A revised version of this paper appears in: *Journal of Logic and Computation* 9(1): 81–103, 1999.

A modal logic for fusing partial belief of multiple reasoners

Luca Boldrin
Dept. of Pure and Applied Math.
University of Padova
Padova, Italy
boldrin@ladseb.pd.cnr.it

Alessandro Saffiotti*
IRIDIA
Université Libre de Bruxelles
Brussels, Belgium
asaffio@ulb.ac.be

Abstract

We present $\mathbf{PL}_n^{\otimes}$, a multi-agent epistemic logic where each agent can perform uncertain (possibilistic) reasoning. The original feature of this logic is the presence of a *distributed belief* operator, with the purpose of merging the belief of different agents. Unlike the corresponding operator in the categorical (non-uncertain) case, our distributed belief operator *accumulates* support for the same fact coming from different agents. This means that opinions shared by different agents can be combined into a stronger distributed belief. This feature is useful in problems like pooling expert opinions and combining information from multiple unreliable sources. We provide a possible worlds semantics and an axiomatic calculus for our logic, and prove soundness, completeness and decidability results. We hint at some possible applications of $\mathbf{PL}_n^{\otimes}$ in the conclusions.

1 Introduction

Consider a community of agents (experts, knowledge bases, processors,...) each holding some belief[1] about a common domain. Distributed belief, introduced in [13], is the belief implicitly held by the agents as a group; we can see distributed belief as the belief held by a (virtual) agent d that collects and merges the beliefs individually held by all the other agents. In Halpern and Moses' example, if Alice believes that ψ and Bob believes that $\psi \to \phi$, then together they have distributed belief that ϕ. Logic-based approaches have been proposed to model multiple agents and their distributed belief (see [9]). These approaches typically assume that each agent's belief is categorical, that is, propositions are either completely believed or completely disbelieved by each agent. So, agents cannot cope with situations in which belief is —as is the case in most real-world domains— a matter of degrees.

*Current address: Dept. of Technology and Science, Örebro University, S-70181 Örebro, Sweden. Email: alessandro.saffiotti@aass.oru.se

[1]All through this paper, we talk about "belief", as opposed to "knowledge", to emphasize that there is no requirement of truth of what is believed.

Several logics have been proposed in the literature that are able to capture partial belief. Most of these logics [22, 7, 11] can be classified as "plausibility logics", since their semantics is based on a unique plausibility measure (see [14, 12, 18] for more about plausibility logics). Other logics [8, 17, 10] are true epistemic logics for uncertainty: Liau and Lin, for example, propose a modal logic where the usual accessibility relation is replaced by a fuzzy one; and Fagin and Halpern's logic can account for multiple (probabilistic) believers. However, we are not aware of any such logic that can perform fusion of multiple agents' beliefs into one distributed belief.

This situation is somehow unfortunate, in that the ability to model intermediate degrees of belief is particularly attractive in a distributed setting. There we can really *fuse* information, since we can use degrees of confidence to weight the opinions of different agents when combining their belief into the distributed belief. So, if the agents agree on some proposition, their individual degrees of belief about it can be accumulated into a stronger degree of distributed belief. One interesting application is the pooling of information coming from several unreliable sources: each source is viewed as an agent, and the degrees of belief attached to the items of information it provides are interpreted as degrees of reliability. Distributed belief represents the result of merging the information coming from all the different sources.

In this paper, we present a modal logic for multiple agents and distributed belief where agents are allowed to hold partial belief. We represent partial belief by using modal operators of the form $\mathbf{B}^i_a$, where i is the index of an agent and a is a number in the $[0,1]$ interval. Given the formula $\mathbf{B}^i_a\phi$, its intended meaning is "Agent i believes, with strength (at least) a, that ϕ is true." We also have a family of modal operators $\mathbf{D}_a, a \in [0,1]$, for representing distributed belief: the intended meaning of $\mathbf{D}_a\phi$ is "The information distributed among the agents, if collected together, would allow to believe with strength (at least) a that ϕ is true." Notice that the index in our modal operators are intended to give a *lower bound* on the degree of belief of the agent.

To see how this works, consider three agents, imaginatively named 1, 2 and 3, and the following facts: (a) agent 1 believes with strength 0.4 that it rains; (b) agent 2 completely believes (*i.e.*, with strength 1) that it is windy; and (c) agent 3 believes with strength 0.7 that whenever it is windy it rains. We represent these statements by the following formulas:

$$\begin{array}{ll} (a) & \mathbf{B}^1_{0.4}\,\mathrm{rain} \\ (b) & \mathbf{B}^2_1\,\mathrm{wind} \\ (c) & \mathbf{B}^3_{0.7}(\mathrm{wind} \rightarrow \mathrm{rain}). \end{array}$$

As we shall show, our logic would deduce from these formulas (for a particular choice of the fusion operator),

$$\mathbf{D}_{0.82}\,\mathrm{rain},$$

that is, by merging the belief of all the agents we can form a distributed belief of strength 0.82 that it rains. Note that this belief is stronger than any belief about rain individually held by any of the agents. In fact, it is even stronger than what you would get if a single agent, say 1, believed all the three facts — in which case, we could only deduce $B^1_{0.7}$rain, *i.e.*, the maximum between (a) and the result of applying *modus ponens* to (b) and (c).

Many choices can be made regarding the numerical behavior of degrees of belief, both regarding how each agent computes the degree of belief in a consequence given the degrees of belief in its premises, and regarding how the degree of distributed belief is computed from the individual degrees of belief. The field of reasoning under uncertainty in AI has seen the use of several mathematical theories for representing partial belief, including subjective probability, Dempster-Shafer's theory of belief functions, and possibility theory. In this paper, we use possibility theory as the basic calculus of degrees.

Two main reasons underlie our choice. First, possibility theory offers a wide range of operators to perform fusion of information: different operators can be chosen depending on the characteristics of the agents and of the information to be combined [6]. In particular, different choices of the fusion operator correspond to assuming different dependence relations among the agents. While we shall discuss this issue in Section 2 below, we emphasize that the logic that we present has the fusion operator as a parameter, and can thus be tailored to the specific application at hand. The second reason for our choice is technical: possibility theory can be integrated into a logic much more easily than other uncertainty calculi (see [7] for an illustration). The study of similar logics based on other calculi is left for future work.

The rest of this paper is organized as follows. In the next section, we recall a few notions about possibility theory. In Sect. 3, we introduce the language, semantics and proof theory of our logic. Soundness, completeness and decidability are proved in Sect. 4. Finally, in Sect. 5, we hint at some possible applications of our logic, place our approach in perspective among the existing literature, and suggest some directions for future work. The proofs of the technical results are collected in an appendix.

2 Possibility theory

Possibility theory [24] is a well known theory for the management of uncertainty, whose aim is the representation of non-statistical uncertainty. We refer the reader to [7] for a comprehensive survey, while we present here just the basic elements of the theory for the sake of completeness. A possibility distribution on a set W is a function $\pi : W \longrightarrow [0,1]$ from the set to the real unitary interval, with the usual normality condition that there exists a $x \in W$ such that $\pi(x) = 1$. In case this condition is not met, we speak of unnormalized possibility distribution. For a given π, we define the associated possibility measure $\Pi : 2^W \longrightarrow [0,1]$ and necessity measure $N : 2^W \longrightarrow [0,1]$ by, respectively:

$$\Pi(X) = \sup_{x \in X} \pi(x)$$

$$N(X) = 1 - \sup_{x \notin X} \pi(x)$$

We shall use the N measure to give semantics to our language: by $\mathbf{B}^i_a \phi$ we will mean that the agent i attributes at least necessity a to the proposition ϕ. As we will see in the following, this is a natural extension of the $\Box$ modality of standard modal logic, and it is also possible to characterize this modality with the requirements that are usually introduced to capture the notions of "knowledge" (S5 axioms) and "belief" (KD45 axioms). This extension, to our knowledge, has been first proposed in [15].

Possibility distributions can be understood as representing uncertain information (or belief) states of an agent, since we can think of W as the set of all possible situations (*worlds*, in a typical modal language), and of $\pi(w)$ as the degree of possibility that we assign to the world w. If we only were to allow π to take values in the extreme points of the interval, *i.e.* in $\{0, 1\}$, each π would represent a (non-empty, under the normality condition) subset of W. The $\subseteq$ relation on subsets of W can be conservatively extended to possibility distributions, by defining: $\pi_1 \leq \pi_2$ iff for any $x \in W$, $\pi_1(x) \leq \pi_2(x)$. If $\pi_1 \leq \pi_2$, the information state associated to π_1 is more informative than the one associated to π_2, since it better constrains the possible elements in W. Note that, if $\pi_1 \leq \pi_2$, then $N_1(X) \geq N_2(X)$ for all $X \subseteq W$.

Going on with the extension of set operations on possibility distributions, let us consider what the corresponding of intersection should be. The first idea is to use the point to point minimum, *i.e.*,

$$(\pi_1 \wedge \pi_2)(w) = \min(\pi_1(w), \pi_2(w)).$$

This choice is not unique, and there are many alternatives (see [6] and [2] for wide analysis). The operators which are most commonly used in the literaturs are *continuous T-norms*, which we generally denote by $\otimes$. A T-norm is any binary operator on $[0, 1]$ that is commutative, associative, non-decreasing in each argument, and has 1 as unit. Given any T-norm $\otimes$, we can define a dual operator $\oplus$, called a T-conorm, by:

$$(\pi_1 \oplus \pi_2)(x) = 1 - ((1 - \pi_1(x)) \otimes (1 - \pi_2(x))).$$

T-conorms have the same properties as T-norms except that they have 0 as unit. For example, the T-conorm associated to $\wedge$ is given by the pointwise maximum:

$$(\pi_1 \vee \pi_2)(w) = \max(\pi_1(w), \pi_2(w)).$$

T-norms and T-conorms are used as a generalization of logical conjunction and disjunction, respectively. The most common among T-norms are, beside the min, the real product:

$$(\pi_1 \times \pi_2)(w) = \pi_1(w) \cdot \pi_2(w),$$

and the Łukasiewicz product:

$$(\pi_1 \times_L \pi_2)(w) = \max(\pi_1(w) + \pi_2(w) - 1, 0).$$

T-norms and T-conorms are the most commonly used operators in the literature on possibility distributions, and we shall base our logic on them (see [23] for a thorough analysis of T-norms and their properties). The adaptation of our framework to include other families of operators (e.g., averaging operators, or OWA operators) is left for future inspection.

