



HAL
open science

A Computationally Grounded Logic of Graded Belief

Emiliano Lorini, François Schwarzentruber

► **To cite this version:**

Emiliano Lorini, François Schwarzentruber. A Computationally Grounded Logic of Graded Belief. 17th European Conference on Logics in Artificial Intelligence (JELIA 2021), May 2021, virtual event, France. pp.245-261, 10.1007/978-3-030-75775-5_17 . hal-03453909

HAL Id: hal-03453909

<https://hal.science/hal-03453909>

Submitted on 30 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Computationally Grounded Logic of Graded Belief

Emiliano Lorini¹ and François Schwarzentruber²

¹ IRIT-CNRS, Toulouse University, France

²ENS Rennes, France

Abstract. We present a logic of graded beliefs with a formal semantics grounded on the notion of belief base. It has modal operators which represent what an agent would believe if she removed k pieces of information from her belief base. We provide a sound and complete axiomatics for our logic as well as an optimal model checking algorithm. To illustrate its expressive power, we apply it to modeling social influence and epistemic explanation.

1 Introduction

Epistemic logic (in the broad sense) captures epistemic attitudes including knowledge and belief. It has been extensively studied by philosophers [27], computer scientists [18,39] and economists [33] and applied to a variety of fields of AI including security protocols [22,9], blockchain protocol [22,38] and epistemic planning [6]. Its language extends that of propositional logic by a modal operator for expressing an agent's knowledge or belief. Multi-agent extensions of epistemic logic have been studied in which modal operators are parameterized by agent names identifying the knower (or the believer) in the system.

Several extensions and variants of epistemic logic dealing with the notion of graded belief have been proposed. This includes logics of probabilistic beliefs [17,50,31] as well as logics of graded belief based on a qualitative or semi-qualitative notion of plausibility [1,32,49,3]. As pointed out by [40,51], while in quantitative approaches belief states are represented by classical probabilistic measures or by alternative numerical accounts, such as lexicographic probabilities or conditional probabilities, in a semi-qualitative setting, such as rank-based systems [44] and possibility theory [16], belief states are represented by qualitative measures assigning orders of magnitude. Finally, qualitative approaches employ a plausibility ordering (also called epistemic entrenchment ordering) on possible worlds. Other approaches use graded modalities whereby the degree of a belief is a function of the number of worlds in which the believed formula is true [48,47,8], or of the number of evidences which support it [2]. All these approaches to graded belief use Kripke semantics in which agents' epistemic states are modeled via accessibility relations over possible worlds.

Kripke semantics have severe limitations in modeling knowledge. First, the number of possible worlds is huge in real applications: for instance, in a classical card game with four players having each 8 cards among a set of 32 cards, the number of possibilities is $\binom{32}{8} \times \binom{24}{8} \times \binom{16}{8} = 9.95 \times 10^{16}$ possible worlds. For this reason, although the model checking problem of epistemic logic based on Kripke semantics is in polynomial time,

in many applications, the model cannot be explicitly constructed. Thus, it is hard to implement graded beliefs through Kripke semantics.

Kripke models and their corresponding accessibility relations can be compactly represented by means of BDDs [46], by Boolean formulas and programs [11,12,13,10] or by the notion of visibility [25]. In these approaches, the size of the models scales up, although the corresponding symbolic model checking is PSPACE-complete. However, they capture restricted notions of belief and knowledge and the formal languages they use to succinctly represent the Kripke model are of different nature than the standard epistemic language. This makes it difficult to define the notion of graded belief.

In this paper, we present a novel account of graded belief exploiting the notion of belief base. The belief base approach to knowledge representation is well-established since at least 30 years in a single-agent setting [30,23,42,5]. In this approach, an agent has non-deductively closed explicit beliefs in her belief base and we call implicit belief a statement that can be inferred from her explicit beliefs.

As we aim at a multi-agent setting, we rely on the idea of using belief bases as a semantics for multi-agent epistemic logic which was put forth in [34] and developed in a series of papers with the aim of capturing multi-agent belief dynamics [36,37] and higher-order epistemic reasoning [35], and of elucidating the connection between distributed belief and belief merging [26]. We enrich the logic presented in [34] with a notion of graded belief. It is expressed by modal operators of the form \Box_i^k , where i is an agent and k is a positive integer capturing the agent's strength of belief. At the semantic level, such operators are interpreted via graded epistemic accessibility relations of type \mathcal{R}_i^k , one per $k \in \mathbb{N}$. Such a relation specifies the set of states that agent i considers possible after having removed at most k pieces of information from her belief base. This means that the higher the value of k of a given state for agent i , the higher the degree of plausibility of that state for agent i .¹ Indeed, states with a high degree of plausibility are states which satisfy a large number of information in the agent's belief base. In the extreme case, maximally plausible states for the agent are states to which value 0 is assigned, since they satisfy all information in the agent's belief base. This is a crucial aspect of our approach which distinguishes it from the standard extensional Kripke-style semantics for epistemic logic. While in the standard semantics an agent's plausibility ordering or the corresponding plausibility measure over states is given as a primitive (see, e.g., [3,1,32]), in our approach they are computed from and grounded on the agents' belief base. The latter provides an advantages for formal verification since the model checking problem is formulated in our logic in a more compact way than in existing logics of graded belief. Also, from the conceptual point view, our logic offers a minimalistic approach to graded belief in which the only primitive concept is belief base, while the concept of graded implicit belief is derived from it.

Our approach provides a succinct semantics of epistemic states exploiting belief bases as well as a conservative extension of standard epistemic logic by the notions of explicit belief and graded belief.

