The Review of Symbolic Logic
Volume 15, Number 1, March 2022

PRIORITY MERGE AND INTERSECTION MODALITIES

ZOÉ CHRISTOFF
University of Groningen
NORBERT GRATZL
Ludwig-Maximilians-Universität München
and
OLIVIER ROY
University of Bayreuth

Abstract. We study the logic of so-called lexicographic or priority merge for multi-agent plausibility models. We start with a systematic comparison between the logical behavior of priority merge and the more standard notion of pooling through intersection, used to define, for instance, distributed knowledge. We then provide a sound and complete axiomatization of the logic of priority merge, as well as a proof theory in labeled sequents that admits cut. We finally study Moorean phenomena and define a dynamic resolution operator for priority merge for which we also provide a complete set of reduction axioms.

In epistemic logic, the standard notion of pooled information is distributed knowledge, defined by taking the intersection of the agents' indistinguishability relations or information cells.[1] It captures what could be known, either by the agents themselves or by an outsider, by pooling or putting together all of the agents' information in a given state. The logic of distributed knowledge is well known [13], and so are its relations, both static [26, 30] and dynamic [1, 31], with individual knowledge and other group notions like shared and common knowledge.

Intersection pooling is well suited to represent the aggregation of true information. However, the model-theoretic intersection of softer types of information, for instance plausibility orderings [4, 18], yields counter-intuitive results. If "distributed beliefs" are defined by taking the intersection of underlying plausibility orderings on possible worlds, they become inconsistent when opinions are radically divergent[2] and, to the

Received: June 27, 2019.
2020 *Mathematics Subject Classification*: 03A05, 03B42, 03B45, 03B70, 91B14.
Key words and phrases: epistemic logic, group knowledge, priority merge, logic of beliefs, dynamic epistemic logic.

[1] The earliest reference to distributed knowledge in epistemic logic, although not using the term explicitly, seems to be [24]. Interestingly, the notion appeared in 1986 in the bibliometric literature [29] under the name of "undiscovered public knowledge." With the notable exception of [19], no contact between these two areas seems to have been made so far.

[2] We use here the intentionally vague "radically divergent" because when one moves from knowledge to attitudes that can be mistaken, like belief, mutually inconsistent beliefs need not result in inconsistent pooled belief. Nor, however, are mutually consistent beliefs sufficient

© Association for Symbolic Logic, 2021
doi:10.1017/S1755020321000058

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

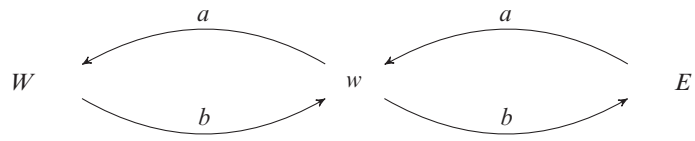


Fig. 1. Opposite plausibility orderings regarding the position of the hiker.

extent that they are modelled by normal modal operators, they "explode." Problems loom, however, even without explosion. Consider the following example.

Example 1. *Two members of a rescue team have radically different opinions regarding the position of a lost hiker. For simplicity let us consider this in a two-dimensional space* (Figure 1).

Agent a believes that the lost hiker is far away to the west of a given area (*relative to the current position w of the team*), *the further the more plausible. Agent b believes that the hiker is far away to the east of that area, the further the more plausible. Taking the intersection of their plausibility orderings yields the conclusion that they, as a group, rule out that the hiker is anywhere except right in the middle, i.e., at w. This seems to be an overly strong conclusion to reach from such diverging evidence. It does not help this rescue team in deciding, for instance, in which direction to send their search helicopter.*

Given the issues raised by applying intersection merge on plausibility orderings, two ways forward present themselves. To avoid explosion, one option is to keep the understanding of pooled information in terms of intersection modalities but to weaken the underlying logic, so that it becomes more tolerant of inconsistencies. That is the route taken in [25]. However, such a choice would not change the result of intersection in cases such as our rescue team example. To do this, one needs to look for alternative notions of pooled information and to compare their logical behavior with standard intersection modalities. This is what we do in this paper. More specifically, we look at the so-called lexicographic or priority merge operations [2] and their corresponding modal logic. This is independently motivated by the facts that, on the one hand, operations akin to priority merge have found a broad array of applications, some of them predating the study by Andréka et al. [2]—c.f. [27] in non-monotonic reasoning and [14] in social choice theory—and some more recent in philosophical logic [6, 8]. On the other hand, however, so far its only logical study has been done in extended modal languages [17].

After a short survey of the relevant literature on priority merge, we introduce priority merge modalities and discuss the expressivity as well as the key validities and invalidities of the resulting logic, especially in comparison with the better-known intersection modalities. We then move on to its axiomatization and proof theory. Finally, we focus on its so-called Moorean phenomena and introduce a dynamic resolution operator, for which we give reduction axioms.

§1. Priority merge: a short overview. Suppose we have a group of agents each equipped with a preorder over a given set of alternatives. This set can be interpreted in

for the pooled belief to be consistent. What guarantees consistency of distributed knowledge in standard epistemic logic is factivity of the knowledge of individuals, syntactically the "T" axiom, or semantically, the fact that the underlying indistinguishability relation is reflexive.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

many ways, for instance as a set of possible worlds, of partial states, or of choice alternatives. Given these interpretations, the preorders can be viewed as ordinal preferences or, in the epistemic case, as relative plausibility orderings. We will use the epistemic interpretation throughout this paper, mainly to fix intuitions. The results we present below, however, are of course not bound to that interpretation. They apply just as well to the logic of priority merges of preferences.

Priority merge is an operation that aggregates or "pools" the *individual* plausibility preorders to construct the *group*'s plausibility ordering. Unlike merging by taking intersection, priority merge makes use of the fact that we might have more information about the relative expertise (or importance) of the group members, i.e., that they can be somehow ordered according to this criterion. Intuitively, the aggregation works as follows. Assume, for now, that the priority ordering is linear. Then priority merge proceeds by lexicographically considering the strict preferences of each agent in order of priority: every pair of states strictly ordered by the agent at the top of the priority order will also be strictly ordered for the group. For the pairs which the topmost agent is indifferent about, we move to the second agent and order them strictly if she does, and so on, until all pairs are strictly ordered or we have gone through all agents. In other words, priority merge reflects the fact that, regarding how much of a say they should have in the group's joint opinion, the agents can be ordered in descending priority.

Consider again the rescue team example depicted in Figure 1. Assume now that agent a has priority over agent b, say, because a is the chief of the rescue team. Applying priority merge in this case would simply result in adopting a's opinion as the one of the team, thus sending the search helicopter towards the west. This illustrates how priority merge differs from intersection merge, by sometimes going towards one extreme: adopting the ordering of the agent with the highest priority as the one of the group. On the other end of the spectrum, when all agents are on a par in terms of expertise (or importance), priority merge boils down to merging by intersection. There are, however, many possibilities in between intersection merge and simply adopting the ordering of the highest-ranked agent, as the following example shows.

Example 2. *Agent a considers it more likely that the lost hiker is in the north than in the south, but equally likely that he is in the west or in the east. Agent b, on the other hand, considers it more likely that the lost hiker is in the south than in the north, and more likely that he is in the west rather than in the east. This is represented in* Figure 2.

Assuming that agent a has priority over b, applying priority merge results in the model pictured in Figure 3. *Collectively, the team considers it more likely that the hiker is north than south, conforming in this respect to the opinion of agent a, and more likely that he is west than east, thanks to the opinion of agent b. Had the team to choose a direction to send their search helicopter in, they would therefore send it towards northwest.*

Priority merge as just described has been generalized by Andréka et al. [2] to arbitrary *priority operators* which poll, using the priority merge rule, any number of relations using priority *graphs* for agents, i.e., relations that need be neither linear nor transitive. Besides providing an axiomatic characterization of priority operators and studying which properties of the merged relations it preserves, Andréka et al. also show that any priority operator can be represented as a combination of two simple operations. This result is fundamental for what follows. To state it precisely, we first introduce more formal details.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

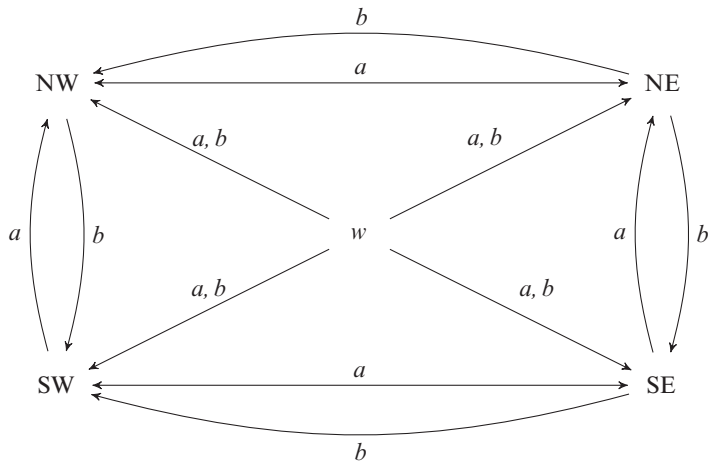


Fig. 2. Partially diverging plausibility orderings regarding the position of the lost hiker.

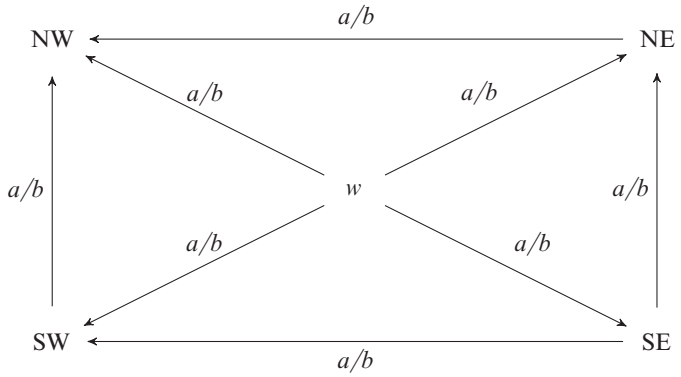


Fig. 3. The group's plausibility ordering obtained by applying priority merge on the model in Figure 2, when a has priority over b (which we denote by a/b).

Definition 1.1 (Plausibility frame). *Let $\mathcal{A}$ be a finite set of agents. A plausibility frame $\mathcal{F}$ is a tuple*

$$\langle W, (R_a)_{a \in \mathcal{A}} \rangle$$

where W is a non-empty set of states and for each $a \in A$, R_a is a preorder, i.e., a transitive and reflexive relation on W.

When wR_av holds, we say that v is at least as plausible for a as w is. Note that the preorders need not be total.[3] As usual we define the strict sub-relation $R_a^<$ of R_a

[3] Incomparability between states should be distinguished from indifference. The former captures the idea that an agent is not able to order two states with respect to their relative plausibility. This can be due to severe uncertainty, for instance. Indifference, on the other hand, captures the idea that an agent considers both states equally plausible. Among the postulates on plausibility or preference orderings, comparability or totality is often viewed as one of the least plausible. C.f. for instance [12]. For that reason we do not assume it here. We come back briefly to the consequences of this assumption in the conclusion.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

as $wR_a^{<}v$ iff wR_av but not vR_aw. Indifference, written $w \equiv_a v$, is defined as wR_av and vR_aw.

The two basic pooling methods that we will study here are then defined as follows:

DEFINITION 1.2 (Basic aggregated relations). *Let a, b be two agents in $\mathcal{A}$. Then*

- *(Basic Intersection) $wR_{a||b}v$ iff both wR_av and wR_bv.*
- *(Basic Priority) $wR_{a/b}v$ iff $wR_a^{<}v$, or $wR_{a||b}v$.*

The first is the intersection operation in its two-agents version. The n-ary version of this operation is the one standardly used to interpret distributed knowledge modalities, in the special case where each agent's preorder is an equivalence relation. The second basic pooling method, the priority merge relation for two agents only, is our main object of study. The notation a/b marks the fact that the relation is obtained by giving a priority over b.[4]

DEFINITION 1.3 (Priority graph). *A priority graph $\mathcal{P}$ is a pair $\langle G, \succ\rangle$ where G is a finite set of agents and $\succ$ is a partial order on G, where $a \succ b$ means that a has priority over b.*[5]

DEFINITION 1.4 (General aggregated relations). *Let $\mathcal{F} = \langle W, (R_a)_{a\in\mathcal{A}}\rangle$ be a plausibility frame, $G \subseteq \mathcal{A}$ a non-empty group of agents and $\langle G, \succ\rangle$ a priority graph on G. Then,*

- *(General Intersection) $R_\cap^G = \bigcap_{a\in G} R_a$.*
- *(General Priority) $wR_\succ^G v$ iff for all $a \in G$ either wR_av or there is some $b \in G$ such that $b \succ a$ and $wR_b^{<}v$.*

Andréka et al.'s [2] main result, for the present paper, shows that any priority operator, i.e., for any number of agents and any given priority graph, can be defined using these two basic operations. This result is best formulated by considering the *preferential algebra M* generated as follows.