T-norm based operations on possibility distributions naturally extend operations on sets. Set intersection can be conceived as a rather rough way to merge information states: if A and B are subsets of W, each of which denotes the set of possible worlds allowed by an agent, we define the merging of the states to be the set of possibilities allowed by both agents; similarly, T-norms perform a merge of two uncertain information states. The fact that there are more choices available for merging possibility distributions corresponds to the unknown dependence relation among the agents. For example, if we can assume independence between agents, then we can use a *reinforcing* operator (e.g., product) that accumulates the support given to the same hypothesis by different agents. If, however, we have no specific information about this relation (e.g., it may even be that two separate items of information originate from the same source) we should perform a cautious fusion, and use an *idempotent* operator, *i.e.*, one such that $\pi_1 \otimes \pi_1 = \pi_1$. This guarantees that the amount of support to an hypothesis is independent on the number of agents supporting it. (Note that this only holds for the minimum T-norm, $\wedge$.)

In general, the "strength" of a combinator depends on the informativeness of the resulting information states; for example we have that

$$\pi_1 \times_L \pi_2 \;\leq\; \pi_1 \times \pi_2 \;\leq\; \pi_1 \wedge \pi_2, \tag{1}$$

which says that the information state resulting from Łukasiewicz combination is stronger than the one obtained by the product combinator which is in turn stronger than the $\wedge$ combinator. (In fact it can be proved that $\times_L$ is the strongest and $\wedge$ is the weakest of all the continuous T-norms).

In our logic, we will allow the use of an arbitrary T-norm for the semantical modeling of the distributed knowledge operator. This means that the virtual agent d can combine the knowledge of the various agents under different assumptions about their mutual dependence, depending on the application domain.

3 The logic $\mathbf{PL}_n^{\otimes}$

We now define in some detail the logic $\mathbf{PL}_n^{\otimes}$. $\mathbf{PL}_n^{\otimes}$ has two parameters: the number n of agents, and the operator $\otimes$ used to combine information into distributed belief. We will comment briefly on the choice of an adequate combination operator in the conclusions. All the results below hold for any $n \geq 2$ (the case when $n = 1$ is trivial) and for any continuous T-norm $\otimes$.

3.1 Language

$\mathbf{PL}_n^{\otimes}$ is defined on the modal propositional language $\mathcal{L}$ built as usual from a countable set of atomic propositional symbols and the logical connectives $\neg$, $\wedge$, $\vee$ and $\rightarrow$, and closed over the modal operators $\mathbf{D}_a$ and $\mathbf{B}_a^i$, with a a rational number in the unit interval $[0, 1]$, and $i \in 1, \ldots, n$. The intended reading of the modal operators $\mathbf{B}_a^i$ and $\mathbf{D}_a$ is the one discussed in the Introduction. The reason for restricting a to be rational is to keep the language countable (more on this at the end of Section 4).

3.2 Semantics

To precisely specify the meaning of the formulas of our language, we give $\mathbf{PL}_n^{\otimes}$ semantics in a Kripke style. First, we define the semantic structures on which we interpret formulas of $\mathcal{L}$.

Definition 1 A $\Pi_n^{\otimes}$*-frame* for $\mathcal{L}$ is a tuple $F = \langle W, \pi_0, \pi_1, \ldots, \pi_n \rangle$ such that: W is a finite non-empty set (of "worlds"); each π_i maps each world w to an unnormalized possibility distribution $\pi_{i,w}$ over W; and, for any $w \in W$,

$$\pi_{0,w} \leq \bigotimes_{i=1}^{n} \pi_{i,w}. \tag{2}$$

A $\Pi_n^{\otimes}$*-structure* for $\mathcal{L}$ is a pair $M = \langle F, V \rangle$, where $F = \langle W, \pi_0, \pi_1, \ldots, \pi_n \rangle$ is a $\Pi_n^{\otimes}$-frame and V maps propositional variables in $\mathcal{L}_0$ to subsets of W.

Intuitively, any world w in a $\Pi_n^{\otimes}$-structure contains two sorts of information: (a) information about which facts are true at that world, embodied in the V truth assignment; and (b) information about which facts each agent i considers possible at that world, embodied in the $\pi_{i,w}$ possibility distribution — the value of $\pi_{i,w}(w')$ measures the extent to which agent i, in w, considers possible the state of affairs represented by world w'.[2] The π_0 function has a special status. At any world w, the $\pi_{0,w}$ possibility distribution represents the state of belief of a virtual agent d, represented by the $\mathbf{D}_a$ operators, that combines the belief of all the other agents. The constraint (2) accounts for this intended interpretation of π_0. Intuitively, agent d considers possible only those worlds that are possible for all the other agents; and the degree of this possibility is computed by the combination operator $\otimes$. Note that the constraint (2) corresponds to the one normally used in categorical logics for distributed belief (e.g., [9]), namely, that the accessibility relation associated to distributed belief be a subset of the intersection of all the accessibility relations associated to the agents.

We now make precise what it means for a formula to be true at a world w of a given $\Pi_n^{\otimes}$-structure M by defining the satisfaction relation $\models$ between pairs M, w and formulas of $\mathcal{L}$.

Definition 2 The satisfaction relation $\models$ for $\Pi_n^{\otimes}$-structures is defined as follows:

$$\begin{array}{lcl}
M,w \models p & \iff & w \in V(p) \\
M,w \models \neg\phi & \iff & \text{it is not the case that } M,w \models \phi \\
M,w \models \phi \wedge \psi & \iff & M,w \models \phi \text{ and } M,w \models \psi \\
M,w \models \mathbf{B}_a^i \phi & \iff & N_{i,w}([\![\phi]\!]_M) \geq a \\
M,w \models \mathbf{D}_a \phi & \iff & N_{0,w}([\![\phi]\!]_M) \geq a
\end{array}$$

where $[\![\phi]\!]_M =_{\text{def}} \{w' \in W \mid M,w' \models \phi\}$ (note that the definition is recursive, but founded), and $N_{i,w}$ is the necessity measure associated to $\pi_{i,w}, i = 0, 1, \ldots, n$. We say that a formula ϕ is valid in the $\Pi_n^{\otimes}$-structure M iff for each world w of M, $M, w \models \phi$. We say that ϕ is $\Pi_n^{\otimes}$-valid iff ϕ is valid in any $\Pi_n^{\otimes}$-structure.

The only peculiar clauses in the above definition are the ones for the modal operators. Intuitively, $\mathbf{B}_a^i \phi$ is true at w if, at that world, agent i regards the set of worlds where ϕ holds as being necessary at the level (at least) a. And $\mathbf{D}_a \phi$ is true at w if the same holds for the pseudo-agent d. The semantics for $\mathbf{B}_a^i$ and $\mathbf{D}_a$ account for the intended reading of the index a as a lower bound on the degree of belief; in particular, note that $\mathbf{B}_a^i \phi$ entails $\mathbf{B}_b^i \phi$ for any $b \leq a$ (similarly for $\mathbf{D}_a$).

[2] As it will become clear in Section 3.4, we could equivalently write our models by using multivalued accessibility relations in place of the π_i distributions (see [17, 10]).

To see how the $\mathbf{D}_a$ operators capture the intended merging of information, note that, by the properties of possibility distributions, the constraint (2) implies that, for any $X \subseteq W$,

$$N_{0,w}(X) \geq \bigoplus_{i=1}^{n} N_{i,w}(X).$$

Thus, the degree of distributed belief in ϕ, measured by $N_{0,w}([\![\phi]\!]_M)$, is (at least) the one obtained by combining, through the $\oplus$ T-conorm, the degrees of belief in ϕ for each individual agent i, measured by $N_{i,w}([\![\phi]\!]_M)$. Note that distributed belief can be strictly richer than the combination of all the agents' belief; in a sense, we allow the combining agent d to have some belief of its own, which do not originate in the other agents' belief. A similar phenomenon can be observed in categorical accounts of distributed belief.[3]

3.3 Proof theory

The semantic characterization gives us some intuitions about the intended meaning of our logic. In order to be able to *use* the logic to perform inferences, we need to find an inference system for it. In this section, we show that $\mathbf{PL}_n^{\otimes}$ can be completely characterized by a Hilbert-style axiomatization. This axiomatization will also give us further insights into the behavior of $\mathbf{PL}_n^{\otimes}$.

Definition 3 The logic $\mathbf{PL}_n^{\otimes}$ is defined by the following axiom schemas and inference rules, where ϕ and ψ range over formulas of $\mathcal{L}$; a, b and c over rational numbers in $[0, 1]$; and $i = 1, \ldots, n$.

A0	All propositional tautologies	
A1	$\mathbf{B}_0^i \bot$	
A2	$\mathbf{B}_a^i(\phi \to \psi) \to (\mathbf{B}_b^i \phi \to \mathbf{B}_c^i \psi)$	$c \leq \min\{a, b\}$
A3	$\mathbf{D}_a(\phi \to \psi) \to (\mathbf{D}_b \phi \to \mathbf{D}_c \psi)$	$c \leq \min\{a, b\}$
A4	$\left(\bigwedge_{i=1}^{n} \mathbf{B}_{a_i}^i \phi\right) \to \mathbf{D}_c \phi$	$c = \bigoplus_{i=1}^{n} a_i$

MP	from ϕ and $\phi \to \psi$ deduce ψ
NEC	from ϕ deduce $\mathbf{B}_1^i \phi$

We say that a formula ϕ is a theorem of $\mathbf{PL}_n^{\otimes}$, written $\vdash \phi$, if ϕ is obtained from A0–A5 by a finite number of applications of MP, NEC and uniform substitutions. If Γ is a finite subset of $\mathcal{L}$, we write $\Gamma \vdash \phi$ to mean $\vdash (\bigwedge_{\psi \in \Gamma} \psi) \to \phi$. Finally, we say that Γ is consistent if $\Gamma \not\vdash \bot$.

