

A Depth-Bounded Semantics for Becoming Informed^{*}

Marco Larotonda¹ and Giuseppe Primiero²

¹Kube Partners

`marcolarotonda@kubepartners.com`

²Logic, Uncertainty, Computation and Information Group

Department of Philosophy, University of Milan

`giuseppe.primiero@unimi.it`

Abstract. We present the three-valued modal logic DBBL-BI_n to formally express information transmission among ordered agents bounded by limited access to repositories and where secrecy is admissible, viz. agents are not forced to transmit every data they possess to everyone else. The language, along with standard formulae for information holding at reachable states, includes formulae for agent and group information transmission, as well as assertion of trustworthy information. The description of information accessibility and transmission among agents is represented by formulae that hold in virtue of two distinct kinds of relations. We illustrate the application of the formal system with some intuitive examples.

Keywords: Logic of Information · Bounded Resources · Trustworthiness

1 Introduction

The development of reliable hybrid AI requires to integrate: (i) resource-bounded rationality, to account for agents who may have limited inferential or informational resources like humans; (ii) dynamic rationality, to model aids to knowledge and computational processes by externally received information; and (iii) models for trustworthy communication, in which information may be considered trustworthy if consensus among (a sufficiently large or relevant set of) sources is reached. A logic that models trustworthy communication within a resource-bounded context is in this light highly desirable. The present work aims at offering a semantics with these features.

Regarding (i), Depth-Bounded Boolean Logic (DBBL) [9] is a logic for single-agent reasoning characterized by an informational semantics that allows to distinguish between actual and virtual information. The former is information actually held by the agent; when an agent limits herself to actual information, she is said to reason at 0-depth. The latter is best explained proof-theoretically as

^{*} This research has been partially funded by the Project PRIN2020 BRIO (2020SSKZ7R) awarded by the Italian Ministry of University and Research (MUR).

the information that an agent might assume and then discharge for the derivation of new knowledge through an application of the rule of bivalence (RB).

$$\text{Rule of Bivalence } \frac{[\phi]^1 \quad [\neg\phi]^1}{\psi} \quad 1$$

When an agent employs k nested instances of RB, she is said to reason at depth k . At each k -depth a tractable inference relation is obtained, and the limit of this sequence is the classical entailment relation.

As for (ii), Multi-Agent Depth-Bounded Boolean Logic (MA-DBBL) [8] is an extension of DBBL modelling a multi-agent setting by shifting the interpretation of bound from the cognitive abilities of the agents to their ability to acquire information through the use of external resources. Under this interpretation, the depth k at which an agent is able to infer measures the number of distinct external sources that offer information necessary for the inference. Accordingly, MA-DBBL accounts for dynamic contexts where agents share information, via a modal operator of “becoming informed” inspired by [14, 15], simulating the epistemic action of a private announcement.

Finally (iii): the notion of trust has received great attention in order to reason about the security of a system [2, 16, 17]. MA-DBBL contains itself a policy of trust: an additional operator for “being informed” inspired by [3, 10] holds when an agent receives the same message by every other agent more informed than herself, thus expressing truthful information as content on which consensus between agents holds [20]. Hence an agent is truthfully informed only in case of trustworthy information. Such distinction seems crucial in contexts of information exchange where possible biases or disinformation campaigns are in place, also through artificial agents.

In this light, MA-DBBL seems well-placed to model systems where the above mentioned requirements are needed. But it suffers from a major limitation, since it allows no secrecy: every agent transmits every information she possesses. This might be a welcomed feature when the aim is to model communication in highly collaborative settings, but it is not realistic in many ordinary contexts. Moreover, MA-DBBL is developed only proof-theoretically. In this paper we present DBBL-BI _{n} , a variant logic equipped with a relational semantics that accounts for a multi-agent system where agents are ordered hierarchically and have access to increasingly extensive information states. In the tradition of role based access control theory [18], we intend such a hierarchy as defined among agents with shared competencies but with different degrees of access to sensitive or relevant information. Agents whose epistemic states do not allow to infer the truth value of given contents can obtain information externally and are free to share it or keep it private. Truthful and trustworthy information for an agent is characterised by content shared by all agents higher in the hierarchy. Given the condition imposed on the hierarchy, trustworthiness does not reduce to a form of democratic consensus. Indeed, when an agent evaluates whether she can trust

some information, she consider only the most reputable sources, i.e. those agents that are higher than herself with respect to the aforementioned hierarchy.

The paper is structured as follows. In Section 2 we draw some comparisons with some related works, in Section 3 we consider some scenarios that our logic models, in Section 4 we introduce the syntax and in Section 5 the semantics. In Section 6 we come back to the examples to show their formalization. Finally, in Section 7 we suggest some possible extensions of the logic.

2 Related Works

Dynamic Epistemic Logic (DEL) is one of the main frameworks that deal with information update. Within DEL, standard epistemic models are updated with action models in order to represent how knowledge changes when a certain action takes place [4]. Despite the success of DEL, there are three reasons why we chose a different framework.

First of all, DBBL-BI_n interprets states differently from DEL. In the latter case, states represent possible ways the world could be according to someone’s knowledge. For the former, states are distinct repositories containing pieces of information. Therefore, we are able to represent within our logic the authorizations that each agent possess for reading the content of a particular source.