The paper is organized as follows. In Section 2, we introduce the language of our multi-agent epistemic logic of explicit belief and graded implicit belief. Section 3 presents its formal semantics exploiting belief bases. Section 4 presents the first appli-

¹ This is in line with the theory of qualitative uncertainty by [44].

cation of our logical framework to modeling the concept of social influence. Section 5 is the core part of the paper and provides an axiomatics proven to be sound and complete relative to the belief base semantics. In Section 6, we extend the base logic by conditional belief operators and generalize the completeness result to it. Section 7 presents a model checking algorithm for the base logic and its extension. Section 8 presents the second application of our logic: we illustrate its expressive power to account for a variety of notions of epistemic explanation. Finally, in Section 9, we conclude. ²

2 Graded doxastic language

Assume a countably infinite set of atomic propositions Atm and a finite set of agents $Agt = \{1, \dots, n\}$. We define the language for representing agents' explicit beliefs and agents' graded implicit beliefs in two steps. First, we define the language $\mathcal{L}_0(Atm, Agt)$ for representing agents' explicit beliefs by the grammar:

$$\alpha ::= p \mid \neg\alpha \mid \alpha_1 \wedge \alpha_2 \mid \Delta_i\alpha,$$

where p ranges over Atm and i ranges over Agt . The formula $\Delta_i\alpha$ is read “agent i explicitly believes that α ”. Second, the language $\mathcal{L}(Atm, Agt)$ extends the language $\mathcal{L}_0(Atm, Agt)$ by graded implicit belief operators. It is defined by:

$$\varphi ::= \alpha \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_i^k\varphi,$$

where α ranges over $\mathcal{L}_0(Atm, Agt)$, i ranges over Agt and k ranges over \mathbb{N} . For notational convenience we write \mathcal{L}_0 instead of $\mathcal{L}_0(Atm, Agt)$ and \mathcal{L} instead of $\mathcal{L}(Atm, Agt)$, when the context is unambiguous. The other Boolean constructions \top , \perp , \rightarrow and \leftrightarrow are defined in the standard way. For every formula $\varphi \in \mathcal{L}$, we write $Atm(\varphi)$ to denote the set of atomic propositions of type p occurring in φ . Moreover, for every set of formulas $X \subseteq \mathcal{L}$, we define $Atm(X) = \bigcup_{\varphi \in X} Atm(\varphi)$.

The formula $\Box_i^k\varphi$ is read “agent i would implicitly believe φ , for every removal of at most k pieces of information from her belief base”. The value k can also be conceived as the extent to which agent i believes that φ . Indeed, the higher the number of information in the belief base that can be removed without affecting the belief, the stronger the belief. Thus, $\Box_i^k\varphi$ can also be read “agent i believes that φ with degree (or strength) at least k ”. The abbreviation $\Diamond_i^k\varphi \stackrel{\text{def}}{=} \neg\Box_i^k\neg\varphi$ defines the concept of belief compatibility. The formula $\Diamond_i^k\varphi$ has to be read “ φ would be compatible with agent i 's explicit beliefs, for some removal of at most k pieces of information from her belief base”.

Example 1. Let us discuss the informal meaning of some formulas by means of an example of a single robot i exploring an area. Formula $\Delta_i\text{fire}$ says that she explicitly believes there is fire. We could also have $\Delta_i(\text{fire} \rightarrow \text{danger})$. Thus, we would have $\Box_i^0\text{danger}$. If now we also have $\Delta_i\text{radiation}$ and $\Delta_i(\text{radiation} \rightarrow \text{danger})$, we would have $\Box_i^1\text{danger}$. In words, the agent would believe that there is a danger at strength 1, because she would still infer danger even if at most one explicit belief is removed.

² Detailed proofs of the results of the paper are given at the following address: <http://people.irisa.fr/Francois.Schwarzentruber/research/jelia2021/jelia2021BelieBases.pdf>

We denote by \mathcal{L}^- the fragment of language \mathcal{L} obtained by the rule: $\varphi ::= \alpha \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_i^0\varphi$. The abbreviations $\Box_i\varphi \stackrel{\text{def}}{=} \Box_i^0\varphi$ and $\Diamond_i\varphi \stackrel{\text{def}}{=} \Diamond_i^0\varphi$ define the concepts of ungraded implicit belief and belief compatibility. Moreover, the abbreviation $\Box_i^k\varphi \stackrel{\text{def}}{=} \Box_i^k\varphi \wedge \neg\Box_i^{k+1}\varphi$ has to be read “agent i believes that φ with degree (or strength) equal to k ”.

3 Belief base semantics

Following [34,36], we now present a formal semantics for the language \mathcal{L} exploiting belief bases. Unlike the standard Kripke semantics for epistemic logic in which the notions of epistemic alternative and plausibility of a world (or state) are given as primitive, in this semantics they are defined from the primitive concept of belief base.

Definition 1 (State). *A state is a tuple $B = (B_1, \dots, B_n, S)$ where for every $i \in \text{Agt}$, $B_i \subseteq \mathcal{L}_0$ is agent i 's finite belief base, and $S \subseteq \text{Atm}$ is the actual environment. The set of all states is denoted by \mathbf{S} .*

The sublanguage $\mathcal{L}_0(\text{Atm}, \text{Agt})$ is interpreted with respect to states, as follows.

Definition 2 (Satisfaction relation). *Let $B = (B_1, \dots, B_n, S) \in \mathbf{S}$ be a state. Then:*

$$\begin{aligned} B \models p &\iff p \in S, \\ B \models \neg\alpha &\iff B \not\models \alpha, \\ B \models \alpha_1 \wedge \alpha_2 &\iff B \models \alpha_1 \text{ and } B \models \alpha_2, \\ B \models \Delta_i\alpha &\iff \alpha \in B_i. \end{aligned}$$

Observe in particular the set-theoretic interpretation of the explicit belief operator: agent i explicitly believes that α if and only if α is included in her belief base.

It is also worth considering belief correct states, according to which every fact explicitly believed by an agent is true.

Definition 3 (Belief correct state). *A state $B = (B_1, \dots, B_n, S)$ is belief correct if and only if, for every agent $i \in \text{Agt}$ and for every $\alpha \in B_i$, we have then $B \models \alpha$. The set of all belief correct states is denoted by \mathbf{S}_{BC} .*

A multi-agent belief model (MAB) is defined to be a state supplemented with a set of states, called *context*. The latter includes all states compatible with the common ground [45], i.e., the body of information that the agents commonly believe to be the case.

Definition 4 (Multi-agent belief model). *A multi-agent belief model (MAB) is a pair (B, Ctx) , where $B \in \mathbf{S}$ and $\text{Ctx} \subseteq \mathbf{S}$. The class of MABs is denoted by \mathbf{M} .*

Note that in Definition 4 we do not require $B \in \text{Ctx}$. The following definition introduces the notion of graded doxastic alternative.