A0 and MP say that $\mathbf{PL}_n^{\otimes}$ is an extension of classical propositional calculus. The task of A2 and NEC is to bring propositional calculus inside the scope of the $\mathbf{B}_a^i$ operators. This means that individual agents can perform propositional deductions inside their belief sets; moreover, degrees of belief propagate across these deductions by using the min rule, as sanctioned by possibility theory. Finally, A1 attributes to any proposition a lower bound of belief equal to 0 (this is because, for all ϕ, $\vdash \bot \to \phi$,

[3] As it turns out, a notion of distributed belief that corresponds *exactly* to the intersection of individual belief states is not definable in a possible world setting [21].

and hence $\vdash \mathbf{B}^i_0\phi$ by A2 and NEC). It is easy to verify that the following patterns of reasoning follow from these axioms:

$$
\begin{array}{ll}
\mathbf{B}^i_a\phi \rightarrow \mathbf{B}^i_b\phi & \text{for any } b \leq a \\
(\mathbf{B}^i_a\phi \wedge \mathbf{B}^i_b\psi) \rightarrow \mathbf{B}^i_c(\phi \wedge \psi) & \text{with } c = \min(a, b) \\
(\mathbf{B}^i_a\phi \wedge \mathbf{B}^i_b\psi) \rightarrow \mathbf{B}^i_c(\phi \vee \psi) & \text{with } c = \max(a, b).
\end{array}
\quad (3)
$$

The axioms for distributed belief are more peculiar. A3 says that distributed belief behaves as any other agent's belief. A4 is the axiom that effectively characterizes distributed belief. It corresponds to the semantic constraint (2), and tells us that the belief distributely held by all the agents about a proposition ϕ must include the fusion of all the individual agents' beliefs about ϕ, and that the corresponding degree is computed via the $\oplus$ T-conorm. (Recall that, by A1, each agent believes any proposition ϕ at a degree of at least 0.)

The following consequences of our axioms can give us further insights into the behavior of distributed belief:

$$
\begin{array}{ll}
\mathbf{B}^i_a\phi \rightarrow \mathbf{D}_a\phi & \\
(\mathbf{B}^i_a\phi \wedge \mathbf{B}^j_b\psi) \rightarrow \mathbf{D}_c(\phi \wedge \psi) & \text{with } c = \min(a, b) \\
(\mathbf{B}^i_a\phi \wedge \mathbf{B}^j_b\psi) \rightarrow \mathbf{D}_c(\phi \vee \psi) & \text{with } c = a \oplus b \text{ provided that } i \neq j.
\end{array}
\quad (4)
$$

So, distributed belief includes any individual belief; in particular, D inherits $\mathbf{D}_0\bot$ and $\mathbf{D}_1\top$, and all the properties (3) of B. Moreover, degrees of belief from different agents are combined using min for conjunctions and using $\oplus$ for disjunctions. The last property shows how $\mathbf{PL}^{\otimes}_n$ can perform strengthening of belief coming from different agents. If two agents 1 and 2 both believe one proposition ϕ with respective strengths a and b, then, as $\phi \vee \phi \equiv \phi$, we can infer $\mathbf{D}_{a\oplus b}\phi$. From (1) above, we have $a \oplus b \geq \max\{a, b\}$ in general, and $a \oplus b > \max\{a, b\}$ if we use a non-idempotent T-norm (e.g., the product). In the latter case, distributed belief effectively *reinforces* the beliefs shared by different agents. Note by contrast that from $\mathbf{B}^i_a\phi \wedge \mathbf{B}^i_b\psi$ we can only deduce $\mathbf{B}^i_c(\phi \vee \psi)$ with $c = \max\{a, b\}$. As $\otimes$ and $\oplus$ are commutative and associative, the above considerations extend to any number of agents and any sequence of combination.

The example shown in the introduction assumes that we chose the product T-norm for $\otimes$; in this case, we have $a \oplus b = a + b - ab$. It is easy to see how this example follows from the axioms of $\mathbf{PL}^{\otimes}_n$:

$\mathbf{B}^1_1(\text{rain} \rightarrow (\text{wind} \rightarrow \text{rain}))$	NEC
$\mathbf{B}^1_{0.4}(\text{wind} \rightarrow \text{rain})$	(a) and A2
$\mathbf{B}^2_0(\text{wind} \rightarrow \text{rain})$	A1
$\mathbf{D}_{0.82}(\text{wind} \rightarrow \text{rain})$	(c) and A4
$\mathbf{D}_1\text{wind}$	(b) and A4
$\mathbf{D}_{0.82}\text{rain}$	A3.

3.4 The axiom schemas T, D, 4 and 5

$\mathbf{PL}^{\otimes}_n$ is an extension of K_n, the minimal epistemic logic for multi-agent belief. In fact, if we only consider the $\mathbf{B}^i_1$ operators, that is, if we only consider categorical belief, then A0, A2, MP and NEC reduce to the usual axioms of K_n. Typical epistemic logic also make use of (some of) the axiom schemas T: $\Box\phi \rightarrow \phi$; D: $\neg\Box\bot$; 4: $\Box\phi \rightarrow \Box\Box\phi$;

and 5: $\neg\Box\phi \to \Box\neg\Box\phi$. In our frame, these axioms are translated as follows, where $i \in \{1, \ldots, n\}$, and α is some fixed threshold in $[0, 1]$.

$$\begin{array}{ll} \alpha\textbf{-T} & \mathbf{B}^i_\alpha\phi \to \phi \\ \alpha\textbf{-D} & \neg\mathbf{B}^i_\alpha\bot \\ \alpha\textbf{-4} & \mathbf{B}^i_\beta\phi \to \mathbf{B}^i_\beta\mathbf{B}^i_\beta\phi \quad \text{for any } \beta \geq \alpha \\ \alpha\textbf{-5} & \neg\mathbf{B}^i_\beta\phi \to \mathbf{B}^i_\beta\neg\mathbf{B}^i_\beta\phi \quad \text{for any } \beta \geq \alpha \end{array} \tag{5}$$

Intuitively, α-**T** says that agents cannot believe (more than α) false propositions — a requirement commonly imposed on knowledge. For instance, we can use

$$1\textbf{-T} \qquad \mathbf{B}^i_1\phi \to \phi$$

to say that facts that are completely believed must be true; note that α-**T** entails β-**T** for any $\beta \geq \alpha$. α-**D** expresses a weaker requirement: it says that agents' belief cannot be (more than α-) inconsistent; again, α-**D** entails β-**D** for any $\beta \geq \alpha$. Finally, α-**4** and α-**5** concern the requirements imposed on introspective (or "higher-order") belief. They say that agents are perfectly introspective about what they believe (4) and what they do not believe (5), provided the degree of belief is at least α. For example,

$$1\textbf{-4} \qquad \mathbf{B}^i_1\phi \to \mathbf{B}^i_1\mathbf{B}^i_1\phi$$

says that if agents completely believe something, they completely believe to believe it; and 0-**4** says that this is true for any degree of belief. Axioms similar to the (5) can be written for the $\mathbf{D}_a$ modality. We can add any subset of these axioms to $\mathbf{PL}^{\otimes}_n$ in order to obtain a specific logic tailored to a specific application. Moreover we can use different α thresholds in different schemes (notice however that α must be strictly positive in α-**T** and α-**D** in order for these axioms to be consistent.)

In standard modal logics, it is well known that the axioms T, D, 4 and 5 correspond to certain restrictions on the accessibility relation: namely, that it be reflexive (T), serial (D), transitive (4), or Euclidean (5) [5]. We can prove a similar result for the axioms (5) with respect to the following conditions (see [16] for a related characterization).

Definition 4 Let $F = \langle W, \pi_0, \pi_1, \ldots, \pi_n \rangle$ be a $\Pi^{\otimes}_n$-frame. We say that F is

$$\begin{array}{lcl} \alpha\text{-reflexive} & \iff & (\forall i, v)(\pi_{i,v}(v) > 1 - \alpha) \\ \alpha\text{-serial} & \iff & (\forall i, v)(\sup_{w \in W} \pi_{i,v}(w) > 1 - \alpha) \\ \alpha\text{-transitive} & \iff & (\forall i, v, z)(\pi_{i,v}(z) \leq 1 - \alpha \Rightarrow \\ & & \pi_{i,v}(z) \geq \sup_{w \in W} \min(\pi_{i,v}(w), \pi_{i,w}(z)) \\ \alpha\text{-euclidean} & \iff & (\forall i, w, z)(\pi_{i,w}(z) \leq 1 - \alpha \Rightarrow \\ & & \pi_{i,w}(z) \geq \sup_{v \in W} \min(\pi_{i,v}(w), \pi_{i,v}(z)) \end{array}$$

The key step to prove the correspondence is to note that we can write $\Pi^{\otimes}_n$-frames as classical Kripke models containing countably many accessibility relations, one for each operator $\mathbf{B}^i_a$. We substitute each π_i by the set of binary relations over W $\{R^i_\alpha \mid \alpha$ a rational number in $[0, 1]\}$, where

$$R^i_\alpha = \{(v, w) \mid \pi_{i,v}(w) > 1 - \alpha\}.$$

Note that, for any i, the R^i_α relations are linearly ordered by inclusion: if $\alpha \leq \beta$, then $vR^i_\alpha w$ implies $vR^i_\beta w$. It is easy to see that the truth condition for each modality $\mathbf{B}^i_a$ is equivalent to a standard Kripke condition using the accessibility relation R^i_a.

Lemma 5 $M, w \models \mathbf{B}^i_a \phi$ iff, for all $v \in W$, $wR^i_a v$ implies $M, v \models \phi$.

Now, each of the axioms (5) determines the usual property on the corresponding accessibility relations. For example, α-**T** determines the reflexivity of R^i_α for any i. And α-**4** determines the transitivity of R^i_β for any i and any $\beta \geq \alpha$. (Note that R^i_0 is always the empty relation, hence it is not reflexive nor serial in any $\Pi^\otimes_n$-frame; this corresponds to the fact that 0-**T** and 0-**D** are inconsistent in $\mathbf{PL}^\otimes_n$.) The following lemma links the properties of the accessibility relations to those in Def. 4.[4]

Lemma 6 Let F be a $\Pi^\otimes_n$-frame. Then, F is

1. α-reflexive iff, for all i and for all $\beta \geq \alpha$, R^i_β is reflexive.
2. α-serial iff, for all i and for all $\beta \geq \alpha$, R^i_β is serial.
3. α-transitive iff, for all i and all $\beta \geq \alpha$, R^i_β is transitive.
4. α-euclidean iff, for all i and all $\beta \geq \alpha$, R^i_β is Euclidean.

Corollary 7 The axiom schemas (5) determine the properties in Def. 4, taken in the order.

Proof By the standard result, α-**T** determines, for any i, the reflexivity of R^i_α (hence of R^i_β for any $\beta \geq \alpha$); α-**D** determines, for any i, the seriality of R^i_α (hence of R^i_β for any $\beta \geq \alpha$); and α-**4** and α-**5** respectively determine, for any i and any $\beta \geq \alpha$, the transitivity and the euclideanicity of R^i_β. Lemma 6 tells us that these properties correspond to those in Def. 4, in the order. □

Again, similar considerations can be made for the $\mathbf{D}_a$ modality.