DEFINITION 1.5 (Preferential Algebra). *Given a plausibility frame $\mathcal{F} = \langle W, (R_a)_{a\in\mathcal{A}}\rangle$, the preferential algebra M is the smallest set generated by*

- *All elements of $(R_a)_{a\in\mathcal{A}}$ are in M.*
- *If $x, y \in M$ then $R_{x||y}$ and $R_{x/y}$ are also in M.*

THEOREM 1 (Theorem 7.1 [2]). *For any general aggregated relation R^G there is a relation R' in M such that $R^G = R'$.*

In particular, this implies that any n-ary priority merge operation is definable using a composition of intersection and strict relations. For a group G of three agents $a \succ b \succ c$, for instance, the relation $R_\succ^G$ is defined as $R_{(a/b)/c}$, which boils down to

$$(R_a^{<} \cup (R_a \cap R_b^{<})) \cup ((R_a^{<} \cup (R_a \cap R_b)) \cap R_c).$$

[4] To avoid confusion, note that in [2], a/b is read instead as "b has priority over a".

[5] Andréka et al. [2] defines a more general notion of priority graph (see their Def. 2.4), possibly infinite, where an agent may be represented by several nodes in the same graph. They show that this possibility of repetition increases greatly the expressivity of their setting. The scope of Theorem 1 is therefore also wider than what we highlight here. For simplicity, the examples in this paper are restricted to the case where agents are not repeated. However, the logical system to be introduced encompasses the general case, too.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

Nonlinear priority graphs may require different combinations of intersection and priority merge. If, for instance, we have $a \succ' c$ and $b \succ' c$ but a and b are incomparable, then the priority merge operation $R^{G}_{\succ'}$ will only consider pairs on which the strict and indifferent parts of a's and b's respective plausibility order agree, before moving on to c. This will correspond to $R_{(a||b)/c}$.

Building on this general result, Girard [16] has studied the modal logic of general aggregated relations. He has shown that this logic can be completely axiomatized in hybrid logic, an extension of the modal language that contains explicit references to states. Girard and Seligman [17] went on to show that this logic can be given a well-behaved proof theory, i.e., one that admits weakening, contraction and, most importantly, analytic cut. They do so, again, by combining the expressive power of hybrid logic and the deductive power of labeled sequents. They leave open whether either of the results (complete axiomatization and well-behaved proof theory) could be obtained *without* the hybrid machinery. One of the main contributions of this paper, besides the comparison with intersection modalities and the study of Moorean phenomena, is to answer this question in the affirmative.

§2. Modal language and semantics. We are working with a modal language containing two families of modalities, weak and strict, for each possible *relation term*. The set of such terms T is defined in the same way as the elements of preferential algebras above (Definition 1.5): each $a \in \mathcal{A}$ is a term, and $x||y$ as well as x/y are terms whenever x and y are. Later on, we will use $T_{||}$ to denote the set of terms restricted to elements of $\mathcal{A}$ and terms generated only by intersection, and $T_{/}$ for the set consisting of elements of $\mathcal{A}$ and the terms generated only by priority merge.

We can now define our modal language:

DEFINITION 2.6. *The language $\mathcal{L}$ is defined as follows*:

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \vee \varphi \mid [t]\varphi \mid [t]^{<}\varphi \mid U\varphi,$$

where p is an element of a given countable set of atomic propositions Φ and $t \in T$ is a term.

U is the universal modality, with E its dual. All the "box" operators are standard Kripke modalities on their respective plausibility orderings. They have been interpreted as "defeasible knowledge", i.e., attitudes which cannot be defeated by any true information, for instance in [3, 5, 23]. We will also use this terminology to refer to them.[6] This language is interpreted in plausibility models:

DEFINITION 2.7. *Let $\mathcal{F} = \langle W, (R_a)_{a\in\mathcal{A}}\rangle$ be a plausibility frame and Φ a countable set of atomic propositions. A plausibility model $\mathcal{M}$ is a pair $\langle \mathcal{F}, V\rangle$ where $V : \Phi \to \mathcal{P}(W)$ is a valuation function.*

[6] Note that the more classical notion of knowledge as truth in all epistemically accessible states, i.e., in the same equivalence class or partition cell, could also have been added to this language. One would then add a third family of operators interpreted over connected components of R_t, i.e., truth in all states that comparable to the current one, that is either more or less plausible. See again [3, 5, 23] for details.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

DEFINITION 2.8. *For any plausibility model* $\mathcal{M} = \langle W, (R_a)_{a\in\mathcal{A}}\rangle, V$, *state* $w \in W$ *and formulas* $p, \varphi, \psi \in \mathcal{L}$:

$$\mathcal{M}, w \vDash p \text{ iff } w \in V(p)$$
$$\mathcal{M}, w \vDash \neg\varphi \text{ iff } \mathcal{M}, w \nvDash \varphi$$
$$\mathcal{M}, w \vDash \varphi \vee \psi \text{ iff } \mathcal{M}, w \vDash \varphi \text{ or } \mathcal{M}, w \vDash \psi$$
$$\mathcal{M}, w \vDash [t]\varphi \text{ iff for all } v \in W \text{ with } wR_t v, \mathcal{M}, v \vDash \varphi$$
$$\mathcal{M}, w \vDash [t]^{<}\varphi \text{ iff for all } v \in W \text{ with } wR_t^{<} v, \mathcal{M}, v \vDash \varphi$$
$$\mathcal{M}, w \vDash U\varphi \text{ iff for all } v \in W, \mathcal{M}, v \vDash \varphi$$

$R_t^{<}$ is defined as the strict version of the R_t relation—$wR_t^{<}v$ iff $wR_t v$ but not $vR_t w$—as opposed to the relation formed by aggregating the strict version of t's sub-terms' relations. In the Appendix, we consider briefly the consequences of using this latter definition. We furthermore write $||\varphi||$ for the "truth set" of a formula φ in a particular model $\mathcal{M}$, i.e., $\{w\colon \mathcal{M}, w \vDash \varphi\}$. Validity for the classes of plausibility models and of frames are defined as usual.

It is well known that this language is strong enough to express, beyond defeasible knowledge ($[a]\varphi$), the more standard notion of belief as truth in all states beyond some plausibility threshold, as well as its conditional variant [9, 11]. Let us write $B_a(\varphi/\psi)$ for "a believes φ conditional on ψ." One can give two different semantics for this operator, which will turn out to behave differently when we consider their aggregated counterparts. The first definition, belief as truth in the *most* plausible worlds, goes as follows. Let $\max_{R_a}(X)$ be defined as $\{w \in X\colon$ there is no $y \in X$ such that $xR_a^{<}y\}$. Then,

$$\mathcal{M}, w \vDash B_a^1(\varphi/\psi) \text{ iff } \max_{R_a}(||\psi||) \subseteq ||\varphi||.$$

Even in serial plausibility frames this notion of belief does not satisfy the usual D axiom, because in infinite models, $\max_{R_a}(||\psi||)$ might be empty in the case where we do not require R_a to be converse well-founded. The standard way to recover D is to use this more general definition:

$$\mathcal{M}, w \vDash B_a^2(\varphi/\psi) \text{ iff } \forall v \geq_a w,\ \exists v' \geq_a v(v' \in ||\psi|| \wedge \forall v'' \geq_a v'(v'' \in ||\psi \rightarrow \varphi||)).$$

Both notions are definable in $\mathcal{L}$ [9] and can be straightforwardly lifted to arbitrary terms to obtain the corresponding notions of conditional and unconditional aggregated beliefs. Unconditional belief in φ, for an arbitrary term t, is then defined as $B_t(\varphi/\top)$. In the remainder of the paper, We will write B_t^i when we do not need to specify B_t^1 or B_t^2.

The notion of beliefs and unconditional belief might, in turn, be used to study the interpretation of $[a]$. It is straightforward to check, for instance, that unconditionally believing φ, in the B_t^2-sense of beliefs, becomes equivalent to defeasibly knowing that one might defeasibly know φ.

$$B_a^2(\varphi) \leftrightarrow [a]\langle a\rangle[a]\varphi$$

From this the D axiom for $B_a^2(\varphi)$ follows directly from T for $[a]$. C.f. [23] for a thorough discussion.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

§3. Comparing priority merge and intersection modalities. Priority merge and intersection are of course different ways of pooling beliefs, but they are not unrelated. This section is meant to highlight their similarities and differences, as reflected in their logical behavior. We focus on three groups of relations. First, we consider questions of expressive power and interdefinability of the different modal operators in $\mathcal{L}$. Second, we compare how a group's pooled beliefs relate to the beliefs of its subgroups. And third, we consider the respective introspective properties of these notions. Figure 6 on page 11 and Table 1 on page 15 summarize the results.

3.1. Expressive power and interdefinability. To compare the expressivity of different fragments of our language $\mathcal{L}$, let us first fix some notation. We will write $\mathcal{L}_{[t],S}$ for the fragment of $\mathcal{L}$ that contains only the weak modalities $[t]$, for all $t \in S \subseteq T$, and $\mathcal{L}_{[t],[t]^{<},S}$ for the one containing both the weak and the strict modalities. The language of multi-agent doxastic logic, for instance, is written $\mathcal{L}_{[t],\mathcal{A}}$, and the extension of this language with an intersection modality is $\mathcal{L}_{[t],T_{||}}$. Note that our full language $\mathcal{L}$ is therefore also denoted by $\mathcal{L}_{[t],[t]^{<},T}$.

As a warm-up result it is good to keep in mind that both the priority merge and the intersection modalities can of course express the individual defeasible knowledge operators. More generally, for any relation term t we indeed have:

$$[t/t]\varphi \leftrightarrow [t]\varphi \leftrightarrow [t||t]\varphi.$$

The second biconditional is well known from modal logic for intersection modalities, and the first is a trivial consequence of the definition of priority merge. The same equivalences hold for the strict versions of these boxes. The languages $\mathcal{L}_{[t],T_{||}}$ and $\mathcal{L}_{[t],T_{||}\backslash\mathcal{A}}$ are thus equally expressive, and so are $\mathcal{L}_{[t]T_{/}}$ and $\mathcal{L}_{[t],T_{/}\backslash\mathcal{A}}$ (and similarly for the languages containing the strict operators).

Group modalities are not, in general, definable in terms of the relations that they aggregate. Again the case of the intersection modalities is well known. To show this for priority merge we can adapt the standard notion of bisimulation to the family of relations in our preferential algebras (Definition 1.5):

DEFINITION 3.9. *Let t be a relation term and $\mathcal{M} = \langle W, (R_a)_{a\in\mathcal{A}}, V\rangle$ and $\mathcal{M}' = \langle W', (R'_a)_{a\in\mathcal{A}}, V'\rangle$ two plausibility models. We say that a relation $Z \subseteq W \times W'$ is a t-bisimulation if and only if for all wZw' we have*:

1. *$\mathcal{M}, w$ and $\mathcal{M}', w'$ make the same atomic propositions true*;
2. *whenever there exists a $v \in W$ such that wR_tv, there is a v' such that $w'R'_tv'$ and vZv'*;
3. *whenever there exists a $v' \in W'$ such that $w'R'_tv'$, there is a v such that wR_tv and vZv'; and*
4. *analogously for the strict relation $R_t^{<}$.*

We say that $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar whenever there is a bisimulation Z between them.

Each t-modality, both in its weak and strict variants, is invariant under t-bisimulation.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

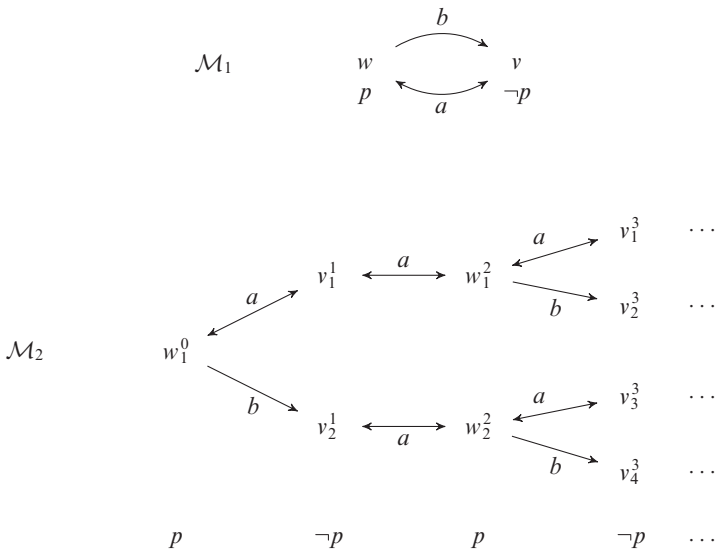


Fig. 4. Two a- and b-bisimilar models.