Definition 5 (Graded doxastic alternatives). Let $i \in \text{Agt}$ and let $k \in \mathbb{N}$. Then, \mathcal{R}_i^k is the binary relation on the set \mathbf{S} such that, for all $B = (B_1, \dots, B_n, S), B' = (B'_1, \dots, B'_n, S') \in \mathbf{S}$:

$$B\mathcal{R}_i^k B' \text{ if and only if } |\{\alpha \in B_i : B' \models \alpha\}| \geq (|B_i| - k),$$

$B\mathcal{R}_i^k B'$ means that B' is a k -level doxastic alternative for agent i at B , that is to say, B' is a state that at B agent i considers possible after having removed at most k pieces of information from her belief base. Graded doxastic accessibility relations induce a plausibility ordering over states, as in [44,32].³ For notational convenience, we write \mathcal{R}_i instead of \mathcal{R}_i^0 . Clearly, $B\mathcal{R}_i B'$ if and only if $B' \models \alpha$, for every $\alpha \in B_i$.

The following definition extends Definition 2 to the full language \mathcal{L} . Its formulas are interpreted with respect to MABs. (We omit Boolean cases, as they are defined in the usual way.)

Definition 6 (Satisfaction relation (cont.)). Let $(B, \text{Cxt}) \in \mathbf{M}$. Then:

$$\begin{aligned} (B, \text{Cxt}) \models \alpha &\iff B \models \alpha, \\ (B, \text{Cxt}) \models \Box_i^k \varphi &\iff \forall B' \in \text{Cxt} : \text{if } B\mathcal{R}_i^k B' \text{ then } (B', \text{Cxt}) \models \varphi. \end{aligned}$$

We consider the subclass of MABs that guarantee correctness of the agents' beliefs.

Definition 7 (Belief correct MAB). The MAB (B, Cxt) is belief correct (BC) if and only if $B \in \text{Cxt}$ and, for every $i \in \text{Agt}$ and for every $B' \in \text{Cxt}$, $B'\mathcal{R}_i B'$. The class of MABs satisfying BC is denoted by \mathbf{M}_{BC} .

Saying that (B, Cxt) satisfies BC is the same thing as saying that $B \in \text{Cxt}$ and, for every $i \in \text{Agt}$, the relation $\mathcal{R}_i \cap (\text{Cxt} \times \text{Cxt})$ is reflexive. The condition $B \in \text{Cxt}$ in Definition 7 is necessary to make the agents' implicit beliefs correct, i.e., to make the formula $\Box_i \varphi \rightarrow \varphi$ valid.

As the following proposition highlights, belief correctness for MABs is completely characterized by the fact that the actual world is included in the agents' common ground and that the agents' explicit beliefs are correct in the sense of Definition 3.

Proposition 1. A MAB (B, Cxt) satisfies BC if and only if $B \in \text{Cxt}$ and $\text{Cxt} \subseteq \mathbf{S}_{BC}$.

Let $\varphi \in \mathcal{L}$. We say that φ is valid relative to the class \mathbf{M} (resp. \mathbf{M}_{BC}), denoted by $\models_{\mathbf{M}} \varphi$ (resp. $\models_{\mathbf{M}_{BC}} \varphi$), if and only if, for every $(B, \text{Cxt}) \in \mathbf{M}$ (resp. $(B, \text{Cxt}) \in \mathbf{M}_{BC}$) we have $(B, \text{Cxt}) \models \varphi$. We say that φ is satisfiable for the class \mathbf{M} (resp. \mathbf{M}_{BC}) if and only if $\neg\varphi$ is not valid for the class \mathbf{M} (resp. \mathbf{M}_{BC}).

As the following theorem indicates, graded belief operators add expressivity to the non-graded language \mathcal{L}^- .

Theorem 1. The language \mathcal{L} is strictly more expressive than the language \mathcal{L}^- .

Proof (sketch). By contradiction, suppose there is a formula φ from the single-agent version of \mathcal{L}^- that is equivalent to $\Box_1^1 p$. Let us consider a formula $\psi = p \wedge \dots \wedge p$ such that $\Delta_1 \psi$ does not appear in φ . We have:

³ Note that a belief base B_i may contain non-independent formulas p and $p \wedge p$ which count twice when computing relations \mathcal{R}_i^k . We could consider non-redundant belief bases in which redundant formulas such as $p \wedge p$ are not allowed. We leave the analysis of the notion of redundancy for future work.

- $((\{p, \psi\}, \emptyset), \mathbf{S}) \models \Box_1^1 p$;
- $((\{p, \psi\}, \emptyset), \mathbf{S}) \models \varphi$;
- $((\{p\}, \emptyset), \mathbf{S}) \not\models \Box_1^1 p$;
- $((\{p\}, \emptyset), \mathbf{S}) \not\models \varphi$

By $((\{p, \psi\}, \emptyset), \mathbf{S}) \models \varphi$, we also have $((\{p\}, \emptyset), \mathbf{S}) \models \varphi$ because φ does not talk about ψ being in the base of 1. \square

4 Social influence

In this section, we apply the language \mathcal{L} and its belief base semantics to the analysis of the concept of social influence. In social sciences [41], social influence is conceived as the causal connection between an agent's belief (or opinion) and other agents' beliefs: an agent (the influencee) believes that α because and as long as she believes that other credible agents (the influencers) believe that α . It has been shown to play a crucial role in information dynamics in multi-agent systems (see, e.g., [21,14,43]).

The belief in the information source's credibility is an essential component of the influence process. Indeed, for a rational agent i to be influenced by another agent j 's opinion, i must believe that j 's opinion is correct and well-founded, that is to say, j must not have wrong beliefs about the subject at matter.

In line with [15], we assume that an agent i 's belief that another agent j 's is credible about α is identified with i 's belief that 'if agent j believes that α , then α is true'. This captures a form of i 's trust in j , namely, i 's trust in j 's credibility about α . We note the latter by $Trust(i, j, \alpha)$ and define it as follows:

$$Trust(i, j, \alpha) \stackrel{\text{def}}{=} \Delta_i(\Delta_j \alpha \rightarrow \alpha).$$

Let $2^{Agt^*} = 2^{Agt} \setminus \{\emptyset\}$ and its elements be denoted by G, G', \dots . As the following proposition indicates if an agent i has trusts in the credibility of each information source in group G and explicitly believes that each of them explicitly believes that α , then she should conclude that α is true with strength at least $|G| - 1$. This means that an agent cumulates information received from different credible information sources to determine her degree of belief: the higher the number of credible sources in support of α , the stronger the influence, the higher the degree of the belief that α . We recall that \Box_i is defined by \Box_i^0 . This is the reason why the resulting degree of belief is $|G| - 1$ instead of $|G|$. For example, if G is a singleton, the resulting degree is 0 which means that agent i believes that α .

