet paradox (cf. Brams and Fishburn's (2002) voting survey).

Many solutions have been proposed for resolving such conflicts – using *Borda counts* or *instant runoff voting* (aka *single transferable vote*) (cf. Center for Voting and Democracy, 2002) are two popular examples. As before, we do not attempt to resolve the conflicts but, instead, make them explicit in a way that supports flexibility in the choice of resolution methodology and allows for semantically well-behaved iteration of aggregation.

Third, the end-point members of the family of voting functions have special significance. The voting function for 0 is equivalent to the union operator we saw earlier that takes all opinions seriously, i.e., is *indiscriminate*:

Proposition 24 *If $S \subseteq \mathcal{S}$, then* $\mathrm{vt}_0(S) = \mathrm{Un}(S)$.

At the other extreme, we have the voting function for 1. In this case, it is equivalent to taking the intersection of the sources' belief states, i.e., only accepting unanimous opinions:

Proposition 25 *If $S \subseteq \mathcal{S}$, then* $\mathrm{vt}_1(S) = \bigcap_{s \in S} <^s$.

In contrast to vt_0 which generates many conflicts, vt_1 generates a lot of agnosticism.

Fourth, voting functions are opinion-centered; that is, if the proportion of agnostic sources is larger than p, the voting function for p will not necessarily reflect this agnosticism as it would in the case of an opinion. If, for example, the belief states of three sources over

11. $\lfloor x \rfloor$ denotes the *floor* of x, i.e., the largest integer less than or equal to x.

$\mathcal{W} = \{a, b\}$ are $\{(a, b)\}$, $\{(b, a)\}$, and $\{\}$, respectively, then the voting function for $p = 1/3$ will produce a conflict with respect to a and b, not agnosticism. However, this is not to say that abstainers have no impact on the final result. The fact that abstainers are counted among the total number of voters has the effect that agnosticism with respect to a pair of worlds counts as a "no" vote for both possible opinions. This issue usually does not arise in standard voting schemes because these usually assume that sources totally rank candidates.

However, the most important observation is that members of the family of voting operators do not produce belief states in general. As we've already shown, vt_0 produces a modular relation that is not necessarily transitive. At the other end of the spectrum, vt_1 produces a transitive relation that is not necessarily modular:

Proposition 26 *Suppose $S \subseteq \mathcal{S}$. $\mathrm{vt}_1(S)$ is transitive but not necessarily modular.*

For the other members of the family, the result may be neither modular nor transitive, as the Condorcet paradox in Example 11 illustrates for $1/3 < p \leq 2/3$. In fact, we can construct such a scenario for every $0 < p < 1$:

Proposition 27 *If $\|\mathcal{W}\| \geq 3$, then for every $p \in (0, 1)$, there exists $\mathcal{S}$ such that $\mathrm{vt}_p(\mathcal{S})$ is neither modular nor transitive.*

Part of the problem is that voting may introduce conflicts which may imply other conflicts. As before, we need to take the transitive closure to infer these implied conflicts. In the Condorcet paradox example, this produces the fully connected belief state as we would hope. Unfortunately, closing under transitivity does not necessarily restore modularity as well, as the following example demonstrates:

Example 12 *Let $\mathcal{W} = \{a, b, c\}$ and S be such that $1/3$ of the sources have belief state $\{(a, b), (b, c), (a, c)\}$, $1/3$ have $\{(b, c), (c, a), (b, a)\}$, and $1/3$ have $\{(b, a), (b, c)\}$. Then, for $p > 2/3$, $\mathrm{vt}_p(S) = \mathrm{vt}_p(S)^+ = \{(b, c)\}$ which is not modular.*

We solve this problem by defining a natural modular closure operation which converts a transitive relation into a belief state. We will then define a modular-transitive closure operation which will take the result of an arbitrary voting function and transform it into a belief state using a transitive closure followed by a modular closure.

6.2 Modular-Transitive Closure

We start by defining a helper function which returns the *level* of a world in a relation, i.e., the length of the longest path (along strict edges) from a world to a member of the choice set of $\mathcal{W}$. For convenience, throughout this modular-transitive closure subsection we will use $\preceq$ to denote an arbitrary relation, $\prec$ and $\bowtie$ to denote its asymmetric and symmetric restrictions, respectively.

Definition 25 *The* level *of $x \in \mathcal{W}$ in a transitive relation $\preceq$ over $\mathcal{W}$, written $\mathrm{lev}_\preceq(x)$, is*

$$\mathrm{lev}_\preceq(x) = \begin{cases} 0 & \text{if } x \in \mathrm{ch}(\mathcal{W}, \preceq) \\ 1 + \max\limits_{y \in \mathcal{W}} (\{\mathrm{lev}_\preceq(y) : y \prec x\}) & \text{otherwise.} \end{cases}$$

(Recall that ch is the choice set function defined in Definition 2.) The following simple properties relating $\preceq$ and $\mathrm{lev}_\preceq$ are immediate:

Proposition 28 *Suppose* $\preceq$ *is a transitive relation over* $\mathcal{W}$ *and* $x, y \in \mathcal{W}$.

1. *If* $x \prec y$*, then* $\mathrm{lev}_{\preceq}(x) < \mathrm{lev}_{\preceq}(y)$.
2. *If* $x \bowtie y$*, then* $\mathrm{lev}_{\preceq}(x) = \mathrm{lev}_{\preceq}(y)$.
3. *If* $\mathrm{lev}_{\preceq}(x) < \mathrm{lev}_{\preceq}(y)$*, then* $\exists z.\ \mathrm{lev}_{\preceq}(z) = \mathrm{lev}_{\preceq}(x) \land z \prec y$.
4. *If* $\mathrm{lev}_{\preceq}(x) = \mathrm{lev}_{\preceq}(y)$*, then* $x \preceq y$ *iff* $y \preceq x$.

We now define the *modular closure* of a relation to be the relation that results from fully connecting all equi-level alternatives unless they are fully disconnected:

Definition 26 *The* modular closure $\mathrm{MC}(\preceq)$ *of a transitive relation* $\preceq$ *over* $\mathcal{W}$ *is the relation such that* $(x, y) \in \mathrm{MC}(\preceq)$ *iff*

1. $\mathrm{lev}_{\preceq}(x) < \mathrm{lev}_{\preceq}(y)$ *or*
2. $\mathrm{lev}_{\preceq}(x) = \mathrm{lev}_{\preceq}(y)$ *and* $\exists x', y'.\ \mathrm{lev}_{\preceq}(x') = \mathrm{lev}_{\preceq}(y') = \mathrm{lev}_{\preceq}(x) \land x' \bowtie y'$.

Intuitively, as long as we have reason to doubt that some pair in a level are interchangeable, we doubt all the pairs at that level, but only then. Note that the definition of MC is similar to one of the equivalent constructions of the rational closure Lehmann and Magidor (1992) describe.

We see that MC indeed makes any transitive relation modular while preserving transitivity:

Proposition 29 *If* $\preceq$ *is a transitive relation over* $\mathcal{W}$*, then* $\mathrm{MC}(\preceq) \in \mathcal{B}$.

MC is an additive process that changes a relation minimally to achieve modularity while preserving transitivity and the levels of the worlds:

Proposition 30 *Suppose* $\preceq$ *is a transitive relation over* $\mathcal{W}$ *and* $\preceq^* = \mathrm{MC}(\preceq)$.

1. $\preceq \subseteq \preceq^*$ *and* $\prec \subseteq \prec^*$.
2. *If* $\preceq$ *is modular, then* $\preceq^* = \preceq$.
3. $\mathrm{lev}_{\preceq^*}(x) = \mathrm{lev}_{\preceq}(x)$ *for all* $x \in \mathcal{W}$.
4. *If* $\preceq' \in \mathcal{B}$ *such that* $\preceq \subseteq \preceq'$ *and* $\mathrm{lev}_{\preceq'}(x) = \mathrm{lev}_{\preceq}(x)$ *for all* $x \in \mathcal{W}$*, then* $\preceq^* \subseteq \preceq'$.

We now define the *modular, transitive closure* of an arbitrary relation $\preceq$ as MC applied to the transitive closure of $\preceq$, and show that the result is a belief state:

Definition 27 *The* modular, transitive closure $\mathrm{MT}(\preceq)$ *of a relation* $\preceq$ *over* $\mathcal{W}$ *is the relation* $\mathrm{MC}\,(\preceq^+)$.

Proposition 31 *If* $\preceq$ *is a relation over* $\mathcal{W}$*, then* $\mathrm{MT}(\preceq) \in \mathcal{B}$.

MT is also a minimally additive operator:

Proposition 32 *Suppose* $\preceq$ *is a relation over* $\mathcal{W}$ *and* $\preceq^* = \mathrm{MT}(\preceq)$.

1. $\preceq \subseteq \preceq^*$.
2. *If* $\preceq$ *is transitive, then* $\preceq^* = \mathrm{MC}(\preceq)$.
3. *If* $\preceq$ *is modular, then* $\preceq^* = \preceq^+$.
4. *If* $\preceq$ *is modular and transitive, then* $\preceq^* = \preceq$.
5. *If* $\preceq$ *has no conflicts, then neither does* $\preceq^*$.

6.3 The Aggregation Family

We are now fully equipped to solve the problem of incorporating voting into aggregation. First, consider the special case where all sources have the same rank. Our aggregation operators will construct an aggregate belief state by first applying voting, then closing under MT:

Definition 28 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$, *then* $\mathrm{AGREq}_p(S) = \mathrm{MT}\,(\mathrm{vt}_p(S))$.

Proposition 33 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$, *then* $\mathrm{AGREq}_p(S) \in \mathcal{B}$.