4 Soundness, Completeness and Decidability

In this section, we consider the problem of proving that the axiomatic system above can deduce all and only the formulas that are valid according to our semantics, that is, that this system actually captures the intended meaning. The proof of the soundness of $\mathbf{PL}^\otimes_n$ (all the consequences of $\mathbf{PL}^\otimes_n$ are $\Pi^\otimes_n$-valid) is a matter of routine.

Theorem 8 (Soundness) $\mathbf{PL}^\otimes_n$is sound with respect to $\Pi^\otimes_n$-structures.

Proof By induction on the length of the derivation. As usual, it suffices to check that the axioms are valid, and that the rules preserve validity. Let M be any $\Pi^\otimes_n$-structure whose set of worlds is W_M, and w be any world in W_M.

[4] We could of course apply the same procedure to import other popular axiom schemas in our frame; for example, the *Brouwersche* axiom B: $\phi \to \mathbf{B}^i_a \neg \mathbf{B}^i_a \neg \phi$ corresponds to the symmetry of R^i_a. Axioms other than T, D, 4 and 5, however, do not have a straightforward epistemic interpretation.

A0. Obvious by the definition of the satisfaction relation over negation and conjunction.

A1. $M, w \models \mathbf{B}^i_0\bot$ iff $1 - \sup_{v \notin [\![\bot]\!]_M} \pi_{i,w}(v) \geq 0$, that is, iff $\sup_{v \in W} \pi_{i,w}(v) \leq 1$. But this is always true because possibility functions take their values in $[0, 1]$.

A2. Assume $M, w \models \mathbf{B}^i_a(\phi \rightarrow \psi)$ and $M, w \models \mathbf{B}^i_b\phi$. Then $N_{i,w}([\![\phi \rightarrow \psi]\!]_M) \geq a$ and $N_{i,w}([\![\phi]\!]_M) \geq b$. Hence, by the properties of necessity functions,

$$N_{i,w}([\![\psi]\!]) \geq N_{i,w}([\![\psi \wedge \phi]\!]_M) = N_{i,w}([\![\phi \rightarrow \psi]\!]_M \cap [\![\phi]\!]_M) \geq \min(a, b)$$

and so $M, v \models \mathbf{B}^i_{\min(a,b)}\psi$.

A3. The same proof as for A2.

A4. Assume $M, w \models \bigwedge_{i=1}^n \mathbf{B}^i_{a_i}\phi$; then $N_{i,w}([\![\psi]\!]) \geq a_i$ for any i. By the properties of necessity functions, since $\pi_{0,w} \leq \bigotimes_{i=1}^n \pi_{i,w}$, it holds

$$N_{0,w}([\![\phi]\!]_M) \geq \bigoplus_{i=1}^n N_{i,w}([\![\phi]\!]_M) \geq \bigoplus_{i=1}^n a_i,$$

hence $M, v \models \mathbf{D}_c\phi$, where $c = \bigoplus_{i=1}^n a_i$.

MP Obvious by the definition of the satisfaction relation.

NEC If for all M and all $w \in W_M$, $M, w \models \phi$, then for all M the set $[\![\phi]\!]_M = W_M$, hence for all $w \in W_M$, $N_{i,w}([\![\phi]\!]_M) = 1$ and so $M, w \models \mathbf{B}^i_1\phi$. □

Proving completeness (all $\Pi^{\otimes}_n$-valid formulas can be derived in $\mathbf{PL}^{\otimes}_n$) is, as it is usually the case, much harder, and we devote most of this section to this. The proof is long and somehow technical, so we start with some comments on the overall schema. Also, to improve readability, we defer to the Appendix the proofs which are not necessary to a first understanding.

One technical difficulty is that our logic is not compact. To see this, note that the infinite set $\{\mathbf{B}^i_{1-\frac{1}{k}}\psi \mid k = 1, 2, \ldots\} \cup \{\neg\mathbf{B}^i_1\psi\}$ does not have a $\Pi^{\otimes}_n$-model, although any finite subset of it does. This rules out the possibility of proving completeness by using a standard canonical model construction (which would imply compactness) [5]. What we shall do instead is to build, for any given formula ψ, a specific $\Pi^{\otimes}_n$-structure M^ψ by using finitary means; this corresponds to building a canonical model and filtering it through ψ in one step (see [8] for the use of a similar technique in the context of probabilistic logics). The aim is to build M^ψ in such a way that, if ψ is not a theorem of $\mathbf{PL}^{\otimes}_n$, then ψ will fail at some world of M^ψ. This means that any formula which cannot be derived in $\mathbf{PL}^{\otimes}_n$ is not $\Pi^{\otimes}_n$-valid; the contrapositive of this is, exactly, completeness.

We will proceed as follows. We will assume for the length of the proof that a formula ψ is given. The first step will be to build, from this formula, the finitary analogous of the maximally consistent sets used in the standard Henkin construction; we will show that these sets behave nicely for our purposes. Second, we will use these sets to build M^ψ, the canonical $\Pi^{\otimes}_n$-structure *for* ψ. Third, we will prove that M^ψ enjoys the fundamental property of canonical models: the formulas true at a world w of M^ψ are exactly those formulas that appear in w. This is the main step in our proof, where possibility theory cames into play. The last step, proving completeness, will follow easily from this result: all we will have to do is to show that, if ψ is not a theorem of $\mathbf{PL}_n^{\otimes}$, then there is a world in M^{ψ} that does not contain ψ, and then ψ is not true at that world.

So, let Gen(ψ) denote the set consisting of all the sub-formulae of ψ and their negations. We call *ψ-full set* any subset Γ of $Gen(\psi)$ that is maximally consistent in $Gen(\psi)$ — that is, Γ is consistent in $\mathbf{PL}_n^{\otimes}$, and any proper extension of it is inconsistent. (Note that this definition is system-relative). It is easy to see that $Gen(\psi)$ is finite, and so is the number of ψ-full sets. The maximally consistent sets used in the standard Henkin construction have the remarkable property of behaving as semantic truth sets: the following lemmas show that ψ-full sets have the same property, provided we restrict attention to the sub-language $Gen(\psi)$.

Lemma 9 Let Γ be any ψ-full set. Then

(i) For any ϕ in $Gen(\psi)$, ϕ is in Γ iff $\Gamma \vdash \phi$;

(ii) For any ϕ in $Gen(\psi)$, exactly one of ϕ and $\neg\phi$ is in Γ;

(iii) For any $(\phi_1 \wedge \phi_2)$ in $Gen(\psi)$, $(\phi_1 \wedge \phi_2)$ is in Γ iff both ϕ_1 and ϕ_2 are in Γ.

Lemma 10 Let Γ be any consistent subset of $Gen(\psi)$. Then there is a ψ-full set Δ that contains Γ.

Lemma 11 Let $\phi \in Gen(\psi)$ and $\Gamma \subseteq Gen(\psi)$. Then, $\Gamma \vdash \phi$ if and only if ϕ belongs to any ψ-full set that contains Γ.

In order to define the canonical model M^{ψ} for ψ, we introduce the notion of possibilistic constraint. We call (possibilistic) constraint over W a couple (A, a), where $A \subseteq W$ and $a \in [0, 1]$; we say that a possibility distribution π satisfies the constraint (A, a) iff $N(A) \geq a$. It follows from the properties of possibility distributions that for any set of constraints $\mathcal{C} = \{(A_i, a_i) \mid i \in I\}$ there exists a unique *least committed possibilistic solution*, *i.e.*, a possibility distribution that satisfies all of the constraints, and is minimal with respect to the order $\leq$ (recall Section 2). For a single constraint, this solution is given by

$$Lcs(A, a)(w) = \begin{cases} 1 & \text{if } w \in A \\ 1 - a & \text{otherwise.} \end{cases}$$

For a set $\mathcal{C}$ of constraints, it is given by

$$LCS(\mathcal{C})(w) = \bigwedge_{(A_i, a_i) \in \mathcal{C}} Lcs(A_i, a_i)(w) = \max_{i \in I, w \notin A_i} (1 - a_i).$$

We are now ready to define our canonical models.

Definition 12 (Canonical model) The *canonical model for ψ* is the $\Pi_n^{\otimes}$-structure $M^{\psi} = \langle W^{\psi}, V^{\psi}, \pi_0^{\psi}, \pi_1^{\psi}, \ldots, \pi_n^{\psi} \rangle$ given by

1. W^{ψ} is the set of all the ψ-full sets;

2. for each propositional variable p, $V^\psi(p) = \{w \in W^\psi \mid p \in w\}$;

3. for each $w \in W^\psi$ and $i = 1, \ldots, n$, $\pi^\psi_{i,w} = LCS(\{(|\phi|^\psi, a) \mid \mathbf{B}^i_a\phi \in w\})$;

4. $\pi^\psi_{0,w} = \pi^{\mathbf{D}}_w \wedge \pi^{\otimes}_w$, where $\pi^{\mathbf{D}}_w = LCS(\{(|\phi|^\psi, a) \mid \mathbf{D}_a\phi \in w\})$, and $\pi^{\otimes}_w = \bigotimes_{i=1}^n \pi^\psi_{i,w}$;

where $|\phi|^\psi =_{\text{def}} \{v \in W^\psi | \phi \in v\}$.

It is easy to see that M^ψ is indeed a $\Pi^\otimes_n$-structure, and that it is finite. Intuitively, each $\pi^\psi_{i,w}$ is so built that agent i believes in w exactly those formulas ϕ for which $\mathbf{B}^i_a\phi$ is in w; $\pi^\psi_{0,w}$ is built in a similar way with respect to $\mathbf{D}_a\phi$, but also includes the $\otimes$ combination of all the other $\pi^\psi_{i,w}$'s. The aim of this definition is clearly to enable us to show that M^ψ enjoys the fundamental property of canonical models, stated in the following lemma.

Lemma 13 (Truth lemma) Let χ be in Gen(ψ) and $w \in W^\psi$; then $\chi \in w$ iff $M^\psi, w \models \chi$.

Proof The proof is by induction on the structure of χ. As usual with this kind of proofs, the only non trivial cases are those relative to the modal operators. The crucial steps for these cases are proved as separate lemmas in the Appendix.

χ is a propositional variable. Then the thesis holds by definition of V^ψ.

$\chi \equiv \neg\phi$. By Lemma 9(ii), we know that $\neg\phi \in w$ iff $\phi \notin w$, and by induction hypotheses $\phi \notin w$ iff $M^\psi, w \not\models \phi$, which means that $M^\psi, w \models \neg\phi$.