The second reason regards how actions are described. Even though DEL is able to account for highly complex and structured actions via appropriate action models, they lack a relevant feature that we are interested in. Indeed, an action model is underspecified with respect to the agent (or group of agents) that is responsible for the occurrence of an action. On the contrary, we include formulae which make explicit the agent involved in the information-changing action. This is particularly relevant when trust assessments enter the picture: it does make a difference whether information is sent by a trusted or by an untrusted source.

The third point regards depth-boundedness. In the present work, differently from DBBL [9], virtual information is interpreted as information received by a distinct agent, and the depth of the reasoning is the distance between the receiver of a content and its original source. Therefore, depth measures how much a content has been shared among the agents. This is another parameter missing in DEL that could be important for trust assessment.

Trust is a central notion for the analysis of secure computational systems. In particular, the cognitive theory of Castelfranchi and Falcone [5], which analyses the notion of trust in terms of goals, beliefs, capabilities, and intentions, has received great attention from logicians. For example, [12] builds BDI-like [13] logical models of trust based on [5]. A distinct approach on trust is encompassed by [1, 11]. Those works are characterized by a “speaks for” modality which expresses delegation among agents. Moreover, [11] formalizes a “says” modality, and in this logic a content is trusted if it is said by an agent who is delegated via the “speaks for” modality. Finally, in [16] trust is conceived as a consistency checking function: an agent trusts an incoming message if it is consistent with her own profile; when this is not the case, two distinct policies may occur. On the

one hand, if the message was issued by a less reputable agent, then that content is distrusted, i.e. rejected. On the other hand, if the message was issued by a more reputable agent, then an operation of mistrust follows, leading to accepting the message and to a contraction of the reader's initial profile. In the present work, we conceive trust as agreement between a relevant set of agents: an agent needs to check if everyone more informed than herself agrees on a content in order to trust it.

3 Examples

We start by providing an example to illustrate the type of situations which our logic aims at modelling. Anne, Bob, and Charles have distinct authorizations to access three servers: s_1 , s_2 , and s_3 . Charles is authorized to read from all of them, Bob from all but s_3 , and Anne only from s_1 . This is reflected by an order $c < b < a$. Each of them is at liberty to share information with the others. In this context the transmission of information is represented by acquiring new authorizations to access more servers.

Example 1. Bob knows that $\neg(p \wedge u)$, since p and u cannot both be true at the same time but he does not know whether $\neg p$ or $\neg u$ is the case. This gap is filled by Charles, who decides to share information p and r stored on s_3 with Bob who can read these contents from its access to s_2 . As no one in the hierarchy above Bob disagrees on those formulae, he decides to write them on s_2 , and now Bob is in a position to determine a previously unnoticed facts, i.e. that $\neg u$ is true. Bob chooses to write p and r on different parts of s_2 based on whether he wants to share or not those contents.

Example 2. Bob decides to share p with Anne who can now read this content from its access to s_1 . But as long as she gets it only from Bob, she might be in doubt whether to trust it. But if Charles shares himself p from s_3 in a way that Anne can read it from s_1 , then every agent above Anne will have shared p with her directly, and now she trusts p , and writes that content on s_1 .

Example 3. Bob shares $\neg q$ from s_2 in a way that Anne can read it from s_1 . But Anne does not trust $\neg q$, because she does not receive it by every agent more informed than herself. Finally, Charles shares also r from s_3 in a way that Bob can read it from s_2 . Bob trusts r . He is therefore allowed to share r on its own, but he chooses to not transmit r to Anne, possibly because r is reserved information. Indeed, Bob gave to Anne only the authorization to access the part of s_2 containing $\neg q$, and not that with r .

4 Syntax

In the present section we specify the elements of the syntax. Firstly, we declare what are the symbols employed within the language, then we explain how formulae are built from those sets of symbols.

For the language of DBBL-BI_n we need the following objects: a finite set \mathcal{A} of indices for agents; a finite set \mathcal{S}^0 of indices for atomic informational states; a symbol for a composition function $+$ which takes as arguments a pair of states, and returns as output a state which is the composition of the two. \mathcal{R} is a set of relation symbols for accessibility relations: R_i is for single agent i accessibility among states and $R_{i,j}$ for information transmission from agent i to agent j . R_G and $R_{G,j}$ are the relations for accessibility and for transmissions from group of agents $G \subseteq \mathcal{A}$. \mathcal{P} is a finite set of propositional variables. Formation of complex formulae is closed under the set \mathcal{C} of classical connectives, and under the set \mathcal{K} of epistemic operators.

Definition 1 (Syntax of DBBL-BI_n). *The following sets constitutes the elements of the syntax.*

$$\begin{aligned}\mathcal{A} &= \{i, j, \dots, l\} \\ \mathcal{S}^0 &= \{s, s', s'', \dots, s^n\} \\ \mathcal{F} &= \{+\} \\ \mathcal{R} &= \{R_i, R_{i,j}, R_G, R_{G,j}\} \\ \mathcal{P} &= \{p, q, \dots, r\} \\ \mathcal{C} &= \{\wedge, \vee, \rightarrow, \neg\} \\ \mathcal{K} &= \{\diamond, BI_i, DBI_i, I_i\}\end{aligned}$$

Definition 2 (Language of DBBL-BI_n). *Formulae of the language of DBBL-BI_n are inductively defined by the following grammar in BNF:*

$$\begin{aligned}s : \phi_j &::= s : p_j \mid s : (\neg\phi)_j \mid s : (\phi \wedge \phi)_j \mid s : (\phi \vee \phi)_j \mid s : (\phi \rightarrow \phi)_j \mid \\ & \quad s : \diamond\phi_j \mid s : BI(\phi_i)_j \mid s : DBI(\phi_G)_j \mid s : I(\phi_i)_j \\ \tau &::= R_j(s, s') \mid R_{i,j}(s, s')\end{aligned}$$

Note that as it is standard in labelled logics [21], also relational formulae are introduced within the syntax. Labelled formulae of the form $s : \phi_j$ are read as “agent j has access to information ϕ at state s ”. Moreover, since labelled formulae could be indexed by groups of agents as well, $s : \phi_G$ are read as “group G distributively holds information ϕ at state s ”. In the following, to aid readability we rewrite respectively $s : BI(\phi_i)_j$ as $s : BI_j\phi_i$, $s : DBI(\phi_G)_j$ as $s : DBI_j\phi_G$ and $s : I(\phi_i)_j$ as $s : I_j\phi_i$.