We can now easily generalize this definition to accommodate a ranking on the sources. We accept an opinion if enough individuals at the highest rank with an opinion support it, then close under MT:

Definition 29 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$, *then* $\mathrm{AGRRf}_p(S)$ *is the relation*

$$\Big\{(x,y) : \exists s \in S.\ x <^s y \wedge (x,y) \in \mathrm{vt}_p(\{s' \in S : s' \equiv_S s\}) \wedge \Big(\forall s' \in S.\ s' \sqsupset_S s \Rightarrow x \approx^{s'} y\Big)\Big\}.$$

Definition 30 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$, *then* $\mathrm{AGR}_p(S) = \mathrm{MT}(\mathrm{AGRRf}_p(S))$.

Proposition 34 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$, *then* $\mathrm{AGR}_p(S) \in \mathcal{B}$.

All the aggregation functions we have encountered so far are special cases of this general family:

Proposition 35 *Suppose* $S \subseteq \mathcal{S}$ *and* $p \in [0,1]$.

1. *If* $\sqsupseteq_S$ *is fully connected, then* $\mathrm{AGR}_p(S) = \mathrm{AGREq}_p(S)$.
2. *If* $\sqsupseteq_S$ *is a total order, then* $\mathrm{AGR}_p(S) = \mathrm{AGRRf}_p(S) = \mathrm{AGRRf}(S) = \mathrm{AGR}(S)$.
3. $\mathrm{AGR}_0(S) = \mathrm{AGR}(S)$.

As an obvious consequence of the last property, AGR_0 satisfies the modified Arrovian conditions.

Corollary 35.1 *Let* $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ *and* $\mathrm{AGR}_f(<^{s_1}, \ldots, <^{s_n}) = \mathrm{AGR}_0(S)$. AGR_f *satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

6.4 Fusion

Fusion is still defined as in Definition 21, i.e., the belief state created by fusion of a set of agents is the aggregate belief state of the agents' cumulative informant sets. However, we now use AGR_p rather than AGR to compute the aggregate belief state.

Once again, we want to compute fusion without storing all the belief states of all the informant sources, if possible. Unfortunately, this is not possible in general for aggregation functions based on voting. The reason is we need to keep track of the actual identity of the sources supporting each opinion so as to avoid "double-counting" sources shared by multiple agents.

However, we can often do better than the $O(n^2m)$ space required to store the full sources, where $n = \|\mathcal{W}\|$ and m is the number of informant sources. We only store those parts that matter. Specifically, for a given source, we only store those opinions for which the source is one of the highest ranked supporting an opinion over the corresponding worlds. We can effectively accomplish this by extending the pedigreed belief state so that we label each opinion not only with the rank of the highest ranking sources supporting an opinion over the corresponding worlds, but also with the set of unique identifiers for the sources supporting the particular opinion. We also maintain a table that stores, for each rank represented in the set of informant sources, the set of identifiers for all sources at that rank. We call the resulting representation a *support pedigreed belief state*.

Definition 31 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$. A's* support pedigreed belief state *is a triple* $(l, \mathrm{sup}, \mathrm{rtab})$ *where*

- $l : \mathcal{W} \times \mathcal{W} \to \mathcal{R} \cup \{\clubsuit\}$ *such that* $l(x, y) = \max(\{\mathrm{rank}(s) : x \not\approx^s y,\ s \in S\} \cup \{\clubsuit\})$ *where* $\clubsuit \notin \mathcal{R}$ *and* $\clubsuit < r$ *for all* $r \in \mathcal{R}$,
- $\mathrm{sup} : \mathcal{W} \times \mathcal{W} \to 2^S$ *such that* $\mathrm{sup}(x, y) = \{s \in S : \mathrm{rank}(s) = l(x, y),\ x <^s y\}$, *and*
- $\mathrm{rtab} : \mathrm{ranks}(S) \to 2^S$ *such that* $\mathrm{rtab}(r) = \{s \in S : \mathrm{rank}(s) = r\}$.

Note that l is symmetric: $l(x, y) = l(y, x)$. On the other hand, sup is not. Now, we can easily compute an agent's belief state from its support pedigreed belief state. To compute the proportion of support for a particular opinion, we simply divide the size of the support set for that opinion by the number of informant sources with the labeled rank.

Proposition 36 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$, with support pedigreed belief state* $(l, \mathrm{sup}, \mathrm{rtab})$*, and using aggregation function AGR_p for $p \in [0, 1]$. A's belief state is the relation*

$$\mathrm{MT}\left(\{(x, y) : \|\mathrm{sup}(x, y)\| > 0 \text{ and } \|\mathrm{sup}(x, y)\| / \|\mathrm{rtab}(l(x, y))\| \geq p\}\right)$$

Observe that, unlike with pedigreed belief states, support pedigreed belief states label *all* possible opinions, not just those appearing in the agent's induced belief state, i.e., whose support falls below the threshhold. The reason is another agent may come along later with enough new votes to cross the threshold, in which case the votes from the earlier sources become relevant. Similarly, support pedigreed belief states maintain rank information even for ranks not appearing as labels. No source of a particular rank may currently support any opinion, but another agent may later bring sources of that rank supporting an opinion hitherto unsupported by any source of equal or higher rank. The correct computation of the proportion of support for this opinion must take into account the earlier sources.

Before we address fusion, let us consider the space required to store a support pedigreed belief state $(l, \mathrm{sup}, \mathrm{rtab})$. l requires $O(n^2)$ space, rtab requires $\Theta(m)$ space, and, if $r_{\max}$ denotes the number of sources of a rank having the most sources with that rank, sup requires $O(n^2 r_{\max})$ space, for a total of $O(n^2 r_{\max} + m)$ space.[12] Thus, in a best-case scenario where, for example, sources are strictly ranked, a support pedigreed belief state only requires $O(n^2 + m)$ space since each opinion has at most one supporter. However, in

12. As before, we assume representing ranks requires constant space. We assume that we can represent each source label with constant space as well.

the worst-case scenario where, for example, sources are all equally ranked so that $r_{\max} = m$, we will still need $O(n^2 m)$ space.

Now, computing the support pedigreed belief state resulting from fusion is straightforward. For each opinion, we set l to be the highest l value for that opinion among the agents and set sup to be the union of sup sets for that opinion of all agents with that l value. And for each rank represented in some agent's rank table, we set rtab to be the union of the rtab sets of all agents for whom it is defined.

Definition 32 *Let S and $\mathcal{A}$ be as in Definition 21, $\sqsupseteq_S$, a total preorder, and $\mathcal{P}_{\mathcal{A}}$, the set of support pedigreed belief states of the agents in $\mathcal{A}$. The* support pedigreed fusion *of $\mathcal{P}_{\mathcal{A}}$, written $\oplus_{\text{sup}}(\mathcal{P}_{\mathcal{A}})$, is $(l, \text{sup}, \text{rtab})$ where*

1. $l : \prec \rightarrow \mathcal{R}$ *such that* $l((x, y)) = \max(\{l'(x, y) : (l', \text{sup}', \text{rtab}') \in \mathcal{P}_{\mathcal{A}}\})$,
2. $\text{sup} : \mathcal{W} \times \mathcal{W} \rightarrow 2^S$ *such that*

$$\text{sup}(x, y) = \bigcup_{(l', \text{sup}', \text{rtab}') \in \mathcal{P}_{\mathcal{A}},\ l'(x,y) = l(x,y)} \text{sup}'(x, y),$$

and

3. $\text{rtab} : \text{ranks}(S) \rightarrow 2^S$ *such that*

$$\text{rtab}(r) = \bigcup_{(l', \text{sup}', \text{rtab}') \in \mathcal{P}_{\mathcal{A}},\ r \in \text{range}(\text{rtab}')} \text{rtab}'(r).$$

Proposition 37 *Let $\mathcal{A}$, $\mathcal{P}_{\mathcal{A}}$, S, and $\sqsupseteq_S$ be as in Definition 32. Then $\oplus_{\text{sup}}(\mathcal{P}_{\mathcal{A}})$ is the support pedigreed belief state of $\oplus(\mathcal{A})$.*

Thus, in addition to the potential savings in space gained by using support pedigreed belief states, we also potentially save in the time needed to compute fusion since, for a given opinion, we do not need to consider the opinions of lesser ranked sources.

7. Related Work

Much of the work in belief aggregation has been geared towards unbiased kinds of belief pooling. Besides the work in social choice we described in Section 3.2, recent attempts from the belief revision community (e.g. Borgida & Imielinski, 1984; Baral, Kraus, Minker, & Subrahmanian, 1992; Liberatore & Schaerf, 1995; Makinson, 1997; Revesz, 1997; Konieczny & Pérez, 1998; Meyer, 2001; Benferhat, Dubois, Kaci, & Prade, 2002) have sought to modify the AGM theory to capture "fair" revisions, that is, revisions where the revisee and reviser's beliefs are treated equally seriously. Like our proposal, Benferhat et al. and Meyer accommodate iterative merging. Benferhat et al.'s proposal is also distinct in that they approach the problem from a possibilistic logic point of view. Besides the restriction to equally-ranked sources, these fairness-based proposals differ from ours in that they are generally syntactic in nature in the sense that sentences are prioritized rather than possible worlds. Meyer's proposal is an exception; his belief states are *epistemic states*, structures in the style of Spohn's (1988) ordinal conditional functions (aka κ-rankings). In fact, Meyer, Ghose, and

Chopra (2001) have shown that a number of simple aggregation operators on epistemic belief states *also* satisfy Arrow's postulates when appropriately modified for this context (unrestricted domain, restricted range, and IIA in particular need modification). Unfortunately, epistemic states are enriched total preorders and, thus, suffer from the problems we described earlier, i.e., the inability to explicitly handle conflicts.

Cantwell's (1998) work is also syntactic in nature, but does allow for sources of differing credibility. Cantwell addresses a complementary problem to our own: deciding what information to reject given the subset of informing sources rejecting the information. He assumes a generalization of our credibility ordering, a partial preorder over sets of sources. He explores ways of inducing a partial preorder over sentences based on this ordering, then uses this ordering to determine a subset (although not all) of the sentences to reject. Another difference from our work is that he only considers the non-counterfactual beliefs of sources.

