



HAL
open science

Exploiting Belief Bases for Building Rich Epistemic Structures

Emiliano Lorini

► **To cite this version:**

Emiliano Lorini. Exploiting Belief Bases for Building Rich Epistemic Structures. Seventeenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2019), Jul 2019, Toulouse, France. hal-02414909

HAL Id: hal-02414909

<https://hal.science/hal-02414909>

Submitted on 16 Dec 2019

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.

Exploiting Belief Bases for Building Rich Epistemic Structures

Emiliano Lorini

IRIT, CNRS, Toulouse University
Toulouse, France

Emiliano.Lorini@irit.fr

We introduce a semantics for epistemic logic exploiting a belief base abstraction. Differently from existing Kripke-style semantics for epistemic logic in which the notions of possible world and epistemic alternative are primitive, in the proposed semantics they are non-primitive but are defined from the concept of belief base. We show that this semantics allows us to define the universal epistemic model in a simpler and more compact way than existing inductive constructions of it. We provide (i) a number of semantic equivalence results for both the basic epistemic language with ‘individual belief’ operators and its extension by the notion of ‘only believing’, and (ii) a lower bound complexity result for epistemic logic model checking relative to the universal epistemic model.

1 Introduction

Type spaces were introduced by Harsanyi [22] in order to formally represent higher-order probabilistic beliefs of rational players in strategic situations. The notion of *universal* type space was studied in detail by Mertens & Zamir [36]: it is defined to be the “largest” type space which contains all possible states of the world as well as all possible belief hierarchies of the players. Mertens & Zamir showed that, under certain topological assumptions, every Harsanyi type space can be mapped to the universal type space by a morphism which preserves the current state of the world as well as the players’ belief hierarchies. Alternative constructions of the universal type space under different topological assumptions can be found in [9, 23]. Heifez & Samet [24] proved a variant of Mertens & Zamir’s result for the general measure-theoretic case. Battigalli & Siniscalchi [2] extended the universal type space construction to conditional probabilistic beliefs. More recently, Bjorndahl & Halpern [7] provided a logical analysis of the probabilistic structure of the universal type space.¹

A qualitative version of the universal type space was introduced by Fagin et al. [12] (see also [13]). Specifically, Fagin et al. provided a construction of the “largest” (or “universal”) epistemic model for the epistemic logic $S5^n$ which contains all possible knowledge hierarchies of the agents in the system. They studied its relationship with the standard multi-relational semantics of epistemic logic [14] both for the basic epistemic language with ‘individual knowledge’ operators and for its extension by common knowledge operators.

Both constructions of the universal epistemic model in its probabilistic version à la Mertens & Zamir and in its qualitative version à la Fagin et al. are inductive. One has to define first the set of possible states of the world. Then, the agents’ first-order beliefs about the states of the world are defined. In the third step, one has to define the agents’ second-order beliefs about the states of nature and the agents’

¹An analysis of the relationship between Harsanyi’s type spaces and the multi-relational Kripke semantics of epistemic logic can be found in [15].

first-order beliefs, and so on so forth. More generally, in order to define an agent's $k + 1$ -order belief one has to define first the agents' k -order beliefs.

The general aim of the present paper is to offer a simple and compact construction of the universal epistemic model that does not require an inductive construction of the agents' belief hierarchies, as in Mertens & Zamir's and Fagin et al.'s definitions. Our construction of the universal epistemic model is based on a semantics for epistemic logic exploiting a belief base abstraction which was recently introduced in [33]. Differently from existing multi-relational Kripke semantics for epistemic logic in which the notions of possible world and epistemic alternative are primitive, in the belief base semantics they are non-primitive but are defined from the concept of belief base. Specifically, in this semantics it is assumed that at a given state s agent i considers state s' possible if and only if s' satisfies all formulas that are included in agent i 's belief base at s .

The initial motivation for introducing such a semantics was to bridge two traditions that have rarely talked to each other in the past. On the one hand, we have epistemic logic: it started in the 60s with the seminal work of Hintikka [26] on the logics of knowledge and belief, it was extended to the multi-agent setting at the end of 80s [14, 37] and then furtherly developed during the last 20 years, the period of the "dynamic turn", with growing research on dynamic epistemic logic [11]. On the other hand, we have syntactic approaches to knowledge representation and reasoning mainly proposed in the area of artificial intelligence (AI). The latter includes, for instance, work on belief base and knowledge base revision [20, 21, 4], belief base merging [29], input-output logic [35], as well as more recent work on the so-called "database perspective" to the theory of intention by [39]. All these approaches defend the idea that right level of abstraction for modeling rational agents is the "belief base" (or "knowledge base") level whereby the agent is identified with the set of facts that she believes (or knows).

The paper is organized as follows. In Section 2, we present the belief base semantics for epistemic logic, as defined in [33, 34]. We show how the basic language with 'individual belief' operators can be interpreted on this semantics. We also show how the semantics can be easily incorporate the belief correctness assumption, thereby allowing us to model knowledge instead of belief. We show that the belief base semantics is equivalent to the standard multi-relational Kripke semantics, in the sense that they lead to the same set of validities. Section 3 is the core of the paper: we define the universal epistemic model with the help of the belief base semantics. As emphasized above, our definition is simpler and more compact than existing definitions of the universal epistemic model, since it does not require an inductive construction of the agents' belief hierarchies. We show that, as far as the basic epistemic language with 'individual belief' operators is concerned, the set of validities relative to the universal epistemic model and the set of validities relative to the generic belief base semantics are the same. In Section 4, we extend the basic epistemic language by the notion of 'only believing'. We show that the equivalence between the two semantics does not hold anymore in the context of this more expressive language: the universal epistemic model has more validities than the generic belief base semantics. Section 5 is devoted to the comparison between Fagin et al.'s inductive construction of the universal epistemic model and the construction of Section 3 exploiting belief bases. In particular, we show that, from the point of the basic epistemic language and of its extension by the notion of 'only believing', there is no difference between the two constructions, as they give rise to the same set of validities. In Section 6, we provide a compact formulation of epistemic logic model checking which exploits the definition of the universal epistemic model given in Section 3. We show that such a compact formulation makes epistemic logic model checking PSPACE-hard, whereas standard epistemic logic model checking relative to the multi-relational Kripke semantics is polynomial in the size of the model and the length of the formula. In Section 7, we conclude.

2 Belief base semantics for epistemic logic

In this section, we present a semantics for epistemic logic (EL) in which the accessibility relations in an epistemic model are not primitive but they are defined from the primitive concept of multi-agent belief base. The semantics was first introduced in [33].

2.1 Multi-agent belief base models

Assume a countably infinite set of atomic propositions $Atm = \{p, q, \dots\}$ and a finite set of agents $Agt = \{1, \dots, n\}$.

We define the language $\mathcal{L}_0(Atm, Agt)$ by the following grammar in Backus-Naur Form (BNF):

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

where p ranges over Atm and i ranges over Agt . $\mathcal{L}_0(Atm, Agt)$ is the language for representing explicit beliefs of multiple agents. For simplicity, we sometimes write \mathcal{L}_0 instead of $\mathcal{L}_0(Atm, Agt)$, when the context is unambiguous. The formula $\Delta_i\alpha$ can be read as “agent i explicitly believes that α is true” or “ α is in agent i ’s belief base”. In this language, we can represent higher-order explicit beliefs, for example $\Delta_i\Delta_j\alpha$ express the fact that agent i explicitly believes that agent j explicitly believes that α is true.

Definition 1 (Multi-agent belief base). *A multi-agent belief base (MBB) is a tuple $B = (B_1, \dots, B_n, V)$ where (i) for every $i \in Agt$, $B_i \subseteq \mathcal{L}_0$ is agent i ’s belief base, and (ii) $V \subseteq Atm$ is the actual state. The class of MBBs is denoted by \mathbf{B} .*

Formulas of the language \mathcal{L}_0 are interpreted relative to MBBs as follows.

Definition 2 (Satisfaction relation for formulas in \mathcal{L}_0). *Let $B = (B_1, \dots, B_n, V) \in \mathbf{B}$. Then, the satisfaction relation \models between B and formulas in \mathcal{L}_0 is defined as follows:*

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

Definition 3 (Correct MBB). *Let $B = (B_1, \dots, B_n, V) \in \mathbf{B}$. We say that B is correct if and only if, for all $i \in Agt$ and for all $\alpha \in \mathcal{L}_0$, if $\alpha \in B_i$ then $B \models \alpha$. The class of correct MBBs is denoted by \mathbf{CB} .*

The following definition introduces the concept of epistemic alternative.

Definition 4 (Epistemic alternatives). *Let $B, B' \in \mathbf{B}$. Then, $B \mathcal{R}_i B'$ if and only if, for every $\alpha \in B_i$, $B' \models \alpha$, where the satisfaction relation \models follows Definition 2.*

$B \mathcal{R}_i B'$ means that B' is an epistemic alternative for agent i at B . The idea of the previous definition is that B' is an epistemic alternative for agent i at B if and only if, B' satisfies all facts that agent i explicitly believes at B .

A multi-agent belief model (MAB) is defined to be a multi-agent belief base supplemented with a set of multi-agent belief bases, called *context*. The latter includes all multi-agent belief bases that are compatible with the agents’ common ground [40], i.e., the body of information that the agents commonly believe to be the case.