Proposition 2. *Let $i \in Agt$ and $G \in 2^{Agt^*}$. Then,*

$$\models_{\mathbf{M}} \bigwedge_{j \in G} (Trust(i, j, \alpha) \wedge \Delta_i \Delta_j \alpha) \rightarrow \Box_i^{|G|-1} \alpha. \quad (1)$$

The following example concretely illustrates the social influence process with the help of the belief base semantics.

Example 2. Suppose agent cb is a chatbot connected to the Internet who has to provide information to a human user about the quality of a certain movie. The chatbot has access

to four recommender systems about movies: Netflix (nf), Rotten Tomatoes (rt), IMDb (im) and Amazon (am). Each recommender system provides an evaluation whether the movie is good or not which is used by the chabot to form an opinion about the movie. Consider an arbitrary MAB (B, Cxt) such that agent cb's belief base in B is:

$$B_{cb} = \bigcup_{j \in \{nf, rt, im, am\}} \{\Delta_j good \rightarrow good, \Delta_j \neg good \rightarrow \neg good\} \cup \{\Delta_{nf} good, \Delta_{rt} good, \Delta_{am} good, \Delta_{im} \neg good\}.$$

We have the following:

$$(B, Cxt) \models \bigwedge_{j \in \{nf, rt, im, am\}} (Trust(cb, j, good) \wedge Trust(cb, j, \neg good)) \wedge \Delta_{cb} \Delta_{nf} good \wedge \Delta_{cb} \Delta_{rt} good \wedge \Delta_{cb} \Delta_{am} good \wedge \Delta_{cb} \Delta_{im} \neg good \wedge \square_{cb}^2 good \wedge \square_{cb}^0 \neg good.$$

This means that in the situation described by the MAB (B, Cxt) , (i) cb trusts the credibility of each of the four recommender systems both about the fact that the movie is good and about the fact that the movie is not good, and (ii) cb believes that Netflix, Rotten Tomatoes and Amazon evaluates it as a good movie, while IMDb evaluates it as a not good movie. Furthermore, thanks to (i) and (ii), we have that (iii) cb believes that the movie is good with strength equal to 2 and believes that the movie is not good with strength equal to 0. This means that cb's degree of belief that the movie is good, is strictly higher than the degree of belief that the movie is not good, since there are more credible sources in support of the former than credible sources in support of the latter.

It is worth noting that in Proposition 2, information sources are supposed to be independent. Indeed, formula $Trust(i, j, \alpha)$ relies on the assumption that i believes that j is credible about α , *regardless of what the other agents believe*. This explains why j 's opinion contributes to increase i 's strength of belief. The situation is different when agent i merely trusts the credibility of a group of information sources including agent j as a whole. In this case, i will not be influenced by j 's beliefs unless the other sources have the same belief. This covers the case in which information sources in the group are dependent so that their opinions jointly contribute to increase i 's strength of belief, but not individually. To see this formally, let us generalize the previous definition to trust in a group's credibility, as follows:

$$Trust(i, G, \alpha) \stackrel{\text{def}}{=} \Delta_i \left(\left(\bigwedge_{j \in G} \Delta_j \alpha \right) \rightarrow \alpha \right) \text{ with } G \in 2^{Agt*}.$$

As the following proposition indicates, trusting a group's credibility about α and believing that each source in the group believes that α is sufficient for forming the belief that α . Nonetheless, it is not sufficient for forming a belief with a strictly higher degree, since the group counts as a single unit of influence.

Proposition 3. *Let $i \in \text{Agt}$, $G \in 2^{\text{Agt}^*}$. Then,*

$$\models_{\mathbf{M}} (\text{Trust}(i, G, \alpha) \wedge \bigwedge_{j \in G} \Delta_i \Delta_j \alpha) \rightarrow \Box_i \alpha, \quad (2)$$

$$\not\models_{\mathbf{M}} (\text{Trust}(i, G, p) \wedge \bigwedge_{j \in G} \Delta_i \Delta_j p) \rightarrow \Box_i^1 p. \quad (3)$$

5 Axiomatics and decidability

This section is devoted to define two logical systems of explicit belief and graded implicit belief. They are called LGDA and $\text{LGDA}_{\mathbf{T}\Box_i}$, where LGDA stands for ‘‘Logic of Graded Doxastic Attitudes’’.

Let us start with the definition of the two logics.

Definition 8 (LGDA). *We define LGDA to be the extension of classical propositional logic by the following axioms and rule of inference:*

$$\begin{aligned} (\Box_i^k \varphi \wedge \Box_i^k (\varphi \rightarrow \psi)) &\rightarrow \Box_i^k \psi && \mathbf{(K}_{\Box_i^k}) \\ \Box_i^k \varphi &\rightarrow \Box_i^{k'} \varphi \text{ if } k' \leq k && \mathbf{(Mon}_{\Box_i^k}) \\ \left(\bigwedge_{\alpha \in X} \Delta_i \alpha \right) &\rightarrow \Box_i^k \bigvee_{\substack{X' \subseteq X: \\ |X'| \geq |X| - k}} \bigwedge_{\beta \in X'} \beta \text{ if } |X| > k && \mathbf{(Int}_{\Delta_i, \Box_i^k}) \\ \frac{\varphi}{\Box_i^k \varphi} &&& \mathbf{(Nec}_{\Box_i^k}) \end{aligned}$$

We define $\text{LGDA}_{\mathbf{T}\Box_i}$ to be the extension of the logic LGDA by the following axiom:

$$\Box_i \varphi \rightarrow \varphi \quad \mathbf{(T}_{\Box_i})$$

Axiom $\mathbf{K}_{\Box_i^k}$ and the rule of inference $\mathbf{Nec}_{\Box_i^k}$ are the basic principles of the normal modal operator \Box_i^k for graded implicit belief. Axiom $\mathbf{Mon}_{\Box_i^k}$ is a monotonicity principle for graded implicit belief: implicitly believing that φ with degree at least k implies implicitly believing that φ with degree at least k' if $k' \leq k$. Finally, Axiom $\mathbf{Int}_{\Delta_i, \Box_i^k}$ is the interaction principle between explicit and graded implicit belief: if an agent explicitly believes every fact in X , then she should implicitly believe with degree at least k that there exists a subset X' of X such that $|X'| \geq |X| - k$ and every fact in X' is true. The reason why we do not consider the case $|X| \leq k$ is that if $|X| \leq k$ then $\Box_i^k (\bigvee_{X' \subseteq X: |X'| \geq |X| - k} \bigwedge_{\beta \in X'} \beta)$ is equivalent to $\Box_i^k \top$ which in turn is equivalent to \top . It is also worth noting that if $|X| = 1$ then Axiom $\mathbf{Int}_{\Delta_i, \Box_i^k}$ acquires the simpler form $\Delta_i \alpha \rightarrow \Box_i \alpha$.