$\chi \equiv (\phi_1 \wedge \phi_2)$. By Lemma 9(iii), $(\phi_1 \wedge \phi_2) \in w$ iff both $\phi_1 \in w$ and $\phi_2 \in w$. By induction hypothesis, the latter is true iff both $M^\psi, w \models \phi_1$ and $M^\psi, w \models \phi_2$, which means that $M^\psi, w \models (\phi_1 \wedge \phi_2)$.

$\chi \equiv \mathbf{B}^i_a\phi$. The *only if* part is easy: suppose that $\mathbf{B}^i_a\phi \in w$; then, by definition of $\pi^\psi_{i,w}$, $N^\psi_{i,w}(|\phi|^\psi) \geq a$, and $M^\psi, w \models \mathbf{B}^i_a\phi$ follows by noticing that $|\phi|^\psi = [\![\phi]\!]_{M^\psi}$ by induction hypothesis. To see that the reverse holds, suppose that $M^\psi, w \models \mathbf{B}^i_a\phi$. Then, by using the induction hypothesis, $N^\psi_{i,w}(|\phi|^\psi) \geq a$. Let $\Gamma = \{\mathbf{B}^i_{a_j}\phi_j \mid \mathbf{B}^i_{a_j}\phi_j \in w\}$, and recall that $\pi^\psi_{i,w}$ is, by definition, the least committed possibilistic solution of

$$\{(|\phi_j|^\psi, a_j) \mid \mathbf{B}^i_{a_j}\phi_j \in \Gamma\}.$$

Now, Lemma A.3 (in the Appendix) tells us that, under these conditions, we have

$$\Gamma \vdash \mathbf{B}^i_a\phi$$

Then, by using Lemma 11, any ψ-full set containing Γ must also contain $\mathbf{B}^i_a\phi$: so, in particular, $\mathbf{B}^i_a\phi \in w$.

$\chi \equiv \mathbf{D}_a\phi$. The *only if* part can be proved as for the B case. For the *if* part, suppose that $M^\psi, w \models \mathbf{D}_a\phi$; then, by using the induction hypothesis, $N^\psi_{0,w}(|\phi|^\psi) \geq a$. Consider the sets $\Gamma_i = \{\mathbf{B}^i_{a_{i,j}}\phi_{i,j} \mid \mathbf{B}^i_{a_{i,j}}\phi_{i,j} \in w\}$, $i = 1, \ldots, n$, and $\Gamma_0 =$

$\{\mathbf{D}_{a_j}\phi_j \mid \mathbf{D}_{a_j}\phi_j \in w\}$, and let $\Gamma = \Gamma_0 \cup \bigcup_{i=1}^n \Gamma_i$. Now $\pi_{0,w}^{\psi} = \pi_w^{\mathbf{D}} \wedge \pi_w^{\otimes}$, so we can apply Lemma A.6 (in the Appendix) to get

$$\Gamma \vdash \mathbf{D}_a\phi.$$

Finally, we use Lemma 11 as above to conclude that $\mathbf{D}_a\phi \in w$. □

We now have all the needed ingredients to establish the completeness of $\mathbf{PL}_n^{\otimes}$.

Theorem 14 (Completeness) $\mathbf{PL}_n^{\otimes}$is complete with respect to $\Pi_n^{\otimes}$-structures.

Proof Consider any $\Pi_n^{\otimes}$-valid formula ψ, and suppose that ψ is not a theorem of $\mathbf{PL}_n^{\otimes}$. Then, $\neg\psi$ is consistent. As $\neg\psi$ is in Gen(ψ), by Lemma 10 it must belong to some world in W^{ψ}, say w_0. Then, by Lemma 13, $M^{\psi}, w_0 \models \neg\psi$, and hence $M^{\psi}, w_0 \not\models \psi$. But this means that ψ fails at the $\Pi_n^{\otimes}$-structure M^{ψ}, and so it cannot be $\Pi_n^{\otimes}$-valid, contradicting our assumption. Hence, ψ must be a theorem of $\mathbf{PL}_n^{\otimes}$. □

This result, together with Theorem 8, gives the following

Corollary 15 The calculus $\mathbf{PL}_n^{\otimes}$ is determined by the class of $\Pi_n^{\otimes}$-structures.

We end this section by considering the issue of the decidability of our logic. A common way to establish decidability of a modal logic is to prove that it enjoys the finite model property: every non-theorem fails in a finite structure, where "finite" means with a finite set W of worlds (this guarantees that the set of non-theorems is recursively enumerable). The situation is a slightly more complex in our case. In fact, even when W is finite, the set of possibility functions over W is not countable, and then the class of finite $\Pi_n^{\otimes}$-structures is not recursively enumerable. To fix this situation, we consider the following class of structures.

Definition 16 A $\Pi_n^{\otimes}$-structure $\langle W, \pi_0, \pi_1, \ldots, \pi_n, V\rangle$ is *rational finite* if W is finite and, for any i and w, the possibility distribution $\pi_{i,w}$ is rational-valued.

It is easy to see that the class of rational finite $\Pi_n^{\otimes}$-structures is recursively enumerable. The following lemma plays the same role, in our case, as the standard finite model property; it holds for some elements of our family of logics only: those based on a rational $\otimes$.

Lemma 17 Let $\otimes$ be any continuous rational T-norm. Then, any non-theorem of $\mathbf{PL}_n^{\otimes}$ fails at a rational finite $\Pi_n^{\otimes}$-structure.

We can now state our decidability result.

Theorem 18 (Decidability) If $\otimes$ is rational, then $\mathbf{PL}_n^{\otimes}$ is decidable.

Proof Let ψ be any formula of $\mathcal{L}$. Lemma 17 guarantees that we have an algorithm that terminates if ψ is not a theorem of $\mathbf{PL}_n^{\otimes}$: we generate all the rational finite $\Pi_n^{\otimes}$-structures until we find one in which ψ fails. On the other hand, the proof theory of $\mathbf{PL}_n^{\otimes}$ gives us an algorithm that terminates if ψ is a theorem of $\mathbf{PL}_n^{\otimes}$: we generate all the theorems of $\mathbf{PL}_n^{\otimes}$ until we find ψ. (We can do this because the indexes of the modal operators are rational, and hence the set of substitution instances of the axiom schemas is countable.) By dove-tailing these two algorithms, we obtain an effective procedure to decide theoremhood in $\mathbf{PL}_n^{\otimes}$. □

The decidability of $\mathbf{PL}_n^{\otimes}$ critically depends on the assumption that degrees of belief are rational numbers, and that a rational-valued T-norm is used for merging belief (the most commonly used T-norms, $\wedge$, $\times$ and $\times_L$, are rational). However, we emphasize that our soundness and completeness results would still hold if we allowed real numbers as indexes, and arbitrary continuous T-norms for combination.

5 Discussion

Practical multi-agent reasoning systems need the ability to combine partial beliefs originating from different agents; at present, this can only be done by means of non-logical tools, like Bayesian nets, ruling out the possibility of symbolic reasoning on uncertain data. We believe that the richness of the merging operations allowed by uncertainty models should be captured inside a formal system. This can become particularly attractive when we want to mix uncertain information with certain, high-level symbolic knowledge.

In this paper, we have presented a logic (in fact, a parametric family of logics) able to represent and merge partial belief of multiple agents. This merging ability can be applied to problems like pooling expert opinions and combining information from multiple unreliable sources. The logical formalism makes it particularly convenient to incorporate prior knowledge (e.g., facts or general rules) in the reasoning process. In a related paper [3], we report an application of $\mathbf{PL}_n^{\otimes}$ to a problem in autonomous robotics involving fusion of sensorial and prior information from multiple sources.

It should be noticed that, although the $\mathbf{PL}_n^{\otimes}$ logic allows us to encode numerical degrees of belief, it does not force us to do so. In particular, we can represent crisp knowledge in $\mathbf{PL}_n^{\otimes}$ by means of the $\mathbf{B}_1^i$ and $\mathbf{D}_1$ operators, as we did for item (b) in the Introduction. It is easy to see that $\mathbf{PL}_n^{\otimes}$ behaves like a standard multi-agent epistemic logic (i.e., $\mathbf{K}_n$ and the stronger related systems) with distributed belief if we only use crisp information. Clearly, however, we are mainly interested in using $\mathbf{PL}_n^{\otimes}$ in applications in which we are able to use numerical degrees to quantify the reliability of items of information.

The behavior of the distributed belief operator provided by $\mathbf{PL}_n^{\otimes}$ depends on the choice of the $\otimes$ T-norm. Using the min T-norm results in simple collection of information: from $\mathbf{B}_a^i\phi$ and $\mathbf{B}_b^j\phi$ we only deduce $\mathbf{D}_{\max\{a,b\}}\phi$, and no accumulation is achieved. In a sense, we obtain the same result as if one single agent believed everything that is believed by any one of the agents. This is perhaps the simplest way to extend the distributed belief operator introduced in classical epistemic modal logics. However, the interest of using degrees of belief becomes more evident if we use a non-idempotent operator for $\otimes$, like the product or the Łukasiewicz T-norms. In this case, the distributed belief operator $\mathbf{D}$ *reinforces* the belief shared by different agents: in the example above, we obtain $\mathbf{D}_c\phi$ where, in general, both $c > a$ and $c > b$.

The use of a reinforcing combination operator $\otimes$ only makes sense when the agents can be regarded as independent sources of information. In such cases, having two or more agents to agree on a proposition legitimately counts as a stronger evidence for the truth of this proposition. Moreover, we can capture different reinforcement behaviors using different operators. For instance, if we use the product T-norm for $\otimes$ in the example above, we get $c = a + b - ab$. This type of combination is most adequate in applications in which the agents can be regarded as independent sources of probabilistic information. If instead we use the Łukasiewicz T-norm, we get $c = \min(a + b, 1)$. This type of combination is adequate in applications in which the agents provide additive degrees of support following a voting metaphor. Finally, if independence cannot be granted, a cautious combinator (the idempotent min T-norm) should be used.

On the theoretical side, the $\mathbf{PL}_n^{\otimes}$ logic contributes to two distinct fields: classical multi-agent epistemic logics [9], which we extend to deal with uncertain information; and epistemic logics for uncertainty [17, 8, 10], that we enrich with a distributed belief operator. In that, $\mathbf{PL}_n^{\otimes}$ is related to other multi-modal systems for partial belief, such as [8, 15, 16, 17] and [10]. The closest relative of $\mathbf{PL}_n^{\otimes}$, however, is the system proposed in [4], where a connective is introduced in the language for merging information from independent sources. In fact, we borrowed from that work the idea of representing the merging of information *inside* the logic using T-norm products.