Definition 5 (Multi-agent belief model). *A multi-agent belief model (MBM) is a pair (B, Cxt) , where $B \in \mathbf{B}$ and $Cxt \subseteq \mathbf{B}$. The class of MBMs is denoted by \mathbf{M} .*

2.2 Interpretation of epistemic logic language

Thanks to the epistemic accessibility relations defined in Definition 4, we are able to interpret the language of epistemic logic relative to MBMs. Such language is denoted by $\mathcal{L}_{\text{EL}}(Atm, Agt)$ and is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_i\varphi$$

where p ranges over Atm and i ranges over Agt . For simplicity, we write \mathcal{L}_{EL} instead of $\mathcal{L}_{\text{EL}}(Atm, Agt)$, when the context is unambiguous. The other Boolean constructions \top , \perp , \vee , \rightarrow and \leftrightarrow are defined from p , \neg and \wedge in the usual way.

The formula $\Box_i\varphi$ has to be read “agent i implicitly (or potentially) believes that φ is true” in the sense that agent i can derive φ from her explicit beliefs (i.e., from the information in her belief base). For the sake of simplicity, we sometimes read the formula $\Box_i\varphi$ as “agent i believes that φ ”. We define the dual operator \Diamond_i as follows:

$$\Diamond_i\varphi \stackrel{\text{def}}{=} \neg\Box_i\neg\varphi.$$

$\Diamond_i\varphi$ has to be read “ φ is compatible (or consistent) with agent i ’s explicit beliefs”.

As usual, for every formula $\varphi \in \mathcal{L}_{\text{EL}}$ we denote by $Atm(\varphi)$ the set of atoms in Atm occurring in φ .

We are ready to define what it means for a multi-agent belief model (MBM) (B, Cxt) to satisfy a formula φ in \mathcal{L}_{EL} , written $(B, Cxt) \models \varphi$.

Definition 6 (Satisfaction relation for formulas in \mathcal{L}_{EL}). *Let $B = (B_1, \dots, B_n, V) \in \mathbf{B}$ and let $(B, Cxt) \in \mathbf{M}$. Then:*

$$\begin{aligned} (B, Cxt) \models p &\iff p \in V \\ (B, Cxt) \models \neg\varphi &\iff (B, Cxt) \not\models \varphi \\ (B, Cxt) \models \varphi \wedge \psi &\iff (B, Cxt) \models \varphi \text{ and } (B, Cxt) \models \psi \\ (B, Cxt) \models \Box_i\varphi &\iff \forall B' \in Cxt : \text{if } B \mathcal{R}_i B' \text{ then } (B', Cxt) \models \varphi \end{aligned}$$

Note that, according to the last clause, an agent i implicitly believes that φ (i.e., $\Box_i\varphi$) if and only if φ is true at all states that are compatible with the information in i ’s belief base.

Figure 1 illustrates the general idea behind the belief base semantics, especially for what concerns the relationship between the agents’ belief bases and the agents’ common ground (or context) and the relationship between the latter and the agents’ implicit beliefs. While an agent’s belief base captures the agent’s private information, the common ground captures the agents’ public information. An agent’s implicit belief corresponds to a fact that the agent can deduce from the public information and her private information.

Given a formula φ in \mathcal{L}_{EL} , we denote by $depth(\varphi)$ its model depth which is defined as follows:

$$\begin{aligned} depth(p) &= 0 \\ depth(\neg\varphi) &= depth(\varphi) \\ depth(\varphi \wedge \psi) &= \max(depth(\varphi), depth(\psi)) \\ depth(\Box_i\varphi) &= depth(\varphi) + 1 \end{aligned}$$

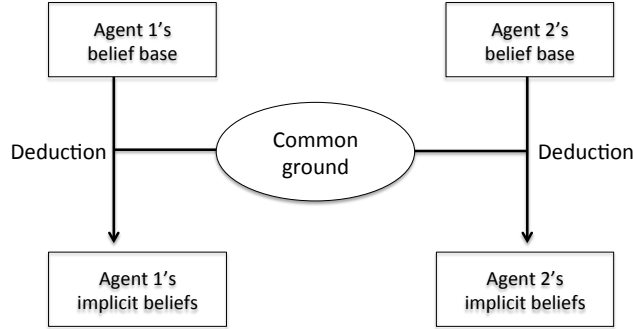


Figure 1: Conceptual framework

2.3 Model classes and validity

In some situations, it may be useful to assume that agents' beliefs are correct, i.e., what an agent believes is true. When talking about correct (or true) explicit and implicit beliefs, it is usual to call them explicit and implicit knowledge. Indeed, we assume that the terms “true belief”, “correct belief” and “knowledge” are synonyms. The following definition introduces belief correctness for multi-agent belief models.

Definition 7 (Belief correctness). *Let $(B, Cxt) \in \mathbf{M}$. We say that (B, Cxt) satisfies belief correctness (BC) if and only if $B \in Cxt$ and, for every $B' \in Cxt$, $B' \mathcal{R}_i B'$.*

As the following proposition highlights, belief correctness for multi-agent belief models 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, i.e., if an agent has α in her belief base then α is true in the actual state of the world.

Proposition 8. *Let $(B, Cxt) \in \mathbf{M}$. Then, (B, Cxt) satisfies BC if and only if $B \in Cxt$ and $B' \in \mathbf{CB}$ for all $B' \in Cxt$.*

We let \mathbf{M}_{BC} denote the class of MBMs satisfying property BC.

Let $\varphi \in \mathcal{L}_{EL}$, we say that φ is valid for the class \mathbf{M} , denoted by $\models_{\mathbf{M}} \varphi$, if and only if, for every $(B, Cxt) \in \mathbf{M}$, we have $(B, Cxt) \models \varphi$. We say that φ is satisfiable for the class \mathbf{M} if and only if $\neg\varphi$ is not valid for the class \mathbf{M} . Satisfiability and validity of a formula φ in \mathcal{L}_{EL} relative to the class \mathbf{M}_{BC} (denoted by $\models_{\mathbf{M}_{BC}} \varphi$) are defined in an analogous way.

2.4 Equivalence with Kripke semantics

In this section, we show that the belief base semantics for the epistemic language \mathcal{L}_{EL} defined in the previous section is equivalent to the traditional multi-relational Kripke semantics for epistemic logic [14].

Definition 9 (Multi-relational Kripke model). *A multi-relational Kripke model is a structure $M = (W, \Rightarrow_1, \dots, \Rightarrow_n, \omega)$ such that W is a set of states, $\Rightarrow_i \subseteq W \times W$ is agent i 's epistemic accessibility relation, and $\omega : Atm \rightarrow 2^W$ is a valuation function. Multi-relational Kripke models satisfying reflexivity are those for which every relation \Rightarrow_i is reflexive (i.e., for every $w \in W$, $w \Rightarrow_i w$).*

The interpretation of formulas in \mathcal{L}_{EL} relative to a multi-relational Kripke model $M = (W, \Rightarrow_1$

, ..., \Rightarrow_n, ω) and a world w in W is defined as follows:

$$\begin{aligned} (M, w) &\models p \text{ iff } w \in \omega(p) \\ (M, w) &\models \neg\varphi \text{ iff } (M, w) \not\models \varphi \\ (M, w) &\models \varphi \wedge \psi \text{ iff } (M, w) \models \varphi \text{ and } (M, w) \models \psi \\ (M, w) &\models \Box_i\varphi \text{ iff } \forall v \in W : \text{if } w \Rightarrow_i v \text{ then } (M, v) \models \varphi \end{aligned}$$

Notions of validity and satisfiability relative to the class of multi-relational Kripke models are defined in the usual way. We denote the fact that φ is valid relative to the class of multi-relational Kripke models (resp. multi-relational Kripke models satisfying reflexivity) by $\models_{\text{KripkeM}} \varphi$ (resp. $\models_{\text{ReflKripkeM}} \varphi$).

The equivalence between the belief base semantics for the language \mathcal{L}_{EL} and the multi-relational Kripke semantics is stated by the following theorem.

Theorem 10. *Let $\varphi \in \mathcal{L}_{EL}$. Then,*

- $\models_{\mathbf{M}} \varphi$ if and only if $\models_{\text{KripkeM}} \varphi$,
- $\models_{\mathbf{M}_{BC}} \varphi$ if and only if $\models_{\text{ReflKripkeM}} \varphi$.

3 Universal context

We here define the notion of *universal context*, as the context containing all possible explicit belief hierarchies for the agents.

Definition 11 (α -context). *Let $\alpha \in \mathcal{L}_0$. We define*

$$\mathbf{B}_\alpha = \{B \in \mathbf{B} : B \models \alpha\}$$

to be the α -context. The universal context is simply \mathbf{B}_\top which is the same thing as \mathbf{B} .

The formula α in the previous definition corresponds to the concept of *integrality constraint*, as traditionally defined in the area of knowledge representation [38, 29].

It is worth to consider a more specific notion of universal context under the assumption of belief correctness.