As the following theorem highlights, the two logics are sound and complete relative to the belief base semantics defined in the previous section.

Theorem 2. *The logic LGDA is sound and complete for the class of MABs, whereas the logic $\text{LGDA}_{\mathbf{T}\Box_i}$ is sound and complete for the class of belief correct MABs.*

Proof (sketch). The theorem relies on the fact that the belief base semantics for the language \mathcal{L} given in Section 2 is equivalent to a “weaker” semantics exploiting enriched Kripke structures of the form $M = (W, \mathcal{D}, \mathcal{N}, \mathcal{V})$ where W is a non-empty set of worlds, $\mathcal{D} : \text{Agt} \times W \rightarrow 2^{\mathcal{L}^0}$ with $\mathcal{D}(i, w)$ finite for every $i \in \text{Agt}$ and $w \in W$, $\mathcal{N} : \text{Agt} \times W \times \mathbb{N} \rightarrow 2^W$ and $\mathcal{V} : \text{Atm} \rightarrow 2^W$ such that for all $i \in \text{Agt}$, for all $w \in W$ and for all $k, k' \in \mathbb{N}$:

$$\begin{aligned} \mathcal{N}(i, w, k) &\subseteq \{v \in W : |\text{Sat}_M(i, w, v)| \geq (|\mathcal{D}(i, w)| - k)\}, \\ \mathcal{N}(i, w, k') &\subseteq \mathcal{N}(i, w, k) \text{ if } k' \leq k, \end{aligned}$$

and with respect to which \mathcal{L} -formulas are interpreted as follows (boolean cases are omitted for simplicity): (i) $(M, w) \models p$ iff $w \in \mathcal{V}(p)$, (ii) $(M, w) \models \Delta_i \alpha$ iff $\alpha \in \mathcal{D}(i, w)$, (iii) $(M, w) \models \Box_i^k \varphi$ iff $\forall v \in \mathcal{N}(i, w, k) : (M, v) \models \varphi$, with $\text{Sat}_M(i, w, v) = \{\alpha \in \mathcal{D}(i, w) : (M, v) \models \alpha\}$. \square

The following decidability result is a consequence of the finite model property for logics LGDA and LGDA_{\Box_i} .

Theorem 3. *The satisfiability problem of LGDA (resp. LGDA_{\Box_i}) relative to the class of MABs (resp. belief correct MABs) is decidable.*

6 Conditional belief operators

In this section, we extend the language \mathcal{L} by conditional belief operators of type \Box_i^{+X} and \Box_i^{-X} . They capture, respectively, what agent i would implicitly believe if she added all information in X to her belief base, and what she would believe if she removed all information in X from her belief base. The new language is denoted by $\mathcal{L}_{\text{cond}}$. The semantic interpretation of these new operators is given in the following definition.

Definition 9. *Let $(B, \text{Cxt}) \in \mathbf{M}$ with $B = (B_1, \dots, B_n, S)$. Then,*

$$\begin{aligned} (B, \text{Cxt}) &\models \Box_i^{+X} \varphi \text{ iff } (B^{i+X}, \text{Cxt}) \models \Box_i \varphi, \\ (B, \text{Cxt}) &\models \Box_i^{-X} \varphi \text{ iff } (B^{i-X}, \text{Cxt}) \models \Box_i \varphi, \end{aligned}$$

where $B^{i+X} = (B_1^{i+X}, \dots, B_n^{i+X}, S^{i+X})$ and $B^{i-X} = (B_1^{i-X}, \dots, B_n^{i-X}, S^{i-X})$ with:

$$\begin{aligned} B_i^{i+X} &= B_i \cup X, & B_i^{i-X} &= B_i \setminus X, \\ B_j^{i+X} &= B_j^{i-X} = B_j \text{ if } i \neq j, & S^{i+X} &= S^{i-X} = S. \end{aligned}$$

Interestingly, the following axioms show that the new operators \Box_i^{+X} and \Box_i^{-X} do not add expressivity to the language.

$$\Box_i^{+X} \varphi \leftrightarrow \Box_i \left(\left(\bigwedge_{\alpha \in X} \alpha \right) \rightarrow \varphi \right) \quad (\text{+X})$$

$$\begin{aligned} \Box_i^{-X} \varphi &\leftrightarrow \bigwedge_{X' \subseteq X} \left(\left(\left(\bigwedge_{\alpha \in X'} \Delta_i \alpha \right) \wedge \left(\bigwedge_{\alpha \in X \setminus X'} \neg \Delta_i \alpha \right) \right) \rightarrow \right. \\ &\quad \left. \bigwedge_{X'' \subseteq X'} \Box_i^{|X''|} \left(\left(\bigwedge_{\alpha \in X''} \neg \alpha \right) \rightarrow \varphi \right) \right) \quad (\text{-X}) \end{aligned}$$

In axiom **+X**, we simply evaluate φ -states, possible for agent i , that satisfy the guard $\bigwedge_{\alpha \in X} \alpha$. Axiom **-X** mimics the removing of X . To do that, we first identify the subset X' of formulas in X that actually appear in the base of agent i . Formulas in X' are the formulas that are indeed removed while formulas in $X \setminus X'$ are not present in the base of agent i . Then we should impose that φ holds in all possible worlds when formulas in X' are not enforced anymore. In particular, the clause for $X'' = X'$ says that if we remove $|X'|$ formulas, and if these $|X'|$ removed formulas are those in X' and are false (the guard $\bigwedge_{\alpha \in X'} \neg \alpha$) then φ holds. The definition of axiom **-X** is more subtle. As some formulas in $\alpha \in X'$ cannot be made false (because they are tautologies), we consider all guards $\bigwedge_{\alpha \in X''} \neg \alpha$ for all subsets X'' of X' .