We are not, of course, the first to consider using the lexicographic ordering for aggregation purposes. Lexicographic operators have long been studied in the fields of management and social science; Fishburn (1974) gives a good survey of much of that work. More recently, researchers in artificial intelligence have taken an interest in these operators; examples include Grosof (1991), Maynard-Reid II and Shoham (2001) and Andréka, Ryan, and Schobbens (2002).

Grosof uses lexicographic aggregation of preorders as a means of tackling the problem of default reasoning in the presence of conflicting defaults. Besides the more general preorders being aggregated, another interesting difference from our work is that although Grosof does not allow for sources of equal rank, he does allow for sources of *incomparable* rank, i.e., the ranking on sources is a strict *partial* order. Thus, in the extreme case where the ordering is completely disconnected, the operator reduces to our Un operator (and, thus, does not necessarily preserve transitivity).

Andréka et al., on the other hand, frame their work in the context of preference aggregation. They go one step further than Grosof and allow input relations to be arbitrary. They prove that the lexicographic operator is the only one that satisfies a variation on Arrow's properties – unanimity, IIA, preservation of transitivity, and a weaker version of non-dictatorship. (We should point out that from the perspective of Arrow's *original* framework, the relation with the highest priority is *always* a dictator.) They describe a collection of other properties besides transitivity preserved by the operator. However, as in our work, they do not preserve totality.

Our work derives much of its inspiration from Maynard-Reid II and Shoham's work. They restrict their attention to total preorders, but this does not create problems because they assume sources to be totally ordered. They focus, instead, on the strong connection between belief aggregation and iterated belief revision. They show that $\oplus$ can be used as an iterated belief operator in an AGM-based setting, then compare its properties as such against those of a representative sampling of well-known iterated belief operator proposals – Boutilier's (1996) natural revision, Darwiche and Pearl's (1997) operators, Spohn's (1988) ordinal conditional function revision, Lehmann's (1995) widening rank model revision, and Williams's (1994) conditionalization and adjustment operators. They show that $\oplus$ is the only operator among them that is semantically well-behaved: the results of all the other operators depend on the order of iteration.

Finally, to our knowledge, none of these related approaches outside of social choice have yet been extended to incorporate voting.

8. Conclusion

We have described a semantically clean representation – the class of modular, transitive relations – for collective qualitative 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 preordered by credibility. We showed that this operator circumvents Arrow's Impossibility result in a satisfactory manner. We also described a mechanism for fusing the belief states of different agents that iterates well and extended the framework to incorporate voting.

We have assumed that all agents share the credibility ranking on sources. In general, these rankings can vary among agents, and even change over time. Furthermore, an agent's ranking function can depend on the context; different sources may have different areas of expertise. Exploring the behavior of fusion in these more general settings is an obvious next step.

Note that although we have described operators to incorporate voting, under no condition will any of these ever side with lower rank sources when they conflict with higher rank sources, no matter how many of these disagreeing lower rank sources there are. An aggregation scheme that behaves differently would have to be built on fundamentally different assumptions than our framework.

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.

Acknowledgments

We want to thank Yoav Shoham and the anonymous referees of this paper for their insightful and comprehensive feedback. A preliminary version of this paper appears in the *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000)* (the first author's last name was Maynard-Reid II at that time). This work was partially supported by a National Physical Science Consortium Fellowship and the Jean et Hélène Alfassa Fund for Research in Artificial Intelligence.

Appendix 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 asymmetric restriction. 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* ch *defines a choice function iff $\leq$ is acyclic.*

Proof: See Sen's (1986) proof. □

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's (1963) proof. □

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

Proof: Let $\preceq$ be a total, quasi-transitive, non-transitive relation. 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 *$\mathcal{T}_< \subset \mathcal{B}$ and is the set of irreflexive relations in $\mathcal{B}$.*

Proof: Let $x, y, z \in \mathcal{W}$. We first show that $\mathcal{T}_< \subset \mathcal{B}$. Let $\prec \in \mathcal{T}_<$. Then there exists $\preceq \in \mathcal{T}$ such that $\prec$ is the asymmetric restriction of $\preceq$. By definition, $\preceq$ is transitive, so by Proposition 1, so is $\prec$. 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 $x \preceq z$. If $y \preceq z$, then $z \not\preceq x$ (otherwise $y \preceq x$ by transitivity, a contradiction), so $x \prec z$. And if, on the other hand, $y \not\preceq z$, then $z \preceq y$ by totality, so $z \prec y$. Suppose, instead, $z \preceq x$. Then $z \preceq y$ by transitivity and $y \not\preceq z$ (otherwise $y \preceq x$ by transitivity, a contradiction), so $z \prec y$. Thus, $x \prec z$ or $z \prec y$, so $\prec$ is modular.

Now we show that $\prec \in \mathcal{B}$ is in $\mathcal{T}_<$ if and only if it is irreflexive. If $\prec \in \mathcal{T}_<$, it is asymmetric, so it is irreflexive. Suppose, instead, $\prec$ is irreflexive. We define a relationship $\preceq$, show that $\prec$

is its asymmetric restriction, and show that $\preceq$ is in $\mathcal{T}$. Let $\preceq$ be defined as $x \preceq y$ iff $y \not\prec x$. We first show that $\prec$ is the asymmetric restriction of $\preceq$. Suppose $\prec'$ is the asymmetric restriction 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 next 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 $x \preceq z$, and, thus $\preceq$ is transitive. □

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 33 *$x \equiv y$ 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 $x \bowtie y$, then $[x] = [y]$.*

Proof: Suppose $x, y \in \mathcal{W}$ and $x \bowtie y$. 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 34 *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 $x' \prec y'$. $\square$

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 $[y] = [z]$, 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 $\leq$. $\square$

We name the members of the partition $W_0, \dots, W_n$ 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 $y \not\prec x$, $[y] \not\leq [x]$ 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 $\mathbf{W} = \langle W_0, \dots, W_n \rangle$ 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. $\square$

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 asymmetric restriction 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 the 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 asymmetric restriction. (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 asymmetric restriction 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 the 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* $\text{Un}(S)$ *is modular but not necessarily transitive.*

Proof: Let $\prec = \text{Un}(S)$. 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 $\text{Un}(S)$, $x \prec z$ 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 $<^{s_2} = \{(b, a), (c, a)\}$. $\text{Un}(S)$ is not transitive. □

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

Proof: The transitive closure of any relation is transitive. Since $\text{Un}(S)$ is modular, the transitive closure of $\text{Un}(S)$ is also modular by Proposition 1. □

Proposition 13 *If* $S \subseteq \mathcal{S}$*, then* $\text{AGRRf}(S)$ *is modular but not necessarily transitive.*

Proof: We first prove modularity. Suppose $x, y, z \in \mathcal{W}$ and $(x, y) \in \text{AGRRf}(S)$. Then there exists $s \in S$ such that $x <^s y$ and for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$. By modularity of $<^s$, either $x <^s z$ or $z <^s y$. Since S is finite, this implies that either there exists $s' \in S$ such that $x <^{s'} z$ and for all $s'' \sqsupset_S s' \in S$, $x \approx^{s''} z$, or there exists $s' \in S$ such that $y <^{s'} z$ and for all $s'' \sqsupset_S s' \in S$, $y \approx^{s''} z$. Thus, $(x, z) \in \text{AGRRf}(S)$ or $(z, y) \in \text{AGRRf}(S)$, so $\text{AGRRf}(S)$ is modular.

Suppose $\mathcal{W} = \{x, y, z\}$ and $S = \{s_1, s_2\}$ such that $s_1 = \{(x, y), (z, y)\}$, $s_2 = \{(y, x), (y, z)\}$, and $s_1 \equiv_S s_2$. Then $\text{AGRRf}(S) = \{(x, y), (z, y), (y, x), (y, z)\}$ which is not transitive. □

Proposition 14 *If $S \subseteq \mathcal{S}$ and $\sqsupset_S$ is a total order, then* $\mathrm{AGRRf}(S) \in \mathcal{B}$.

Proof: We've already proven in Proposition 13 that $\mathrm{AGRRf}(S)$ is modular. Let $\prec = \mathrm{AGRRf}(S)$ and suppose $x, y, z \in \mathcal{W}$. 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 $s_1' \in S$, $s_1' \sqsupset_S s$ 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 $s_2' \in S$, $s_2' \sqsupset_S s$ implies $y \not<^{s_2'} z$ and $z \not<^{s_2'} y$. Suppose $s_1 \sqsupset_S s_2$ (the case $s_2 \sqsupset_S s_1$ 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 $s' \sqsupset_S s_1$. Then $x \not<^{s'} y$ and $y \not<^{s'} x$. And, since $s_1 \sqsupset_S s_2$, $s' \sqsupset_S s_2$, 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 15 *If $S \subseteq \mathcal{S}$, then* $\mathrm{AGR}^*(S) \in \mathcal{B}$.

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

$$
\begin{aligned}
\mathrm{AGR}^*(S) &= \big\{(x,y) : \exists r \in \mathcal{R}.\ x <_r y \wedge \big(\forall r' \in \mathrm{ranks}(S).\ r' > r \Rightarrow x \approx_{r'} y\big)\big\} \\
&= \Big\{(x,y) : \exists s \in S'.\ x <^s y \wedge \Big(\forall s' \in S'.\ s' \sqsupset_{S'} s \Rightarrow x \approx^{s'} y\Big)\Big\} \\
&= \mathrm{AGRRf}(S').
\end{aligned}
$$

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 14. □

Proposition 16 *If $S \subseteq \mathcal{S}$, then* $\mathrm{AGR}(S) \in \mathcal{B}$.