Definition 12 (α -context with belief correctness). *Let $\alpha \in \mathcal{L}_0$. We define*

$$\mathbf{CB}_\alpha = \{B \in \mathbf{CB} : B \models \alpha\}$$

to be the α -context with belief correctness. The universal context with belief correctness is simply \mathbf{CB}_\top which is the same thing as \mathbf{CB} .

Let $\varphi \in \mathcal{L}_{EL}$, we say that φ is valid relative to the universal context \mathbf{B}_\top , denoted by $\models_{\mathbf{B}_\top} \varphi$, if and only if, for every $B \in \mathbf{B}_\top$, we have $(B, \mathbf{B}_\top) \models \varphi$. Analogously, we say that φ is valid relative to the universal context satisfying BC, denoted by $\models_{\mathbf{CB}_\top} \varphi$, if and only if, for every $B \in \mathbf{CB}_\top$, we have $(B, \mathbf{CB}_\top) \models \varphi$.

We say that φ is satisfiable relative to the universal context \mathbf{B}_\top (resp. \mathbf{CB}_\top) if and only if $\neg\varphi$ is not valid relative to the universal context \mathbf{B}_\top (resp. \mathbf{CB}_\top).

The following theorem highlights that the set of validities for the language \mathcal{L}_{EL} relative to the universal context \mathbf{B}_\top (resp. \mathbf{CB}_\top) and the set of validities relative to the class of models \mathbf{M} (resp. \mathbf{M}_{BC}) are the same. In other words, the language \mathcal{L}_{EL} is not able to distinguish the universal context including all possible hierarchies of the agents' beliefs from incomplete models in which some belief hierarchy may be missing.

Theorem 13. Let $\varphi \in \mathcal{L}_{EL}$. Then,

- $\models_{\mathbf{M}} \varphi$ if and only if $\models_{\mathbf{B}_{\top}} \varphi$,
- $\models_{\mathbf{M}_{BC}} \varphi$ if and only if $\models_{\mathbf{CB}_{\top}} \varphi$.

The following theorem is a corollary of Theorems 10 and 13.

Theorem 14. Let $\varphi \in \mathcal{L}_{EL}$. Then,

- $\models_{\mathbf{B}_{\top}} \varphi$ if and only if $\models_{\mathbf{KripkeM}} \varphi$,
- $\models_{\mathbf{CB}_{\top}} \varphi$ if and only if $\models_{\mathbf{ReflKripkeM}} \varphi$.

The notion of universal context allows us to easily define the notion of model with *maximal general uncertainty*, that is, to say a model in which the agents have maximal uncertainty about the state of the world as well as maximal uncertainty about the agents' k -order beliefs for every $k \geq 1$, where the order of a belief is defined inductively as follows: (i) an agent's belief is order 1 if and only if its content is a propositional formula that does not mention beliefs; (ii) an agent's belief is order $k + 1$ if and only if it is a belief about the agents' k -order beliefs.

Definition 15 (MBM with maximal general uncertainty). Let $B = (B_1, \dots, B_n, V) \in \mathbf{B}$ such that $B_i = \emptyset$ for every $i \in \text{Agt}$. Then, (B, \mathbf{B}_{\top}) is called *MBM with maximal general uncertainty*, while (B, \mathbf{CB}_{\top}) is called *MBM with maximal general uncertainty and belief correctness*.

Note that in our semantics maximal general uncertainty coincides with the fact that the agents' belief bases are empty (i.e., the agents do not know anything). The following proposition is a direct consequence of Theorem 13.

Proposition 16. Let $\varphi \in \mathcal{L}_{EL}$, let $i \in \text{Agt}$ and let (B, \mathbf{B}_{\top}) (resp. (B, \mathbf{CB}_{\top})) be a MBM with maximal general uncertainty (resp. a MBM with maximal general uncertainty and belief correctness). Then,

- φ is satisfiable for the class \mathbf{M} if and only if $(B, \mathbf{B}_{\top}) \models \diamond_i \varphi$,
- φ is satisfiable for the class \mathbf{M}_{BC} if and only if $(B, \mathbf{CB}_{\top}) \models \diamond_i \varphi$.

It highlights the essential aspects of models with maximal general uncertainty. To see this, suppose φ and $\neg\varphi$ are both satisfiable for the class \mathbf{M} (resp. for the class \mathbf{M}_{BC}). Then, by Proposition 16, we have $(B, \mathbf{B}_{\top}) \models \diamond_i \varphi \wedge \diamond_i \neg\varphi$ (resp. $(B, \mathbf{CB}_{\top}) \models \diamond_i \varphi \wedge \diamond_i \neg\varphi$), where $\diamond_i \varphi \wedge \diamond_i \neg\varphi$ captures agent i 's uncertainty about φ . More generally, for every formula φ yielding information either about the state of the world or about the agents' higher-order beliefs, if φ and $\neg\varphi$ are both satisfiable, then in a model with maximal uncertainty every agent i has uncertainty about φ .

4 Only believing

The aim of this section is to show that if we increase the expressive power of our multimodal epistemic language by the notions of 'believing at most' and 'only believing', then the equivalence result between the generic semantics in terms of MBMs and the universal context semantics of Section 3 does not hold anymore.

Let $\mathcal{L}_{EEL}(Atm, Agt)$ be the language which extends language \mathcal{L}_{EL} by modal operators of implicitly believing 'at most' of the form ∇_i , where EEL stands for "Extended Epistemic Logic". It is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_i \varphi \mid \nabla_i \varphi$$

where p ranges over Atm and i ranges over Agt . For simplicity, we write \mathcal{L}_{EEL} instead of $\mathcal{L}_{EEL}(Atm, Agt)$, when the context is unambiguous.

The formula $\nabla_i\varphi$ has to be read “agent i believes at most that $\neg\varphi$ ” and has the following interpretation relative to MBMs.

Definition 17 (Satisfaction relation (cont.)). *Let $B = (B_1, \dots, B_n, V) \in \mathbf{B}$ and let $(B, Cxt) \in \mathbf{M}$. Then:*

$$(B, Cxt) \models \nabla_i\varphi \iff \forall B' \in (Cxt \setminus \mathcal{R}_i(B)) : (B', Cxt) \models \varphi$$

where $\mathcal{R}_i(B) = \{B' : B\mathcal{R}_iB'\}$.

The definition of modal depth of a formula in \mathcal{L}_{EEL} extends the definition of the modal depth of formulas in \mathcal{L}_{EL} by the following additional clause:

$$depth(\nabla_i\varphi) = depth(\varphi) + 1$$

By combining the operators \Box_i and ∇_i in the appropriate way, we can reconstruct the universal modality [25, 16]:

$$U\varphi \stackrel{\text{def}}{=} \Box_i\varphi \wedge \nabla_i\varphi$$

where $U\varphi$ has to be read “ φ is universally true”. We can moreover reconstruct the “only believing” modality [31, 18, 30]:

$$O_i\varphi \stackrel{\text{def}}{=} \Box_i\varphi \wedge \nabla_i\neg\varphi$$

where $O_i\varphi$ has to be read “all that agent i believes is φ ”. From the perspective of the logic of only believing, all that agent i believes is φ if and only if agent i believes at least that φ is true (i.e., $\Box_i\varphi$) and she believes at most that φ is true (i.e., $\nabla_i\neg\varphi$).

Definitions of validity and satisfiability for formulas in \mathcal{L}_{EEL} relative to the class \mathbf{M} (resp. \mathbf{M}_{BC}) and relative to the universal context \mathbf{B}_\top (resp. \mathbf{CB}_\top) coincide with those for formulas in \mathcal{L}_{EL} given above.

The following theorem highlights that the semantics for the language \mathcal{L}_{EEL} based on the universal context contains more validities than the semantics for the language \mathcal{L}_{EEL} based on the generic class \mathbf{M} .

Theorem 18. *We have the following relationship between sets of validities:*

- $\{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{M}} \varphi\} \subset \{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{B}_\top} \varphi\}$,
- $\{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{M}_{BC}} \varphi\} \subset \{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{CB}_\top} \varphi\}$.

Showing that $\{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{M}} \varphi\} \subseteq \{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{B}_\top} \varphi\}$ is trivial since $(B, \mathbf{B}_\top) \in \mathbf{M}$ for every $B \in \mathbf{B}_\top$. Furthermore, it is easy to find a formula χ such that $\not\models_{\mathbf{M}} \chi$ and $\models_{\mathbf{B}_\top} \chi$. The following formula is an example:

$$\chi \stackrel{\text{def}}{=} E\left(\bigwedge_{p \in X} p \wedge \bigwedge_{q \in Y} \neg q\right)$$

with $E\psi \stackrel{\text{def}}{=} \neg U\neg\psi$ and finite $X, Y \subseteq Atm$ such that $X \cap Y = \emptyset$.

Note that, we can easily adapt the proof of Theorem 10 to show that the set of \mathcal{L}_{EEL} -validities relative to the class of multi-relational Kripke models and the set of validities relative to the generic belief base semantics are the same.

Theorem 19. *Let $\varphi \in \mathcal{L}_{EEL}$. Then,*

- $\models_{\mathbf{M}} \varphi$ if and only if $\models_{\mathbf{KripkeM}} \varphi$,
- $\models_{\mathbf{M}_{BC}} \varphi$ if and only if $\models_{\mathbf{ReflKripkeM}} \varphi$.