OBSERVATION 1. *Neither are the modalities* $[t/t']$ *definable in terms of* $[t]$ *and* $[t']$ *nor, mutatis mutandis, their strict variant. Furthermore, the intersection modalities* $[t||t']$ *are not definable in terms of* $[t]$, $[t']$, $[t/t']$ *and* $[t'/t]$ *together.*

Proof. For the first part consider the two models in Figure 4. We claim that the relation Z, connecting w in $\mathcal{M}_1$ to all the w_j^i in $\mathcal{M}_2$, and the same for v and all the v_j^i, is an a- and b-bisimulation. To show this we have to argue that the back and forth conditions between $\mathcal{M}_1$ and $\mathcal{M}_2$ hold. So starting from the direction from $\mathcal{M}_1$ to $\mathcal{M}_2$:

- wR_av and wZw_j^i. We have that $w_j^iR'_av_{2j-1}^{i+1}$ and vZv_{2j-1}^{i+1}.
- $wR_b^<v$ and wZw_j^i. We have that $w_j^iR_b^{<'}v_{2j}^{i+1}$ and vZv_{2j}^{i+1}.
- vR_aw and vZv_j^i. We have that $v_j^iR'_aw_j^{i+1}$ and wZw_j^{i+1}.

The argument from $\mathcal{M}_2$ to $\mathcal{M}_1$ is entirely similar except that it involves an additional case when $v_{2j-1}^{i+1}R_aw_j^i$. From this we can conclude that w and w_1^0 satisfy exactly the same formulas built using only $[a]$ and $[b]$ and their strict variant.

We have, however, that $\mathcal{M}_1, w \vDash \langle a/b\rangle\neg p$ while $\mathcal{M}_2, w_1^0 \vDash [a/b]p$. The first should be obvious. For the second, observe that, at state w_1^0, only the reflexive arrows survive priority merge. The strict arrow from w_1^0 to v_2^1 gets eliminated because we have neither $w_1^0R_a^<v_2^1$ nor $w_1^0R_av_2^1$. The R_a (double-)arrow from w_1^0 to v_1^1 gets eliminated this time because these states are incomparable for b. Note finally that $\mathcal{M}_1, w \vDash \langle a/b\rangle^<\top$, while $\mathcal{M}_2, w_1^0 \vDash [a/b]^<\bot$.

For the second claim, consider the pair of a- and b-bisimilar models in Figure 5, standardly used to show that intersection modalities are not definable in terms of the modalities they intersect. The new observation here is that these two models are also a/b- and b/a-bisimilar. Indeed, even if priority merge removes the strict arrow of the

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press



Fig. 5. Two a-, b-, a/b-, and b/a-bisimilar models.

low-priority agent in $\mathcal{M}_2$, the remaining arrow is sufficient to match the one in $\mathcal{M}_1$. However, $\mathcal{M}_1, w \vDash \langle a||b\rangle p$, while $\mathcal{M}_2, w' \nvDash \langle a||b\rangle p$. □

While it is not the case that even given the individual weak and strict modalities, priority merge can express intersection, the converse is actually true: priority merge is definable in $\mathcal{L}_{[t],[t]^<,T_{||}}$:

Observation 2. *The following are valid in the class of plausibility models*:

$$[t/t']\varphi \leftrightarrow [t]^<\varphi \wedge [t||t']\varphi.$$

$$[t/t']^<\varphi \leftrightarrow [t]^<\varphi \wedge [t||t']^<\varphi.$$

Proof. This observation is a direct corollary of Definition 1.2 and Theorem 1. But since this result is fundamental for the axiomatization and proof theory that we give later on, we provide the argument from first principles nonetheless.

We show that the equivalence holds in "diamond" form.

$$\langle t/t'\rangle\varphi \leftrightarrow \langle t\rangle^<\varphi \vee \langle t||t'\rangle\varphi.$$

From left to right: Suppose there is a v such that $wR_{t/t'}v$ and $\mathcal{M}, v \vDash \varphi$. Given the definition of $wR_{t/t'}v$ this can happen in two cases. First $wR_t^<v$, which makes the first disjunct true. In the second case we have $w \equiv_t v$ and $wR_{t'}^<v$. Since the latter implies $wR_{t'}v$ we get $wR_{t||t'}v$. For the right-to-left direction, suppose that $\mathcal{M}, w \vDash \langle t\rangle^<\varphi \vee \langle t||t'\rangle\varphi$. If the first disjunct is true we know that there is a v such that $wR_t^<v$ and $\mathcal{M}, v \vDash \varphi$. From $wR_t^<v$ we obtain $wR_{t/t'}v$ directly. Now if, instead, the second disjunct is true, we know that there is a v' is such that $\mathcal{M}, v' \vDash \varphi$ and $wR_{t||t'}v'$. Now there are three cases to consider: either one of these two relations between w and v' is strict, or they are both symmetric. In each of these cases we obtain $wR_{t/t'}v'$.

Now consider the second equivalence, again in diamond form:

$$\langle t/t'\rangle^<\varphi \leftrightarrow \langle t\rangle^<\varphi \vee \langle t||t'\rangle^<\varphi.$$

Suppose that there is a v such that $wR_{t/t'}^<v$ and $\mathcal{M}, v \vDash \varphi$. This means, in particular, that $wR_{t/t'}v$ but not the other way around. By unpacking the definition of $R_{t/t'}$, we know that this can happen in only two cases: either $wR_t^<v$, or $wR_t^{\equiv}v$ and $wR_{t'}^<v$. The first case gives us the first disjunct directly, and the second ensures that v is accessible from w by the strict version of the intersection of R_t and $R_{t'}$. For the other direction, suppose that either there is a v such that $wR_t^<v$ and $\mathcal{M}, v \vDash \varphi$, or a v' such that $wR_{t||t'}^<v'$ and $\mathcal{M}, v' \vDash \varphi$. The first case, as before, gives, $wR_{t/t'}^<v$ directly. The second case gives us that both wR_tv' and $wR_{t'}v'$, but that the converse of either of these fail. Either way we get $wR_{t/t'}^<v'$. □

To summarize: Adding pooling modalities, either intersection or priority merge, to the language of multi-agent individual defeasible knowledge increases expressive

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

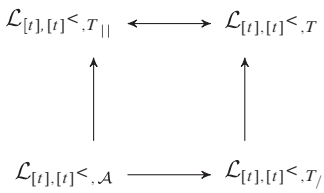


Fig. 6. Overview of the relative expressive powers of the languages of intersection and priority merge modalities. Unidirectional arrows point towards strictly more expressive languages, bidirectional arrows link equally expressive ones.

power, and so does adding intersection to the language of priority merge. As we have just shown, however, the converse is not true. Adding priority merge to the language containing strict individual and intersection modalities does not increase expressive power. This is summarized in Figure 6.

3.2. The relation between beliefs of subgroups and groups. Let us first recall some of the salient properties of distributed knowledge. They will serve as a benchmark for comparison. We write K_a for an individual's knowledge operator and D_G for distributed knowledge in group G.

$$K_a\varphi \to D_G\varphi, \text{ whenever } a \in G. \qquad (\text{incl–}D_G)$$

$$D_G\varphi \to D_{G'}\varphi, \text{ whenever } G \subseteq G'. \qquad (\text{mon–}D_G)$$

$$D_G\varphi \to \varphi. \qquad (\text{T–}D_G)$$

(incl–D_G) expresses that distributed knowledge is a sub-relation of the epistemic equivalence relations of each agent. Group monotonicity (mon–D_G) is a more general version of that axiom: whatever is distributed knowledge for any of its subgroups is also distributed knowledge for the group. (T–D_G) is the well-known "truth axiom." Since not all of the belief modalities we work with validate T, the "consistency axiom" D, which follows from T, will be more prominent:

$$D_G\varphi \to \neg D_G\neg\varphi. \qquad (\text{D–}D_G)$$

Which of these logical properties still hold in the case of our subgroup- and group-pooled belief operators depends on whether one considers belief as $[t]$, B^1_t, or B^2_t. We look at them in turn.

Consider first the case of the defeasible knowledge operators $[t]$. Recall that any term t is of the form a, $t||t'$, or t/t'. The logical behavior of the intersection modalities $[t||t']$ only differs from the one of distributed knowledge by their introspective properties (c.f. Section 3.3). Both are intersections of reflexive relations. Indeed we get:

$$\vDash [t]\varphi \to [t||t']\varphi, \text{ for any } t', \qquad (\text{incl–}[t||t'])$$

$$\vDash [t||t']\varphi \to \varphi. \qquad (\text{T–}[t||t'])$$

Observe that, since intersection between more than two relations is defined as the composition of binary intersections, (incl–$[t||t']$) is also a defeasible knowledge counterpart of (mon–D_G). Recall also that the order in which the terms appear in $[t||t']$ does not matter. (T–$[t||t']$) is valid, because $[t||t']$ is interpreted over weak plausibility orderings, which are reflexive. The formula is not valid in its strict variant, of course.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

Defeasible knowledge operators behave differently in the case of priority merge. The implication from subgroup to group belief obtained via priority merge is order-dependent. That is, on the one hand, we do get

$$\vDash [t]\varphi \to [t/t']\varphi, \text{ for any } t' \qquad \text{(incl–}[t/t']\text{)}$$

but, on the other hand,

$$\nvDash [t]\varphi \to [t'/t]\varphi.$$

Validity of the first formula follows from the fact that $R_{t/t'} \subseteq R_t$, by definition. For a counter-example to the second just consider a model with two states, only one of which makes p true, where two agents a and b order these states in strictly opposite orders.

Since reflexivity is preserved under priority merge as well, we do however keep the T axiom, and thus consistency, for priority merge operators.

$$\vDash [t/t']\varphi \to \varphi. \qquad \text{(T–}[t/t']\text{)}$$

Finally, the (un-)definability results presented above give us a clear picture of the relation between priority merge and intersection modalities for defeasible knowledge. The former implies the latter, but not the other way around:

$$\vDash [t/t']\varphi \to [t||t']\varphi,$$
$$\nvDash [t||t']\varphi \to [t/t']\varphi.$$

If we now turn to the belief operators B^1_t and B^2_t, the picture changes again. Already for intersection modalities $B^i_{t||t'}$ we lose the inclusion property, with the model in the previous paragraph providing a counter-example for both B^1_a and B^2_a.

$$\nvDash B^i_t\varphi \to B^i_{t||t'}\varphi.$$

The individual version of these belief operators do not validate T, and neither do their merged modalities. The situation for D is more interesting. Recall that, while the operator B^2_t does validate D, B^1_t does not. The consistency of group beliefs and that of individual beliefs are, however, independent from one another.

OBSERVATION 3. *Both* $B^1_a\bot \wedge B^1_b\bot \wedge \neg B^1_{a||b}\bot$ *and* $\neg B^1_a\bot \wedge \neg B^1_b\bot \wedge B^1_{a||b}\bot$ *are consistent.*

Proof. Consider the two models in Figure 7. First, in $\mathcal{M}_1$, $max_{R_a} = \emptyset$ and $max_{R_b} = \emptyset$, but $max_{R_{a||b}} \neq \emptyset$. Therefore, $\mathcal{M}_1, 0 \vDash B^1_a\bot \wedge B^1_b\bot \wedge \neg B^1_{a||b}\bot$. And second, in $\mathcal{M}_2$, $max_{R_a} \neq \emptyset$ and $max_{R_b} \neq \emptyset$ but $max_{R_{a||b}} = \emptyset$. Therefore, $\mathcal{M}_2, 0 \vDash \neg B^1_a\bot \wedge \neg B^1_b\bot \wedge B^1_{a||b}\bot$. □

We already saw that priority-merged defeasible knowledge invalidates inclusion for subgroups which are not at the top of the priority ordering, and the same holds for belief as maximum plausibility, under both semantics:

$$\vDash B^i_t\varphi \to B^i_{t/t'}\varphi, \qquad \text{(incl–}B^i_{t/t'}\text{)}$$
$$\nvDash B^i_t\varphi \to B^i_{t'/t}\varphi.$$

The situation regarding consistency, i.e., axiom D, is similar as for intersection. For $B^2_{t/t'}$, consistency is inherited directly from reflexivity, which is also preserved by priority

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

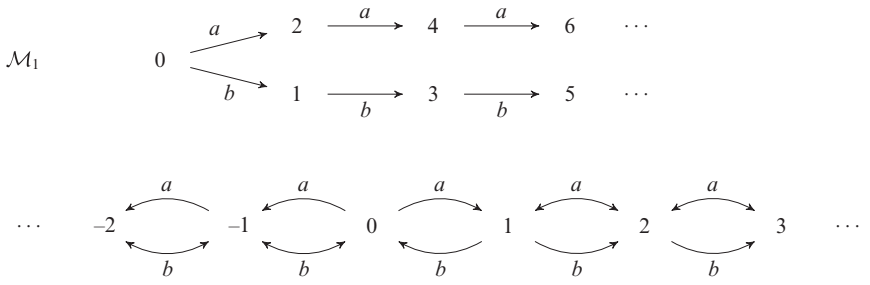


Fig. 7. Consistency of B^1_a and B^1_b is independent from consistency of $B^1_{a||b}$. $\mathcal{M}_1, 0 \vDash B^1_a\bot \wedge B^1_b\bot \wedge \neg B^1_{a||b}\bot$, while $\mathcal{M}_2, 0 \vDash \neg B^1_a\bot \wedge \neg B^1_b\bot \wedge B^1_{a||b}\bot$.