Relational formulae express accessibility between states:

- we read $R_j(s, s')$ as “state s' is accessible from state s for agent j ”;
- we read $R_{i,j}(s', s)$ as “agent i gives the authorization to agent j to access state s' from state s ”.

Modal formulae allow to reason about processes of information access and transmission:

- we read $s : \Diamond\phi_j$ as “agent j can access from state s a state that contains information ϕ ”;
- we read $s : BI_j\phi_i$ as “agent j becomes informed at state s of information ϕ by agent i ”;
- we read $s : DBI_j\phi_G$ as “agent j becomes distributively informed at state s of information ϕ by group G ”;
- we read $s : I_j\phi$ as “agent j is informed at state s of information ϕ ”.

The meanings of the last three sentences have substantial differences. If $s : BI_j\phi_i$ holds, then there is a channel across which agent i makes available to agent j information ϕ . The *DBI* operator is analogous to the notion of distributed knowledge in standard epistemic logic: in this case a channel is established by possibly many agents in the group G for a single agent j to access information available to them. Note that in the case of *DBI* is the group G who possesses distributed information about ϕ . The difference between becoming informed and being informed lies in the degree of warranty that an agent has towards some information. If $s : BI_j\phi_i$ holds, then agent j has just received access to a piece of information ϕ from i at state s . If $s : I_j\phi$, additionally the agent possesses a sufficient amount of warranty that the information ϕ she received access to at state s is trustworthy in terms of the same information becoming accessible from all agents that stand in a certain relation with her, namely those higher in a shared hierarchy. For this reason, if agent j is informed of ϕ , then she might make that content available to other agents. This does not happen if $s : BI_j\phi_i$ but not $s : I_j\phi$.

5 Semantics

In the semantics the elements of the syntax are interpreted as agents, states, functions, and accessibility relations. We don't use another notation to distinguish the symbols of the syntax from their interpretation in the semantics in order to avoid unnecessary burden on the reader.

Definition 3 (Model). *A model for DBBL-BI_n is a tuple:*

$$\mathcal{M} = ((\mathcal{S}^0, +), \mathcal{A}, \{R_i\}_{i \in \mathcal{A}}, \{R_{i,j}\}_{\{i,j\} \subseteq \mathcal{A}}, \preceq, \mathcal{P}, v)$$

where:

- $+$ is a commutative, associative, and idempotent dyadic function such that:

$$\begin{aligned} \mathcal{S}^{n+1} &= \mathcal{S}^n \cup \{s' + s'' \mid s', s'' \in \mathcal{S}^n\}; \\ \mathcal{S} &= \bigcup_{n \in \mathbb{N}} \mathcal{S}^n; \end{aligned}$$

i.e. \mathcal{S}^0 is the set of atomic states; \mathcal{S}^{n+1} is the union of \mathcal{S}^n and of the set of all states composed from elements of \mathcal{S}^n ; \mathcal{S} is the union of all \mathcal{S}^n ;

– \mathcal{A} is a finite set of agents;

– $R_i \subseteq \mathcal{S} \times \mathcal{S}$ is a preorder such that:

if $(s, s' + s'') \in R_i$, then $(s' + s'', s') \in R_i$ and $(s' + s'', s'') \in R_i$;

if $(s, s') \in R_i$ and $(s, s'') \in R_i$, then $(s, s' + s'') \in R_i$;

i.e. if an agent is able to access a composite state $s' + s''$, then she can access also to its parts s' and s'' . Moreover, if an agent is able to access to two distinct states s' and s'' , then she can access also to their composition $s' + s''$.

– $R_{i,j} \subseteq \{(s', s) \mid (s', s') \in R_i, (s, s) \in R_j\}$ such that:

if $(s, s) \in R_i$, then $(s, s) \in R_{i,i}$;

i.e. $R_{i,j}$ satisfies a trivial condition for self-information: if an agent is authorized to access s , then she receives by herself all information stored at s .

– $\preceq \subseteq \mathcal{A} \times \mathcal{A}$ is a preorder;

– $\mathcal{P} = \{p, q, \dots, r\}$;

– $v : \mathcal{S} \mapsto (\mathcal{P} \rightarrow \{1, 0\})$ is the valuation function (with \rightarrow denoting a partial function) s.t. for all $s, s', s'' \in \mathcal{S}$:

- if $(p, 1) \in v(s)$, then $(p, 0) \notin v(s')$;
- if $(p, 0) \in v(s)$, then $(p, 1) \notin v(s')$;
- if $s' + s'' = s$, then $v(s') \subseteq v(s)$ and $v(s'') \subseteq v(s)$.