The following is a direct consequence of Theorems 18 and 19.

Theorem 20. *We have the following relationship between sets of validities:*

- $\{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{KripkeM}} \varphi\} \subset \{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{B}_\top} \varphi\}$,
- $\{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{ReflKripkeM}} \varphi\} \subset \{\varphi \in \mathcal{L}_{EEL} : \models_{\mathbf{CB}_\top} \varphi\}$.

5 Qualitative belief structures

In this section, we consider a different representation of the universal epistemic model containing all possible belief hierarchies for the agents in the system first proposed by [12] (see also [13]). Differently from the definition of the universal context given in Section 3, which does not require any inductive construction of belief hierarchies, Fagin et al.'s definition of the universal epistemic model is inductive. As emphasized by [6], Fagin et al.'s universal epistemic model can be seen as qualitative counterpart of the notion of probabilistic universal type space by [36]. A similar inductive construction of the universal epistemic model was proposed more recently by [3] (see also [1]), as a semantics for the logic of multi-agent only knowing.

Fagin et al. study the universal epistemic model for the multimodal logic $S5^n$, including both positive and negative introspection principles for knowledge. We here consider variants of Fagin et al.'s construction for the multimodal logics K^n and KT^n , the latter including the veracity principle for knowledge according to which what an agent knows cannot be false.²

Given an arbitrary set Y , let $F(Y)$ be the set of functions with domain Y and codomain $\{0, 1\}$. We define the set Z_k in an inductive way as follows:

$$\begin{aligned} Z_0 &= F(Atm) \\ Z_{k+1} &= Z_k \times F(Z_k)^n. \end{aligned}$$

Thus, we have

$$Z_{k+1} = Z_0 \times F(Z_0)^n \times \dots \times F(Z_k)^n.$$

Elements of Z_0 are denoted by f_0 , while elements of $F(Z_k)^n$ are denoted by f_{k+1}, f'_{k+1}, \dots . Moreover, for every $i \in Agt$ and for every $f_{k+1} \in F(Z_k)^n$, $f_{k+1}(i)$ denotes the i -th component in the tuple f_{k+1} .

Elements of Z_k are called $k+1$ -ary worlds, or simply $k+1$ -worlds. Elements of $\bigcup_{k \in \mathbb{N}_0} Z_k$ are called worlds. The set of worlds is denoted by \mathbf{W} .

The set of belief structures is defined as follows:

$$\mathbf{BS} = Z_0 \times F(Z_0)^n \times F(Z_1)^n \times \dots$$

Thus, a belief structure is a countably infinite sequence $f = (f_0, f_1, \dots)$ such that $f_0 \in Z_0$ and, for every $k \in \mathbb{N}_0$, $f_{k+1} \in F(Z_k)^n$. It follows that $f = (f_0, f_1, \dots)$ is a belief structure if and only if, for every $k \in \mathbb{N}_0$, (f_0, \dots, f_k) is a $k+1$ -world.

²Belle & Lakemeyer [3] offer an inductive construction of the universal epistemic model for the multimodal logic $K45^n$.

Definition 21 (Coherent world and belief structure). A world (f_0, \dots, f_k) is said to be coherent if and only if, the following conditions hold:

- for every $i \in \text{Agt}$ and for every $2 \leq h \leq k$:
if $f_h(i)(g_0, \dots, g_{h-1}) = 1$ then $f_{h-1}(i)(g_0, \dots, g_{h-2}) = 1$;
- for every $i \in \text{Agt}$ and for every $1 \leq h \leq k-1$:
if $f_h(i)(g_0, \dots, g_{h-1}) = 1$ then there exists $g_h \in F(Z_{h-1})^n$ such that $f_{h+1}(i)(g_0, \dots, g_{h-1}, g_h) = 1$.

The set of coherent worlds is denoted by **CW**. We say that the belief structure $f = (f_0, f_1, \dots)$ is coherent if, for every $k \in \mathbb{N}_0$, (f_0, \dots, f_k) is a coherent world. The set of coherent belief structures is denoted by **CBS**.

The following definition introduces the property of correctness for worlds and belief structures.

Definition 22 (Correct world and belief structure). A world (f_0, \dots, f_k) with $k \geq 1$ satisfies belief correctness if and only if, for every $i \in \text{Agt}$, $f_k(i)(f_0, \dots, f_{k-1}) = 1$. The set of coherent worlds satisfying belief correctness (BC) is denoted by **CW_{BC}**. We say that the belief structure $f = (f_0, f_1, \dots)$ is correct if, for every $k \geq 1$, (f_0, \dots, f_k) is a correct world. The set of coherent belief structures satisfying belief correctness (BC) is denoted by **CBS_{BC}**.

A formula φ of the language \mathcal{L}_{EEL} is interpreted relative to a coherent $k+1$ -world $(f_0, \dots, f_k) \in \text{CW}$ such that $k \geq \text{depth}(\varphi)$, as follows:

$$\begin{aligned}
(f_0, \dots, f_k) \models p &\iff f_0(p) = 1 \\
(f_0, \dots, f_k) \models \neg\varphi &\iff (f_0, \dots, f_k) \not\models \varphi \\
(f_0, \dots, f_k) \models \varphi \wedge \psi &\iff (f_0, \dots, f_k) \models \varphi \text{ and } (f_0, \dots, f_k) \models \psi \\
(f_0, \dots, f_k) \models \Box_i\varphi &\iff \forall (g_0, \dots, g_{k-1}) \in Z_{k-1} : \\
&\quad \text{if } f_k(i)(g_0, \dots, g_{k-1}) = 1 \\
&\quad \text{then } (g_0, \dots, g_{k-1}) \models \varphi \\
(f_0, \dots, f_k) \models \nabla_i\varphi &\iff \forall (g_0, \dots, g_{k-1}) \in Z_{k-1} : \\
&\quad \text{if } f_k(i)(g_0, \dots, g_{k-1}) = 0 \\
&\quad \text{then } (g_0, \dots, g_{k-1}) \models \varphi
\end{aligned}$$

The following proposition is a generalization of [12, Lemma 2.5] to the language \mathcal{L}_{EEL} .³

Proposition 23. Let $\text{depth}(\varphi) = h$, $k \geq h$, and $(f_0, \dots, f_h), (f_0, \dots, f_k) \in \text{CW}$. Then, $(f_0, \dots, f_h) \models \varphi$ if and only if $(f_0, \dots, f_k) \models \varphi$.

Let $\varphi \in \mathcal{L}_{\text{EEL}}$ and let $f = (f_0, f_1, \dots) \in \text{CBS}$ be a coherent belief structure. We say that f satisfies φ , denoted by $f \models \varphi$ if $(f_0, \dots, f_h) \models \varphi$ where $h = \text{depth}(\varphi)$. It is worth noting that, by Proposition 23, if $k \geq h$ then $f \models \varphi$ if and only if $(f_0, \dots, f_k) \models \varphi$. We moreover say that φ is valid relative to the class **CBS**, denoted by $\models_{\text{CBS}} \varphi$, if $f \models \varphi$ for every $f \in \text{CBS}$. We say that φ is satisfiable relative to the class **CBS** if $\not\models_{\text{CBS}} \neg\varphi$. Definitions of validity relative to the class **CBS_{BC}** (denoted by $\models_{\text{CBS}_{\text{BC}}} \varphi$) and satisfiability are defined analogously.

It is straightforward to adapt the proof of [12][Theorem 3.1] in order to prove the following equivalence result between the coherent belief structure semantics and the Kripke semantics relative to the language \mathcal{L}_{EL} .

³[12, Lemma 2.5] only applies to the standard epistemic language with epistemic operators \Box_i and without operators ∇_i .

Theorem 24. *Let $\varphi \in \mathcal{L}_{EL}$. Then,*

- $\models_{\mathbf{CBS}} \varphi$ if and only if $\models_{\mathbf{KripkeM}} \varphi$,
- $\models_{\mathbf{CBS}_{BC}} \varphi$ if and only if $\models_{\mathbf{RefKripkeM}} \varphi$.

The following theorem is a corollary of Theorems 14 and 24.

Theorem 25. *Let $\varphi \in \mathcal{L}_{EL}$. Then,*

- $\models_{\mathbf{B}_\top} \varphi$ if and only if $\models_{\mathbf{CBS}} \varphi$,
- $\models_{\mathbf{CB}_\top} \varphi$ if and only if $\models_{\mathbf{CBS}_{BC}} \varphi$.

The following theorem strengthens Theorem 25 by stating that the set of \mathcal{L}_{EEL} -validities relative to the universal context \mathbf{B}_\top (resp. \mathbf{CB}_\top) is the same as the set of \mathcal{L}_{EEL} -validities relative to the class \mathbf{CBS} (resp. \mathbf{CBS}_{BC}).

Theorem 26. *Let $\varphi \in \mathcal{L}_{EEL}$. Then,*

- $\models_{\mathbf{B}_\top} \varphi$ if and only if $\models_{\mathbf{CBS}} \varphi$,
- $\models_{\mathbf{CB}_\top} \varphi$ if and only if $\models_{\mathbf{CBS}_{BC}} \varphi$.