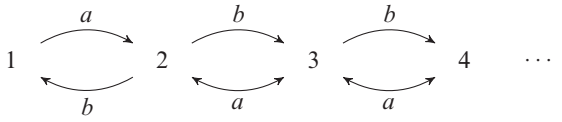


Fig. 8. State 1 satisfies $\neg B^1_a\bot$ and $\neg B^1_b\bot$, but also $B^1_{a/b}\bot$.

merge. For $B^1_{t/t'}$, neither the individual nor the priority merged modalities validate D. Furthermore, we also get that consistency at the subgroup level does not entail consistency at the group level.

Observation 4. *The following is consistent*: $\neg B^1_a\bot \wedge \neg B^1_b\bot \wedge B^1_{a/b}\bot$.

Proof. See Figure 8.

□

The converse, however, turns out to depend again on the priority order. Consistency of B^1_t is necessary for $B^1_{t/t'}$ to be consistent as well. Indeed, given that R_t is serial, inconsistency in B^1_t can only arise if all maximal linear chains starting at a given state are infinite and strict for t. Those then get copied directly into $R_{t/t'}$. However, consistency of $B^1_{t'}$ is not necessary for the consistency of $B^1_{t/t'}$.

Finally, when moving from defeasible knowledge operators to B^i_t we also loose the dependency between our two aggregated-belief modalities:

$$\nvDash B^i_{t/t'}\varphi \to B^i_{t||t'}\varphi,$$
$$\nvDash B^i_{t||t'}\varphi \to B^i_{t/t'}\varphi.$$

3.3. Introspection. Again, for the purpose of comparison, it is useful to start with a reminder of the key introspective properties of distributed knowledge:

$$D_G\varphi \to D_G D_G\varphi, \tag{4–D_G}$$

$$\neg D_G\varphi \to D_G \neg D_G\varphi. \tag{5–D_G}$$

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

These properties are directly inherited, semantically,[7] from the transitivity and euclidicity of the underlying individual relations.

The defeasible knowledge modalities are of course S4, and so they validate positive introspection ("4"), too, but not negative introspection ("5"). The counter-examples for the latter are standard [2, 3, 5].

$$\models [t||t']\varphi \rightarrow [t||t'][t||t']\varphi, \qquad (4\text{–}[t||t'])$$
$$\not\models \neg[t||t']\varphi \rightarrow [t||t']\neg[t||t']\varphi.$$

Similarly as for distributed knowledge, (4–$[t||t']$) is valid but not syntactically derivable, in the axiom system to be introduced below. It will therefore need to be added as an axiom. Positive introspection is also valid for the priority merge modality, although as we will see (in Section 4.1), it is derivable in that axiom system.

It is well known that, unlike for instance with common knowledge, intersection modalities are not individually introspective. It is possible for φ to be distributed knowledge without any of the group members knowing that it is. Since this is a special case of our intersection modalities, we get

$$\not\models [t||t']\varphi \rightarrow [t][t||t']\varphi.$$

This formula is not valid either for priority merge, both for the first and the second agent in the priority order. The reason for that is that the usual counter-examples to these validities for intersection modalities apply as well for priority merge.

$$\not\models [t/t']\varphi \rightarrow [t][t/t']\varphi,$$
$$\not\models [t/t']\varphi \rightarrow [t'][t/t']\varphi.$$

However, defeasible knowledge aggregated by priority merge are "distributively introspective":

$$\models [t/t']\varphi \rightarrow [t||t'][t/t']\varphi.$$

This follows from the facts that $[t/t']$ validates 4 and that $[t/t'][t/t']\varphi$ implies $[t||t'][t/t']\varphi$. The other way around, however, is not valid; defeasible knowledge aggregated by intersection are not "introspective" through priority merge. This comes from the fact that all states strictly more plausible for the term t with highest priority are considered in the evaluation of $[t/t']$, irrespective of whether t' can access them or not, i.e. whether they are taken into account in evaluating $[t||t']$.

$$\not\models [t||t']\varphi \rightarrow [t/t'][t||t']\varphi.$$

In comparison with defeasible knowledge, belief operators B_t^i validate stronger introspection principles: just as the individual modalities do, the intersection and priority merge operators satisfy both 4 and 5.

$$\models B_t^i\varphi \rightarrow B_t^i B_t^i\varphi, \qquad (4\text{–}B_t^i)$$
$$\models \neg B_t^i\varphi \rightarrow B_t^i\neg B_t^i\varphi. \qquad (5\text{–}B_t^i)$$

[7] Because of the standard modal language's inability to fully express the fact that D_G corresponds to intersection of individual modalities, these formulas are not derivable in S5 augmented with distributed knowledge. They have to be taken as axioms.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

"Subgroup introspection" about intersected belief $B^i_{t||t'}$ also fails. In fact, since modalities B^i_t do not satisfy T, it is even possible for the agents to be completely mistaken about what is "distributed belief" in the group. For instance, $B^2_{a||b}\varphi \wedge B^2_a B^2_{a||b}\neg\varphi$ is consistent. The same holds for beliefs aggregated with priority merge $B^i_{t/t'}$:

$$\not\models B^i_{t||t'}\varphi \to B^i_t B^i_{t||t'}\varphi,$$
$$\not\models B^i_{t/t'}\varphi \to B^i_t B^i_{t/t'}\varphi.$$

The converse principle does hold for priority merge but not for intersection merge. That is:

$$\models B^i_t B^i_{t/t'}\varphi \to B^i_{t/t'}\varphi, \qquad \text{(conv.sub.intro–}B^i_{t/t'}\text{)}$$
$$\not\models B^i_t B^i_{t||t'}\varphi \to B^i_{t||t'}\varphi.$$

Finally, while "distributed introspection" is valid for priority-merged defeasible knowledge, it fails for B^i_t:

$$\not\models B^i_{t/t'}\varphi \to B^i_{t||t'}B^i_{t/t'}\varphi.$$

The source of this failure is analogous to the one that makes priority-merge introspection fail for defeasible knowledge. The states that are considered in the evaluation of $[t/t']$ might be different than those that are considered in the evaluation of $[t||t']$.

§4. Axiomatization. We now turn to the axiomatization of the set of validities in $\mathcal{L}$ over the class of plausibility frames. Recall that Girard and Seligman [17] have provided such an axiomatization using the additional expressive power of hybrid logic. They left open the question whether this set could also be axiomatized *without* nominals. We show here that it can. The key to this result is Observation 2, the

Table 1. *Summary of the comparison between the standard group operators and ours*

	T	D	4	5	Incl	Mon	Sub.intro	Conv.sub.intro	Dis.intro
C_G	✓	✓	✓	✓			✓	✓	✓
D_G	✓	✓	✓	✓	✓	✓		✓	✓(4)
$[t\|\|t']$	✓	✓	✓		$[t]\varphi \to [t\|\|t']\varphi$				✓(4)
$[t/t']$	✓	✓	✓		$[t]\varphi \to [t/t']\varphi$				$[t/t']\varphi \to [t\|\|t'][t/t']\varphi$
$B^1_{t\|\|t'}$			✓	✓					
$B^1_{t/t'}$			✓	✓	$B^1_t\varphi \to B^1_{t/t'}\varphi$			$B^1_t B^1_{t/t'}\varphi \to B^1_{t/t'}\varphi$	
$B^2_{t\|\|t'}$		✓	✓	✓					
$B^2_{t/t'}$		✓	✓	✓	$B^2_t\varphi \to B^2_{t/t'}\varphi$			$B^2_t B^2_{t/t'}\varphi \to B^2_{t/t'}\varphi$	

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

definability of the priority-merge modality $[t/t']$ in terms of $[t]$, $[t]^<$ and $[t||t']$. As a result, the task of axiomatizing validities for the whole language $\mathcal{L}$ boils down to axiomatizing the simpler language $\mathcal{L}_{[t],[t]^<,T_{||}}$. There are well-known techniques for axiomatizing each of its components, i.e., weak/strict preferential modalities and intersection modalities. Therefore, the main challenge of the completeness result is to show how these techniques can be consistently combined.

The axiom system that we use contains standard axioms to capture basic properties of operators U, $[t]$ (for all $t \in T_{||}$), as well as of their strict variants:

1. S5 for the universal modality U.
2. The necessitation rule, T, and 4 for $[t]$.
3. K and the necessitation rule for $[t]^<$.

The K axiom for the individual "weak" modalities is derivable from the following axiom, which comes from [15], together with the axiom (red–$[t||t']$) that we introduce below.[8]

$$([t]\varphi \wedge [t'](\varphi \to \psi)) \to [t||t']\psi. \tag{dis}$$

Observe that, using necessitation, (dis) also entails the more standard expression of the "inclusion" property for intersection modalities mentioned in Section 3.2:

$$[t]\varphi \to [t||t']\varphi.$$

This property does *not* hold for the strict version of the intersection modalities. It would hold if the latter were defined on the basis of the intersection of the strict underlying relations. Recall, however, that instead we define them using the strict version of the intersection of the (weak) relations. Therefore, the corresponding inclusion validity for the strict modality is rather:

$$\bigwedge_{a \in t} [a]^<\varphi \to [t]^<\psi, \tag{wincl–$[t]^<$}$$

with $a \in t$ a shorthand for "$a \in \mathcal{A}$ occurs in t."

Our syntax requires adding a set of axioms that allows us to manipulate the terms inside modalities for intersection. In (red–$[t||t']$) the notation $\sigma(t : t'||t' \to t')$ stands for the result of substituting in t all instances of $t'||t'$ with t'. We use a similar set of axioms for the strict variant of the intersection modalities.

$$[t||t']\varphi \leftrightarrow [t'||t]\varphi, \tag{com–$[t||t']$}$$

$$[t||(t'||t'')]\varphi \leftrightarrow [(t||t')||t'']\varphi, \tag{asso–$[t||t']$}$$

$$[t]\varphi \leftrightarrow [\sigma(t : t'||t' \to t')]\varphi, \tag{red–$[t||t']$}$$

$$[t||t']^<\varphi \leftrightarrow [t'||t]^<\varphi, \tag{com–$[t||t']^<$}$$

$$[t||(t'||t'')]^<\varphi \leftrightarrow [(t||t')||t'']^<\varphi, \tag{asso–$[t||t']^<$}$$

$$[t]^<\varphi \leftrightarrow [\sigma(t : t'||t' \to t')]^<\varphi. \tag{red–$[t||t']^<$}$$

[8] Note that in the presence of K and necessitation the rule of substitution of logically equivalent formulas is also derivable. That is: if $\varphi \leftrightarrow \psi$ is derivable then $[t]\varphi \leftrightarrow [t]\psi$ is derivable as well. This applies also to the strict modalities, since for them K and necessitation are parts of the basic axiomatization.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

Observe that, given (com–$[t||t']$), (asso–$[t||t']$), and (red–$[t||t']$), every formula of the form $[t]\varphi$ where t is a complex relational term is provably equivalent to a formula $[t']\varphi$ where each individual agent $a \in \mathcal{A}$ occurs at most once in t'. Since $\mathcal{A}$ is finite there can be only finitely many such formulas. The same holds for formulas of the form $[t]^{<}\varphi$.

OBSERVATION 5. *For any $\varphi \in \mathcal{L}_{[t],[t]^{<},T_{||}}$, there are only finitely many non-equivalent formulas of the form $[t||t']\varphi$, and only finitely many non-equivalent formulas of the form $[t||t']^{<}\varphi$.*

This implies furthermore that, up to logical equivalence, we can treat any intersection modality based on a complex term t as an intersection modality over the *set* of agents which occur in t. Neither duplicates nor the order between agents matter. Moreover, we can measure the complexity of a term t by simply counting the number of different agents occurring in it. Hence the notation $a \in t$ used already for "$a \in \mathcal{A}$ occurs in t."

Our final set of axioms regiments the interaction between the weak and the strict variant of the different modalities in our language. These are again standard axioms, this time from preference logic [7].

$$[t]\varphi \to [t]^{<}\varphi, \qquad \text{(incl–}[t]^{<}\text{)}$$

$$U\varphi \to [t]\varphi, \qquad \text{(incl–}[t]\text{)}$$

$$[t]^{<}\varphi \to [t][t]^{<}\varphi, \qquad (\text{inter}_1)$$

$$[t]^{<}\varphi \to [t]^{<}[t]\varphi, \qquad (\text{inter}_2)$$

$$(\neg\varphi \wedge [t]^{<}\psi \wedge [t](\neg\psi \to [t]\varphi)) \to [t]\psi. \qquad (\text{inter}_3)$$

Axiom inter_3 is perhaps easier to read in contrapositive:

$$\varphi \wedge \langle t\rangle\psi \to (\langle t\rangle^{<}\psi \vee \langle t\rangle(\psi \wedge \langle t\rangle\varphi)).$$

In words: suppose that φ is true at a state w, and that that same state can access, through the weak relation R_t, a state v where ψ is true. Then two things can happen: either v can also access w, or it does not. In the first case this means that $\langle t\rangle(\psi \wedge \langle t\rangle\varphi)$ is true at w itself. In the second case, precisely because of the definition of the strict relation, we get $\langle t\rangle^{<}\psi$ true at w.