$\mathbf{PL}_n^{\otimes}$ departs from the logic in [4] in that use modal operators to distinguish distinct sources, and adopt a possible world semantics instead of an algebraic one. These facts have two main consequences. First, our logic allows for nested epistemic reasoning: we can say, for instance, that $\mathbf{B}_a^1 \mathbf{D}_b \phi$, meaning that agent 1 believes at level a that the formula ϕ is "distributively believed" at level b. Second, negation behaves differently in the two logics: while our negation is a typical modal negation (which means absence of information), negation in the algebraic setting represents positive information on some "orthogonal" formula.

The logic in [4] has one more degree of expressiveness than $\mathbf{PL}_n^{\otimes}$, in that it allows for many-level merging of information. For instance, by the formula $((A \otimes B) \& C) \otimes D$ they express the total information collected by agent 1 who receives from agent 2 the token of information $(A \otimes B) \& C$, and from an independent agent 3 the token D; agent 2 has, in turn, collected the two tokens A and B from the independent agents 4 and 5, and added his own token C. We could incorporate a similar mechanism in $\mathbf{PL}_n^{\otimes}$ by defining a set of epistemic operators D^G, where G is a subset of the agents; the intended meaning of D^G is to only combine the evidence provided by the agents in G. The introduction of operators for sub-group distributed belief is left for future work. Other issues that remain to be explored in future work include:

- A detailed analysis of the complexity of $\mathbf{PL}_n^{\otimes}$;
- The definition of the uncertain counterpart of the operator of *common knowledge* [9]; and
- A similar construction for a logic based on a more general representation of the information states (e.g., belief functions [1], [19]).

Acknowledgments This work has benefit from discussions with, and comments by, Natasha Alechina, Bruno Marchal, Antonia Sinachopoulos, Philippe Smets and Claudio Sossai. Work by Luca Boldrin was performed while he was visiting ILLC at the University of Amsterdam. Alessandro Saffiotti was supported by the BELON project, founded by the *Communauté Française de Belgique*.

References

[1] N. Alechina and P. Smets. A note on modal logics for partial belief. ILLC Technical Note X-94-06, University of Amsterdam, NL, 1994.

[2] I. Bloch. Information combination operators for data fusion: a comparative review with classification. IEEE Trans. on Systems, Man and Cybernetics, Part A, 26 (1), 1996, pp. 52–67.

[3] L. Boldrin and A. Saffiotti A modal logic for merging information from unreliable sources. Available by anonymous FTP from `iridia.ulb.ac.be`, file /pub/saffiotti/uncertainty/isis96.ps.gz.

[4] L. Boldrin and C. Sossai. Local possibilistic logic. J. of Applied Non-Classical Logic 7(3), 1997.

[5] B.F. Chellas. Modal Logic: an introduction. Cambridge University Press, 1980.

[6] D. Dubois and H. Prade. A review of fuzzy set aggregation connectives. Information Sciences 36, 1985, pp. 85-121.

[7] D. Dubois, J. Lang and H. Prade. Possibilistic logic In: D. Gabbay, C. Hogger and J. Robinson (eds.) Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3. Clarendon Press 1994.

[8] R. Fagin and J.Y. Halpern. Reasoning about knowledge and probability. J. of the ACM, 41, 1994, pp. 340-367.

[9] R. Fagin, J.Y. Halpern, Y. Moses and M.Y. Vardi. Reasoning about knowledge. The MIT Press 1995.

[10] F. Esteva, P. García, L. Godo and R. Rodriguez. A modal account of similarity-based reasoning. Technical Report 96/08. IIIA, Coseil Superior d'Investigacions Científiques, Barcelona, Spain, 1996.

[11] P. Hajek, D. Harmancova and R. Verbrugge. A qualitative fuzzy possibilistic logic. Int. J. of Approximate Reasoning 7, 1994.

[12] P. Hajek. On logics of approximate reasoning. In: M. Masuch and L. Pólos (eds.), Knowledge representation and reasoning under uncertainty. LNAI 808, Springer-Verlag 1994, pp. 17-29.

[13] J.Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. Procs. of the ACM Symp. on Principles of Distributed Computing, 1984, pp. 50-61.

[14] U. Höhle. A survey on the fundamentals of fuzzy set theory. To appear in Kreith (ed.) Handbook of Mechanical Engineering. GRC Press.

[15] C.J. Liau and B.I-P. Lin. Quantitative modal logic and possibilistic reasoning. Procs. of ECAI 92, 1992, pp. 43-47.

[16] C.J. Liau and B.I-P. Lin. Reasoning about higher order uncertainty in possibilistic logic. In: J. Komorowski and Z.W. Raś (Eds.) Methodologies for Intelligent Systems. LNAI 689, Springer-Verlag 1983, pp. 316-325.

[17] C.J. Liau and I-P. Lin. A theoretical investigation into quantitative modal logic. Fuzzy sets and systems 75, 1995, 355–363.

[18] T. Murai, M. Miyakoshi and M. Shimbo. Measure-based semantics for modal logic. In: R. Lowen and M. Roubens (eds.) Fuzzy logic. Kluwer Academic Publishers, 1993, pp. 395-405.

[19] A. Saffiotti. Partial epistemic logics based on belief functions — preliminary report. Iridia Technical Report 95-21. Université Libre de Bruxelles, Belgium, 1995.

[20] A. Saffiotti and L. Wesley. Perception-based self-localization using fuzzy locations. In L. Dorst, M. van Lambalgen and F. Voorbraak (eds.) Reasoning with Uncertainty in Robotics. LNAI 1093, Springer, Berlin, DE, 1996, pp. 368–385.

[21] W. van der Hoek. Modalities for reasoning about knowledge and quantities. PhD dissertation, Free University of Amsterdam, 1992.

[22] F. Voorbraak. As far as I know. Epistemic logic and uncertainty. PhD dissertation, Dept. of Philosophy, Utrecht University, 1993.

[23] S. Weber. A general concept of fuzzy connectives, negations and implications based on T-norms and T-conorms. Fuzzy sets and systems, 11, 1983, pp. 115-134.

[24] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. Fuzzy sets and systems 1, 1978, pp. 3-28

Appendix: Proofs of technical lemmas

Lemma 6 $M, w \models \mathbf{B}_a^i \phi$ iff, for all $v \in W$, $wR_a^i v$ implies $M, v \models \phi$.

Proof We have that

$$M, w \models \mathbf{B}_a^i \phi \text{ iff } 1 - \sup_{v \notin [\![\phi]\!]_M} \pi_{i,w}(v) \geq a \text{ iff } \sup_{v \notin [\![\phi]\!]_M} \pi_{i,w}(v) \leq 1 - a.$$

The last condition is equivalent to

$$\text{For all } v \in W, \quad M, v \not\models \phi \text{ implies } \pi_{i,w}(v) \leq 1 - a.$$

By taking the contrapositive, we have

$$\text{For all } v \in W, \quad \pi_{i,w}(v) > 1 - a \text{ implies } M, v \models \phi$$

which gives our thesis by recalling the definition of R_a^i. □

Lemma 7 Let F be a $\Pi_n^{\otimes}$-frame. Then, F is

1. α-reflexive iff, for all i and for all $\beta \geq \alpha$, R_β^i is reflexive.
2. α-serial iff, for all i and for all $\beta \geq \alpha$, R_β^i is serial.
3. α-transitive iff, for all i and all $\beta \geq \alpha$, R_β^i is transitive.
4. α-euclidean iff, for all i and all $\beta \geq \alpha$, R_β^i is euclidean.

Proof 1. Note that R_α^i reflexive implies R_β^i reflexive for any $\beta \geq \alpha$. R_α^i is reflexive iff, for all v, $vR_\alpha^i v$ iff, for all v, $\pi_{i,v}(v) > 1 - \alpha$, that is, iff F is α-reflexive.

2. Note that R_α^i serial implies R_β^i serial for any $\beta \geq \alpha$. R_α^i serial iff for all v there is a w such that $vR_\alpha^i w$, that is, there is a w such that $\pi_{i,v}(w) > 1 - \alpha$. But this is true iff $\sup_{w \in W} \pi_{i,v}(w) > 1 - \alpha$, that is, iff F is α-serial.

3. Assume R_β^i transitive for all i and for all $\beta \geq \alpha$. For some fixed v, w, z, take $x =_{\text{def}} \min(\pi_{i,v}(w), \pi_{i,w}(z))$. By transitivity of R_β^i, we have:

$$\text{For all } \beta \geq \alpha, \quad vR_\beta^i w \text{ and } wR_\beta^i z \text{ implies } vR_\beta^i z$$

hence, by definition of R_β^i,

$$\text{For all } \beta \geq \alpha, \quad x > 1 - \beta \text{ implies } \pi_{i,v}(z) > 1 - \beta$$

which can also be written as

$$\text{For all } \beta' \leq 1 - \alpha, \quad x > \beta' \text{ implies } \pi_{i,v}(z) > \beta' \tag{6}$$

Assume now $\pi_{i,v}(z) \leq 1 - \alpha$; we want to prove that $\pi_{i,v}(z) \geq x$. In fact, suppose $\pi_{i,v}(z) < x$; by density of rational numbers, there exists a β' such that

$\pi_{i,v}(z) < \beta' < x \leq 1-\alpha$. It is $\beta' \leq 1-\alpha$, hence, using 6, since $x > \beta'$, it is also $\pi_{i,v}(z) > \beta'$, contraddicting the hypothesis. Hence we have proved that if $\pi_{i,v}(z) \leq 1-\alpha$, then $\pi_{i,v}(z) \geq x$, i.e. that F is α-transitive.

For the reverse, assume $x > 1-\beta$, for some $\beta \geq \alpha$. By α-transitivity of the frame, if $\pi_{i,v}(z) \leq 1-\alpha$ it is $\pi_{i,v}(z) \geq x > 1-\beta$. If $\pi_{i,v}(z) > 1-\alpha$, we are done, since $1-\alpha \geq 1-\beta$.

4. Assume R^i_β euclidean for all i and for all $\beta \geq \alpha$; take $x =_{\text{def}} \min(\pi_{i,v}(w), \pi_{i,v}(z))$. For some fixed v, w, z, similarly to point 3, we have:

$$\text{For all } \beta \geq \alpha, \quad x > 1-\beta \text{ implies } \pi_{i,w}(z) > 1-\beta$$

Reasoning as in point 3, we conclude (using density of rationals) that if $\pi_{i,w}(z) \leq 1-\alpha$, then $\pi_{i,w}(z) \geq x$. Hence F is α-euclidean.