The function v associates to each state a partial valuation over \mathcal{P} . The valuation function satisfies three constraints. The first two impose monotonicity: if a proposition is true (resp. false) at some state, then it cannot be false (resp. true) elsewhere. Accordingly, the present work does not consider the transmission of contradictory information. The third condition is needed in order to correctly represent the fact that some states are part of other states.

Via these elements of the semantics, we are now in a position to define new objects: the hierarchy $\preceq \subseteq \mathcal{A} \times \mathcal{A}$, and the accessibility relations for groups of agents. We start with the former.

Definition 4. *Agent i is informed at least as much as agent j if and only if i has access to every state at which j has access: $i \preceq j$ iff $S_j \subseteq S_i$, where $S_i = \{s \in \mathcal{S} \mid R_i(s, s)\}$.*

Both R_i and $R_{i,j}$ could be extended for groups of agents. $(s, s') \in R_G$ means that group G is authorized to access state s' from s ; $(s, s') \in R_{G,j}$ means that the group G gives the authorization to j to access the composite set s from s' . The extension to groups of agents is obtained as follows.

Definition 5 (Accessibility relations for groups of agents).

$(s, s') \in R_G$ iff $(s, s') \in R_i$ for some $i \in G$;
 $(s_1 + \dots + s_n, s) \in R_{G,j}$ iff for every $1 \leq m \leq n$ there is $i \in G$ s.t. $(s_m, s) \in R_{i,j}$.

These definitions say that a group G can access a state if and only if at least one agent $i \in G$ can. Moreover, a group G gives the authorization to access a composite state if and only if every component of that state is made accessible to j by some agent $i \in G$. Note that when a group is formed by a single agent we treat $R_{\{i\},j}$ and $R_{i,j}$ as equivalent.

Below we use A as a meta-variable for both labelled and relational formulae.

5.1 Satisfiability relations

In this subsection we introduce the satisfiability relations. $\mathcal{M} \Vdash^k A$ means that model \mathcal{M} makes A true at depth k . Falsity is standard by negation. A model makes a labelled formula undetermined ($*$) just in case it makes it neither true nor false, i.e. $\mathcal{M} \Vdash_*^k s : \phi_i$ iff $\mathcal{M} \not\Vdash^k s : \phi_i$ and $\mathcal{M} \not\Vdash^k s : \neg\phi_i$, where $s : \phi_i$ means that i holds ϕ true at s , and $s : \neg\phi_i$ means that i holds $\neg\phi$ true at s , i.e. ϕ is false at s . When a model \mathcal{M} does not satisfy at a depth k either of these two formulae, agent i lacks any information about the truth-value of ϕ at s , remaining undetermined for her.

Recall that informally the depth at which a formula is validated is a parameter that measures the distance between the agent who evaluates a formula and the original source. Hence, for example, $\mathcal{M} \Vdash^k s : \phi_i$ means that ϕ is true at state s for agent i after that formula went through at most k many informational channels. We consider this depth as a meta-information not available to the agents, but known to the modeller. Nonetheless, agents are conscious of the lowest k -bound by counting the nested operators in formulae in which their index occurs as the outermost one for a BI operator. For example, if an agent h holds at some state that $BI_h BI_j BI_l p_i$, she knows that the distance between herself and the agent who issued p , i.e. agent i , is at least 3. Satisfaction of relational formulae is not qualified by a depth, since they do not express epistemic states of the agents but properties of the model.

In the following we employ two special function symbols, \mathcal{F}_\neg^v and \mathcal{F}_\bullet^v , with \bullet ranging over $\{\wedge, \vee, \rightarrow\}$. \mathcal{F}_\neg^v is the deterministic function that computes the truth-value of the negation of formulae given valuation v , and \mathcal{F}_\bullet^v is the non-deterministic function that computes the truth-value of formulae whose main connective is one of $\{\wedge, \vee, \rightarrow\}$ given valuation v . Those functions agree with the truth-tables of Table 1. Note that in the following clauses, s and s' ranges over the full set of states \mathcal{S} .

Definition 6 (Satisfaction of formulae).

1. $\mathcal{M} \Vdash R_i(s, s')$ iff $(s, s') \in R_i$
2. $\mathcal{M} \Vdash R_{i,j}(s, s')$ iff $(s, s') \in R_{i,j}$

Table 1. Informational truth-tables.

\wedge	1 0 *	\vee	1 0 *	\rightarrow	1 0 *	\neg	
1	1 0 *	1	1 1 1	1	1 0 *	1	0
0	0 0 0	0	1 0 *	0	1 1 1	0	1
*	* 0 *,0	*	1 * *,1	*	1 * *,1	*	*

3. $\mathcal{M} \Vdash^0 s : p_i$ iff $(p, 1) \in v(s)$ and $\mathcal{M} \Vdash R_i(s, s)$
4. $\mathcal{M} \Vdash^k s : \neg\phi_i$ iff $\mathcal{F}_-^v(s : \phi_i) = 1$
5. $\mathcal{M} \Vdash^k s : (s : \phi_i \bullet s : \psi_i)$ iff $\mathcal{F}_\bullet^v(s : \phi_i, s : \psi_i) = 1$ with $\bullet \in \{\wedge, \vee, \rightarrow\}$
6. $\mathcal{M} \Vdash^k s : \diamond\phi_j$ iff $\mathcal{M} \Vdash^k s' : \phi_j$ for some s' s.t. $\mathcal{M} \Vdash R_j(s, s')$
7. $\mathcal{M} \Vdash^{k+1} s : BI_j\phi_i$ iff $\mathcal{M} \Vdash^k s' : \phi_i$ for some s' s.t. $\mathcal{M} \Vdash R_{i,j}(s', s)$
8. $\mathcal{M} \Vdash^{k+1} s : DBI_j\phi_G$ iff $\mathcal{M} \Vdash^k s' : \phi_G$ for some s' s.t. $\mathcal{M} \Vdash R_{G,j}(s', s)$
9. $\mathcal{M} \Vdash^{k+1} s : I_j\phi_i$ iff $\mathcal{M} \Vdash^{k'+1} s : BI_j\phi_i$ for all (at least one) $i \prec j$, and $k' \leq k$