THEOREM 2. *The set of axioms and rules just listed is sound and complete with respect to the class of plausibility frames.*

Proof. The proof follows standard "bulldozing" ([10], pp. 221–222)[9] and "splitting" techniques except for three crucial places: First we need to carefully choose the strict representative sub-relations of $<_t$-clusters, which ensures that the usual canonical model construction can be bulldozed in a way that leaves the intersection relations sufficiently intact. Those can then be fixed by making copies of states, as in the

[9] The so-called "bulldozing" technique is used to remedy the fact that, in the canonical model, the strict relations $R_t^{<}$ may not be irreflexive and assymetric. Some maximally consistent sets might form what have been called "strict clusters", i.e., sets of MCS where the strict relation $R_t^{<}$ is actually reflexive, transitive, and symmetric. Figuratively, [28], who developed the technique, thought of them as chunks or balls that need to be *flattened*, that is, where $R_t^{<}$ has to be transformed in the largest assymetric, transitive and irreflexive sub-relation of R_t. Hence the name "bulldozing."

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

standard completeness proof for distributed knowledge. Here complex terms need to be decomposed step-by-step, while being careful of also copying symmetric R_t, i.e., non-strict links, in order not to undo what has been achieved by bulldozing.

As usual, we need to show that every consistent set Σ is satisfiable. A standard Lindenbaum lemma shows that any such set can be extended to a maximally consistent set Σ^+. In order for the universal modality U to be indeed universal in the model that we will construct, we need to restrict the set W of maximally consistent sets to those that contain exactly the same formulas of the form $U\varphi$ as Σ^+. This is standard, so for readability we will not adopt any explicit notation for it.

DEFINITION 4.10 (canonical model). *Let the canonical model* $\mathcal{M} = \langle W, \{R_t, R_t^<\}_{t \in T_{||}}, V\rangle$ *be such that*:

- *W is the set of all maximally consistent sets as described above*;
- *the canonical relations are defined as usual, for each* $t \in T_{||}$:
  - $\Gamma R_t \Delta$ *iff for all* $\varphi \in \Delta$, *we have* $\langle t\rangle\varphi \in \Gamma$, *and*
  - $\Gamma R_t^< \Delta$ *iff for all* $\varphi \in \Delta$, *we have* $\langle t\rangle^<\varphi \in \Gamma$;
- $V(p) = \{\Gamma\colon p \in \Gamma\}$. □

The required existence and truth lemmas are standard. In this model we can have that the relations $R_t^<$ are not irreflexive, as well as that $R_{t||t'} \subsetneq R_t \cap R_{t'}$. We address these problems in that order, i.e., for the first defect, we "bulldoze" the canonical model, and then for the second, we make some copies of states to "split" the canonical relations.

An $R_t^<$-cluster is a set C of elements of W such that $\Gamma R_t^< \Delta$ for all $\Gamma, \Delta \in C$. These need to be "flattened out" in a way that does not disrupt the existing intersections of relations. We do this as follows. Start by picking an arbitrary strict ordering $<^*$ on W. For any given agent a, consider the set of all $R_a^<$-clusters and index them with a sufficiently large set I_a. For each such cluster C_i^a, the strict ordering $<_i^a$ is defined as $<^* \cap (C_i^a \times C_i^a)$.

Define C as the set of those $\Gamma \in W$ for which there is an agent a such that Γ is in an $R_a^<$-cluster, and W^- as the complement of C. Then let B be $C \times \mathbb{N}$. We are now ready to define the new bulldozed model M'. We take $W' = W^- \cup B$. A projection function β will be useful for what follows, and will be the basis for the key bisimulation argument later on. So we define $\beta\colon W' \to W$ as $\beta(x) = x$ for all $x \in W^-$ and otherwise $\beta((\Gamma, n)) = \Gamma$. The valuation $V'(x)$ is set to $V(\beta(x))$, for all $x \in W'$.

Now for the relations. What we do is bulldoze all the $R_a^<$ and take care of complex terms on the basis of that. Take any pair $x, y \in W'$,

1. If x or y is in W^-, or $\beta(x)$ and $\beta(y)$ are in different a-clusters, then $xR_a^{<'}y$ iff $\beta(x)R_a^<\beta(y)$.
2. If $\beta(x)$ and $\beta(y)$ are in the same $R_a^<$-cluster we make a further distinction by cases (note that in this case $x = (\Gamma, m)$ and $y = (\Delta, n)$ for some $\Gamma, \Delta \in W$ and $m, n \in \mathbb{N}$):
   - $m \neq n$: then $xR_a^{<'}y$ iff $m < n$;
   - $m = n$: then we use the strict order $<_i^a$ chosen for the cluster C_i^a to which $\beta(x)$ and $\beta(y)$ belong: $xR_t^{<'}y$ iff $\beta(x) <_i^a \beta(y)$.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

The weak relation R'_a is then defined as follows:

$$xR'_a y \text{ iff } \begin{cases} \beta(x)R_a\beta(y), & \text{if } x \text{ or } y \text{ or both are not in any } a\text{-cluster;} \\ \text{either } x = y \text{ or } xR_a^{<'} y, & \text{otherwise.} \end{cases}$$

Now we consider complex terms. The first thing to observe is that (wincl–$[t]^<$) ensures that if a state y is accessible from x through the strict version of the intersection of the relations of individual agents appearing in a complex term t, then y must also be accessible through the strict relation of *some* a appearing in t. The converse direction should not be expected. It states that if a state y is accessible from x through the strict relation of some a appearing in t, then y must also be accessible through the strict version of the intersection of the relations of individual agents appearing in t. This boils down to say that $R^<_{t||t}$ is defined as the intersection of $R_{t||t}$ and $R^<_a$ for some a. Claim 1 below shows that bulldozing takes care of this problem.

Observation 6. *For all* $\Delta, \Gamma \in W$, $\Gamma R^<_{t||t'}\Delta$ *only if* $/\Gamma R_{t||t'}\Delta$ *and* $\Gamma R^<_a\Delta$ *for some* $a \in t||t'$.

Proof. Suppose that $\Gamma R^<_{t||t'}\Delta$ and take $\varphi \in \Delta$. We have to start by showing that both $\langle t||t'\rangle\varphi$ and, for some $a \in t||t'$, $\langle a\rangle^<\varphi$ are in Γ. The first follows directly from our axiom (incl–$[t]^<$) and the fact that $\langle t||t'\rangle^<\varphi \in \Gamma$. For the second, suppose that this is not the case. So for all $a \in t||t'$, we have $[a]\neg\varphi \in \Gamma$. But then (wincl–$[t]^<$) gives directly $[t||t']\neg\varphi \in \Gamma$, which would entail $\neg\varphi \in \Delta$, against our assumption. □

So, for the new version of the strict relation for complex terms, we take:

$$xR^{<'}_{t||t'} y \text{ iff } \begin{cases} xR_a^{<'} y & \exists a \in t||t' \text{ s.t. } \beta(x), \beta(y) \text{ are in the same } a\text{-cluster;} \\ \beta(x)R^<_{t||t}\beta(y), & \text{otherwise.} \end{cases}$$

The idea behind this definition is this: If there is an $a \in t||t'$ such that both $\beta(x)$ and $\beta(y)$ are in a a-cluster, then we fix $xR^{<'}_{t||t'} y$ if and only if $xR_a^{<'} y$. This is well-defined for the following reason. Recall that in this case x and y are pairs (Γ, m) and (Δ, n). Then $xR_a^{<'} y$ obtain if either $m < n$ or $\Gamma <_i^a \Delta$. Now take any other agent $b \in t||t'$ such that $\beta(x)$ and $\beta(y)$ are in a b-cluster. The definition $xR_b^{<'} y$ gives us, again, that this obtains if and only if either $m < n$, in which case the ordering of x and y according to $R_a^{<'}$ and $R_b^{<'}$ agree, or $\Gamma <_i^b \Delta$. In that latter case, since both $<_i^a$ and $<_i^b$ are just different restrictions of the same strict ordering $<^*$ to a- and b-clusters, respectively, we get that $x <_i^a y$ if and only if $x <_i^b y$. Well-definedness is immediate in the second case of the definition.

The first thing to observe is that this construction does not invalidate the existence lemma:

Lemma 1. *For all* $x \in W'$, *if* $\langle t||t'\rangle^<\varphi \in \beta(x)$ *then there is a* y *such that* $xR^<_{t||t'}y$ *and* $\varphi \in \beta(y)$.

Proof. The existence lemma in the canonical model gives us a Δ such that $\beta(x)R^<_{t||t'}\Delta$ and $\varphi \in \Delta$. Now, if there is no $a \in t||t'$ such that $\beta(x)$ and Δ are part of the same a-cluster then any y such that $\beta(y) = \Delta$ will do. If there is such an a then take $y = (\Delta, n)$ with $m < n$. □

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

Finally, for $R'_{t||t'}$, we proceed as we did for R_a, except that we use the new $R^{<'}_{t||t'}$:

$$xR_{t||t'}y \text{ iff } \begin{cases} \beta(x)R_{t||t'}\beta(y), & \text{if } x \text{ or } y \text{ is in } W^-; \\ \text{either } x = y \text{ or } xR^{<'}_{t||t}y, & \text{otherwise.} \end{cases}$$

The next step is making sure that all the $R^<_t$ are adequate w.r.t their corresponding weak R_a [7]. That is, we need to show that each $R^<_t$ is the largest irreflexive and transitive sub-relation of R_t. Recall that (inter$_2$) gives us that $\Gamma R_t \Delta$ and not $\Delta R_t \Gamma$ together imply $\Gamma R^<_t \Delta$ already in the canonical model. What remains to be shown is that bulldozing ensures the converse.

Claim 1. *For all terms t, $R^{<'}_t$ is adequate.*

Proof. For $R^{<'}_a$ with $a \in \mathcal{A}$ this can be shown by the standard argument. For complex terms, suppose $R^{<'}_{t||t'}$ is not adequate. Then there is a pair x, y such that $xR^{<'}_{t||t'}y$ and $yR^{<'}_{t||t'}x$. Observe that this case is ruled out by definition if there is no $a \in t||t'$ for which $\beta(x)$ and $\beta(y)$ are in the same a-cluster. So suppose there is such an a. Then the definition of $R^{<'}_t$ gives us that there are b and c both in $t||t'$ such that $xR^{<'}_b y$ and $yR^{<'}_c x$ and, furthermore, $\beta(x)$ and $\beta(y)$ are in the same b-cluster and the same c-cluster. So we know that $x = (\Gamma, n)$ and $y = (\Delta, m)$ for some $\Gamma, \Delta \in W$ and $n, m \in \mathbb{N}$. If $n \neq m$ then $xR^{<'}_b y$ iff $n < m$ iff $xR^{<'}_c y$, contradicting the adequacy of $R^{<'}_b$ and $R^{<'}_c$. If $n = m$ then $xR^{<'}_b y$ iff $\Gamma <^b_i \Delta$ according to the arbitrary $<^b_i$ that we have chosen early on in the construction, and similarly for $<^c_i$. But observe that we choose these two strict orderings in such a way that they agree on Γ and Δ, contradicting again the adequacy of $R^{<'}_b$ and $R^{<'}_c$. □

Claim 2. *M and M' are modally equivalent.*

Proof. For all $x \in W'$,

$$M', x \vDash \varphi \quad \Leftrightarrow \quad M, \beta(x) \vDash \varphi.$$

We show this by induction. The bisimulation argument already given in [7] gives us all the cases except for modalities based on complex terms. So we show first that, for all $x \in W^-$,

$$M', x \vDash \langle t||t' \rangle \varphi \quad \Leftrightarrow \quad M, \beta(x) \vDash \langle t||t' \rangle \varphi.$$

($\Rightarrow$) $M', x \vDash \langle t||t' \rangle \varphi$ implies that there is a $y \in W'$ such that $xR'_{t||t'}y$ and $M', y \vDash \varphi$. The latter gives us $M, \beta(y) \vDash \varphi$ by our inductive hypothesis. It remains to show that $\beta(x)R_{t||t'}\beta(y)$. If x or y are in W^- this follows by definition. Otherwise either $x = y$ or $xR^{<'}_{t||t'}y$. Since $R_{t||t'}$ is reflexive, the first case is taken care of. For the second case, observe that $xR^{<'}_{t||t'}y$ entails $\beta(x)R^<_{t||t'}\beta(y)$ in both cases of the definition, which in turns entails $\beta(x)R_{t||t'}\beta(y)$ by (incl–$[t]^<$).

($\Leftarrow$) $M, \beta(x) \vDash \langle t||t' \rangle \varphi$ implies that there is a $\Delta \in W$ such that $\beta(x)R_{t||t'}\Delta$ and $M, \Delta \vDash \varphi$. Again, the induction hypothesis gives us $M', y \vDash \varphi$ for any y such that $\beta(y) = \Delta$. The rest of the argument follows the steps in Lemma (1): If we can show that $xR'_{t||t'}y$ for some of those y's we are done. We get this by construction if $\beta(x)$ or

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

Δ are in W^-, and for the case of $x = y$. Otherwise $x = (\Gamma, m)$. Pick, as in Lemma (1), $y = (\Delta, n)$ with $n > m$.