Proof: By Proposition 13, $\mathrm{AGRRf}(S)$ is modular. $\mathrm{AGRRf}(S)^+$ is obviously transitive, and, by Proposition 1, it is modular as well. □

Proposition 17 *Suppose $S \subseteq \mathcal{S}$.*

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

Proof:

1. Suppose $\sqsupset_S$ is fully connected. Then the second half of the definition of AGRRf is vacuously true so that $\mathrm{AGRRf}(S)$ simplifies to $\{(x,y) : \exists s \in S.\ x <^s y\}$. But this is exactly $\bigcup_{s \in S} <^s$, i.e., $\mathrm{Un}(S)$, so $\mathrm{AGR}(S) = \mathrm{AGRRf}(S)^+ = \mathrm{Un}(S)^+ = \mathrm{AGRUn}(S)$.
2. Suppose $\sqsupset_S$ is an total order. By Proposition 14, $\mathrm{AGRRf}(S)$ is transitive, so $\mathrm{AGR}(S) = \mathrm{AGRRf}(S)^+ = \mathrm{AGRRf}(S)$.

□

Proposition 18 *Suppose $S \subseteq \mathcal{S}$, $\prec_* = \mathrm{AGRRf}(S)$, $\prec = \mathrm{AGR}(S)$, and $x \not\prec_* y$ for $x, y \in \mathcal{W}$. If $x \prec y$, then $x \bowtie y$.*

Proof: We first show the following lemma:

Lemma 18.1 *Suppose $S \subseteq \mathcal{S}$ and $\prec_* = \mathrm{AGRRf}(S)$. 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 $s' \sqsupset_S s_1 \in S$, $x_{i-1} \approx^{s'} x_i$. Similarly, there exists $s_2 \in S$ such that $x_i <^{s_2} x_{i+1}$ and, for all $s' \sqsupset_S s_2 \in S$, $x_i \approx^{s'} x_{i+1}$. 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 $s_1 \sqsupset_S s_2$. Then $x_i \approx^{s_1} x_{i+1}$ so, by modularity, $x_{i-1} <^{s_1} x_{i+1}$. But then $x_{i-1} \prec_* x_{i+1}$, a contradiction. Similarly, we derive a contradiction if $s_2 \sqsupset s_1$. Thus, $s_1 \equiv_S s_2$.

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 $s_1 \equiv_S s_2$, 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. □

Now 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 18.1, $y = x_n \prec_* \cdots \prec_* x_0 = x$, so $y \prec x$ and $x \bowtie y$. □

Proposition 19 *Let $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ and $\mathrm{AGR}_f(<^{s_1}, \ldots, <^{s_n}) = \mathrm{AGR}(S)$. AGR_f satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

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

Restricted range: AGR_f satisfies restricted range by Proposition 16.

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 $s' \sqsupset_S s \in S$, $x \not<^s y$ 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 19.1 *Suppose $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$, $S' = \{s'_1, \ldots, s'_n\} \subseteq \mathcal{S}$, $s_i \equiv_S s'_i$ for all i, $\prec_* = \mathrm{AGRRf}(S)$, and $\prec'_* = \mathrm{AGRRf}(S')$. 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 $s_i \equiv_S s'_i$, 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 17, the closure step introduces no new opinions under these conditions.)

Now, suppose $x, y \in \mathcal{W}$, $x <^{s_i} y$ iff $x <^{s'_i} y$ for all i, $x \not\bowtie y$, and $x \not\bowtie' y$. We show that $x \prec y$ implies $x \prec' y$ (the other direction is identical). Suppose $x \prec y$. Let $\prec_* = \text{AGRRf}(S)$ and $\prec'_* = \text{AGRRf}(S')$. Since $x \not\bowtie y$, $x \prec_* y$ by Proposition 18. But then $x \prec'_* y$ by Lemma 19.1, 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 19 PROOF)

Proposition 20 *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 $A = \oplus(\mathcal{A})$, 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 20.1 *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 $\preceq = \left(\bigcup_{\leq \in \Pi} \leq^+\right)^+$, $\preceq' = \left(\bigcup_{\leq \in \Pi} \leq\right)^+$, and $a, b \in \Omega$. Suppose $a \preceq b$. 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

$$a = w_0 = x_{00} \leq_0 \cdots \leq_0 x_{0m_0} = w_1 = \cdots = w_{n-1} = x_{(n-1)0} \leq_{n-1} \cdots \leq_{n-1} x_{(n-1)m_{n-1}} = w_n$$

and $w_n = b$, so $a \preceq' b$.

Now suppose $a \preceq' b$. 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 $a \preceq b$. □

Now, let $\prec$ be the belief state induced by $\oplus(\mathcal{A})$. Then $\prec = \text{AGR}(S)$. By Proposition 17, $\prec = \text{AGRUn}(S)$, so

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

By the lemma,

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

□

Proposition 21 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$ and with pedigreed belief state $(\prec, l)$. Then $\prec_r^A$ is the relation*

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

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 $s' \sqsupset_S s \in S$, $x \approx^{s'} y$. In particular, if $x <^{s'} y$ for some $s' \in S$, then $s \sqsupseteq_S s'$, so $\text{rank}(s) \geq \text{rank}(s')$. Thus,

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

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

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

Therefore, $x \prec_r^A y$. □

Proposition 22 *Let $\mathcal{A}$, $\mathcal{P}_{\mathcal{A}}$, S, and $\sqsupseteq_S$ be as in Definition 23. Then $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}})$ is the pedigreed belief state of $\oplus(\mathcal{A})$.*

Proof: Let $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}}) = (\prec, l)$, $\prec' = \text{AGRRf}(S)$, and $l' : \prec' \to \mathcal{R}$ 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 $s' \sqsupset_S s \in S$, $x \not<^{s'} y$ 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 $x <^{s} y$, $\text{rank}(s) = r$, and, for every $s_1 \sqsupset_S s \in S_i$, $x \not<^{s_1} y$ 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 $\sqsupseteq_S$ is a total preorder, it suffices to show that $s \sqsupseteq_S s'$. 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

$x \prec^{A_j}_{\text{rank}(s')} y$ or $y \prec^{A_j}_{\text{rank}(s')} x$. But since $x \prec^{A_i}_r y$, $r = \text{rank}(s) \geq \text{rank}(s')$, so $s \sqsupseteq_S s'$. Furthermore, $l'((x,y)) = \text{rank}(s) = r = l((x,y))$.

Now suppose $x \prec' y$. We show that $x \prec y$, i.e., there exists A_i and r such that $x \prec^{A_i}_r y$ and, for every $A_j \in \mathcal{A}$ and $r' > r \in \mathcal{R}$, $x \not\prec^{A_j}_{r'} y$ and $y \not\prec^{A_j}_{r'} x$, 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 $s' \sqsupset_S s \in S$, $x \not<^{s'} y$ and $y \not<^{s'} x$. Suppose $s \in S_i$. Since $S_i \subseteq S$, it is also the case that for every $s' \sqsupset_S s \in S_i$, $x \not<^{s'} y$ and $y \not<^{s'} x$, so $x \prec^{A_i}_{\text{rank}(s)} y$. Now, let A_j and r' be such that $x \prec^{A_j}_{r'} y$ or $y \prec^{A_j}_{r'} x$. It suffices to show that $\text{rank}(s) \geq r'$. By Proposition 21, there exists $s' \in S_j$ such that $x <^{s'} y$ or $y <^{s'} x$ and $\text{rank}(s') = r'$. But then $s \sqsupseteq_S s'$, so $\text{rank}(s) \geq \text{rank}(s') = r'$. Furthermore, $l((x,y)) = \text{rank}(s) = l'((x,y))$. $\square$

Proposition 23 *If $\mathcal{A}$, $\mathcal{P}_{\mathcal{A}}$, and S are as in Definition 23, $\sqsupseteq_S$ is a total order, and $\oplus_{\text{ped}}(\mathcal{P}_{\mathcal{A}}) = (\prec, l)$, then $\prec^+ = \prec$.*

Proof: Since $\sqsupseteq_S$ is a total order, $\text{AGR}(S) = \text{AGRRf}(S)$ by Proposition 17. Thus, $\prec = \text{AGRRf}(S) = \text{AGR}(S) = \text{AGRRf}(S)^+ = \prec^+$. $\square$

Proposition 24 *If $S \subseteq \mathcal{S}$, then $\text{vt}_0(S) = \text{Un}(S)$.*

Proof: Suppose $(x,y) \in \text{Un}(S)$. Then $S \neq \emptyset$ and $x <^s y$ for some $s \in S$. Thus, $\text{count}_S(x,y) > 0$ and $\text{count}_S(x,y)/\|S\| > 0$, so $(x,y) \in \text{vt}_0(S)$. Suppose, instead, $(x,y) \notin \text{Un}(S)$. Then $x \not<^s y$ for all $s \in S$, so $\text{count}_S(x,y) = 0$, so $(x,y) \notin \text{vt}_0(S)$. $\square$

Proposition 25 *If $S \subseteq \mathcal{S}$, then $\text{vt}_1(S) = \bigcap_{s \in S} <^s$.*

Proof: Suppose $(x,y) \in \bigcap_{s\in S} <^s$. Then $S \neq \emptyset$ and $x <^s y$ for all $s \in S$. Thus, $\text{count}_S(x,y) > 0$ and $\text{count}_S(x,y)/\|S\| \geq 1$, so $(x,y) \in \text{vt}_1(S)$. Suppose, instead, $(x,y) \notin \bigcap_{s\in S} <^s$. Then there exists $s \in S$ such that $x \not<^s y$, so $\text{count}_S(x,y) < \|S\|$. Thus, $\text{count}_S(x,y)/\|S\| < 1$, so $(x,y) \notin \text{vt}_1(S)$. $\square$

Proposition 26 *Suppose $S \subseteq \mathcal{S}$. $\text{vt}_1(S)$ is transitive but not necessarily modular.*

Proof: Let $\mathcal{W} = \{x,y,z\}$ and $S = \{s_1, s_2\}$ where $<^{s_1} = \{(x,y),(y,z),(x,z)\}$ and $<^{s_2} = \{(x,y),(z,y)\}$. Then $\text{vt}_1(S) = \{(x,y)\}$ which is not modular. $\square$

Proposition 27 *If $\|\mathcal{W}\| \geq 3$, then for every $p \in (0,1)$, there exists $\mathcal{S}$ such that $\text{vt}_p(\mathcal{S})$ is neither modular nor transitive.*

Proof: Note that if $\|\mathcal{W}\| = 2$, every relation over $\mathcal{W}$ is either transitive or modular (but not necessarily both), and if $\|\mathcal{W}\| = 1$ every relation over $\mathcal{W}$ is both modular and transitive. Let $\mathcal{W}$ be a set of worlds such that $\|\mathcal{W}\| \geq 3$ and let x, y, and z denote three distinct members of $\mathcal{W}$. We will define $\mathcal{S}$ parameterized by p such that $\text{vt}_p(\mathcal{S})$ is the relation $\{(x,y),(y,z)\}$ which is neither transitive or modular.