The formula $R_i(s, s')$ is true in a model \mathcal{M} if an access relation for agent i holds in \mathcal{M} from state s to state s' in \mathcal{S} . The semantic clause for $R_{i,j}$ is similar.

The formula $s : p_i$ is true at depth 0 in \mathcal{M} iff $(p, 1)$ is in the valuations at s and the agent i has access to s . The negation and other connectives are as by the Table 1.

Clause 6 is for the standard modal operator \diamond . Informally, if $\diamond\phi_i$ holds at s , then agent i has access to a state where ϕ_i holds.

Clause 7 introduces the *BI* operator. Agent j becomes informed at depth $k + 1$ and at state s of ϕ from agent i iff: at the lower depth k agent i gives the authorization to j to access a state s' where ϕ_i is true. By this definition and clause 4, the interpretation of $\mathcal{M} \Vdash^{k+1} s : \neg BI_j\phi_i$ is that $\mathcal{F}_-^v(s : BI_j\phi_i) = 1$ and this holds iff for all s' s.t. $\mathcal{M} \Vdash R_{i,j}(s', s)$, then $\mathcal{M} \Vdash^k s' : \neg\phi_i$ holds. The same reasoning holds for the other modal operators. Note that redundant and trivial information transmissions are allowed: an agent might become informed of a formula she already holds, and since $R_i(s, s)$ implies $R_{i,i}(s, s)$ (see Def. 3), then every agent becomes informed by herself of every formula she holds. Moreover, this clause accounts also for satisfaction of formulae with nested *BI* operators: for example $\mathcal{M} \Vdash^{k+2} s : BI_h BI_j p_i$ is satisfied when h becomes informed at s by j that j becomes informed by i that p . As before, an analogous reasoning holds for the other modal operators.

Clause 8 introduces the *DBI* operator for distributed becoming informed. This operator works as a closure under connectives for *BI* formulae. Suppose $s : BI_h\phi_i$ and $s : BI_h\psi_j$. It seems reasonable to hold also $s : BI_h(\phi \wedge \psi)_{i,j}$. However, the semantics of *BI* forbids this inference, because *BI* represents the transmission of information as a one-to-one relation between agents: exactly one agent is the access provider and exactly one other agent is the access recipient. On the contrary, *DBI* represents a many-to-one transmission of information between agents: there is exactly one agent who is the recipient of access authorization, but

there are possibly many providers in \mathcal{A} . In other words, an agent j is distributively informed of ϕ by a group G when G is distributively informed that ϕ is true, and ϕ is sent by G to j .

Finally, clause 9 says that an agent is informed that ϕ at s and at depth $k + 1$ iff she becomes informed at s and at a maximum depth $k + 1$ that ϕ by every agent (at least one) higher than herself in the hierarchy imposed by \preceq . For conceptual clarity, note that we assume that each agent is aware of this hierarchy.

Definition 7 (Structural conditions).

10. $\mathcal{M} \Vdash^k s : \phi_i$ implies $\mathcal{M} \Vdash^k s + s' : \phi_i$ (Composition)
11. $\mathcal{M} \Vdash^k s : \phi_i$ implies $\mathcal{M} \Vdash^k s : \phi_{i,j}$ (Grouping)
12. $\mathcal{M} \Vdash^k s : \phi_i$ implies $\mathcal{M} \Vdash^{k+1} s : \phi_i$ (Depth-Monotonicity)
13. $\mathcal{M} \Vdash^{k+1} s : I_j \phi_i$ implies $\mathcal{M} \Vdash^{k+1} s : \phi_j$ (Trust)
14. If $\mathcal{M} \Vdash^{k+1} s : I_j \phi_i$ for all ϕ_i s.t. there is $s' \mathcal{M} \Vdash^{k'} s' : \phi_i$ with $k' \leq k$,
and if $\mathcal{M} \Vdash R_{i,j}(s', s)$, and $\mathcal{M} \Vdash R_{j,h}(s, s'')$,
then $\mathcal{M} \Vdash R_{i,h}(s', s'')$ (New Channel)

The clause of state composition says that if an arbitrary ϕ_i is true at state s , then it is also true also at $s + s'$.


The grouping clause says that if agent i holds that ϕ , then also any group $\{i, j\}$ including i distributively holds that ϕ .

It is worth highlighting the importance of depth-monotonicity: if a formula is determined after at most k steps of information transmission, it remains determined even after $k + 1$ processes. What this conditions says is that the transmission of information is conservative (no information is lost), and that it is cumulative (the indeterminacy may be eventually reduced). Moreover, it produces a desirable side-effect: it permits to manipulate formulae that hold at different depths. For example, suppose $\mathcal{M} \Vdash^0 s : p_i$ and $\mathcal{M} \Vdash^1 s : q_i$. By Depth-Monotonicity, $\mathcal{M} \Vdash^1 s : p_i$, and finally $\mathcal{M} \Vdash^1 s : (p \wedge q)_i$. Without the help of Depth-Monotonicity, this kind of inference would require a more complex reasoning.