The final case is for the strict modality. We show that, for all $x \in W'$,

$$M', x \vDash \langle t||t'\rangle^{<}\varphi \quad \Leftrightarrow \quad M, \beta(x) \vDash \langle t||t'\rangle^{<}\varphi.$$

($\Rightarrow$) $M', x \vDash \langle t||t'\rangle^{<}\varphi$ implies that there is a $y \in W'$ such that $xR^{<'}_{t||t'}y$ and $M', y \vDash \varphi$. Again, $M, \beta(y) \vDash \varphi$ by induction hypothesis. As already observed, the definition of $R^{<'}_{t||t'}$ gives us $\beta(x)R^{<}_{t||t'}\beta(y)$ in both cases.

($\Leftarrow$) $M, \beta(x) \vDash \langle t||t'\rangle^{<}\varphi$ implies that there is a $\Delta \in W$ such that $\beta(x)R^{<}_{t||t'}\Delta$ and $M, \Delta \vDash \varphi$. The result follows from Lemma 1 and the induction hypothesis. □

We now turn to the second problem the original canonical model had, namely terms t, t' for which $R_{t||t'} \subsetneq R_t \cap R_{t'}$, and similarly for the strict version of this relation. It should be clear that this is still the case in the bulldozed model M'. In effect, there are two types of defects that we need to consider:

D1 We have $xR'_a y$ and $yR'_a x$ for all $a \in t||t'$, but neither $xR'_{t||t'}y$ nor $yR'_{t||t'}x$.

D2 We have $xR^{<'}_a y$ for all $a \in t||t'$, but not $xR^{<'}_{t||t'}y$.

The third option— $xR'_a y$ and $yR'_a x$ for all $a \in t||t'$, but $xR^{<'}_{t||t'}y$—is ruled out by Observation 6.

So, take any point $x \in W'$ and consider its generated sub-models $G(x)$. We fix the defects in $G(x)$ of type D1 and D2 together, top-down from the most complex to the most simple terms. For each $x \in W'$ we define $G^1(x)$ as follows. Consider the set of pairs (x, y) such that $xR'_t y$ and $xR'_{t'}y$, where $t||t'$ is of the highest non-equivalent term complexity. For each such pair, if it exhibits a defect of type D1, then make two copies of the generated sub-model from y, and set both $xR'_t y$ and $yR'_t x$ for one copy of y, and the same with t' for the other copy. Take the transitive closure of both relations. If, on the other hand, a pair (x, y) exhibits a defect of type D2, then do the same but put $xR^{<'}_t y$ for one copy, and $xR^{<'}_{t'}y$. Repeat that procedure for all y that are successors of x in this new model, and so on for all step n from x, and then again successively for all terms of lower complexity. Call the resulting generated submodel $G^+(x)$. Our new model is the disjoint union of the $G^+(x)$ for all $x \in W'$.

Claim 3. *For all x, the models M', x and $G^+(x), x$ are modally equivalent.*

Proof. The two pointed models are bisimilar, which takes care of modal equivalence for all the operators except the intersection modalities. But observe that the only places where these models differ are where D1 or D2 defects occur. In both cases, the relation corresponding to the intersection modality remains unchanged, because it does not connect defective pairs. As a result, the truth of formulas built on the corresponding modalities is unaffected by the construction. □

Claim 4. *For all x, in $G^+(x)$ we have that $R'_{t||t'} = R'_t \cap R'_t$ and $R^{<'}_{t||t'} = R^{<'}_t \cap R^{<'}_{t'}$.*

Proof. Clear by construction. □

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

The proof can now be finished. By claims (1) and (2), every consistent set of formulas is satisfiable in a model where the strict relations are adequate for all terms. The construction that fixes D1 and D2 types of defects preserves adequacy. So, by claims (3) and (4), every consistent set is satisfiable in a model where not only the strict relations are adequate for all terms, but also the intersection relation coincides with the intersection of its constituting relations.

4.1. Derivable properties of priority merge. We illustrate how the definition of the priority-merge operator $[t/t']$ in terms of $[t||t']$ and $[t]^{<}$ allows us to axiomatically derive some of the key properties of priority merge which we introduced in Section 3 First, the distribution principle K for priority-merged defeasible knowledge:

$$\vdash [t/t'](\varphi \to \psi) \to ([t/t']\varphi \to [t/t']\psi). \qquad (\text{K}\text{–}[t/t'])$$

This follows directly from the definition of $[t/t']$ and the fact that $[t||t']$ and $[t]^{<}$ are themselves normal modalities. We also obtain positive introspection and truthfulness:

$$\vdash [t/t']\varphi \to [t/t'][t/t']\varphi; \qquad (4\text{–}[t/t'])$$

$$\vdash [t/t']\varphi \to \varphi. \qquad (\text{T}\text{–}[t/t'])$$

The former is a direct application of the definition of $[t/t']$ and the respective axioms 4 for both defining modalities; the latter follows directly from (T–$[t||t']$). Finally, we derive the following reduction principle:

$$\vdash [t/t]\varphi \leftrightarrow [t]\varphi.$$

Here, the formula $[t/t]\varphi$ is a shorthand for $[t||t]\varphi \wedge [t]^{<}\varphi$. Left to right: using (red–$[t||t']$) and (dis) we get that the first conjunct is provably equivalent to $[t]\varphi$. Right to left: $[t]\varphi$ entails $[t]^{<}\varphi$ by (incl–$[t]^{<}$), and then we use the same equivalence as in the other direction.

§5. Proof theory. Girard and Seligman [17] have provided a proof system for a hybrid extension of our language $\mathcal{L}$. In addition to this increased expressive power, they also make use of the additional deductive power that is provided by labeled sequents. They show that, in that framework, the set of validities in $\mathcal{L}$ over the class of plausibility frames can be given a sequent calculus that admits analytic cut.

What we show here is that, just as for the axiomatization, the hybrid machinery is not necessary for the proof theory either. We provide a labeled sequent calculus in the framework developed by [20]. This framework contains labels for states as well as relational atoms of the form $xR_t y$. As it turns out, admissibility of cut for our logic, as well as weakening and invertibility of the rules, are corollaries of a general result by Negri regarding frame conditions definable by so-called "geometric conditions" [22].

We use the axioms $\Gamma, x : \varphi \Longrightarrow x : \varphi, \Delta$ and $\Gamma, xR_t y \Longrightarrow xR_t y, \Delta$, as well as the standard rules for the Boolean connectives. The rest of our set of rules is listed below. The rule schema $\mathbf{R}[t]$, marked in bold, has the usual variable condition that y does not appear in the conclusion. We also use an equivalent of L$[t]$ and $\mathbf{R}[t]$ for the strict modalities $[t]^{<}$.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

$$\frac{y : \varphi, x : [t]\varphi, xR_t y, \Gamma \Longrightarrow \Delta}{x : [t]\varphi, xR_t y, \Gamma \Longrightarrow \Delta}\ (\mathrm{L}[t]) \qquad \frac{xR_t y, \Gamma \Longrightarrow y : \varphi, \Delta}{\Gamma \Longrightarrow x : [t]\varphi, \Delta}\ (\mathbf{R}[t])$$

$$\frac{xR_t y, xR_{t||t'} y, \Gamma \Longrightarrow \Delta}{xR_{t||t'} y, \Gamma \Longrightarrow \Delta}\ (\text{incl-int})$$

$$\frac{xR^{<}_{a_1} y, xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta \quad \ldots \quad xR^{<}_{a_n} y, xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta \quad \text{for all } a_1, \ldots, a_n \in t||t'}{xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta}\ (\text{incl-int}<)$$

$$\frac{xR_t y, xR_{t||t} y, \Gamma \Longrightarrow \Delta}{xR_t y, \Gamma \Longrightarrow \Delta}\ (\text{red}) \qquad \frac{xR^{<}_t y, xR^{<}_{t||t} y, \Gamma \Longrightarrow \Delta}{xR^{<}_t y, \Gamma \Longrightarrow \Delta}\ (\text{red}<)$$

$$\frac{xR_t y, xR^{<}_t y, \Gamma \Longrightarrow \Delta}{xR^{<}_t y, \Gamma \Longrightarrow \Delta}\ (\text{incl}<) \qquad \frac{xR^{<}_t z, xR_t y, yR^{<}_t z, \Gamma \Longrightarrow \Delta}{xR_t y, yR^{<}_t z, \Gamma \Longrightarrow \Delta}\ (\text{intr}_2 <)$$

$$\frac{xR^{<}_t z, yR_t z, xR^{<}_t y, \Gamma \Longrightarrow \Delta}{yR_t z, xR^{<}_t y, \Gamma \Longrightarrow \Delta}\ (\text{intr}'_2 <)$$

$$\frac{xR_t y, yR_t x, \Gamma \Longrightarrow \Delta \qquad xR_t y, xR^{<}_t y, \Gamma \Longrightarrow \Delta}{xR_t y, \Gamma \Longrightarrow \Delta}\ (\text{intr}_3)$$

Write $\Vdash \varphi$ for a formula provable in the axiom system presented in Section 4 and $\vdash \Gamma \Longrightarrow \Delta$ for a sequent provable in the system just presented.

Observation 7 (Deductive Equivalence). *For every formula φ, we have*

$$\Vdash \varphi \quad \textit{iff} \quad \vdash \Longrightarrow \varphi$$

Proof. The right-to-left direction follows standardly by soundness of the rules and completeness of the axiom system. For the converse, we simply illustrate the derivation of some of the key axioms using the rules above.

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{xR_a y, xR_b y, xR_{a||b} y, \ldots, y : \psi \Longrightarrow y : \psi}{xR_a y, xR_b y, xR_{a||b} y, \ldots, y : \varphi, y : \varphi \to \psi \Longrightarrow y : \psi}\ (\mathrm{PL})}{xR_a y, xR_b y, xR_{a||b} y, x : [a]\varphi, x : [b](\varphi \to \psi) \Longrightarrow y : \psi}\ \text{twice } (\mathrm{L}[*]))}{xR_{a||b} y, x : [a]\varphi, x : [b](\varphi \to \psi) \Longrightarrow y : \psi}\ \text{twice (incl-int)}}{x : [a]\varphi, x : [b](\varphi \to \psi) \Longrightarrow x : [a||b]\psi}\ (\mathrm{R}[a||b])$$

$$\cfrac{\cfrac{\cfrac{xR_a y, xR_{a||a} y, \ldots, y : \varphi \Longrightarrow y : \varphi}{xR_a y, xR_{a||a} y, x : [a||a]\varphi \Longrightarrow y : \varphi}\ \mathrm{L}[*])}{xR_a y, x : [a||a]\varphi \Longrightarrow y : \varphi}\ \text{red}}{x : [a||a]\varphi \Longrightarrow x : [a]\varphi}\ (\mathrm{R}[a])$$

$$\cfrac{\cfrac{\cfrac{\cfrac{xR_a y, yR^{<}_a z, xR^{<}_a z, x : [a]^{<}\varphi, z : \varphi \Longrightarrow z : \varphi}{xR_a y, yR^{<}_a z, xR^{<}_a z, x : [a]^{<}\varphi \Longrightarrow z : \varphi}\ (\mathrm{L}[a])}{xR_a y, yR^{<}_a z, x : [a]^{<}\varphi \Longrightarrow z : \varphi}\ (\text{intr}_2)}{xR_a y, x : [a]^{<}\varphi \Longrightarrow y : [a]^{<}\varphi}\ (\mathrm{R}[a]^{<})}{x : [a]^{<}\varphi \Longrightarrow x : [a][a]^{<}\varphi}\ (\mathrm{R}[a]^{<})$$

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

$$\dfrac{\dfrac{xR_a^{<}y,\ldots,x:[<_a]\psi,y:\psi\ldots\Longrightarrow x:\varphi,y:\psi}{xR_a^{<}y,\ldots,x:[<_a]\psi,\ldots\Longrightarrow x:\varphi,y:\psi}\,(\mathrm{L}[<_a])\qquad \dfrac{\dfrac{\dfrac{yR_ax,xR_ay,\ldots,x:\varphi\Longrightarrow x:\varphi}{yR_ax,xR_ay,\ldots,y:[R_a]\varphi\Longrightarrow x:\varphi}\,(\mathrm{L}[R_a])}{yR_ax,xR_ay,\ldots,y:\neg\psi\to[R_a]\varphi,y:\neg\psi\Longrightarrow x:\varphi}\,(\mathrm{PL})}{yR_ax,xR_ay,\ldots,y:\neg\psi\to[R_a]\varphi\Longrightarrow x:\varphi,y:\psi}\,(\mathrm{PL})}{yR_ax,xR_ay,\ldots,x:[R_a](\neg\psi\to[R_a]\varphi)\Longrightarrow x:\varphi,y:\psi}\,(\mathrm{L}[R_a])$$

$$\dfrac{xR_ay,x:[R_a^{<}]\psi,x:[R_a](\neg\psi\to[R_a]\varphi)\Longrightarrow x:\varphi,y:\psi}{x:[R_a^{<}]\psi,x:[R_a](\neg\psi\to[R_a]\varphi)\Longrightarrow x:\varphi,x:[R_a]\psi}\,(\mathrm{R}[R_a])$$

□

THEOREM 3. *In the system above, Weakening is height-preserving admissible. All the rules are furthermore height-preserving invertible, and Contration is height-preserving admissible. Finally, Cut is admissible.*

Proof. All four results follow directly from results in [20, 21]. They show that any axiom of the form $A_1 \wedge \cdots \wedge A_n \to B_1 \vee \cdots \vee B_m$ can be translated using the following *regular rule scheme*:

$$\dfrac{B_1, A_1, \ldots A_n, \Gamma \Longrightarrow \Delta \qquad \ldots \qquad B_m, A_1, \ldots A_n, \Gamma \Longrightarrow \Delta}{A_1, \ldots A_n, \Gamma \Longrightarrow \Delta}$$

The four results in the statement of the theorem hold for labeled sequent calculi extended with sets of such regular rule schemes. The axioms in Section 4 have the required syntactic form, and the rules above are their translation in regular schemes. So the results follow here are well. □

§6. Moorean phenomena and a resolution operator. Something can be distributed knowledge among agents of a group while being unknowable by some of its members [1]. Consider again our initial example of a lost hiker. Assume now that a third agent c, say, the search helicopter pilot, knows that φ, where φ is "the position of the hiker is x but a does not know it." In this situation, φ *is* distributed knowledge among a, b and c. Yet, a can never know φ. Even if c announced φ to a and b via her radio, this would not result in a situation where a knows that a does not know that the position of the hiker is x. After receiving the information, a would instead know that the hiker's position is x, and φ would just have become false.