For the reverse, assume $x > 1-\beta$, for some $\beta \geq \alpha$. By α-euclidicity of the frame, if $\pi_{i,w}(z) \leq 1-\alpha$ it is $\pi_{i,w}(z) \geq x > 1-\beta$. If $\pi_{i,w}(z) > 1-\alpha$, we are done, since $1-\alpha \geq 1-\beta$. □

Lemma 10 Let Γ be any ψ-full set. Then

(i) For any ϕ in $Gen(\psi)$, ϕ is in Γ iff $\Gamma \vdash \phi$;

(ii) For any ϕ in $Gen(\psi)$, exactly one of ϕ and $\neg\phi$ is in Γ;

(iii) For any $(\phi_1 \wedge \phi_2)$ in $Gen(\psi)$, $(\phi_1 \wedge \phi_2)$ is in Γ iff both ϕ_1 and ϕ_2 are in Γ.

Proof First of all, note that $\phi \in Gen(\psi)$ implies that $\neg\phi \in Gen(\psi)$; and that $(\phi_1 \wedge \phi_2) \in Gen(\psi)$ implies that both ϕ_1 and ϕ_2 are in $Gen(\psi)$. The results follow easily from the fact that $\mathbf{PL}_n^{\otimes}$ includes propositional calculus by virtue of A0 and MP.

(i) is immediate. Suppose $\phi \in \Gamma$. By propositional reasoning, $\phi \vdash \phi$, and then $\Gamma \vdash \phi$. Vice-versa, suppose $\phi \notin \Gamma$. By maximality of Γ, $\Gamma \cup \{\phi\}$ is inconsistent, and then $\Gamma \not\vdash \phi$.

To prove (ii), note that ϕ and $\neg\phi$ cannot both be in Γ by consistency of Γ. We show that at least one of them must be in Γ. Suppose not. Then, from item (i), we have neither $\Gamma \vdash \phi$ nor $\Gamma \vdash \neg\phi$. But then both $\Gamma \cup \{\neg\phi\}$ and $\Gamma \cup \{\phi\}$ are consistent, and then we must have both $\phi \in \Gamma$ and $\neg\phi \in \Gamma$, contradicting the hypothesis that Γ is consistent.

As for (iii), assume that $(\phi_1 \wedge \phi_2) \in \Gamma$. We show that both ϕ_1 and ϕ_2 must be in Γ. For suppose that $\phi_1 \notin \Gamma$. Then, by item (ii), $\neg\phi_1 \in \Gamma$. But then Γ would be (propositionally) inconsistent, which is impossible. A similar argument goes for ϕ_2. To see the reverse, assume that both ϕ_1 and ϕ_2 are in Γ, and suppose that $(\phi_1 \wedge \phi_2) \notin \Gamma$. Then, again by item (ii), $\neg(\phi_1 \wedge \phi_2) \in \Gamma$, and then Γ would be inconsistent. □

Lemma 11 Let Γ be any consistent subset of $Gen(\psi)$. Then there is a ψ-full set Δ that contains Γ.

Proof Let r be the cardinality of Γ, and let $E = \{\phi_1, \phi_2, \dots, \phi_k\}$ be an enumeration of the formulae in $Gen(\psi)$ such that the first r formulae are all in Γ (we can always do this, as $\Gamma \subseteq Gen(\psi)$). Define the following sequence of sets:

$$\begin{array}{rcl} \Delta_1 & = & \{\phi_1\} \\ \Delta_{i+1} & = & \left\{ \begin{array}{ll} \Delta_i \cup \{\phi_i\} & \text{if } \Delta_i \cup \{\phi_i\} \text{ is consistent} \\ \Delta_i & \text{otherwise} \end{array} \right. \end{array}$$

and let $\Delta = \Delta_k$. As Γ is consistent, we must have $\Delta_r = \Gamma$. Moreover, as $\Delta_i \subseteq \Delta$ for all $1 \leq i \leq k$, we also have $\Delta \supseteq \Gamma$. To prove our thesis, we have to show that Δ is ψ-full. Now, obviously, Δ is a subset of $Gen(\psi)$. That Δ is consistent is also obvious, because Δ_1 is consistent, and all additions made to it are consistent. So, the only non obvious property is that Δ is maximal in $Gen(\psi)$. Suppose it is not. Then there would be a formula ϕ_s of $Gen(\psi)$ such that $\phi_s \notin \Delta$ and $\Delta \cup \{\phi_s\}$ is consistent. Then, also $\Delta_{s-1} \cup \{\phi_s\}$ is consistent, and so we must have $\Delta_s = \Delta_{s-1} \cup \{\phi_s\}$ by construction. But then ϕ_s would be in Δ, contradicting our supposition. Thus, there is no such ϕ_s, and Δ is ψ-full. □

Lemma 12 Let $\phi \in Gen(\psi)$ and $\Gamma \subseteq Gen(\psi)$. Then, $\Gamma \vdash \phi$ if and only if ϕ belongs to any ψ-full set that contains Γ.

Proof Left to right, assume that $\Gamma \vdash \phi$, and suppose that there is a ψ-full set Δ such that $\Gamma \subseteq \Delta$ but $\phi \notin \Delta$. By Lemma 9(ii), $\neg\phi \in \Delta$. Then, $\Gamma \cup \{\neg\phi\}$ is a subset of Δ, and hence consistent. But then $\Gamma \not\vdash \phi$, contradicting our assumption. So, there is no such Δ.

Right to left, suppose that $\Gamma \not\vdash \phi$. Then, $\Gamma \cup \{\neg\phi\}$ is consistent. As $\neg\phi$ is in $Gen(\psi)$, we can use Lemma 10 to conclude that there exists a ψ-full set Δ that contains both Γ and $\neg\phi$. But then, by Lemma 9(ii), $\phi \notin \Delta$, and still Δ contains Γ. □

We now go into the proof of the two main properties needed to prove the truth lemma (Lemma 13), namely, Lemma A.3 and Lemma A.6. It is convenient to introduce the following definitions:

Definition A.1 Let $I = \{1, \dots, n\}$ to represent a set of agents; for any agent $i \in I$ let $\Gamma_i = \{\mathbf{B}^i_{a_{i,j}} \phi_{i,j} \mid j \in J_i\}$ be the set of beliefs associated with that agent. Similarly, let $\Gamma_0 = \{\mathbf{D}_{a_j} \phi_j \mid j \in J_0\}$ be the set of distributed beliefs. Then define:

1. let $\mathbf{B}_a$ be a fixed arbitrary modality, *i.e.*$\mathbf{B}_a = \mathbf{B}^i_a$ or $\mathbf{B}_a = \mathbf{D}_a$. We will denote by $LCS(\Gamma)$ the solution of the constraints $\{(|\phi_j|^\psi, a_j) \mid \mathbf{B}_{a_j}\phi_j \in \Gamma\}$, and by $Lcs(\mathbf{B}_a\phi)$ the solution of the single constraint $(|\phi|^\psi, a)$. (Note that $LCS(\Gamma)$ and $Lcs(\mathbf{B}_a\phi)$ are both possibility functions on W^ψ)
2. $\Gamma = \Gamma_0 \cup \bigcup_{i \in I} \Gamma_i$.
3. Let $\pi^{\otimes} = \bigotimes_{i=1}^n LCS(\Gamma_i)$.
4. $\mathcal{K} = \{(i, j) \mid i \in I, j \in J_i\}$

5. for any $K \subseteq \mathcal{K}$,

$$\Phi_K = \neg \left(\bigwedge_{(i,j)\in K} \phi_{i,j} \ \wedge \bigwedge_{(i,j)\in \bar{K}} \neg\phi_{i,j} \right)$$

where $\bar{K} = \mathcal{K} \setminus K$.

6. for any $K \subseteq \mathcal{K}$, for any $i \in I$

$$\mathbf{B}^i_K = \max_{j:(i,j)\in \bar{K}} a_{i,j}$$

(where we stipulate that $\max_{j\in\emptyset} c_j = 0$)

7. for any $K \subseteq \mathcal{K}$,

$$b_K = 1 - \bigotimes_{i=1}^{n} (1 - \mathbf{B}^i_K)$$

By a property of T-conorms, we also have

$$b_K = \bigoplus_{i=1}^{n} \mathbf{B}^i_K$$

We state some useful facts about least committed possibilistic solutions:

1. $Lcs(\mathbf{B}^1_a(\phi \vee \chi) = Lcs(\mathbf{B}^1_a\phi) \vee Lcs(\mathbf{B}^1_a\chi)$
2. $LCS(\Gamma \cup \Delta) = LCS(\Gamma) \wedge LCS(\Delta)$

The following simple lemma will be useful in the following:

Lemma A.2 For any given π on the domain W^ψ and any formula ϕ, $N_\pi(|\phi|^\psi) \geq a$ iff $\pi \leq Lcs(\mathbf{B}^i_a\phi) \quad (= Lcs(\mathbf{D}_a\phi))$.

Proof Immediate:

$$N_\pi(|\phi|^\psi) \geq a \text{ iff } 1 - \bigvee_{w \notin |\phi|^\psi} \pi(w) \geq a \text{ iff } \bigvee_{w \notin |\phi|^\psi} \pi(w) \leq 1 - a \text{ iff } \pi \leq Lcs(\mathbf{B}^i_a\phi).$$

□

We can now prove the following property used in the proof of the truth lemma to deal with the $\mathbf{B}^i_a$ case.

Lemma A.3 Let $\mathbf{B}_a$ be a fixed arbibtrary modality, *i.e.*$\mathbf{B}_a = \mathbf{B}^i_a$ or $\mathbf{B}_a = \mathbf{D}_a$. Let $\Gamma = \{\mathbf{B}_{a_1}\phi_1, \ldots, \mathbf{B}_{a_m}\phi_m\}$ be any finite set of wwfs, and let $\mathbf{B}_a\phi$ be a wff. If $\pi^* = LCS(\Gamma)$ is such that $N^*(|\phi|^\psi) \geq a$, (where N^* is the necessity measure associated to π^*), then $\Gamma \vdash \mathbf{B}_a\phi$.

Proof Note that, by lemma A.2, the antecedent in the lemma can be replaced by: $LCS(\Gamma) \leq Lcs(\mathbf{B}_a\phi)$. Then we re-state the lemma as:

$$LCS(\Gamma) \leq Lcs(\mathbf{B}_a\phi) \text{ implies } \Gamma \vdash \mathbf{B}_a\phi$$

The proof is by induction on m (the number of contraints):

$m=1$ **:** We have then $Lcs(\mathbf{B}_{a_1}\phi_1) \le Lcs(\mathbf{B}_a\phi)$. There are just two cases in which this can happen:

case 1: $|\phi_1|^\psi \subseteq |\phi|^\psi$ and $a \le a_1$. Then by classical propositional completeness we have $\phi_1 \vdash \phi$, hence, in a few steps, $\mathbf{B}_{a_1}\phi_1 \vdash \mathbf{B}_a\phi$

case 2: $a = 0$. In this case by axiom A1, $\vdash \mathbf{B}_0\phi$.