The I operator yields a policy of trust: when an agent is informed that ϕ then she can write within her state that ϕ .

Finally, clause 14 produces a kind of restricted transitivity for $R_{i,j}$ relations. It says that when an agent j is informed at s of every formula ϕ_i satisfied at a state s'' , then $R_{i,j}(s'', s)$ and $R_{j,h}(s, s')$ entail $R_{i,h}(s'', s')$. Informally, if there is a channel from agent i to j and one from j to h , and if j checked that every content from the former channel is trustworthy, then there is also an indirect channel from i to h . We give a simple example in order to make the idea clear.

In model \mathcal{M}_1 (see Figure 1) there are two channels represented by the following statements: $\mathcal{M}_1 \Vdash R_{i,j}(s_3, s_2)$ and $\mathcal{M}_1 \Vdash R_{j,h}(s_2, s_1)$ (these relations are in red in the model). Therefore, at s_2 agent j receives p from i , and at s_1 agent h receives $BI_j p_i$ from j . Hence, $\mathcal{M}_1 \Vdash^1 s_2 : BI_j p_i$ and $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_j p_i$. Suppose that we are interested in the satisfaction of the formula $s_1 : I_h BI_j p_i$. By clause 9, $\mathcal{M}_1 \Vdash^2 s_1 : I_h BI_x p_i$ iff $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_x p_i$ for all $x \prec h$, i.e. iff $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_j p_i$ and $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_i p_i$. $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_j p_i$ holds



Fig. 1. Model \mathcal{M}_1

just by clause 7. But clause 14 is needed for the satisfaction of $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_i p_i$. Indeed, in this case h must receive $BI_i p_i$ from i , and by clause 7 this means that model \mathcal{M}_1 should satisfy $R_{i,h}(s_3, s_1)$. We know that agent j considers s_3 a trusted source, i.e. $\mathcal{M}_1 \Vdash^1 s_2 : I_j \phi_i$ for every ϕ_i that holds at s_3 , and that s_2 is accessed by h via the authorization granted by j ($\mathcal{M}_1 \Vdash R_{j,h}(s_2, s_1)$). These conditions are sufficient to establish a new indirect informational channel from i to h through the mediation of j . Therefore, clause 14 entails $\mathcal{M}_1 \Vdash R_{i,h}(s_3, s_1)$. Now, both $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_j p_i$ and $\mathcal{M}_1 \Vdash^2 s_1 : BI_h BI_i p_i$ hold. Then it is also the case that $\mathcal{M} \Vdash^2 s_1 : I_h BI_j p_i$, and by *Trust* (clause 13) $\mathcal{M} \Vdash^2 s_1 : BI_h p_i$. This conclusion is perfectly consistent with the semantic clause for BI , since clause 14 entails $\mathcal{M}_1 \Vdash R_{i,h}(s_3, s_1)$.

Note that when two agents are unrelated, e.g. $j \not\preceq i$ and $i \not\preceq j$ there is no propagation of trust. Indeed, according to clause 14 this may occur only when a hierarchy can be established. Consider for example the variant model \mathcal{M}_{1b} (see Figure 2). In this example i and j are unrelated by \preceq . It is easy to check that in this case there is no propagation of trust from i to j to h as it occurs in \mathcal{M}_1 because there is no trust at all between j and i . In order to trust a formula issued by i , agent j needs to be lower than i in the hierarchy imposed by \prec , i.e. it is required that $i \prec j$. Since they are unrelated, j is not able to trust any information coming from agent i , i.e. she is not able to infer any formula $s_2 : I_j \phi_i$, for any $\mathcal{M}_{1b} \Vdash^k s_3 : \phi_i$.

6 Back to the Example

We now provide a detailed analysis of Model \mathcal{M}_2 which represents the transmission of information between Anne, Bob, and Charles as for the Examples in Section 3, see Figure 3. Dashed lines in the model represent state composition, e.g. $s_{2.1}$ and $s_{2.2}$ jointly compose s_2 .