Theorem 4. *Axioms **+X** and **-X** are valid.*

We call LGCDA (Logic of Graded and Conditional Doxastic Attitudes) the extension of logic LGDA by the previous Axioms **-X** and **-X**, and $\text{LGCDA}_{\mathbb{T}\square_i}$ the corresponding extension of the logic $\text{LGDA}_{\mathbb{T}\square_i}$. It is routine exercise to check that these axioms are valid relative to the class of MABs. Thus, by Corollary 2, we have the following completeness result for the logics LGCDA and $\text{LGCDA}_{\mathbb{T}\square_i}$.

Theorem 5. *The logic LGCDA is sound and complete for the class of MABs, whereas the logic $\text{LGCDA}_{\mathbb{T}\square_i}$ is sound and complete for the class of belief correct MABs.*

The following theorem is a direct consequence of Theorem 3.

Theorem 6. *The satisfiability problem of LGCDA (resp. $\text{LGCDA}_{\mathbb{T}\square_i}$) relative to the class of MABs (resp. belief correct MABs) is decidable.*

7 Model checking

Consider these compact formulations of the model checking problems for the language \mathcal{L}_{cond} .

Model checking
Given: $\varphi \in \mathcal{L}_{cond}$, $\alpha \in \mathcal{L}_0$ and a finite $B \in \mathbf{S}$.
Question: Do we have $(B, \mathbf{S}(\alpha)) \models \varphi$?
 with $\mathbf{S}(\alpha) = \{B \in \mathbf{S} : B \models \alpha\}$.

Belief correct model checking
Given: $\varphi \in \mathcal{L}_{cond}$, $\alpha \in \mathcal{L}_0$ and a finite $B \in \mathbf{S}_{BC}$ with $B \models \alpha$.
Question: Do we have $(B, \mathbf{S}_{BC}(\alpha)) \models \varphi$?
 with $\mathbf{S}_{BC}(\alpha) = \{B \in \mathbf{S}_{BC} : B \models \alpha\}$.

where the state $B = (B_1, \dots, B_n, S)$ is said to be finite if S and every B_i are finite. Note that, thanks to Proposition 1 and the fact that $B \models \alpha$, the MAB $(B, \mathbf{S}_{BC}(\alpha))$ in the belief correct variant of model checking belongs to the model class \mathbf{M}_{BC} , as expected.

In [35], it is proved that the previous two problems are PSPACE-hard, already for the fragment of \mathcal{L} with only implicit belief operators of type \square_i .

We are going to focus on the complexity upper bound. To this aim, we follow the approach given in [26]. The algorithm given in Figure 1 checks that a formula φ is true in a given finite state B . Checking $B' \models \alpha$ (see Definition 5) can be done in polynomial time because α does not contain any implicit belief operator; thus it is reducible to the propositional problem by stating any explicit belief as a fresh proposition. Once $|\{\alpha \in B_i : B' \models \alpha\}|$ is computed, the comparison $|\{\alpha \in B_i : B' \models \alpha\}| \geq (|B_i| - k)$, can be done in polynomial time. Furthermore checking that B' is in $\mathbf{S}(\alpha)$ (or $\mathbf{S}_{BC}(\alpha)$) can be done in polynomial time. So checking $B\mathcal{R}_i^k B'$ can be done in polynomial time and space. We are now ready to establish the PSPACE upper bound for the two model checking problems.

Theorem 7. *Both the model checking problem and the belief correct model checking problem are in PSPACE.*

Proof (sketch). The number of nested calls in $mc(B, \varphi)$ is bounded by the size of φ . The local memory used by the recursive call is polynomial in the size of the initial B and the size of φ . Loops “for all $B' \dots$ ” are performed by enumerating the B' containing correct subformulas of formulas in the initial B and in the initial formula φ . Despite there is an exponential number of such B' , storing the current B' only requires a polynomial amount of space.

The algorithm for the belief correct model checking is similar: we just check each time that the states B' under consideration are correct (Definition 3). \square

```

procedure  $mc(B, \varphi)$ 
  match  $\varphi$  do
    case  $p$ : return  $B \models p$ 
    case  $\neg\psi$ : return not  $mc(B, \psi)$ 
    case  $\psi_1 \wedge \psi_2$ : return  $mc(B, \psi_1)$  and  $mc(B, \psi_2)$ 
    case  $\Delta_i \alpha$ : return  $\alpha \in B_i$ 
    case  $\Box_i^k \psi$ :
      for all  $B'$  such that  $B\mathcal{R}_i^k B'$  do
        if not  $mc(B', \psi)$  return false
    case  $\Box_i^{+X} \psi$ :
      for all  $B'$  such that  $B^{+X}\mathcal{R}_i^0 B'$  do
        if not  $mc(B', \psi)$  return false
    case  $\Box_i^{-X} \psi$ :
      for all  $B'$  such that  $B^{-X}\mathcal{R}_i^0 B'$  do
        if not  $mc(B', \psi)$  return false
  return true // when all tests in for loops failed

```

Fig. 1. Generic algorithm for model checking.

8 Epistemic Explanation

In this section, we leverage the language \mathcal{L}_{cond} to model a variety of notions of epistemic explanation. The standard notion of explanation [29,24] is relative to some background theory which together with the *explanans* is used to explain the *explanandum*.

As emphasized by [7], epistemic explanation is relative to an agent’s epistemic state: the agent explains a given fact or observation in the light of her background knowledge. Existing formal models of epistemic explanation including [7,19,20] focus on the single agent case. We generalize the analysis of epistemic explanation to the multi-agent case in which (i) different agents may have diverging explanations of the same fact, and (ii) an agent may include other agents’ beliefs in the explanation of a given fact.

Following [7], we distinguish *factual* explanation from *hypothetical* explanation. In factual explanation both the explanans and the explanandum are believed by the explaining agent, while they are not in hypothetical explanation. Specifically, a factual explanation is a body of information in the agent’s belief base which supports an actual belief of the agent. A hypothetical explanation is relative to a fact that is not actually believed by the agent but that the agent would have believed, if she had believed that the explanans is true. Let us first define factual explanation:

$$FactExpl_i(X, \varphi) \stackrel{\text{def}}{=} \left(\bigwedge_{\alpha \in X} \Delta_i \alpha \right) \wedge \Box_i \varphi \wedge \neg \Box_i^{-X} \varphi.$$

Formula $FactExpl_i(X, \varphi)$ has to be read “according to agent i , X is a factual explanation of φ ”, where X is the explanans and φ is the explanandum. This means that (i) i explicitly believes all facts in X , (ii) i implicitly believes φ , and (iii) if i removed all information in X from her belief base, she would not believe φ anymore. In other words, a factual explanation is a subset of the agent’s belief base which is necessary for the agent to derive φ . Note that this notion of factual explanation can be used by the agent to detect and explain inconsistency of her belief base. In particular, $FactExpl_i(X, \perp)$ means that, according to agent i , inconsistency of her belief base depends on the body of information X .