Let $\mathcal{S} = \{s_1, \ldots, s_n\}$ satisfying the following conditions:

1. $n = \lceil 1/p + 1 \rceil$ if $p \leq 1/3$, $\lceil 2/(1-p) + 1 \rceil$ otherwise.[13]

13. $\lceil x \rceil$ denotes the *ceiling* of x, i.e., the smallest integer greater than or equal to x.

2. $<^{s_1} = \{(w, y) | w \in \mathcal{W}, w \neq y\}$.
3. $<^{s_2} = \{(y, w) | w \in \mathcal{W}, w \neq y\}$.
4. $\lceil pn - 1 \rceil$ of the remaining sources have belief state $\{(x, w) | w \in \mathcal{W}, w \neq x\} \cup \{(w, z) | w \in \mathcal{W}, w \neq z\}$.
5. The remaining sources have fully disconnected belief states.

It is clear that $<^{s_i} \in \mathcal{B}$ for all i. We make two observations: First, observe that $pn > 1$. If $p \leq 1/3$,

$$pn = p\lceil 1/p + 1 \rceil \geq 1 + p > 1.$$

If $p > 1/3$,

$$pn = p\lceil 2/(1 - p) + 1 \rceil \geq 2p/(1 - p) + p$$

which is a monotonically increasing function, so

$$pn > 2(1/3)/(1 - 1/3) + 1/3 = 4/3 > 1.$$

Second, note that the set described in the fifth condition is non-empty since the number of sources described in conditions 2-4 is $2 + \lceil pn - 1 \rceil < 1 + pn$ which is less than n if $n > 1/(1 - p)$. This is true if $p \leq 1/3$ since for these values

$$n = \lceil 1/p + 1 \rceil > 1/p > 1/(1 - p).$$

And it is also true if $p > 1/3$ since

$$n = \lceil 2/(1 - p) + 1 \rceil > 2/(1 - p) > 1/(1 - p).$$

(x, y) appears only in $<^{s_1}$ and each of the belief states described in condition 3, so

$$\text{count}_{\mathcal{S}}(x, y) = 1 + \lceil pn - 1 \rceil \geq 1 + pn - 1 = pn$$

so $(x, y) \in \text{vt}_p(\mathcal{S})$. Similarly, (y, z) appears only in $<^{s_2}$ and each of the belief states described in condition 3, so

$$\text{count}_{\mathcal{S}}(y, z) = 1 + \lceil pn - 1 \rceil \geq 1 + pn - 1 = pn$$

so $(y, z) \in \text{vt}_p(\mathcal{S})$. It remains to show that $\text{vt}_p(\mathcal{S})$ has no other members. We show that the count of each pair is less than pn. $\text{count}_{\mathcal{S}}(w, w) = 0 < pn$ for all $w \in \mathcal{W}$. $\text{count}_{\mathcal{S}}(z, w) = \text{count}_{\mathcal{S}}(w, x) = 0$ for all $w \neq y \in \mathcal{W}$. For all $w \in \mathcal{W} - \{x, y\}$, (w, y) only appears in $<^{s_1}$, so $\text{count}_{\mathcal{S}}(w, y) = 1 < pn$ from our first observation above. For all $w \in \mathcal{W} - \{y, z\}$, (y, w) only appears in $<^{s_2}$, so $\text{count}_{\mathcal{S}}(y, w) = 1 < pn$. Finally, for all $w \in \mathcal{W} - \{x, y, z\}$, (x, z), (x, w), and (w, z) only appear in the belief states described in condition 3, so

$$\text{count}_{\mathcal{S}}(x, z) = \text{count}_{\mathcal{S}}(x, w) = \text{count}_{\mathcal{S}}(w, z) = \lceil pn - 1 \rceil < pn.$$

□

Proposition 28 *Suppose $\preceq$ is a transitive relation over $\mathcal{W}$ and $x, y \in \mathcal{W}$.*

1. *If* $x \prec y$, *then* $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$.
2. *If* $x \bowtie y$, *then* $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$.
3. *If* $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$, *then* $\exists z.\ \text{lev}_{\preceq}(z) = \text{lev}_{\preceq}(x) \wedge z \prec y$.
4. *If* $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$, *then* $x \preceq y$ *iff* $y \preceq x$.

Proof:

1. Suppose $x \prec y$. Then

$$\begin{aligned}\text{lev}_{\preceq}(y) &= 1 + \max_{y' \in \mathcal{W}} \left(\left\{\text{lev}_{\preceq}(y') : y' \prec y\right\}\right) \\ &\geq 1 + \text{lev}_{\preceq}(x) \\ &> \text{lev}_{\preceq}(x).\end{aligned}$$

2. Suppose $x \bowtie y$. If $x \in \text{ch}(\mathcal{W}, \preceq)$ then $y \in \text{ch}(\mathcal{W}, \preceq)$, so $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y) = 0$. Suppose $x \notin \text{ch}(\mathcal{W}, \preceq)$. Then $y \notin \text{ch}(\mathcal{W}, \preceq)$ and $\text{lev}_{\preceq}(x) = 1 + \max_{y' \in \mathcal{W}} (\{\text{lev}_{\preceq}(y') : y' \prec x\})$. If y' is one such element, then $y' \prec y$ by transitivity, so

$$\text{lev}_{\preceq}(y) = 1 + \max_{y'' \in \mathcal{W}} \left(\left\{\text{lev}_{\preceq}(y'') : y'' \prec y\right\}\right) \geq \text{lev}_{\preceq}(x).$$

   By an identical argument, $\text{lev}_{\preceq}(x) \geq \text{lev}_{\preceq}(y)$. Thus, $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$.

3. Suppose $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$. It is sufficient to prove the following: For every non-negative integer $l < \text{lev}_{\preceq}(y)$, there exists $z \prec y$ such that $\text{lev}_{\preceq}(z) = l$. We prove by induction on l. If $l = \text{lev}_{\preceq}(y) - 1$, there must exist $z \prec y$ and $\text{lev}_{\preceq}(z) = l$ by definition. Assume there exists $z \prec y$ for $0 < l < \text{lev}_{\preceq}(y) - 1$ such that $\text{lev}_{\preceq}(z) = l$. Since $l > 0$, there exists $z' \prec z$ such that $\text{lev}_{\preceq}(z') = l - 1$. By transitivity, $z' \prec y$.

4. Suppose $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$. If $x \preceq y$ then $x \not\prec y$ from the first part of this proposition, otherwise $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(x)$, so $y \preceq x$. Similarly, if $y \preceq x$ then $y \not\prec x$, so $x \preceq y$.

□

Proposition 29 *If* $\preceq$ *is a transitive relation over* $\mathcal{W}$, *then* $\text{MC}(\preceq) \in \mathcal{B}$.

Proof: Let $\preceq^* = \text{MC}(\preceq)$ and $x, y, z \in \mathcal{W}$. Suppose $x \preceq^* y$. Then $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$ or $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$ and $\exists x', y' \in \mathcal{W}.\ (\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x) \wedge x' \bowtie y')$. If $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(z)$ or $\text{lev}_{\preceq}(z) < \text{lev}_{\preceq}(y)$, then $x \preceq^* z$ or $z \preceq^* y$. Otherwise, $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y) = \text{lev}_{\preceq}(z)$ and $\exists x', y' \in \mathcal{W}.\ (\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x) \wedge x' \bowtie y')$, so $x \preceq^* z$. Thus, $\preceq^*$ is modular.

Now also suppose $y \preceq^* z$. Then $\text{lev}_{\preceq}(y) < \text{lev}_{\preceq}(z)$ or $\text{lev}_{\preceq}(y) = \text{lev}_{\preceq}(z)$ and $\exists y', z' \in \mathcal{W}.\ (\text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(z') = \text{lev}_{\preceq}(y) \wedge y' \bowtie z')$. If $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$ or $\text{lev}_{\preceq}(y) < \text{lev}_{\preceq}(z)$, then $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(z)$ by transitivity of $<$, so $x \preceq^* z$. Otherwise, $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y) = \text{lev}_{\preceq}(z)$ and $\exists x', y' \in \mathcal{W}.\ (\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x) \wedge x' \bowtie y')$, so $x \preceq^* z$. Thus, $\preceq^*$ is transitive. □

Proposition 30 *Suppose* $\preceq$ *is a transitive relation over* $\mathcal{W}$ *and* $\preceq^* = \text{MC}(\preceq)$.

1. *$\preceq \subseteq \preceq^*$ and $\prec \subseteq \prec^*$.*

2. *If $\preceq$ is modular, then $\preceq^* = \preceq$.*

3. *$\text{lev}_{\preceq^*}(x) = \text{lev}_{\preceq}(x)$ for all $x \in \mathcal{W}$.*

4. *If $\preceq' \in \mathcal{B}$ such that $\preceq \subseteq \preceq'$ and $\text{lev}_{\preceq'}(x) = \text{lev}_{\preceq}(x)$ for all $x \in \mathcal{W}$, then $\preceq^* \subseteq \preceq'$.*

Proof: Let $x, y \in \mathcal{W}$.

1. Suppose $x \preceq y$. Then $\text{lev}_{\preceq}(x) \leq \text{lev}_{\preceq}(y)$. If $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$, $x \preceq^* y$. Suppose $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$. Then $x \bowtie y$, so there exist x', y' such that $\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x)$ and $x' \bowtie y'$, i.e., $x' = x$ and $y' = y$, so $x \preceq^* y$.