More generally, not even the epistemic language \mathcal{L}_{EEL} — which extends the language \mathcal{L}_{EL} by the notions of ‘believing at most’ and ‘only believing’ — can distinguish the belief structure semantics à la Fagin et al. from the universal context semantics exploiting the belief base abstraction.

6 Model checking

The notions of α -context and universal context defined in Section 3 (Definition 11) allow us to offer a compact formulation of the model checking problem for the formulas in the language \mathcal{L}_{EL} .⁴

α -context model checking
Given: $\varphi \in \mathcal{L}_{EL}$, $\alpha \in \mathcal{L}_0$ and a finite $B \in \mathbf{B}$.
Question: Do we have $(B, \mathbf{B}_\alpha) \models \varphi$?

where the multi-agent belief base (MBB) $B = (B_1, \dots, B_n, V)$ is said to be finite if V and every B_i are finite.

Note that, if $\alpha = \top$, then model checking consists in verifying whether φ is true at a given finite MBB of the universal context, that is, in verifying whether $(B, \mathbf{B}_\top) \models \varphi$ for a specific finite $B \in \mathbf{B}_\top$.

The model checking problem for multi-agent epistemic logic \mathbf{K}^n interpreted relative to the standard multi-relational Kripke semantics of Definition 9 is known to be P-complete with respect to the size of the input formula to be checked and the size of the model [17]. In this section, we are going to show that the previous compact formulation of the model checking problem for the formulas in \mathcal{L}_{EL} is PSPACE-hard, the same complexity as the corresponding satisfiability problem. Our result highlights that the gains in terms of compactness of the model checking problem are counterbalanced by losses on the complexity side.

We provide a polynomial reduction of true quantified boolean formulas (TQBF) to model checking for formulas in \mathcal{L}_{EL} .

⁴ Other compact formulations of the epistemic logic model checking have been proposed in the literature. For instance [32] provide a semantics for epistemic logic model checking based on the concept of interpreted system, while the approach by [5, 27, 10] builds on propositional observability.

Let us assume that a quantified boolean formula χ is in (not necessarily closed) prenex normal form and can be written as follows:

$$\lambda.\pi$$

where π is a propositional formula and λ is a (possibly empty) sequence $Q_0p_0\dots Q_m p_m$ such that $\{p_0, \dots, p_m\} \subseteq \text{Atm}(\pi)$ and $Q_0, \dots, Q_m \in \{\exists, \forall\}$, with $\text{Atm}(\pi)$ denoting the set of atoms occurring in π . If $\{p_0, \dots, p_m\} = \text{Atm}(\pi)$ then χ is said to be *closed*, as it has no free variables. The *length* of χ is defined to be the length of the sequence $p_0\dots p_m$. For notational convenience, when λ is non empty, we write $\lambda[k]$ to denote the propositional variable in the k -th position in the sequence $p_0\dots p_m$.

Let us suppose that the language of quantified boolean formulas is built over the set of atomic variables Atm . Moreover, let us define the following translation from the language of quantified boolean formulas to the language $\mathcal{L}_{\text{EL}}(\text{Atm}, \{1\})$:

$$\begin{aligned} \text{tr}(p) &= \text{for } p \in \text{Atm} \\ \text{tr}(\forall p \lambda.\pi) &= \llbracket \lambda \rrbracket \text{tr}(\lambda.\pi) \end{aligned}$$

where

$$\llbracket \lambda \rrbracket \psi \stackrel{\text{def}}{=} \Box_1(\bigcirc_1 \lambda[0] \rightarrow \psi)$$

and

$$\begin{aligned} \bigcirc_1 \lambda[0] &\stackrel{\text{def}}{=} \Diamond_1 \lambda[0] \wedge \Diamond_1 \neg \lambda[0] && \text{if } \lambda \text{ is non-empty} \\ \bigcirc_1 \lambda[0] &\stackrel{\text{def}}{=} \top && \text{otherwise} \end{aligned}$$

and where the Boolean connectives are translated homomorphically. We define the dual operator of $\llbracket \lambda \rrbracket$ as follows:

$$\langle\langle \lambda \rangle\rangle \psi \stackrel{\text{def}}{=} \neg \llbracket \lambda \rrbracket \neg \psi.$$

For every quantified boolean formula $\chi = \lambda.\pi = Q_0p_0\dots Q_m p_m.\pi$ and for every $V \subseteq \text{Atm}$, we define:

$$\begin{aligned} \Sigma_{\chi, V} &= ((\text{Atm}(\pi) \setminus \{\lambda[0]\}) \cap V) \cup \\ &\quad \{\neg q : q \in (\text{Atm}(\pi) \setminus \{\lambda[0]\}) \setminus V\} \cup \\ &\quad \bigcup_{1 \leq k \leq m} \{\Delta_1^{k-1} \bigwedge_{q \in \text{Atm}(\pi) \setminus \{\lambda[k]\}} (\Delta_1 q \vee \Delta_1 \neg q)\} \end{aligned}$$

where $\Delta_1^{k-1} \alpha$ is inductively defined as follows:

$$\begin{aligned} \Delta_1^{k-1} \alpha &\stackrel{\text{def}}{=} \alpha && \text{if } k = 1 \\ \Delta_1^{k-1} \alpha &\stackrel{\text{def}}{=} \Delta_1 \Delta_1^{k-2} \alpha && \text{if } k > 1 \end{aligned}$$

Moreover, let us define

$$B^{(\chi, V)} = (B_1^{(\chi, V)}, V)$$

where

$$\begin{aligned} B_1^{(\chi, V)} &= \Sigma_{\chi, V} && \text{if } \lambda \text{ is non-empty} \\ B_1^{(\chi, V)} &= \emptyset && \text{otherwise} \end{aligned}$$

The following lemma is crucial for proving the main result of this section about complexity of model checking for formulas in \mathcal{L}_{EL} .

Lemma 1. *Let $V \subseteq Atm$ and let $\chi = \lambda.\pi$ be a quantified boolean formula. Then, $V \models \chi$ if and only if $(B^{(\chi, V)}, \mathbf{B}_\top) \models tr(\chi)$.*

Thanks to Lemma 1 and the fact that the TQBF problem (i.e., the problem of checking whether a quantified boolean formula is true in a given valuation) is PSPACE-hard [41], we have a PSPACE-hardness result for model checking of \mathcal{L}_{EL} -formulas. Indeed, let $\chi = Q_0 p_0 \dots Q_m p_m.\pi$ be a quantified boolean formula. Then, the sizes of $B_1^{(\chi, V)}$ and $tr(\chi)$ are polynomial in m . Therefore, by Lemma 1 we have a polynomial reduction of the TQBF problem into EL-model checking. Since the former is PSPACE-hard, then the latter is PSPACE-hard too.

Theorem 27. *The α -context model checking problem for formulas in \mathcal{L}_{EL} is PSPACE-hard.*

From [19] we know that satisfiability checking for \mathcal{L}_{EL} -formulas relative to multi-relational Kripke models of Definition 9 is PSPACE-complete. Thus, thanks to Theorem 14, we can conclude that satisfiability checking for \mathcal{L}_{EL} -formulas relative to the universal context \mathbf{B}_\top is PSPACE-complete as well. Consequently, Theorem 27 also shows that the compact version of model checking for \mathcal{L}_{EL} -formulas exploiting the notion of α -context is at least as complex as satisfiability checking for \mathcal{L}_{EL} -formulas.

7 Conclusion

We have offered a novel representation of the universal epistemic model which does not require an inductive construction of the agents' belief hierarchies. Our solution exploits the notion of belief base, a natural abstraction for representing epistemic attitudes of rational players in interactive situations, widely used in the area of knowledge representation and reasoning (KR). Our lower bound complexity result for model checking relative to the universal epistemic model clearly displays the trade-off between the compactness of the model checking problem representation and its computational complexity.

Directions of future research are manifold. As emphasized in Section 5, Fagin et al.'s original definition of the universal epistemic model applies to the epistemic logic $S5^n$ which includes principles of positive and negative introspection for knowledge (i.e., if an agent knows that φ /does not know that φ , then she knows that she knows that φ /does not know that φ). The logics we presented in this paper do not make any assumption about introspection for belief or knowledge. Future work will be devoted to extend our analysis of the universal epistemic model to variants of epistemic logic with introspection. We also plan to extend our comparative analysis of the different semantics for epistemic logic to more expressive epistemic languages including common belief and distributed belief operators. Last but not least, we plan to explore the connection between the representation of agents' belief hierarchies using the belief base abstraction and type spaces with finite depth of reasoning as defined in [28]. Indeed, the belief base abstraction allows to naturally define the notion of k -level (resource-bounded) reasoner, as an agent whose belief base contains explicit beliefs of at most order k . This is clearly related to the idea of type spaces with finite depth of reasoning à la Kets.

Acknowledgements

Support from the ANR project CoPains and the ANR-3IA Artificial and Natural Intelligence Toulouse Institute is gratefully acknowledged.