Most formal theories of explanation [4,28] agree on a minimality requirement. In order to account for minimality of factual explanation, we need to assume that, for every strict subset X' of X , removing all information in X' from the belief base does not affect i ’s belief that φ . That is, we define:

$$MinFactExpl_i(X, \varphi) \stackrel{\text{def}}{=} FactExpl_i(X, \varphi) \wedge \bigwedge_{\alpha \in X} \Box_i^{-(X \setminus \{\alpha\})} \varphi,$$

where $MinFactExpl_i(X, \varphi)$ has to be read “according to agent i , X is a minimal factual explanation of φ ”. Like $FactExpl_i(X, \varphi)$ the size of $MinFactExpl_i(X, \varphi)$ is polynomial in the size of X . Note that since $\Box_i^{-X} \varphi$ implies $\Box_i^{-X'} \varphi$ for $X' \subset X$, the second conjunct in the definition of $MinFactExpl_i(X, \varphi)$ is equivalent to $\bigwedge_{X' \subset X} \Box_i^{-X'} \varphi$.

The following example illustrates the notion of minimal factual explanation.

Example 3. Let us go back to the example of Section 4. By the model checking algorithm of Figure 1, we can verify that:

$$(B, Cxt) \models MinFactExpl_{cb}(\{\Delta_{nf} good, \Delta_{rt} good, \Delta_{am} good\}, good).$$

This means that the fact that every agent in $\{nf, rt, am\}$ explicitly believes the movie is good is for agent cb a minimal factual explanation that the movie is good. Thus, the body of information $\{\Delta_{nf} good, \Delta_{rt} good, \Delta_{am} good\}$ is necessary for cb to conclude *good*.

We end this section with a definition of hypothetical explanation:

$$\text{HypExpl}_i(X, \varphi) \stackrel{\text{def}}{=} \left(\bigwedge_{\alpha \in X} \neg \Delta_i \alpha \right) \wedge \Box_i^{+X} \varphi \wedge \neg \Box_i^{+X} \perp.$$

$\text{HypExpl}_i(X, \varphi)$ has to be read “according to agent i , X is a hypothetical explanation of φ ”, in the sense that: (i) no piece of information in X is included in i ’s belief base, (ii) i does not believe that φ , (iii) the body of information X would be sufficient for agent i to consistently conclude φ . Similarly to factual explanation, minimality is captured by assuming that there is no strict subset X' of X that is sufficient for agent i to conclude φ :

$$\text{MinHypExpl}_i(X, \varphi) \stackrel{\text{def}}{=} \text{HypExpl}_i(X, \varphi) \wedge \bigwedge_{\alpha \in X} \neg \Box_i^{+(X \setminus \{\alpha\})} \varphi.$$

$\text{MinHypExpl}_i(X, \varphi)$ has to be read “according to agent i , X is a minimal hypothetical explanation of φ ”. Note that since $\Box_i^{+X'} \varphi$ implies $\Box_i^{+X} \varphi$ for $X' \subset X$, the second conjunct in the definition of $\text{MinHypExpl}_i(X, \varphi)$ is equivalent to $\bigwedge_{X' \subset X} \neg \Box_i^{+X'} \varphi$.

9 Conclusion

We defined a graded doxastic language \mathcal{L} to reason about an agent’s implicit beliefs, when a given *number* of explicit beliefs are removed from her belief base.

Our approach could be relevant in many applications: the agents receive beliefs (including higher-order beliefs, e.g., agent 1 knows that agent 2 knows that p) from different sources, or different reasoners. Moreover, the model checking procedure presented in the paper helps to understand what an agents still believes if some beliefs are removed from her belief base. Also, non-AI experts do not need to learn several languages: in our approach, we emphasize that the query language (e.g., does drone 1 believes that drone 2 believes the area is safe?), the inner state description (formulas in bases), and the language for explanation are the *same*.

Concerning the contributions, we showed that the language \mathcal{L} is strictly more expressive than the ungraded doxastic language \mathcal{L}^- . We also introduced the language $\mathcal{L}_{\text{cond}}$ to reason about an agent’s implicit beliefs when a given *set* of explicit beliefs are added/removed to/from her belief base. The languages \mathcal{L} and $\mathcal{L}_{\text{cond}}$ are equally expressive and we provided a sound and complete axiomatization for both of them.

Directions of future work are manifold. On the theoretical side, we plan to study complexity of the satisfiability checking problem for \mathcal{L} and $\mathcal{L}_{\text{cond}}$. On the practical side, we plan to implement the model checking algorithm for \mathcal{L} and $\mathcal{L}_{\text{cond}}$. We also plan to propose a tool for automatic verification and generation of epistemic explanations in multi-agent scenarios involving autonomous agents endowed with epistemic states.

Acknowledgements

Support from the ANR-3IA Artificial and Natural Intelligence Toulouse Institute and from the ANR project CoPains (grant number ANR-18-CE33-0012) is gratefully acknowledged.