   Now suppose $x \prec y$. Then $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$, so $x \preceq^* y$ and $y \not\preceq^* x$, so $x \prec^* y$.

2. From the first part of this proposition, $\preceq \subseteq \preceq^*$, so it suffices to show $\preceq^* \subseteq \preceq$. Suppose $x \preceq^* y$.

   Case 1: $\text{lev}_{\preceq}(x) < \text{lev}_{\preceq}(y)$. Then, by Proposition 28, there exists z such that $\text{lev}_{\preceq}(z) = \text{lev}_{\preceq}(x)$ and $z \prec y$. By modularity, $z \preceq x$ or $x \preceq y$. In the latter case, we're done. In the former case, $z \bowtie x$ otherwise $\text{lev}_{\preceq}(z) \neq \text{lev}_{\preceq}(x)$. Thus, $x \preceq y$ by transitivity.

   Case 2: $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$. Then there exist x', y' such that $\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x)$ and $x' \bowtie y'$. By modularity, $x' \preceq x$ or $x \preceq y'$. In the former case $x' \bowtie x$ by Proposition 28, in the latter $x' \bowtie x$ by Proposition 28 and transitivity. By modularity, $x' \preceq y$ or $y \preceq x$. Again applying Proposition 28 and transitivity, we have $x \bowtie y$, so $x \preceq y$.

3. We prove by induction on the level of x in $\preceq$.

   Base case: $\text{lev}_{\preceq}(x) = 0$. Then $x \in \text{ch}(\mathcal{W}, \preceq)$. Suppose $y \preceq^* x$. Then $\text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$ and there exist x' and y' such that $\text{lev}_{\preceq}(x') = \text{lev}_{\preceq}(y') = \text{lev}_{\preceq}(x) = \text{lev}_{\preceq}(y)$ and $x' \bowtie y'$, so $x \preceq^* y$. Thus, $x \in \text{ch}(\mathcal{W}, \preceq^*)$, so $\text{lev}_{\preceq^*}(x) = 0 = \text{lev}_{\preceq}(x)$.

   Inductive case: Assume that $\text{lev}_{\preceq^*}(x') = \text{lev}_{\preceq}(x')$ for all x' such that $\text{lev}_{\preceq}(x') < \text{lev}_{\preceq}(x)$; we show that $\text{lev}_{\preceq^*}(x) = \text{lev}_{\preceq}(x)$. Let $z = \arg\max_{y' \in \mathcal{W}} (\{\text{lev}_{\preceq^*}(y') : y' \prec^* x\})$; then $\text{lev}_{\preceq^*}(x) = 1 + \text{lev}_{\preceq^*}(z)$. Also, let $y = \arg\max_{y' \in \mathcal{W}} (\{\text{lev}_{\preceq}(y') : y' \prec x\})$; then $\text{lev}_{\preceq}(x) = 1 + \text{lev}_{\preceq}(y)$. By the first part of this proposition, $\prec \subseteq \prec^*$, so $y \prec^* x$. Thus, $\text{lev}_{\preceq^*}(z) \geq \text{lev}_{\preceq^*}(y)$. Furthermore, $\text{lev}_{\preceq^*}(y) = \text{lev}_{\preceq}(y)$ by the inductive hypothesis, so $\text{lev}_{\preceq^*}(z) \geq \text{lev}_{\preceq}(y)$. Now $\text{lev}_{\preceq}(z) < \text{lev}_{\preceq}(x)$ by Definition 26 (otherwise $x \preceq^* z$, a contradiction). By the inductive hypothesis, $\text{lev}_{\preceq^*}(z) = \text{lev}_{\preceq}(z)$, so $\text{lev}_{\preceq^*}(z) < \text{lev}_{\preceq}(x) = 1 + \text{lev}_{\preceq}(y)$. Since levels are integral, $\text{lev}_{\preceq^*}(z) \leq \text{lev}_{\preceq}(y)$, so $\text{lev}_{\preceq^*}(z) = \text{lev}_{\preceq}(y)$. Thus, $\text{lev}_{\preceq^*}(x) = 1 + \text{lev}_{\preceq^*}(z) = 1 + \text{lev}_{\preceq}(y) = \text{lev}_{\preceq}(x)$.

4. Suppose $\preceq' \in \mathcal{B}$, $\preceq \subseteq \preceq'$, and $\text{lev}_{\preceq'}(z) = \text{lev}_{\preceq}(z)$ for all $z \in \mathcal{W}$. It is clear that for any $\preceq'' \in \mathcal{B}$ and $x \in \mathcal{W}$, x is in the partition corresponding to its level, i.e., $W_{\text{lev}_{\preceq''}(x)}$. Since both $\preceq^*$ and $\preceq'$ are both in $\mathcal{B}$ and both preserve the levels of all worlds in $\preceq$, the must have identical partitions. Suppose $x, y \in \mathcal{W}$ and W_i and W_j are the partitions (for both $\preceq^*$ and $\preceq'$) such that $x \in W_i$ and $y \in W_j$.

Suppose $x \preceq^* y$. If $W_i \neq W_j$ then $i < j$ by Proposition 7, so $x \preceq' y$, again by Proposition 7. If, instead $W_i = W_j$, then $\text{lev}_{\preceq^*}(x) = \text{lev}_{\preceq^*}(y)$. By Definition 26, there exist $x', y' \in \mathcal{W}$ such that $\text{lev}_{\preceq^*}(x') = \text{lev}_{\preceq^*}(y') = \text{lev}_{\preceq^*}(x)$ and $x' \bowtie y'$. Thus, $x', y' \in W_i$ and, since $\preceq \subseteq \preceq'$, $x' \bowtie' y'$. By Proposition 7, W_i is either fully connected or fully disconnected in $\preceq'$, so W_i must be fully connected in $\preceq'$. In particular, $x \preceq' y$.

□

Proposition 31 *If $\preceq$ is a relation over $\mathcal{W}$, then* $\text{MT}(\preceq) \in \mathcal{B}$.

Proof: Since $\preceq^+$ is transitive, $\text{MT}(\preceq) = \text{MC}(\preceq^+) \in \mathcal{B}$ by Proposition 29. □

Proposition 32 *Suppose $\preceq$ is a relation over $\mathcal{W}$ and* $\preceq^* = \text{MT}(\preceq)$.

1. $\preceq \subseteq \preceq^*$.
2. *If $\preceq$ is transitive, then* $\preceq^* = \text{MC}(\preceq)$.
3. *If $\preceq$ is modular, then $\preceq^* = \preceq^+$.*
4. *If $\preceq$ is modular and transitive, then $\preceq^* = \preceq$.*
5. *If $\preceq$ has no conflicts, then neither does $\preceq^*$.*

Proof: Let $x, y \in \mathcal{W}$.

1. $\preceq \subseteq \preceq^+$ and, by the first property of Proposition 30, $\preceq^+ \subseteq \text{MC}(\preceq^+) = \preceq^*$, so $\preceq \subseteq \preceq^*$.
2. Since $\preceq$ is transitive, $\preceq^+ = \preceq$. Thus, $\preceq^* = \text{MC}(\preceq^+) = MC(\preceq)$.
3. Since $\preceq$ is modular, $\preceq^+$ is modular by Proposition 1 so, by Proposition 30, $\text{MC}(\preceq^+) = \preceq^+$. Thus, $\preceq^* = \preceq^+$.
4. Since $\preceq$ is transitive, $\preceq^* = \text{MC}(\preceq)$ and since $\preceq$ is modular, $\text{MC}(\preceq) = \preceq$ by Proposition 30, so $\preceq^* = \preceq$.
5. We first prove the following lemma:

   Lemma 32.1 *x and y are in conflict wrt $\preceq^+$ iff they are in conflict wrt $\preceq$.*

   Proof: The "if" direction is obvious since the transitive closure is a monotonically additive operation. For the "only if" direction, suppose x and y are in conflict wrt $\preceq^+$. Then there exist $w_0, \ldots, w_n, z_0, \ldots, z_m \in \mathcal{W}$ such that

   $$x = w_0 \preceq^+ \cdots \preceq^+ w_n = y = z_0 \preceq^+ \cdots \preceq^+ z_m = x.$$

   But then, for each $0 \leq i \leq n-1$, there exist $w_{i0}, \ldots, w_{ip_i} \in \mathcal{W}$ such that

   $$w_i = w_{i0} \preceq \cdots \preceq w_{ip_i} = w_{i+1}.$$

   Similarly, for each $0 \leq j \leq m-1$, there exist $z_{j0}, \ldots, z_{jq_j} \in \mathcal{W}$ such that

   $$z_j = z_{j0} \preceq \cdots \preceq z_{jq_j} = z_{j+1}.$$

   Thus,

   $$x = w_0 \preceq \cdots \preceq w_n = y = z_0 \preceq \cdots \preceq z_m = x,$$

   so x and y are in conflict wrt $\preceq$. □

Now, suppose x and y are in conflict wrt $\preceq^*$. Then $x \bowtie^* y$ since $\preceq^* \in \mathcal{B}$ by Proposition 31. By Propositions 28 and 30,

$$\text{lev}_{\preceq^+}(x) = \text{lev}_{\preceq^*}(x) = \text{lev}_{\preceq^*}(y) = \text{lev}_{\preceq^+}(y)$$

So, since $x \preceq^* y$, there exist $x', y' \in \mathcal{W}$ such that $\text{lev}_{\preceq^+}(x') = \text{lev}_{\preceq^+}(y') = \text{lev}_{\preceq^+}(x)$ and $x' \bowtie^+ y'$ by Definition 26. Thus, $\preceq^+$ has a conflict. By the lemma above, $\preceq$ must also have a conflict.