References

- [1] G. Aucher & V. Belle (2015): *Multi-Agent Only Knowing on Planet Kripke*. In: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, AAAI Press, pp. 2713–2719.
- [2] P. Battigalli & M. Siniscalchi (1999): *Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games*. *Journal of Economic Theory* 88(1), pp. 188–230, doi:10.1006/jeth.1999.2555.
- [3] V. Belle & G. Lakemeyer (2010): *Multi-Agent Only-Knowing Revisited*. In: *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR 2010)*, AAAI Press, pp. 49–59.
- [4] S. Benferhat, D. Dubois, H. Prade & M.-A. Williams (2002): *A practical approach to revising prioritized knowledge bases*. *Studia Logica* 70(1), pp. 105–130.
- [5] J. van Benthem, J. van Eijck, M. Gattinger & K. Su (2015): *Symbolic Model Checking for Dynamic Epistemic Logic*. In: *Proceedings of the 5th International Workshop on Logic, Rationality and Interaction (LORI 2015)*, LNCS 9394, Springer-Verlag, pp. 366–378.
- [6] J. van Benthem & D. Klein (2019): *Logics for Analyzing Games*. In: *Stanford Encyclopedia of Philosophy*.
- [7] A. Bjorndahl & J. Y. Halpern (2017): *From Type Spaces to Probability Frames and Back, via Language*. In: *Proceedings of the Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017)*, pp. 75–87.
- [8] P. Blackburn, M. de Rijke & Y. Venema: *Modal Logic*. Cambridge University Press, Cambridge, doi:https://doi.org/10.1017/CBO9781107050884.
- [9] A. Brandenburger & E. Dekel (1993): *Hierarchies of Beliefs and Common Knowledge*. *Journal of Economic Theory* 59, pp. 189–198, doi:https://doi.org/10.1006/jeth.1993.1012.
- [10] T. Charrier, A. Herzig, E. Lorini, F. Maffre & F. Schwarzentruber (2016): *Building Epistemic Logic from Observations and Public Announcements*. In: *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, AAAI Press, pp. 268–277.
- [11] H. P. van Ditmarsch, W. van der Hoek & B. Kooi (2007): *Dynamic Epistemic Logic*. Kluwer Academic Publishers, doi:10.1007/978-1-4020-5839-4.
- [12] R. Fagin, J. Y. Halpern, & M. Y. Vardi (1991): *A Model-Theoretic Analysis of Knowledge*. *Journal of the ACM* 38(2), pp. 382–428, doi:10.1145/103516.128680.
- [13] R. Fagin, J. Geanakoplos, J. Y. Halpern, & M. Y. Vardi (1999): *The Hierarchical Approach to Modeling Knowledge and Common Knowledge*. *International Journal of Game Theory* 28(3), pp. 331–365.
- [14] R. Fagin, J. Halpern, Y. Moses & M. Vardi (1995): *Reasoning about Knowledge*. MIT Press, Cambridge, doi:10.7551/mitpress/5803.001.0001.
- [15] P. Galeazzi & E. Lorini (2016): *Epistemic logic meets epistemic game theory: a comparison between multi-agent Kripke models and type spaces*. *Synthese* 193(7), pp. 2097–2127, doi:https://doi.org/10.1007/s11229-015-0834-x.
- [16] V. Goranko & S. Passy (1992): *Using the universal modality: gains and questions*. *Journal of Logic and Computation* 2(1), pp. 5–30, doi:https://doi.org/10.1093/logcom/2.1.5.
- [17] E. Grädel & M. Otto (1999): *On logics with two variables*. *Theoretical Computer Science* 224, pp. 73–113, doi:https://doi.org/10.1016/S0304-3975(98)00308-9.

- [18] J. Y. Halpern & G. Lakemeyer (2001): *Multi-agent only knowing*. *Journal of Logic and Computation* 11(1), pp. 41–70, doi:10.1093/logcom/11.1.41.
- [19] J. Y. Halpern & Y. Moses (1992): *A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief*. *Artificial Intelligence* 54(2), pp. 319–379, doi:https://doi.org/10.1016/0004-3702(92)90049-4.
- [20] S. O. Hansson (1993): *Theory contraction and base contraction unified*. *Journal of Symbolic Logic* 58(2), pp. 602–625, doi:10.1007/978-3-319-20451-2_14.
- [21] S. O. Hansson (1999): *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer, Dordrecht.
- [22] J. C. Harsanyi (1967): *Games with incomplete information played by ‘Bayesian’ players*. *Management Science* 14, pp. 159–182, doi:https://doi.org/10.1287/mnsc.1040.0270.
- [23] A. Heifetz (1993): *The Bayesian formulation of incomplete information: the non-compact case*. *International Journal of Game Theory* 21, pp. 329–338, doi:https://doi.org/10.1007/BF01240148.
- [24] A. Heifetz & D. Samet (1998): *Topology-Free Typology of Beliefs*. *Journal of Economic Theory* 82, pp. 324–341, doi:https://doi.org/10.1006/jeth.1998.2435.
- [25] E. Hemaspaandra (1996): *The Price of Universality*. *Notre Dame Journal of Formal Logic* 37(2), pp. 174–203, doi:10.1305/ndjfl/1040046086.
- [26] J. Hintikka (1962): *Knowledge and Belief. An introduction to the logic of the two notions*. Cornell University Press, New York.
- [27] W. van der Hoek, P. Iliev & M. Wooldridge (2012): *A logic of revelation and concealment*. In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, IFAAMAS, pp. 1115–1122.
- [28] W. Kets (2014): *Finite Depth of Reasoning and Equilibrium Play in Games with Incomplete Information*. Northwestern University, Center for Mathematical Studies in Economics and Management Science Discussion Papers 1569.
- [29] S. Konieczny & R. Pino Pérez (2002): *Merging information under constraints: a logical framework*. *Journal of Logic and Computation* 12(5), pp. 773–808, doi:https://doi.org/10.1093/logcom/12.5.773.
- [30] G. Lakemeyer (1993): *All they know: a study in multi-agent autoepistemic reasoning*. In: *Proceedings of the 13th International Joint Conference on Artificial intelligence (IJCAI’93)*, Morgan Kaufmann, pp. 376–381.
- [31] H. J. Levesque (1990): *All I know: a study in autoepistemic logic*. *Artificial Intelligence* 42(2-3), pp. 263–309, doi:https://doi.org/10.1016/0004-3702(90)90056-6.
- [32] A. Lomuscio, H. Qu & F. Raimondi (2015): *MCMAS: an open-source model checker for the verification of multi-agent systems*. *International Journal on Software Tools for Technology Transfer* 19, pp. 1–22, doi:https://doi.org/10.1007/s10009-015-0378-x.
- [33] E. Lorini (2018): *In Praise of Belief Bases: Doing Epistemic Logic Without Possible Worlds*. In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, AAAI Press, pp. 1915–1922.
- [34] E. Lorini & F. Romero (2019): *Decision procedures for epistemic logic exploiting belief bases*. In: *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, ACM, pp. 944–952.
- [35] D. Makinson & L. van der Torre (2000): *Input/output logics*. *Journal of Philosophical Logic* 29, pp. 383–408, doi:https://doi.org/10.1023/A:1004748624537.
- [36] J. F. Mertens & S. Zamir (1985): *Formulation of Bayesian analysis for games with incomplete information*. *International Journal of Game Theory* 14, pp. 1–29, doi:https://doi.org/10.1007/BF01770224.
- [37] J.-J. C. Meyer & W. van der Hoek (1995): *Epistemic Logic for AI and Theoretical Computer Science*. Cambridge University Press, Oxford, doi:https://doi.org/10.1017/CBO9780511569852.
- [38] R. Reiter (1988): *On integrity constraints*. In: *Proceedings of the 2nd Conference on Theoretical aspects of Reasoning about Knowledge (TARK’88)*, Morgan Kaufmann Publishers, pp. 97–111.

- [39] Y. Shoham (2009): *Logical Theories of Intention and the Database Perspective*. *Journal of Philosophical Logic* 38(6), pp. 633–648, doi:https://doi.org/10.1007/s10992-009-9116-8.
- [40] R. Stalnaker (2002): *Common ground*. *Linguistics and Philosophy* 25(5-6), pp. 701–721, doi:10.1023/A:1020867916902.
- [41] L. J. Stockmeyer (1973): *Word problems requiring exponential time (Preliminary Report)*. In: *Proceedings of the Fifth Annual ACM symposium on Theory of Computing (STOC '73)*, ACM, pp. 1–9.

A Proofs

This appendix presents a selection of the proofs of the technical results presented in the paper.

A.1 Proof of Theorem 10

Proof. We only prove the first item. The proof of the second item is similar.

In order to prove the right-to-left direction we show that if φ is satisfiable for \mathbf{M} then it is satisfiable for the class of multi-relational Kripke models. Let $(B, Cxt) \in \mathbf{M}$ such that $(B, Cxt) \models \varphi$. We build the multi-relational Kripke model $M = (W, \Rightarrow_1, \dots, \Rightarrow_n, \omega)$ corresponding to (B, Cxt) as follows:

- $W = \{w_{B'} : B' \in \{B\} \cup Cxt\}$,
- for every $i \in Agt$ and for every $w_{B'}, w_{B''} \in W$, $w_{B'} \Rightarrow_i w_{B''}$ iff $B' \mathcal{R}_i B''$,
- for every $p \in Atm$, $\omega(p) = \{w_{B'} \in W : B' \models p\}$,

Clearly, we have $(M, w_B) \models \varphi$ iff $(B, Cxt) \models \varphi$. Thus, $(M, w_B) \models \varphi$ since we supposed that $(B, Cxt) \models \varphi$.