References

1. G. Aucher. A combined system for update logic and belief revision. In *Proceedings of PRIMA 2004*, volume 3371 of *LNAI*, pages 1–18. Springer-Verlag, 2005.
2. P. Balbiani, D. Fernández-Duque, A. Herzig, and E. Lorini. Stratified evidence logics. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI 2019)*, pages 1523–1529, 2019.
3. A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press, 2008.
4. S. Benferhat, D. Dubois, and H. Prade. Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study part 1: The flat case. *Studia Logica*, 58:17–45, 1997.
5. S. Benferhat, D. Dubois, H. Prade, and M.-A. Williams. A practical approach to revising prioritized knowledge bases. *Stud Logica*, 70(1):105–130, 2002.
6. T. Bolander and M. B. Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):656–680, 2011.
7. C. Boutilier and V. Beche. Abduction as belief revision. *Artificial Intelligence*, 77(1):43–94, 1995.
8. K. Budzynska and M. Kacprzak. A logic for reasoning about persuasion. *Fundamenta Informaticae*, 85:51–65, 2008.
9. M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, 1990.
10. T. Charrier, S. Pinchinat, and F. Schwarzenruber. Symbolic model checking of public announcement protocols. *Journal of Logic and Computation*, 29(8):1211–1249, 2019.
11. T. Charrier and F. Schwarzenruber. Arbitrary public announcement logic with mental programs. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 1471–1479. ACM, 2015.
12. T. Charrier and F. Schwarzenruber. A succinct language for dynamic epistemic logic. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2017)*, pages 123–131. ACM, 2017.
13. T. Charrier and F. Schwarzenruber. Complexity of dynamic epistemic logic with common knowledge. In *Proceedings of the 12th conference on Advances in Modal Logic 12*, pages 103–122. College Publications, 2018.
14. Z. Christoff. A logic for social influence through communication. In *Proceedings of the Eleventh European Workshop on Multi-Agent Systems (EUMAS 2013)*, volume 1113 of *CEUR Workshop Proceedings*, pages 31–39. CEUR-WS.org, 2013.
15. R. Demolombe. Reasoning about trust: A formal logical framework. In *Proceedings of the Second International Conference on Trust Management (iTrust 2004)*, volume 2995 of *LNCS*, pages 291–303. Springer, 2004.
16. D. Dubois and H. Prade. *Possibility theory: an approach to computerized processing of uncertainty*. Plenum Press, 1988.
17. R. Fagin and J. Halpern. Reasoning about knowledge and probability. *Journal of the Association for Computing Machinery*, 41(2):340–367, 1994.
18. R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, 1995.
19. M. A. Falappa, Gabriele Kern-Isberner, and G. R. Simari. Explanations, belief revision and defeasible reasoning. *Artificial Intelligence*, 141:1–28, 2002.

20. P. Gärdenfors. A pragmatic approach to explanations. *Philosophy of Science*, 47(3):404–423, 1980.
21. U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 989–997. ACM, 2015.
22. J. Y. Halpern and R. Pass. A knowledge-based analysis of the blockchain protocol. In *Proceedings of the Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017)*, pages 324–335. EPTCS, 2017.
23. S. O. Hansson. Theory contraction and base contraction unified. *Journal of Symbolic Logic*, 58(2):602–625, 1993.
24. C. Hempel and P. Oppenheim. Studies in the logic of explanation. *Philosophy of Science*, 15:135–175, 1948.
25. A. Herzig, E. Lorini, and F. Maffre. A poor man’s epistemic logic based on propositional assignment and higher-order observation. In *Proceedings of the 5th International Workshop on Logic, Rationality, and Interaction (LORI 2015)*, volume 9394 of *Lecture Notes in Computer Science*, pages 156–168. Springer, 2015.
26. A. Herzig, E. Lorini, E. Perrotin, F. Romero, and F. Schwarzentruher. A logic of explicit and implicit distributed belief. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, volume 325, pages 753–760. IOS Press, 2020.
27. J. Hintikka. *Knowledge and Belief*. Cornell University Press, New York, 1962.
28. J. De Kleer. An assumption-based tms. *Artificial Intelligence*, 28:127–162, 1986.
29. B. Kment. Counterfactuals and explanation. *Mind*, 115(458):261–310, 2006.
30. K. Konolige. What awareness isn’t: A sentential view of implicit and explicit belief. In *Proceedings of the 1st Conference on Theoretical Aspects of Reasoning about Knowledge, Monterey, CA, USA, March 1986*, pages 241–250. Morgan Kaufmann, 1986.
31. B. P. Kooi. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12:381–408, 2003.
32. N. Lavreny and J. Lang. From knowledge-based programs to graded belief-based programs part I: on-line reasoning. In *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI’2004)*, pages 368–372. IOS Press, 2004.
33. L. Lismont and P. Mongin. On the logic of common belief and common knowledge. *Theory and Decision*, 37:75–106, 1994.
34. E. Lorini. In praise of belief bases: Doing epistemic logic without possible worlds. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 1915–1922. AAAI Press, 2018.
35. E. Lorini. Exploiting belief bases for building rich epistemic structures. In L. S. Moss, editor, *Proceedings of the Seventeenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2019)*, volume 297 of *EPTCS*, pages 332–353, 2019.
36. E. Lorini. Rethinking epistemic logic with belief bases. *Artificial Intelligence*, 282, 2020.
37. E. Lorini and F. Romero. Decision procedures for epistemic logic exploiting belief bases. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2019)*, pages 944–952. IFAAMAS, 2019.
38. B. Marinkovic, P. Glavan, Z. Ognjanovic, and T. Studer. A temporal epistemic logic with a non-rigid set of agents for analyzing the blockchain protocol. *Journal of Logic and Computation*, 29(5):803–830, 2019.
39. J. J. Meyer and W. van der Hoek. *Epistemic logic for AI and computer science*. Cambridge University Press, 1995.
40. J. Pearl. From conditional oughts to qualitative decision theory. In D. Heckerman and E. H. Mamdani, editors, *Proceedings of UAI’93*, pages 12–22. Morgan Kaufmann, 1993.
41. L. Rashotte. Social influence. In G. Ritzer and J. M. Ryan, editors, *Concise Blackwell Encyclopedia of Sociology*. Blackwell, 2009.

42. H. Rott. Just because”: Taking belief bases seriously. In *Logic Colloquium*, volume 98, pages 387–408, 1998.
43. N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 1590–1596. AAAI Press, 2015.
44. W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in decision, belief change and statistics*, pages 105–134. Kluwer, 1988.
45. R. Stalnaker. Common ground. *Linguistics and Philosophy*, 25(5-6):701–721, 2002.
46. J. van Benthem, J. van Eijck, M. Gattinger, and K. Su. Symbolic model checking for dynamic epistemic logic - S5 and beyond. *Journal of Logic and Computation*, 28(2):367–402, 2018.
47. W. van der Hoek. *Modalities for Reasoning about Knowledge and Quantities*. PhD thesis, Free University of Amsterdam, 1992.
48. W. van der Hoek and J.-J. Ch. Meyer. Graded modalities in epistemic logic. *Logique & Analyse*, 133-134:251–270, 1991.
49. H. van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese*, 147(2):229–275, 2005.
50. J. van Eijck and F. Schwarzentruher. Epistemic probability logic simplified. In *Advances in Modal Logic 10*, pages 158–177. College Publications, 2014.
51. E. Weydert. General belief measures. In R. L. de Mántaras and D. Poole, editors, *Proceedings of UAI’94*, pages 575–582. Morgan Kaufmann, 1994.