□

Proposition 33 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0, 1]$*, then* $\text{AGREq}_p(S) \in \mathcal{B}$*.*

Proof: Follows immediately from the definition of AGREq_p and Proposition 31. □

Proposition 34 *If* $S \subseteq \mathcal{S}$ *and* $p \in [0, 1]$*, then* $\text{AGR}_p(S) \in \mathcal{B}$*.*

Proof: Again, this follows immediately from Proposition 31. □

Proposition 35 *Suppose* $S \subseteq \mathcal{S}$ *and* $p \in [0, 1]$*.*

1. *If* $\sqsupseteq_S$ *is fully connected, then* $\text{AGR}_p(S) = \text{AGREq}_p(S)$*.*
2. *If* $\sqsupseteq_S$ *is a total order, then* $\text{AGR}_p(S) = \text{AGRRf}_p(S) = \text{AGRRf}(S) = \text{AGR}(S)$*.*
3. $\text{AGR}_0(S) = \text{AGR}(S)$*.*

Proof: Assume $x, y \in \mathcal{W}$.

1. It suffices to show that $\text{AGRRf}_p(S) = \text{vt}_p(S)$ when $\sqsupseteq_S$ is fully connected. Suppose $(x, y) \in \text{AGRRf}_p(S)$. Then, by the definition of AGRRf_p, there exists $s \in S$ such that $(x, y) \in \text{vt}_p(\{s' \in S : s' \equiv_S s\})$. Since $\sqsupseteq_S$ is fully connected, $\{s' \in S : s' \equiv_S s\} = S$, so $(x, y) \in \text{vt}_p(S)$.

   Suppose, instead, $(x, y) \in \text{vt}_p(S)$. By the definition of vt_p, $\text{count}_S(x, y) > 0$ so there exists $s \in S$ such that $x <^s y$. Pick one such s. Again, $\{s' \in S : s' \equiv_S s\} = S$ since $\sqsupseteq_S$ is fully connected, so $(x, y) \in \text{vt}_p(\{s' \in S : s' \equiv_S s\})$. Finally, $\forall s' \in S.\ s' \sqsupset_S s \Rightarrow x \approx^{s'} y$ holds vacuously, so $(x, y) \in \text{AGRRf}_p(S)$.

2. Suppose $\sqsupseteq_S$ is a total order. We have already shown in Proposition 17 that $\text{AGRRf}(S) = \text{AGR}(S)$.

   Next we show that $\text{AGRRf}_p(S) = \text{AGRRf}(S)$. $\text{AGRRf}_p(S)$ is the set (x, y) such that there exists $s \in S$ such that $x <^s y$, $(x, y) \in \text{vt}_p(\{s' \in S : s' \equiv_S s\})$, and, for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$. $\sqsupseteq$ is a total order, $\{s' \in S : s' \equiv_S s\} = \{s\}$. Since $x <^s y$, so $\text{count}_{\{s\}}(x, y) = 1 > 0$ and $\text{count}_{\{s\}}(x, y)/\|\{s\}\| = 1 \geq p$. Consequently, $(x, y) \in \text{vt}_p(\{s' \in S : s' \equiv_S s\})$, proving that this requirement is redundant when $\sqsupseteq_S$ is a total order. Thus, $\text{AGRRf}_p(S)$ is the set (x, y) such that $x <^s y$ and, for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$, i.e., $\text{AGRRf}_p(S) = \text{AGRRf}(S)$.

   Finally, $\text{AGR}_p(S) = \text{MT}(\text{AGRRf}_p(S)) = \text{MT}(\text{AGRRf}(S))$. By Proposiion 14, $\text{AGRRf}(S)$ is modular and transitive, so $\text{AGR}_p(S) = \text{AGRRf}(S)$ by Proposition 32.

3. It suffices to show that $\mathrm{AGRRf}_0(S) = \mathrm{AGRRf}(S)$, since then

$$\mathrm{AGR}_0(S) = \mathrm{MT}(\mathrm{AGRRf}_0(S)) = \mathrm{MT}(\mathrm{AGRRf}(S)) = \mathrm{AGRRf}(S)^+$$

by Propositions 13 and 32, so $\mathrm{AGR}_0(S) = \mathrm{AGR}(S)$.

Suppose $(x, y) \in \mathrm{AGRRf}_0(S)$. Then $x <^s y$ and, for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$, so $(x, y) \in \mathrm{AGRRf}(S)$. Suppose, instead, $(x, y) \in \mathrm{AGRRf}(S)$. Then $x <^s y$ and, for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$. Let $S' = \{s' \in S : s' \equiv_S s\}$. Since $x <^s y$ and $s \in S'$, $\mathrm{count}_{S'}(x, y) > 0$ and $\mathrm{count}_{S'}(x, y)/\|S'\| \geq 0$, so $(x, y) \in \mathrm{vt}_0(S')$. Therefore, $(x, y) \in \mathrm{AGRRf}_0(S)$.

□

Corollary 35.1 *Let $S = \{s_1, \ldots, s_n\} \subseteq \mathcal{S}$ and $\mathrm{AGR}_f(<^{s_1}, \ldots, <^{s_n}) = \mathrm{AGR}_0(S)$. AGR_f satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.*

Proof: Follows immediately from Propositions 35 and 19. □

Proposition 36 *Let A be an agent informed by a set of sources $S \subseteq \mathcal{S}$, with support pedigreed belief state $(l, \mathrm{sup}, \mathrm{rtab})$, and using aggregation function AGR_p for $p \in [0, 1]$. A's belief state is the relation*

$$\mathrm{MT}\left(\{(x, y) : \|\mathrm{sup}(x, y)\| > 0 \textit{ and } \|\mathrm{sup}(x, y)\|/\|\mathrm{rtab}(l(x, y))\| \geq p\}\right)$$

Proof: Let

$$R = \{(x, y) : \|\mathrm{sup}(x, y)\| > 0 \text{ and } \|\mathrm{sup}(x, y)\|/\|\mathrm{rtab}(l(x, y))\| \geq p\}$$

It suffices to show $\mathrm{AGRRf}_p(S) = R$. Suppose $(x, y) \in \mathrm{AGRRf}_p(S)$. Then there exists $s \in S$ such that (a) $x <^s y$, (b) $(x, y) \in \mathrm{vt}_p(\{s' \in S : s' \equiv_S s\})$, and (c) for all $s' \sqsupset_S s \in S$, $x \approx^{s'} y$. By the (a) and (c), $\mathrm{rank}(s) = \max(\{\mathrm{rank}(s) : x \not\approx^s y,\ s \in S\} \cup \{\clubsuit\})$, so $l(x, y) = \mathrm{rank}(s)$. Thus, $\{s' \in S : s' \equiv_S s\} = \{s' \in S : \mathrm{rank}(s') = l(x, y)\} = \mathrm{rtab}(l(x, y))$, so $(x, y) \in \mathrm{vt}_p(\mathrm{rtab}(l(x, y)))$ by (b). By the definition of vt_p, $\mathrm{count}_{\mathrm{rtab}(l(x,y))}(x, y) > 0$ and $\mathrm{count}_{\mathrm{rtab}(l(x,y))}(x, y)/\|\mathrm{rtab}(l(x, y))\| \geq p$. But $\mathrm{sup}(x, y) = \{s' \in S : \mathrm{rank}(s') = l(x, y),\ x <^{s'} y\} = \{s' \in \mathrm{rtab}(l(x, y)) : x <^{s'} y\}$, so $\|\mathrm{sup}(x, y)\| = \mathrm{count}_{\mathrm{rtab}(l(x,y))}(x, y)$. Thus, $\|\mathrm{sup}(x, y)\| > 0$ and $\|\mathrm{sup}(x, y)\|/\|\mathrm{rtab}(l(x, y))\| \geq p$, so $(x, y) \in R$.

Now suppose $(x, y) \in R$. Then (a) $\|\mathrm{sup}(x, y)\| > 0$ and (b) $\|\mathrm{sup}(x, y)\|/\|\mathrm{rtab}(l(x, y))\| \geq p$. Suppose $s \in \mathrm{sup}(x, y)$; by (a), at least one such s exists. By the definition of sup, $x <^s y$, satisfying the first condition of AGRRf_p, and $\mathrm{rank}(s) = l(x, y)$. By the definition of $l(x, y)$, $\mathrm{rank}(s) = \max(\{\mathrm{rank}(s) : x \not\approx^s y,\ s \in S\} \cup \{\clubsuit\})$, so for all $s' \in S$ such that $\mathrm{rank}(s') > \mathrm{rank}(s)$ (i.e., $s' \sqsupset_S s$), $x \approx^{s'} y$. It only remains to show that $(x, y) \in \mathrm{vt}_p(\{s' \in S : s' \equiv_S s\})$. Since $\mathrm{rank}(s) = l(x, y)$, $\{s' \in S : s' \equiv_S s\} = \mathrm{rtab}(l(x, y))$ as we showed above, so

$$\begin{aligned}
&\mathrm{vt}_p(\{s' \in S : s' \equiv_S s\}) \\
&\quad = \mathrm{vt}_p(\mathrm{rtab}(l(x, y))) \\
&\quad = \{(x', y') : \mathrm{count}_{\mathrm{rtab}(l(x,y))}(x', y') > 0,\ \mathrm{count}_{\mathrm{rtab}(l(x,y))}(x', y')/\|\mathrm{rtab}(l(x, y))\| \geq p\}.
\end{aligned}$$

As we showed above, $\|\mathrm{sup}(x, y)\| = \mathrm{count}_{\mathrm{rtab}(l(x,y))}(x, y)$. Making this substitution into (a) and (b), we see that $(x, y) \in \mathrm{vt}_p(\{s' \in S : s' \equiv_S s\})$, so $(x, y) \in \mathrm{AGRRf}_p(S)$. □

Proposition 37 *Let $\mathcal{A}$, $\mathcal{P}_\mathcal{A}$, S, and $\sqsupseteq_S$ be as in Definition 32. Then $\oplus_{\sup}(\mathcal{P}_\mathcal{A})$ is the support pedigreed belief state of $\oplus(\mathcal{A})$.*

Proof: Let $\oplus_{\sup}(\mathcal{P}_\mathcal{A}) = (l, \sup, \mathrm{rtab})$, $l' : \mathcal{W} \times \mathcal{W} \to \mathcal{R} \cup \{\clubsuit\}$ such that

$$l'((x,y)) = \max(\{l''(x,y) : (l'', \sup'', \mathrm{rtab}'') \in \mathcal{P}_\mathcal{A}\}),$$

$\sup' : \mathcal{W} \times \mathcal{W} \to 2^S$ such that

$$\sup'(x,y) = \bigcup_{(l'',\sup'',\mathrm{rtab}'') \in \mathcal{P}_\mathcal{A},\ l''(x,y) = l'(x,y)} \sup''(x,y),$$

and $\mathrm{rtab}' : \mathrm{ranks}(S) \to \mathbb{R}$ such that

$$\mathrm{rtab}'(r) = \bigcup_{(l'',\sup'',\mathrm{rtab}'') \in \mathcal{P}_\mathcal{A},\ r \in \mathrm{range}(\mathrm{rtab}'')} \mathrm{rtab}''(r).$$

It suffices to show that $l = l'$, $\sup = \sup'$, and $\mathrm{rtab} = \mathrm{rtab}'$.