As for the left-to-right direction suppose φ is satisfiable for the class of multi-relational Kripke models. We know that the multimodal logic K^n interpreted relative to multi-relational Kripke models has the finite model property. Consequently, there exists a *finite* multi-relational Kripke model $M = (W, \Rightarrow_1, \dots, \Rightarrow_n, \omega)$ and $w \in W$ such that $(M, w) \models \varphi$. Let $name : W \rightarrow Atm \setminus Atm(\varphi)$ be an injective function. Such an injection exists since the set W is finite and the the set Atm is assumed to be infinite. We define the context $Cxt = \{B^v : v \in W\}$ where, for every $v \in W$ and for every $i \in Agt$,

$$B_i^v = \left\{ \bigvee_{u \in W: v \Rightarrow_i u} name(u) \right\},$$

and $V^v = (\omega(v) \cap Atm(\varphi)) \cup \{name(v)\}$. By induction on the structure of φ , it is routine task to show that $(B^w, Cxt) \models \varphi$ iff $(M, w) \models \varphi$. Thus, $(B^w, Cxt) \models \varphi$ since $(M, w) \models \varphi$. \square

A.2 Proof of Theorem 13

Proof. We prove the first item, as the proof of the second item is analogous.

Left-to-right direction is obvious. As for the right-to-left direction, we prove that if φ is satisfiable relative to the class \mathbf{M} , then it is satisfiable relative to the universal context.

Let $(B, Cxt) \in \mathbf{M}$ such that $(B, Cxt) \models \varphi$. We build the multi-relational Kripke model $M = (W, \Rightarrow_1, \dots, \Rightarrow_n, \omega)$ corresponding to (B, Cxt) as follows:

- $W = \{w_{B'} : B' \in \{B\} \cup Cxt\}$,
- for every $i \in Agt$ and for every $w_{B'}, w_{B''} \in W$, $w_{B'} \Rightarrow_i w_{B''}$ iff $B' \mathcal{R}_i B''$,
- for every $p \in Atm$, $\omega(p) = \{w_{B'} \in W : B' \models p\}$,

where the interpretation of formulas relative to a multi-relational Kripke model $M = (W, \Rightarrow_1, \dots, \Rightarrow_n, \omega)$ and a world w in W was defined in Section 2.4. Clearly, we have $(M, w_B) \models \varphi$ iff $(B, Cxt) \models \varphi$. Thus, we have $(M, w_B) \models \varphi$ since $(B, Cxt) \models \varphi$.

In what follows, for notational convenience we denote elements of W by w, v, u, \dots

We define the filtration of M . Let $\Sigma \subseteq \mathcal{L}_{EL}$ be an arbitrary finite set of formulas which is closed under subformulas. (Cf. Definition 2.35 in [8] for a definition of subformulas closed set of formulas.) Let the equivalence relation \equiv_Σ on W be defined as follows. For all $w, v \in W$:

$$w \equiv_\Sigma v \text{ iff } \forall \varphi \in \Sigma : (M, w) \models \varphi \text{ iff } (M, v) \models \varphi.$$

Let $|w|_\Sigma$ be the equivalence class of the world w with respect to the equivalence relation \equiv_Σ .

We define W_Σ to be the filtrated set of worlds with respect to Σ :

$$W_\Sigma = \{|w|_\Sigma : w \in W\}.$$

Clearly, W_Σ is a finite set.

Let us define the filtrated valuation function ω_Σ . For every $p \in \text{Atm}$, we define:

$$\begin{aligned} \omega_\Sigma(p) &= \{|w|_\Sigma : (M, w) \models p\} && \text{if } p \in \text{Atm}(\Sigma) \\ \omega_\Sigma(p) &= \emptyset && \text{otherwise} \end{aligned}$$

with $\text{Atm}(\Sigma) = \bigcup_{\psi \in \Sigma} \text{Atm}(\psi)$.

Finally, for every $i \in \text{Agt}$, we define agent i 's accessibility relation as follows:

$$\Rightarrow_{i,\Sigma} = \{(|w|_\Sigma, |v|_\Sigma) : |w|_\Sigma, |v|_\Sigma \in W_\Sigma \text{ and } w \Rightarrow_i v\}.$$

Let us define:

$$\Rightarrow_\Sigma = \bigcup_{i \in \text{Agt}} \Rightarrow_{i,\Sigma}.$$

The model $(W_\Sigma, \Rightarrow_{1,\Sigma}, \dots, \Rightarrow_{n,\Sigma}, \omega_\Sigma)$ is the smallest filtration of M under Σ . Let $M_{\{\varphi\}} = (W_{\{\varphi\}}, \Rightarrow_{1,\{\varphi\}}, \dots, \Rightarrow_{n,\{\varphi\}}, \omega_{\{\varphi\}})$. Clearly, we have $(M_{\{\varphi\}}, |w_B|_{\{\varphi\}}) \models \varphi$ since $(M, w_B) \models \varphi$.

In what follows, for notational convenience we denote elements of $W_{\{\varphi\}}$ by x, y, \dots

The next step of the proof consists in unraveling the finite model $M_{\{\varphi\}}$ up to the modal depth of φ .

Let

$$\text{Seq}(\varphi) = \{(x_0, \dots, x_k) : k \leq \text{depth}(\varphi) \text{ and } x_0, \dots, x_k \in W_{\{\varphi\}}\}$$

be the set of sequences of worlds in $W_{\{\varphi\}}$ of length at most $\text{depth}(\varphi)$. Elements of $\text{Seq}(\varphi)$ are denoted by \vec{x}, \vec{y}, \dots . The length of the sequence \vec{x} is denoted by $\text{length}(\vec{x})$. For every $0 \leq k \leq \text{length}(\vec{x})$, $\vec{x}[k]$ denotes the k -th element in \vec{x} while $\vec{x}[\text{last}]$ denotes the last element in \vec{x} . Let $\vec{x} = (x_0, \dots, x_k)$, we write $\vec{x}.y$ to denote the sequence (x_0, \dots, x_k, y) .

We define the tree-like multi-relational Kripke model $M' = (W', \Rightarrow'_1, \dots, \Rightarrow'_n, \omega')$ as follows:

- $W' = \{\vec{x} \in \text{Seq}(\varphi) : \vec{x}[0] = |w_B|_{\{\varphi\}} \text{ and } \forall 0 \leq k < \text{length}(\vec{x}), \vec{x}[k] \Rightarrow_{\{\varphi\}} \vec{x}[k+1]\}$,
- for every $i \in \text{Agt}$ and for every $\vec{x}, \vec{y} \in W'$, $\vec{x} \Rightarrow'_i \vec{y}$ iff $\exists y \in W_{\{\varphi\}}$ such that $\vec{y} = \vec{x}.y$ and $\vec{x}[\text{last}] \Rightarrow_{i,\{\varphi\}} y$,

- for every $p \in \text{Atm}$, $\omega'(p) = \{\vec{x} \in W' : \vec{x}[\text{last}] \in \omega_{\{\varphi\}}(p)\}$.

Clearly, the tree-like model M' is finite.

By induction on the structure of φ , it is routine exercise to check that $(M', \vec{x}_0) \models \varphi$ iff $(M_{\{\varphi\}}, |w_B|_{\{\varphi\}}) \models \varphi$, with $\vec{x}_0 = (|w_B|_{\{\varphi\}})$. Thus, $(M', \vec{x}_0) \models \varphi$ since $(M_{\{\varphi\}}, |w_B|_{\{\varphi\}}) \models \varphi$.

Now, let us denote by $\text{Term}_{M'}$ the set of terminal nodes in the tree M' . That is, $\text{Term}_{M'} = \{\vec{x} \in W' : \exists \vec{y} \in W' \text{ such that } \vec{x} \Rightarrow' \vec{y}\}$ with $\Rightarrow' = \bigcup_{i \in \text{Agt}} \Rightarrow'_i$.

We define k -level explicit mutual belief that α in an inductive way as follows:

$$\begin{aligned} \text{EB}^0 \alpha &\stackrel{\text{def}}{=} \alpha \\ \text{EB}^{k+1} \alpha &\stackrel{\text{def}}{=} \text{EB} \text{EB}^k \alpha \\ \text{MB}^k \alpha &\stackrel{\text{def}}{=} \bigwedge_{0 \leq h \leq k} \text{EB}^h \alpha \end{aligned}$$

where $\text{EB} \alpha \stackrel{\text{def}}{=} \bigwedge_{i \in \text{Agt}} \Delta_i \alpha$.