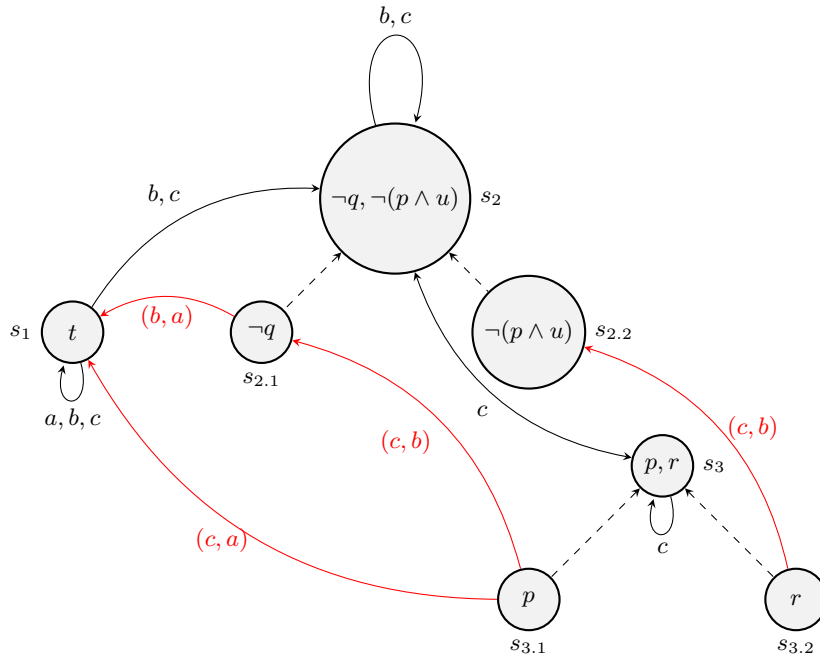

Fig. 2. Model \mathcal{M}_{1b}

Consider Example 1. Bob knows that $\neg(p \wedge u)$, but he does not know whether $\neg p$ or $\neg u$ is the case (this fact is legitimate since propositional connectives have a non-deterministic semantics). Charles gives the authorization to Bob to access s_3 from his access to s_2 , thus receiving both p and r . However, Bob has distinct plans for those two pieces of information: he is prepared to share p , but not r . Therefore, he decides to store incoming information on different parts of s_2 . This is reflected by the satisfaction of the following relational formulae:

- $\mathcal{M}_2 \Vdash R_{c,b}(s_{3.2}, s_{2.2})$, says that information stored at $s_{3.2}$ i.e. r is made available for access at $s_{2.2}$;
- $\mathcal{M}_2 \Vdash R_{c,b}(s_{3.2}, s_{2.1})$, says that information stored at $s_{3.1}$ i.e. p is made available for access at $s_{2.1}$.

Since p is information that Charles owns on its own, then it holds at depth 0 for him, i.e. $\mathcal{M}_2 \Vdash^0 s_{3.1} : p_c$. Bob receives that formula at depth 1: $\mathcal{M}_2 \Vdash^1 s_{2.1} : BI_b p_c$. As Bob receives p from Charles, and there is no one else in the hierarchy above Bob who disagrees with p , then Bob is informed of p at depth 1: $\mathcal{M}_2 \Vdash^1 s_2 : I_b p$. Now Bob satisfies the constraint to trust p , therefore he writes that content on s_2 . Hence, $\mathcal{M}_2 \Vdash^1 s_2 : p_b$. An analogous analysis holds with respect to r . Before receiving information from Charles, Bob knew that $\neg(p \wedge u)$ but he lacked any information about the truth-value of $\neg p$ and $\neg u$, i.e. $\mathcal{M}_2 \Vdash^0 s_2 : \neg(p \wedge u)_b$, $\mathcal{M}_2 \Vdash^0_* s_2 : \neg p_b$, and $\mathcal{M}_2 \Vdash^0_* s_2 : \neg u_b$. But at depth 1 Bob is able to fill these truth-value gaps concluding that $\neg u$ is the case: $\mathcal{M}_2 \Vdash^1 s_2 : \neg u_b$

Consider now Example 2. After having written p on $s_{2.1}$, Bob shares that content with Anne. Therefore, $\mathcal{M}_2 \Vdash R_{b,a}(s_{2.1}, s_1)$. Since the original source of p is Charles, it means that Anne receives p after it went through 2 channels (the first is from Charles to Bob, and the second from Bob to Anne). Therefore: $\mathcal{M}_2 \Vdash^2 s_1 : BI_a p_b$. Additionally, Anne receives p at depth 1 from Charles, i.e. $\mathcal{M}_2 \Vdash R_{c,a}(s_{3.1}, s_1)$. Then, every agent above Anne has shared p with her. We


Fig. 3. Model \mathcal{M}_2

can conclude that Anne is informed that p at depth 2 and that she trusts p at the same depth: i.e. $\mathcal{M}_2 \Vdash^2 s_1 : I_a p$ and $\mathcal{M}_2 \Vdash^2 s_1 : p_a$.

Examples 3 is now straightforward. Anne receives $\neg q$ from Bob, but not from Charles. This is reflected in the relational formulae satisfied by the model: $\mathcal{M}_2 \Vdash R_{b,a}(s_{2.1}, s_1)$ and $\mathcal{M}_2 \not\Vdash R_{c,a}(s_{3.1}, s_1)$. For this reason, $\mathcal{M}_2 \not\Vdash^1 s_1 : BI_a p_c$, and therefore Anne is not able to trust p .

Finally, we highlight two more facts. The first is about *DBI*. At s_1 , Anne receives $\neg q$ from Bob, and p by Charles. So, $\mathcal{M}_2 \Vdash^1 s_1 : BI_a \neg q_b$ and $\mathcal{M}_2 \Vdash^1 s_1 : BI_a p_c$. Thanks to the structural rules and by the definition of transmission by a group (Def. 5), we infer that the group formed by Bob and Charles has distributed information that $(p \wedge \neg q)$, and they transmit that information from the composite state $s_{2.1} + s_{3.1}$ to Anne, who can read that information at s_1 : $\mathcal{M}_2 \Vdash R_{\{b,c\},a}(s_{2.1} + s_{3.1}, s_1)$. Hence, Anne receives distributed information at s_1 that $(p \wedge \neg q)$ by Bob and Charles: $\mathcal{M}_2 \Vdash^1 s_1 : DBI_a(p \wedge \neg q)_{b,c}$. As for the second fact, note that information might or might not be preserved at reachable states. For example, Bob can read t at s_1 , but that information is not preserved when Bob reaches s_2 . Therefore, s_1 might be a state that is only temporarily accessible, and Bob loses the authorization to read into s_1 when he access to s_2 . On the contrary, Charles does not lose any piece of information when he reaches s_3 from s_2 , because in this case the accessibility relation from these two states is symmetric for Charles.