This example illustrates why distributed knowledge cannot be equated with what agents in a group would come to know if they would share everything they know with each other. Ågotnes et Wáng have used these types of Moorean phenomena to argue that the standard definition of distributed knowledge as the "potential knowledge" of a group is not entirely accurate. Indeed, distributed knowledge corresponds to the result of pooling the agents' information, but not necessarily (and sometimes necessarily not) by the agents' themselves. This is not equivalent to the knowledge members would hold *themselves* after sharing all their knowledge. To capture the latter, Ågotnes and Wáng [1] introduces a "resolution" operator.

The Moorean phenomena that Ågotnes and Wáng have observed also appear with priority merge. Indeed, while $[a](p \wedge \neg[a]p)$ is not consistent, the following hold:

$$\not\models \neg[a/b](p \wedge \neg[a]p),$$
$$\not\models \neg[a/b](p \wedge \neg[b]p).$$

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

And similarly, of course, for, respectively, $B^2_a(p \wedge \neg B^2_a p)$ and both $B^2_{a/b}(p \wedge \neg B^2_a p)$ and $B^2_{a/b}(p \wedge \neg B^2_b p)$. For an example, consider world w of model $\mathcal{M}_1$ in Figure 4: $w \vDash [a/b](p \wedge \neg[a]p)$. So, again, to capture the idea of what would be believed after the merge by the very agents whose beliefs are being merged, we need to move from our static merged-belief operators to dynamic resolution ones. The resulting definition will cover both priority-merge and intersection modalities, thus generalizing the resolution operator of [1].

In order to define resolution in full generality over all the terms of our preferential algebra (Definition 1.5), we need some additional notation:

DEFINITION 6.11 (Substitutions). *Let $t \in T$ be a term. We write $at(t)$ for the set of subterms of t that are elements of our basic set of agents $\mathcal{A}$. For any terms t, t', we write $\sigma(at(t') \to t)$ to denote the result of substituting, in t', all occurrences of agents $a \in t$ with t.*

The general resolution operator is then defined in analogy with the resolution operator for distributed knowledge in [1]. The operator $[t!]$ reads "after all agents $a \in t$ publicly announce that they have merged their plausibility orderings according to the operation t." The corresponding model transformation boils down to replacing the plausibility orderings of all agents $a \in t$ by R_t:

DEFINITION 6.12. *Given a model $\mathcal{M} = \langle W, (R_i)_{i\in\mathcal{A}}, V\rangle$, and $t \in T$, the updated model $\mathcal{M}_{t!}$ is defined as $\mathcal{M}_{t!} = \langle W, (R^{t!}_i)_{i\in\mathcal{A}}, V\rangle$, where*

$$R^{t!}_i = \begin{cases} R_t, & \text{if } i \in t, \\ R_i, & \text{otherwise.} \end{cases}$$

The semantics of the dynamic resolution operator is then defined in the standard way:

$$M, w \vDash [t!]\varphi \quad \text{iff} \quad M_{t!}, w \vDash \varphi.$$

We obtain the following reduction axioms:

$$\begin{aligned} [t!]p &\leftrightarrow p, \\ [t!]\varphi \wedge \psi &\leftrightarrow [t!]\varphi \wedge [t!]\psi, \\ [t!]\neg\varphi &\leftrightarrow \neg[t!]\varphi, \\ [t!][t']\varphi &\leftrightarrow [\sigma(at(t') \to t)][t!]\varphi. \end{aligned}$$

As usual, soundness of these rules together with their addition to our axiomatic system give us soundness and completeness for $\mathcal{L}$ extended with dynamic resolution modalities. C.f. for instance [31] for details.

§7. Conclusion. Let us summarize the contributions of this paper. We have introduced a new (non-hybrid) modal logic for group beliefs obtained via lexicographic merge of plausibility orderings. This has allowed us to discuss the main properties of the priority-merged belief modalities, for different notions of belief, in comparison with the better-known intersection operators for distributed knowledge. We have then given a sound and complete axiomatization of the logic as well as a labeled sequents system admitting cut. Finally, we have extended the logic with a dynamic resolution operator

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

to capture precisely what the agents themselves would know, were they to share all their information with each other and aggregate it using priority merge, taking into account Moorean phenomena. All in all, the paper provides a general toolbox for logicians who are interested in group attitudes obtained via lexicographic aggregation.

Note also that our system is more general than it might seem at first sight. Recall that, for simplicity, we have explicitly mentioned only a subclass of the priority graphs as defined in [2], and have not made use of the possibility of repeating agents within one priority graph. As a result, we have left "undercover" some of the expressivity of priority graphs. Indeed, recall that Andréka et al. [2] have shown that any priority merge operation, using their more general notion of priority graphs, is definable using combinations of what we called earlier basic priority merge, and basic intersection. The language we have studied here contains modalities for every relational term definable using these two notions, and so covers aggregation with general priority graphs as well.

On the other hand, we have left aside several questions, which we plan to address in further research. In particular, we have made some assumptions which we could in principle relax. First of all, we have assumed that the individual orderings of alternatives are transitive, since this is usually taken for granted in their doxastic plausibility interpretation. However, transitivity is not always a desideratum, especially when one adopts the preference interpretation of the orderings. By contrast, unlike in most work using preferences and plausibility orderings, we have *not* assumed that these are connected. Many of the results obtained in Section 3 rely on that fact. Therefore, imposing a restriction to connected preorders might result in a quite different picture of the logical relationships between the individual modalities and the different pooled modalities. We leave this comparison for further work.

Second, we have chosen to consider group belief as belief based on the merged plausibility orderings, where the belief operator is defined in the same way for groups and for individuals. This is the most obvious way to go, but it is of course not the only way. One could have various reasons to wish for group attitudes to be more, or less, permissive than ours. We have ourselves started this paper by mentioning that intersection seems too restrictive for group belief in the case where individual opinions diverge, and we have used this remark as a justification of the use of lexicographic merge instead. In the same spirit, one could now counter-argue that group belief, as we have defined it, is too permissive, and/or that it gives too much power to some agents. Indeed, in our sense, whatever the topmost agent believes, the group believes too. One could dislike this result and therefore favor a different notion of group belief, with additional criteria in mind, such as fairness or accuracy.

For instance, one could define a less permissive notion by requiring that, for a group to jointly accept φ, the group believes φ (in our sense) *and* would still believe it despite some reshuffling of the priority ranking: for every possible such ranking, or for a majority of them, or for some specific variation of the current ranking.[10] Concretely, one could claim, for instance, that the group should accept φ if and only if it believes φ (in our sense) with the current ranking *and* when switching the two topmost agents in the ranking. More generally, this suggest exploring the space of aggregation rules that seem plausible without falling prey to the classical, Arrowian impossibilities.[11]

[10] We thank Frederik Van De Putte for suggesting some of these interesting variations.

[11] We thank one anonymous reviewer for pointing out this implication of the suggestion above.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

This discussion only hints toward one possible direction for further research, and there is of course much more to discuss as to whether priority merge is a good candidate for the notion of group belief, or even a useful building block of such a notion. For now, we consider that the purpose of the present paper lies somewhere else, beyond those considerations. Until further research, our aspirations are restricted to offering a formal stepping stone for defining more complex operators and for comparing them in terms of any definable additional desiderata.

§8. Appendix 1: Axiomatization and proof theory for the intersection of strict relations. In this Appendix we show that the second understanding of the strict version of the intersection modalities is also completely axiomatizable and that it also has a well-behaved proof theory. As alluded to in Section 1, the definition we are working on here is

$$wR^{<}_{t||t'}v \quad \Leftrightarrow \quad wR^{<}_{t}v \text{ and } wR^{<}_{t'}v.$$

Observe first that by decomposing the terms t and t' themselves we arrive at the following equivalent definition:

$$wR^{<}_{t||t'}v \quad \Leftrightarrow \quad wR^{<}_{a}v \text{ for all } a \in t||t'.$$

This highlights the fact that the present definition entails the one we have used throughout this paper. As also alluded to earlier, the strict modality based on this definition validates the same, stronger version of the K-axiom as we used for the weak one:

$$([t]^{<}\varphi \wedge [t']^{<}(\varphi \to \psi)) \to [t||t']^{<}\psi. \tag{dis$^{<}$}$$

This is in fact the only change in the axiomatization, but this is a crucial one as it indeed allows to derive:

$$[t]^{<}\varphi \to [t||t']^{<}\varphi. \tag{incl–$[t||t']^{<}$}$$

THEOREM 4. *The set of axioms and rules listed in Section 4, but with* (dis) *replaced with* (dis$_{<}$) *is sound and complete with respect to the class of plausibility frames with the new definition of the strict intersection relation.*

Proof. The proof uses the same construction as the proof of Theorem 2. The difference only shows in the definition of the relations for the complex terms in the bulldozed model. This time we set:

$$xR^{<'}_{t||t'}y \quad \text{iff} \quad \beta(x)R^{<}_{t||t}\beta(y) \text{ and for all } a \text{ occurring in } t||t', xR^{<'}_{a}y.$$

Like before, the existence lemma still holds:

LEMMA 2. *For all $x \in W'$, if $\langle t||t'\rangle^{<}\varphi \in \beta(x)$ then there is a y such that $xR^{<}_{t||t'}y$ and $\varphi \in \beta(y)$.* □

Proof. If x is in W^{-} then this follows directly from the existence lemma for the canonical model. If not, i.e., if $x = (\Gamma, m)$ for some Γ and n, then the same existence lemma gives us a Δ such that $\beta(x)R^{<}_{t||t'}\Delta$ and $\varphi \in \Delta$. By (incl–$[t||t']_{<}$) we know that $\beta(x)R^{<}_{a}\Delta$ for all a. Take $y = (\Delta, n)$ with $m < n$. For any agent a for which Δ is either

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

not part of a cluster or part of a different clusters than Γ, we get by definition that $xR_a^{<}y$. For all others the latter follows from the fact that we chose $m > n$. □

The rest of the proof (i.e., adequacy for the new $R_t^{<'}$, modal equivalence, and splitting) works in essentially the same way as before.

For the proof theory, we only have to replace the rule (incl-int $<$) with the weaker:

$$\frac{xR_a^{<}y, xR_{t||t'}^{<}y, \Gamma \Longrightarrow \Delta \text{ for some } a \in t||t'}{xR_{t||t'}^{<}y, \Gamma \Longrightarrow \Delta} \text{ (incl'-int}<\text{)}$$

THEOREM 5. *The system of Section 5, with (incl-int $<$) replaced by (incl'-int $<$), is deductively equivalent to the axiom system just presented. Furthermore, in this proof system weakening and contraction are height-preserving admissible. All the rules are furthermore height-preserving invertible, and cut is admissible.*

§9. Appendix 2: Details of proof of Theorem 3.