$m>1$ **:** Let $\Gamma' = \Gamma \setminus \{\mathbf{B}_{a_n}\phi_n\}$. We have

$$LCS(\Gamma) = LCS(\Gamma') \wedge Lcs(\mathbf{B}_{a_n}\phi_n) = \begin{cases} LCS(\Gamma') & \text{iff } w \in |\phi_m|^\psi \\ LCS(\Gamma') \wedge (1-a_m) & \text{iff } w \notin |\phi_m|^\psi \end{cases}$$

while $Lcs(\mathbf{B}_a\phi)(w) = \begin{cases} 1 & \text{iff } w \in |\phi|^\psi \\ 1-a & \text{iff } w \notin |\phi|^\psi \end{cases}$

Since $LCS(\Gamma) \le Lcs(\mathbf{B}_a\phi)$, there are two cases:

- $\bigwedge_{j \le m-1, w \notin |\phi_j|^\psi}(1-a_j) \le 1-a$ for any w such that $w \notin |\phi_n|^\psi$ and $w \notin |\phi|^\psi$. In this case it happens that $LCS(\Gamma') \le Lcs(\mathbf{B}_a\phi)$. Then, by inductive hyp., we have $\Gamma' \vdash \mathbf{B}_a\phi$ and hence $\Gamma \vdash \mathbf{B}_a\phi$.
- If the first possibility does not hold, then it must be the case that $1-a_n \le 1-a$, *i.e.a* $\le a_n$, and hence $\mathbf{B}_{a_n}\phi_n \vdash \mathbf{B}_a\phi_n$; we also have (this is true in any case):

$$Lcs(\mathbf{B}_a(\phi \vee \neg\phi_n)) = Lcs(\mathbf{B}_a\phi) \vee Lcs(\mathbf{B}_a\neg\phi_n)$$

Now, we can see that $LCS(\Gamma') \le LCS(\mathbf{B}_a(\phi \vee \neg\phi_n))$; in fact:

  – if $w \in |\phi_n|^\psi$, it is:

$$LCS(\Gamma')(w) \le LCS(\Gamma)(w) \le Lcs(\mathbf{B}_a\phi) \le L(\mathbf{B}_a(\phi \vee \neg\phi_n))$$

  – if $w \notin |\phi_n|^\psi$, then

$$LCS(\Gamma') \le 1 = Lcs(\mathbf{B}_a\neg\phi_n) \le Lcs(\mathbf{B}_a(\phi \vee \neg\phi_n))$$

By inductive hypothesis we have $\Gamma' \vdash \mathbf{B}_a(\phi \vee \neg\phi_n)$, hence, since we also have $\mathbf{B}_{a_n}\phi_n \vdash \mathbf{B}_a\phi_n$, we can prove $\Gamma \vdash \mathbf{B}_a(\phi \vee \neg\phi_n) \wedge \mathbf{B}_a\phi_n$ and, from there, $\Gamma \vdash \mathbf{B}_a\phi$. □

Lemma A.4 Let $\pi^* = LCS(\{\mathbf{D}_{b_K}\Phi_K \mid K \subseteq \mathcal{K}\})$; then $\pi^\otimes = \pi^*$.

Proof The proof consists in a direct checking.

Consider that $\{|\neg\mathbf{B}_K|^\psi \mid K \in \mathcal{K}\}$ is a partition of the set W^ψ, and so for any $w \in W^\psi$ there is a unique $|\neg\mathbf{B}_{K_w}|^\psi$ which contains it. So, if the value of π^* on w only depends on the constraint $(|\mathbf{B}_{K_w}|^\psi, b_{K_w})$, since

$$\pi^*(w) = \bigwedge_{w \notin |\mathbf{B}_K|^\psi}(1-b_K) = \bigwedge_{w \in |\neg\mathbf{B}_K|^\psi}(1-b_K) = 1 - b_{K_w} =$$

$$\bigotimes_{i\in I}(1-\bigvee_{(i,j)\notin K_w} a_{i,j}) = \bigotimes_{i\in I}\bigwedge_{(i,j)\notin K_w}(1-a_{i,j})$$

where K_w is such that $w \in |\neg \mathbf{B}_{K_w}|^{\psi}$. Note that, by the definition of $\mathbf{B}_{K_w}$, $w \in |\neg \mathbf{B}_{K_w}|^{\psi}$ iff $w \in |\phi_{i,j}|^{\psi}$ for any $(i,j) \in K_w$ and $w \in \neg|\phi_{i,j}|^{\psi}$ for any $(i,j) \notin K_w$. Hence, $w \in |\phi_{i,j}|^{\psi}$ iff $(i,j) \in K_w$. This allows us to conclude that

$$\pi^*(w) = \bigotimes_{i\in I}\bigwedge_{w\notin|\phi_{i,j}|^{\psi}}(1-a_{i,j})$$

From the other side,

$$LCS(\Gamma_i)(w) = \bigwedge_{w\notin|\phi_{i,j}|^{\psi}}(1-a_{i,j})$$

and so

$$\pi^{\otimes}(w) = \bigotimes_{i\in I}\bigwedge_{w\notin|\phi_{i,j}|^{\psi}}(1-a_{i,j})$$

Since the two are identical, the proof is concluded. □

Lemma A.5 For any $K \subseteq \mathcal{K}$, $\Gamma \vdash \mathbf{D}_{b_K}\Phi_K$, hence $\Gamma \vdash \bigwedge_{K\subseteq\mathcal{K}} \mathbf{D}_{b_K}\Phi_K \wedge \Gamma_0$. (Where we use Γ_0 to actually denote the formula $\bigwedge_{\phi\in\Gamma_0}\phi$.)

Proof All through this proof we consider a given $K \subseteq \mathcal{K}$.

First, we show that $\Gamma \vdash \mathbf{B}^i_{\mathbf{B}^i_K}\Phi_K$ for any $i = 1, \dots, n$. Consider that, by applying De Morgan's laws,

$$\Phi_K = \bigvee_{(i,j)\in K}\neg\phi_{i,j} \ \vee \bigvee_{(i,j)\in\bar{K}}\phi_{i,j}.$$

We have, for any $(i,j) \in \mathcal{K}$,

$$\Gamma \vdash \mathbf{B}^i_{a_{i,j}}\phi_{i,j},$$

and this holds, in particular, for any $(i,j) \in \bar{K}$. Remember that, for any fixed i, we have defined $\mathbf{B}^i_K = \max_{j:(i,j)\in\bar{K}} a_{i,j}$; take then the element j_i such that $a_{i,j_i} = \mathbf{B}^i_K$, if there is one. Of course, we have

$$\Gamma \vdash \mathbf{B}^i_{a_{i,j_i}}\phi_{i,j_i},$$

and so, since $\phi_{i,j_i} \vdash \Phi_K$, by Nec we have

$$\mathbf{B}^i_1(\phi_{i,j_i} \vdash \Phi_K)$$

and by using A2 we conclude

$$\Gamma \vdash \mathbf{B}^i_{\mathbf{B}^i_K}\Phi_K, \tag{7}$$

If there are no j such that $(i,j) \in \bar{K}$, we have $\mathbf{B}^i_K = 0$, and axiom A1 allows to reach the same conclusion. From (7) we have

$$\Gamma \vdash \bigwedge_{i \in I} \mathbf{B}^i_{\mathbf{B}^i_K} \Phi_K,$$

from which,

$$\Gamma \vdash \mathbf{D}_{b_K} \Phi_K \qquad \text{with } b_K = \oplus_{i \in I} \mathbf{B}^i_K$$

follows by applying A4 (recall that 0 is the unit of $\oplus$).

Since this fact has been proved for an arbitrary $K \subseteq \mathcal{K}$, and, evidently, $\Gamma \vdash \Gamma_0$, we also have

$$\Gamma \vdash \bigwedge_{K \subseteq \mathcal{K}} \mathbf{D}_{b_K} \Phi_K \ \wedge \Gamma_0$$

which concludes our proof. □

We finally state the lemma which is needed in the truth lemma to deal with the $\mathbf{D}_a$ case.

Lemma A.6 Let Γ, π_0, $\pi^{\otimes}$ be as in definition A.1. If $\pi^* = \pi_0 \wedge \pi^{\otimes}$ is such that $N^*(|\phi|^{\psi}) \geq a$, then $\Gamma \vdash \mathbf{D}_a \phi$.

Proof Again, by lemma A.2 we can restate the lemma as

$$\pi_0 \wedge \pi^{\otimes} \leq Lcs(\mathbf{D}_a \phi) \text{ implies } \Gamma \vdash \mathbf{D}_a \phi$$

The proof is just a collection of the lemmas we have proved; it can be divided in four steps as follows:

1. Prove that $\pi^{\otimes}$ as defined above cprresponds to the solution of the set of contraints $\{(\mathbf{B}_K, b_K) \mid K \in \mathcal{K}\}$. This task is performed by lemma A.4

2. By lemma A.3, conclude that

$$\pi_0 \wedge \pi^{\otimes} \leq Lcs(\mathbf{D}_a \phi) \text{ implies } \{\mathbf{D}_{b_K} \mathbf{B}_K \mid K \in \mathcal{K}\} \cup \Gamma_0 \vdash \mathbf{D}_a \phi$$

3. Show that, for any $K \in \mathcal{K}$, $\Gamma \vdash \mathbf{D}_{b_K} \mathbf{B}_K$, hence

$$\Gamma \vdash \bigwedge_{K \in \mathcal{K}} \mathbf{D}_{b_K} \mathbf{B}_K \wedge \Gamma_0$$

   This task is performed by lemma A.5.

4. By the last two points, conclude that $\Gamma \vdash \mathbf{D}_a \phi$. □

Lemma 17 Let $\otimes$ be any continuous rational T-norm. Then, any non-theorem of $\mathbf{PL}^{\otimes}_n$ fails at a rational finite $\Pi^{\otimes}_n$-structure.

Proof The argument goes by noticing that the canonical model M^{ψ} for ψ has a finite set of worlds, and rational valued π_i; to see this second fact, consider that for any i and w, $\pi^{\psi}_{i,w}$ is built from the indexes of modal operators (which are rational) by using the operators max, min and $\oplus$, which all preserve rationality if $\otimes$ is rational. □