Suppose $x, y \in \mathcal{W}$ and agent A_i's support pedigreed belief state is $(l_i, \sup_i, \mathrm{rtab}_i)$.

$$\begin{aligned}
l(x,y) &= \max(\{\mathrm{rank}(s) : x <^s y,\ s \in S\} \cup \{\clubsuit\}) \\
&= \max\left(\bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A}} (\{\mathrm{rank}(s) : x <^s y,\ s \in S_i\} \cup \{\clubsuit\})\right) \\
&= \max\left(\bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A}} \{\max(\{\mathrm{rank}(s) : x <^s y,\ s \in S_i\} \cup \{\clubsuit\})\}\right) \\
&= \max(\{l''(x,y) : (l'', \sup'', \mathrm{rtab}'') \in \mathcal{P}_\mathcal{A}\}) \\
&= l'(x,y).
\end{aligned}$$

Also,

$$\begin{aligned}
\sup(x,y) &= \{s \in S : \mathrm{rank}(s) = l(x,y),\ x <^s y\} \\
&= \bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A}} \{s \in S_i : \mathrm{rank}(s) = l(x,y),\ x <^s y\} \\
&= \bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A}} \{s \in S_i : \mathrm{rank}(s) = l'(x,y),\ x <^s y\} \\
&= \bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A},\ l_i(x,y) = l'(x,y)} \{s \in S_i : \mathrm{rank}(s) = l_i(x,y),\ x <^s y\} \\
&= \bigcup_{(l'',\sup'',\mathrm{rtab}'') \in \mathcal{P}_\mathcal{A},\ l''(x,y) = l'(x,y)} \sup''(x,y) \\
&= \sup'(x,y).
\end{aligned}$$

Finally,

$$\begin{aligned}
\mathrm{rtab}(r) &= \{s \in S : \mathrm{rank}(s) = r\} \\
&= \bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A}} \{s \in S_i : \mathrm{rank}(s) = r\} \\
&= \bigcup_{S_i \text{ informs } A_i,\ A_i \in \mathcal{A},\ r \in \mathrm{range}(\mathrm{rtab}_i)} \mathrm{rtab}_i(r) \\
&= \bigcup_{(l'', \mathrm{sup}'', \mathrm{rtab}'') \in \mathcal{P}_{\mathcal{A}},\ r \in \mathrm{range}(\mathrm{rtab}'')} \mathrm{rtab}''(r) \\
&= \mathrm{rtab}'(x, y).
\end{aligned}$$

□

Appendix 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
$\leq^+$: transitive closure of $\leq$
$\mathrm{ch}(X, \leq)$: choice set of X wrt $\leq$
$\|X\|$: cardinality of set X

$\mathcal{W}$: finite set of possible worlds
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
$\preceq$: weak likelihood
$\sim$: equal likelihood, agnosticism
$\bowtie$: conflict
Bel: belief of a conditional statement
Agn: agnosticism over a conditional statement
Con: conflict over a conditional statement

$\mathcal{T}$: set of total preorders
$\mathcal{T}_<$: strict versions of total preorders
$\mathcal{Q}$: 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
$\approx^s$: source agnosticism
$\bowtie^s$: source conflict
$\mathcal{R}$: set of ranks
r: element of $\mathcal{R}$
rank(s): rank of source s
ranks(S): set of ranks of sources in S
$\sqsupseteq$, $\sqsupseteq_S$: credibility ordering over $\mathcal{S}$, $S \subseteq \mathcal{S}$, respectively

Un: union of a set of belief states
AGRUn: aggregation via union
AGRRf: aggregation via refinement
AGR: general aggregation

$\mathcal{A}$: set of agents
A: element of $\mathcal{A}$
$\prec^A$: A's induced belief state
$(\prec, l)$: pedigreed belief state
$\prec^A_r$: restriction of A's pedigreed belief state to rank r
$\oplus$: fusion
$\oplus_{\text{ped}}$: pedigreed fusion

vt_p: voting function for p
lev: level of a world in a transitive relation
MC: modular closure
MT: modular, transitive closure
AGREq_p: aggregation with voting without refinement
AGRRf_p: aggregation with voting via refinement
AGR_p: general aggregation with voting
$(l, \text{sup}, \text{rtab})$: support pedigreed belief state
$\oplus_{\text{sup}}$: support pedigreed fusion

References

Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, *50*, 510–530.

Andréka, H., Ryan, M., & Schobbens, P.-Y. (2002). Operators and laws for combining preference relations. *Journal of Logic and Computation*, *12*(1), 13–53.

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

Baral, C., Kraus, S., Minker, J., & Subrahmanian, V. S. (1992). Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, *8*(1), 45–71.

Benferhat, S., Dubois, D., Kaci, S., & Prade, H. (2002). Possibilistic merging and distance-based fusion of propositional information. *Annals of Mathematics and Artificial In-*

telligence, *34*(1–3), 217–252.

Black, D. (1958). *The Theory of Committees and Elections*. Cambridge University Press, Cambridge.

Borgida, A., & Imielinski, T. (1984). Decision making in committees: A framework for dealing with inconsistency and non-monotonicity. In *Proceedings of the Workshop on Nonmonotonic Reasoning*, pp. 21–32.

Boutilier, C. (1996). Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, *25*, 263–305.

Brams, S. J., & Fishburn, P. C. (2002). Voting procedures. In Arrow, K. J., Sen, A. K., & Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare*, Vol. 1 of *Handbooks in Economics*, chap. 4, pp. 173–236. Elsevier Science.

Cantwell, J. (1998). Resolving conflicting information. *Journal of Logic, Language, and Information*, *7*, 191–220.

Center for Voting and Democracy (2002). Instant runoff voting. http://www.fairvote.org/irv/.

Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence*, *89*(1–2), 1–29.

Fishburn, P. C. (1974). Lexicographic orders, utilities and decision rules: A survey. *Management Science*, *20*(11), 1442–1471.

Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Gärdenfors, P., & Makinson, D. (1994). Nonmonotonic inference based on expectations. *Artificial Intelligence*, *65*(1), 197–245.

Gärdenfors, P., & Rott, H. (1995). Belief revision. In Gabbay, D. M., Hogger, C. J., & Robinson, J. A. (Eds.), *Epistemic and Temporal Reasoning*, Vol. 4 of *Handbook of Logic in Artificial Intelligence and Logic Programming*, pp. 35–132. Oxford University Press, Oxford.

Grosof, B. (1991). Generalizing prioritization. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR '91)*, pp. 289–300.

Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, *17*, 157–170.

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

Katsuno, H., & Mendelzon, A. O. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, *52*(3), 263–294.

Konieczny, S., & Pérez, R. P. (1998). On the logic of merging. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pp. 488–498.

Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, *44*(1–2), 167–207. CoRR: cs.AI/0202021.

Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton University Press.

Lehmann, D. (1995). Belief revision, revised. In *Proceedings of the Fourteenth International Joint Conference of Artificial Intelligence (IJCAI '95)*, pp. 1534–1540.

Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail?. *Artificial Intelligence*, *55*(1), 1–60. CoRR: cs.AI/0202022.

Liberatore, P., & Schaerf, M. (1995). Arbitration: A commutative operator for belief revision. In *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI '95)*, pp. 217–228.

Luce, D. R. (1956). Semiorders and a theory of utility discrimination. *Econometrica*, *24*(2), 178–191.

Makinson, D. C. (1997). Screened revision. *Theoria*, *63*(1–2), 14–23. Special issue on non-prioritized belief revision.

Maynard-Reid II, P., & Shoham, Y. (2001). Belief fusion: Aggregating pedigreed belief states. *Journal of Logic, Language, and Information*, *10*(2), 183–209.

Meyer, T. (2001). On the semantics of combination operations. *Journal of Applied Non-Classical Logics*, *11*(1–2), 59–84.

Meyer, T., Ghose, A., & Chopra, S. (2001). Social choice, merging and elections. In Benferhat, S., & Besnard, P. (Eds.), *Proceedings of the Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU2001)*, pp. 466–477.

Ramsey, F. P. (1931). *Foundations of Mathematics and Other Logical Essays*. Routledge and Kegan Paul, New York.

Revesz, P. Z. (1997). On the semantics of arbitration. *International Journal of Algebra and Computation*, *7*(2), 133–160.

Sen, A. (1986). Social choice theory. In Arrow, K. J., & Intriligator, M. D. (Eds.), *Handbook of Mathematical Economics*, Vol. III, chap. 22, pp. 1073–1181. Elsevier Science Publishers.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W. L., & Skyrms, B. (Eds.), *Causation in Decision, Belief Change, and Statistics, II*, pp. 105–134. Kluwer Academic Publishers.

Williams, M.-A. (1994). Transmutations of knowledge systems. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR '94)*, pp. 619–629.