LEMMA 3. *$x : A, \Gamma \Longrightarrow \Delta, x : A$ is provable for arbitrary A.*

Proof. By induction on the complexity of A. The basic cases, both for atomic propositions and relational atoms, hold by definition. The induction step for the propositional connectives is trivial. The modal cases follow [20]. Here is one example:

$$\frac{\dfrac{y : \varphi, x : [t]\varphi, xR_t y, \Gamma \Longrightarrow \Delta, y : \varphi}{x : [t]\varphi, xR_t y, \Gamma \Longrightarrow \Delta, y : \varphi}\,(\mathrm{L}[t])}{x : [t]\varphi, \Gamma \Longrightarrow \Delta, x : [t]\varphi}\,(\mathrm{R}[t])$$

□

For the next lemma, the Substitution Lemma, we need the following notation, again taken from [20]: Let $xRy(z/x)$ the substitution of all instances of x in xRy, and similarly for $x : A(z/x)$. If Γ is a multi-set of formulas then $\Gamma(z/x)$ is the applications of the substitution of all elements of Γ.

LEMMA 4. *If $\Gamma \Longrightarrow \Delta$ is derivable then $\Gamma(z/x) \Longrightarrow \Delta(z/x)$ is derivable in a derivation of the same length.*

Proof. Again by induction on the length of proof. The basic case is trivial, since of $\Gamma \Longrightarrow \Delta$ is an axiom, then so is $\Gamma(z/x) \Longrightarrow \Delta(z/x)$. The induction step is also completely trivial given our inductive hypothesis for all the rules that are specific to this paper, since none of them have a variable condition. The only subtlety are for applications of the two R rules, but these are handled exactly as Negri [20] does: we use our inductive hypothesis twice, once to replace the variable of interest, and one for the other variable that gets introduced in the relational atom, in order to avoid potential clashes. □

LEMMA 5. *Weakening is height-preserving admissible.*

Proof. There are six cases to check: labeled formulas, strict and weak relational atoms, on each side. We show two of them for weakening on the right. The left side proceeds similarly. Again this argument is the same as in [20].

$$\frac{\Gamma \Longrightarrow \Delta}{\Gamma \Longrightarrow \Delta, xR_t y} \qquad \frac{\Gamma \Longrightarrow \Delta}{\Gamma \Longrightarrow \Delta, x : \varphi}.$$

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

For both cases the argument is similar. If $\Gamma \Longrightarrow \Delta$ has been obtained by any rule except $R[t]$ or its strict counterpart, weakening can be applied to the premise of that rule, and the rule itself applied afterward. For instance:

$$\dfrac{\dfrac{xR_t y, xR_{t||t'} y, \Gamma \Longrightarrow \Delta}{xR_{t||t'} y, \Gamma \Longrightarrow \Delta}\,(\text{incl'-int})}{xR_{t||t'} y, \Gamma \Longrightarrow \Delta, x : \varphi} \quad \rightsquigarrow \quad \dfrac{\dfrac{xR_t y, xR_{t||t'} y, \Gamma \Longrightarrow \Delta}{xR_t y, xR_{t||t'} y, \Gamma \Longrightarrow \Delta, x : \varphi}\,(\text{incl'-int})}{xR_{t||t'} y, \Gamma \Longrightarrow \Delta, x : \varphi}$$

In case the premise has been obtained from $R[t]$ or its strict counterpart, the variable x should be first substituted for a fresh variable, in order to avoid clashes. We know by the previous Lemma that this can be done in an height-preserving way. The argument then proceeds as above. □

LEMMA 6. *All rules are height-preserving invertible.*

Proof. We have to show that if the conclusion of a rule is derivable in n steps, so is the premise of that rule. As usual, the only relevant cases are where the formulas or relational atoms in the conclusion of the rule are not principal at the last step. The argument for that is generic, and rest on the fact that all the rules specific to our system allow for arbitrary contexts Γ and Δ to the left and to the right. For illustration we present the case of the rule (incl-int $<$). From a derivation of $xR^{<}_{t||t'} y, \Gamma' \Longrightarrow \Delta'$, a simple application of weakening, which we just shown to be height-preserving admissible, gives us $xR^{<}_{a} y, xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta$. This case also illustrates the usefulness of having all our new rules in the regular rule scheme identified by Negri Von Plato [21].

$$\dfrac{xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta}{xR^{<}_{t||t'} y, \Gamma' \Longrightarrow \Delta'}\,\text{Some rule } \mathcal{R} \quad \rightsquigarrow \quad \dfrac{xR^{<}_{a} y, xR^{<}_{t||t'} y, \Gamma \Longrightarrow \Delta}{xR^{<}_{a} y, xR^{<}_{t||t'} y, \Gamma' \Longrightarrow \Delta'}\,\mathcal{R}\cdot$$ □

LEMMA 7. *Contraction is height-preserving admissible.*

Proof. The argument is routine in the case where the contraction formula is not principal. For the other cases it is also straightforward, this time since none of the rules specific to this system requires the use of copies of labelled formulas or relational atoms. Here is an illustration, involving contraction on one side of $(intr_3)$:

$$\dfrac{\dfrac{xR_t y, xR_t y, yR_t x, \Gamma \Longrightarrow \Delta \qquad xR_t y, xR^{<}_t y, \Gamma \Longrightarrow \Delta}{xR_t y, xR_t y, \Gamma \Longrightarrow \Delta}\,(\text{intr}_3)}{xR_t y, \Gamma \Longrightarrow \Delta}$$

This transforms to the following.

$$\dfrac{\dfrac{xR_t y, xR_t y, yR_t x, \Gamma \Longrightarrow \Delta}{xR_t y, yR_t x, \Gamma \Longrightarrow \Delta} \qquad xR_t y, xR^{<}_t y, \Gamma \Longrightarrow \Delta}{xR_t y, \Gamma \Longrightarrow \Delta}\,(\text{intr}_3)$$ □

THEOREM 6. *Cut is admissible.*

Proof. We show this by induction on the length of derivations of the premises of cut plus main induction on the complexity of the cut-formula. The base case is again routine. For the induction step, observe that, with the exception of both R-rules, our rules only operate on relational atoms on the left. The former can be handled as in [20]. Otherwise this restricts the cases to be checked to those where the cut relational atom on the right is non-principal at the last step. Here is one illustration. All other

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

cases work similarly.

$$\dfrac{\dfrac{\Gamma'' \Longrightarrow \Delta'', xR^{<}_{t||t'}y}{\Gamma' \Longrightarrow \Delta', xR^{<}_{t||t'}y} \quad \dfrac{xR^{<}_{a_1}y, xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta \quad \ldots \quad xR^{<}_{a_n}y, xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta}{xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta}\ \text{incl-int}<}{\Gamma, \Gamma' \Longrightarrow \Delta, \Delta'}\ \text{Cut}$$

This transforms to the following.

$$\dfrac{\Gamma'' \Longrightarrow \Delta'', xR^{<}_{t||t'}y \quad \dfrac{xR^{<}_{a_1}y, xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta \quad \ldots \quad xR^{<}_{a_n}y, xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta}{xR^{<}_{t||t'}y, \Gamma \Longrightarrow \Delta}\ \text{incl-int}<}{\dfrac{\Gamma, \Gamma'' \Longrightarrow \Delta, \Delta''}{\Gamma, \Gamma' \Longrightarrow \Delta, \Delta'}}\ \text{Cut}$$

□

Acknowledgments. The authors would like to thank Marta Bilkova, Dominik Klein, Frederik Van De Putte, Ondrej Majer, and Igor Sedlar for helpful comments and suggestions. Fiancial support of the Deutsche Forschungsgemeinschaft (DFG) and the Grantová agentura České republiky (GAČR), joint project RO 4548/6–1 "From Shared Evidence to Group Attitudes," as well as, for Gratzl, research project no. 390218268 of the DFG, is gratefully acknowledged.

BIBLIOGRAPHY

[1] Ågotnes, T. & Wáng, Y. (2017). Resolving distributed knowledge. *Artificial Intelligence*, **252**,1–21.

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

[3] Baltag, A., Bezhanishvili, N., Özgün, A., & Smets, S. (2013). The topology of belief, belief revision and defeasible knowledge. In *International Workshop on Logic, Rationality and Interaction*. Berlin: Springer, pp. 27–40.

[4] Baltag, A. & Renne, B. (2016). Dynamic epistemic logic. In Zalta, E. N., editor. *The Stanford Encyclopedia of Philosophy*, Winter 2016 edition. Stanford, CA: Metaphysics Research Lab, Stanford University.

[5] Baltag, A. & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors. *Logic and the Foundations of Game and Decision Theory*, Texts in Logic and Games, Vol. 3. Amsterdam: Amsterdam University Press, pp. 9–58.

[6] ———. (2013). Protocols for belief merge: Reaching agreement via communication. *Logic Journal of IGPL*, **21**(3), 468–487.

[7] Benthem, J., Girard, P., & Roy, O. (2009). Everything else being equal: A modal logic for ceteris paribus preferences. *Journal of Philosophical Logic*, **38**(1), 83–125.

[8] Benthem, J., Grossi, D., & Liu, F. (2014). Priority structures in deontic logic. *Theoria*, **80**(2), 116–152.

[9] Benthem, J., van Otterloo, S., & Roy, O. (2005). Preference logic, conditionals and solution concepts in games. In Lagerlund, L. and Sliwinski, editors, *Modality Matters*: *Twenty-Five Essays in Honour of Krister Segerberg*. Uppsala: Department of Philosophy, Uppsala Philosophical Studies.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

[10] Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal logic*. In *Cambridge Tracts in Theoretical Computer Science*, Vol. **53**. United Kingdom: Cambridge University Press.

[11] Boutilier, C. (1994). Conditional logics of normality: A modal approach. *Artificial Intelligence*, **68**(1), 87–154.

[12] Bradley, R. (2017). *Decision Theory With a Human Face*. Cambridge: Cambridge University Press.

[13] Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. (2004). *Reasoning About Knowledge*. Cambridge, MA: MIT Press.

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

[15] Gerbrandy, J. (1998). Distributed knowledge. In *Proceedings of the Second Workshop on the Semantics and Pragmatics of Dialogue*. Twente: University of Twente, pp. 111–124.

[16] Girard, P. (2011). Modal logic for lexicographic preference aggregation. In *Games, Norms and Reasons*. Netherlands: Springer, pp. 97–117.

[17] Girard, P. & Seligman, J. (2009). An analytic logic of aggregation. In *Proceedings of the* 3*rd Indian Conference on Logic and Its Applications*. Berlin, Heidelberg: Springer-Verlag, pp. 146–161.

[18] Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, **17**(2), 157–170.

[19] Kazmer, M. M., Lustria, M. L. A., Cortese, J., Burnett, G., Kim, J.-H., Ma, J., & Frost, J. (2014). Distributed knowledge in an online patient support community: Authority and discovery. *Journal of the Association for Information Science and Technology*, **65**(7), 1319–1334.

[20] Negri, S. (2005). Proof analysis in modal logic. *Journal of Philosophical Logic*, **34**(5–6), 507.

[21] Negri, S. & Von Plato, J. (1998). Cut elimination in the presence of axioms. *Bulletin of Symbolic Logic*, **4**(4),418–435.

[22] Negri, S., von Plato, J., & Ranta, A. (2008). *Structural Proof Theory*. Cambridge: Cambridge University Press.

[23] Özgün, A. (2017). Evidence in Epistemic Logic: A Topological Perspective. Ph. D. thesis.

[24] Parikh, R. & Ramanujam, R. (1985). Distributed processes and the logic of knowledge. In *Workshop on Logic of Programs*. Berlin: Springer, pp. 256–268.

[25] Punčochář, V. & Sedlár, I. (2017). Substructural logics for pooling information. In *International Workshop on Logic, Rationality and Interaction*. Berlin: Springer, pp. 407–421.

[26] Roelofsen, F. (2007). Distributed knowledge. *Journal of Applied Non-Classical Logics*, **17**(2), 255–273.

[27] Ryan, M. (1993). Defaults in specifications. In [1993] *Proceedings of the IEEE International Symposium on Requirements Engineering*. IEEE, pp. 142–149.

[28] Segerberg, K. (1971). *An Essay in Classical Modal Logic*. Uppsala: Filosofiska Föreningen Och Filosofiska Institutionen Vid Uppsala Universitet.

[29] Swanson, D. R. (1986). Undiscovered public knowledge. *The Library Quarterly*, **56**(2), 103–118.

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press

[30] van der Hoek, W., van Linder, B., & Meyer, J.-J. (1999). Group knowledge is not always distributed (neither is it always implicit). *Mathematical Social Sciences*, **38**(2), 215–240.

[31] van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic Epistemic Logic*, Vol. 337. Berlin: Springer Science & Business Media.

DEPARTMENT OF ARTIFICIAL INTELLIGENCE
UNIVERSITY OF GRONINGEN
GRONINGEN, NETHERLANDS
E-mail: z.l.christoff@rug.nl

MUNICH CENTER FOR MATHEMATICAL PHILOSOPHY
LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN
MÜNCHEN, GERMANY
E-mail: norbert.gratzl@lrz.uni-muenchen.de

DEPARTMENT OF PHILOSOPHY
UNIVERSITY OF BAYREUTH
BAYREUTH, GERMANY
E-mail: olivier.roy@uni-bayreuth.de

https://doi.org/10.1017/S1755020321000058 Published online by Cambridge University Press