7 Conclusions and Future Works

DBBL-BI_n models information transmission by agents through access authorization to parts of available memory and preserving secrecy when required. Agents can receive private communications from an agent, or from a group of agents by the operator of becoming informed (*BI*), and distributed becoming informed (*DBI*) respectively. The *I* operator models a policy of trust: when the same information is received by all agents with more access, the receiver is safe in trusting the message.

Several extensions are foreseen. Firstly, DBBL-BI_n has an appropriate proof-theory formulated in natural deduction style, and standard soundness and completeness results. These results are not included here for reasons of space. Since the aim of depth-bounded logics is to account for computationally tractable consequence relations, then it is desirable to study computational complexity for DBBL-BI_n, devising a decision procedure working in polynomial time.

Secondly, DBBL-BI_n can be extended with an additional parameter expressing different degrees of inferential ability as standardly understood in DBBL [9], to complement the measure of the distance between source and receiver presented here.

Finally, we aim at enriching DBBL-BI_n with a suitable way to compute trustworthiness by means of a threshold function and of degrees of beliefs as in [2, 6, 7, 19]. Moreover, it would be desirable to model updates with contradictory information, and have a method to eliminate inconsistencies by means of operations of negative trust as in [16].

References

1. Abadi, M., Burrows, M., Lampson, B., Plotkin, G.: A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems (TOPLAS)* **15**(4), 706–734 (1993)
2. Aldini, A., Tagliaferri, M.: Logics to reason formally about trust computation and manipulation. In: *International Workshop on Emerging Technologies for Authorization and Authentication*. pp. 1–15. Springer (2019)
3. Allo, P.: The logic of ‘being informed’ revisited and revised. *Philosophical Studies* **153**(3), 417–434 (2011)
4. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: *Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*. pp. 43–56 (1998)
5. Castelfranchi, C., Falcone, R.: *Trust theory: A socio-cognitive and computational model*. John Wiley & Sons (2010)
6. Ceolin, D., Primiero, G.: A granular approach to source trustworthiness for negative trust assessment. In: *IFIP International Conference on Trust Management*. pp. 108–121. Springer (2019)
7. Chen, T., Primiero, G., Raimondi, F., Rungta, N.: A computationally grounded, weighted doxastic logic. *Stud Logica* **104**(4), 679–703 (2016). <https://doi.org/10.1007/s11225-015-9621-4>, <https://doi.org/10.1007/s11225-015-9621-4>

8. Cignarale, G., Primiero, G.: A multi-agent depth bounded boolean logic. In: International Conference on Software Engineering and Formal Methods. pp. 176–191. Springer (2020)
9. D’Agostino, M.: An informational view of classical logic. *Theoretical Computer Science* **606**, 79–97 (2015)
10. Floridi, L.: The logic of being informed. *Logique et analyse* pp. 433–460 (2006)
11. Genovese, V., Rispoli, D., Gabbay, D.M., Van Der Torre, L.: Modal access control logic-axiomatization, semantics and fol theorem proving. In: STAIRS 2010, pp. 114–126. IOS Press (2010)
12. Herzig, A., Lorini, E., Hübner, J.F., Vercouter, L.: A logic of trust and reputation. *Logic Journal of the IGPL* **18**(1), 214–244 (2010)
13. Meyer, J.J., Broersen, J., Herzig, A.: Bdi logics. In: van Ditmarsch, H., Halpern, J., van der Hoek, W., Kooi, B., et al. (eds.) *Handbook of Epistemic Logic*. College Publications (2015)
14. Primiero, G.: An epistemic constructive definition of information. *Logique et Analyse* pp. 391–416 (2007)
15. Primiero, G.: An epistemic logic for becoming informed. *Synthese* **167**(2), 363–389 (2009)
16. Primiero, G.: A logic of negative trust. *J. Appl. Non Class. Logics* **30**(3), 193–222 (2020). <https://doi.org/10.1080/11663081.2020.1789404>, <https://doi.org/10.1080/11663081.2020.1789404>
17. Primiero, G., Raimondi, F.: A typed natural deduction calculus to reason about secure trust. In: 2014 Twelfth Annual International Conference on Privacy, Security and Trust. pp. 379–382. IEEE (2014)
18. Sandhu, R.S., Ferraiolo, D.F., Kuhn, D.R.: The NIST model for role-based access control: towards a unified standard. In: Rebersburg, K., Youman, C.E., Atluri, V. (eds.) *Fifth ACM Workshop on Role-Based Access Control, RBAC 2000*, Berlin, Germany, July 26-27, 2000. pp. 47–63. ACM (2000). <https://doi.org/10.1145/344287.344301>, <https://doi.org/10.1145/344287.344301>
19. Termine, A., Primiero, G., D’Asaro, F.A.: Modelling accuracy and trustworthiness of explaining agents. In: Ghosh, S., Icard, T. (eds.) *Logic, Rationality, and Interaction - 8th International Workshop, LORI 2021, Xi’an, China, October 16-18, 2021, Proceedings*. *Lecture Notes in Computer Science*, vol. 13039, pp. 232–245. Springer (2021). https://doi.org/10.1007/978-3-030-88708-7_19, https://doi.org/10.1007/978-3-030-88708-7_19
20. Van Ditmarsch, H., van Der Hoek, W., Kooi, B.: *Dynamic epistemic logic*, vol. 337. Springer Science & Business Media (2007)
21. Viganò, L.: *Labelled Non-Classical Logics*. Springer, Boston (2000)