We define the labelling function L over nodes in W' as follows:

$$\begin{aligned} L(\vec{x}) &= \bigwedge_{p \in \text{Val}(\varphi, \vec{x})} p \wedge \bigwedge_{p \in \text{Atm}(\varphi) \setminus \text{Val}(\varphi, \vec{x})} \neg p \wedge \bigwedge_{i \in \text{Agt}} \Delta_i \perp \\ &\text{if } \vec{x} \in \text{Term}_{M'} \\ L(\vec{x}) &= \bigwedge_{p \in \text{Val}(\varphi, \vec{x})} p \wedge \bigwedge_{p \in \text{Atm}(\varphi) \setminus \text{Val}(\varphi, \vec{x})} \neg p \wedge \bigwedge_{i \in \text{Agt}} \Delta_i \bigvee_{\vec{y}: \vec{x} \Rightarrow'_i \vec{y}} L(\vec{y}) \\ &\text{if } \vec{x} \in W' \setminus \text{Term}_{M'} \end{aligned}$$

where $\text{Val}(\varphi, \vec{x}) = \{p \in \text{Atm}(\varphi) : (M', \vec{x}) \models p\}$ is the set of atoms in φ which are true at \vec{x} . Note that every $L(\vec{x})$ is a (finitary) formula of the language \mathcal{L}_0 .

The labelling function is used to construct, for every $i \in \text{Agt}$, the \mathcal{L}_0 -formula which corresponds to the “tree” of valuations rooted in $(|w_B|_{\{\varphi\}})$.

Let $\vec{x}_0 = (|w_B|_{\{\varphi\}})$ and let $B' = (B'_1, \dots, B'_n, V')$ such that:

- for every $i \in \text{Agt}$, $B'_i = \bigcup_{p \in \text{Atm} \setminus \text{Atm}(\varphi)} \{\text{MB}^{\text{depth}(\varphi)} \neg p\} \cup \{\bigvee_{\vec{y} \in W': \vec{x}_0 \Rightarrow'_i \vec{y}} L(\vec{y})\}$,
- $V' = \{p : \vec{x}_0 \in \omega'(p)\}$.

Furthermore, let us define the (universal) multi-relational Kripke model $M'' = (W'', \Rightarrow''_1, \dots, \Rightarrow''_n, \omega'')$ corresponding to the universal context \mathbf{B}_\top as follows: (i) $W'' = \{w_{B''} : B'' \in \mathbf{B}_\top\}$, (ii) for every $i \in \text{Agt}$ and for every $w_{B''}, w_{B'''}$ in W , $w_{B''} \Rightarrow_i w_{B'''}$ iff $B'' \mathcal{R}_i B'''$, and (iii) for every $p \in \text{Atm}$, $\omega(p) = \{w_{B''} \in W : B'' \models p\}$.

By construction of B' , we can show that $(M'', w_{B'})$ and (M', \vec{x}_0) with $\vec{x}_0 = (|w_B|_{\{\varphi\}})$ are $\text{depth}(\varphi)$ -bisimilar [8, Definition 2.30]. Thus, $(M'', w_{B'}) \models \varphi$ since $(M', \vec{x}_0) \models \varphi$. It follows that $(B', \mathbf{B}_\top) \models \varphi$, since clearly $(M'', w_{B'}) \models \varphi$ iff $(B', \mathbf{B}_\top) \models \varphi$. \square

A.3 Proof of Proposition 23

Proof. The proof is by induction on the structure of the formulas. Boolean cases are trivial. The case $\varphi = \Box_i \psi$ is proved in the same way as [12, Lemma 2.5]. Let us prove the case $\varphi = \nabla_i \psi$. Assume that $k \geq h = \text{depth}(\varphi)$ and $(f_0, \dots, f_k) \models \nabla_i \psi$. Since $\varphi = \Box_i \psi$, we have $h \geq 1$. Moreover, let $(g_0, \dots, g_{h-1}) \in Z_{h-1}$ such that $f_h(i)(g_0, \dots, g_{h-1}) = 0$. By Definition 21, for all $(g_0, \dots, g_{k-1}) \in Z_{k-1}$, we have $f_k(i)(g_0, \dots, g_{k-1}) =$

0. Since $(f_0, \dots, f_k) \models \nabla_i \psi$, it follows by definition that $(g_0, \dots, g_{k-1}) \models \psi$ for all $(g_0, \dots, g_{k-1}) \in Z_{k-1}$. Thus, by induction hypothesis, $(g_0, \dots, g_{h-1}) \models \psi$. It follows that, for every $(g_0, \dots, g_{h-1}) \in Z_{h-1}$ such that $f_h(i)(g_0, \dots, g_{h-1}) = 0$, $(g_0, \dots, g_{h-1}) \models \psi$. Thus, $(f_0, \dots, f_h) \models \nabla_i \psi$. The proof of the other direction is similar. \square

A.4 Proof of Theorem 26

Proof. We only prove the first item. The second item can be proved in a similar way.

Let $\varphi \in \mathcal{L}_{\text{EEL}}$. We will prove that φ is satisfiable relative to \mathbf{B}_\top iff φ is satisfiable relative to the class **CBS**.

We build the multi-relational Kripke model $M = (W, \mathcal{R}_1, \dots, \mathcal{R}_n, \mathcal{R}_1^c, \dots, \mathcal{R}_n^c, \omega)$ corresponding to the $\text{Atm}(\varphi)$ -restriction of the universal context \mathbf{B}_\top as follows:

- $W = \{B' \in \mathbf{B}_\top : \forall q \in \text{Atm} \setminus \text{Atm}(\varphi) \text{ and } \forall k \geq 0 : (B', \mathbf{B}_\top) \models \text{MB}^k \neg q\}$,
- for all $i \in \text{Agt}$ and for all $B', B'' \in W$, $B' \mathcal{R}_i^c B''$ iff $B'' \notin \mathcal{R}_i(B')$,
- for all $p \in \text{Atm}$, $\omega(p) = \{B' \in W : B' \models p\}$,

where the k -level mutual belief operator MB^k is defined as in the proof of Theorem 13 (Section A.2). In what follows, for notational convenience we denote elements of W by w, v, u, \dots . Interpretation of \mathcal{L}_{EEL} -formulas relative to M and to a world $w \in W$ is defined as in Section 2.4 for boolean formulas and for formula $\Box_i \varphi$. We add the following clause for the ∇_i -operator:

$$(M, w) \models \nabla_i \varphi \text{ iff } \forall v \notin \mathcal{R}_i(w) : (M, v) \models \varphi$$

We leave to the reader the task of checking that φ is satisfiable relative to \mathbf{B}_\top iff M satisfies φ .

Similarly, we build the multi-relational Kripke model $M' = (W', \mathcal{T}_1, \dots, \mathcal{T}_n, \mathcal{T}_1^c, \dots, \mathcal{T}_n^c, \omega')$ corresponding to the $\text{Atm}(\varphi)$ -restriction of the set of coherent belief structures as follows:

- $W' = \{f \in \mathbf{CBS} : \forall q \in \text{Atm} \setminus \text{Atm}(\varphi), f_0(q) = 0 \text{ and if } g_0(q) = 1 \text{ then } f_k(i)(g_0, \dots, g_{k-1}) = 0, \forall k \geq 0, \forall i \in \text{Agt}, \forall g \in \mathbf{CBS}\}$,
- for all $i \in \text{Agt}$ and for all $f, g \in W'$, $f \mathcal{T}_i g$ iff $f_k(i)(g_0, \dots, g_{k-1}) = 1$ for all $k > 1$,
- for all $i \in \text{Agt}$ and for all $f, g \in W'$, $f \mathcal{T}_i^c g$ iff $g \notin \mathcal{T}_i(f)$,
- for all $p \in \text{Atm}$, $\omega'(p) = \{f \in W' : f_0(p) = 1\}$.

Interpretation of \mathcal{L}_{EEL} -formulas relative to M' and to a world $f \in W'$ is defined as usual. We have the following truth conditions for the operators \Box_i and ∇_i :

$$(M, w) \models \Box_i \varphi \text{ iff } \forall v \in W : \text{if } w \mathcal{T}_i v \text{ then } (M, v) \models \varphi$$

$$(M, w) \models \nabla_i \varphi \text{ iff } \forall v \in W : \text{if } w \mathcal{T}_i^c v \text{ then } (M, v) \models \varphi$$

We leave to the reader the task of checking that φ is satisfiable relative to **CBS** iff M' satisfies φ .

For every world $w \in W$ in the Kripke model M , we build a coherent belief structure $f^w = (f_0^w, f_1^w, \dots)$. We define f_0^w to be the function such that, for every $p \in \text{Atm}$, $f_0^w(p) = 1$ iff $w \in \omega(p)$. Moreover, suppose f_0^w, \dots, f_k^w have been defined for each $w \in W$. Then, we define f_{k+1}^w to be the function such that, for every $i \in \text{Agt}$, $f_{k+1}^w(i)^{-1}(\{1\}) = \{(g_0^v, \dots, g_k^v) : w \Rightarrow_i v\}$ where $f_{k+1}^w(i)^{-1}(\{1\})$ is the inverse image by $f_{k+1}^w(i)$ of the subset $\{1\}$ of the codomain $\{0, 1\}$.

Let us define the mapping $\tau : w \mapsto f^w$ from W to W' . The following intermediate proposition will be useful for the rest of the proof.

Proposition 28. *Let $f \in W'$. Then, there exists $v \in W$ such that $\tau(v) = f$.*

Proof. Let $f = (f_0, f_1, \dots) \in W'$. We show how to construct $v \in W$ such that $\tau(v) = f$.

Let us define:

$$\beta_{i,(f_0,f_1)} \stackrel{\text{def}}{=} \bigvee_{g_0 \in Z_0: f_1(i)(g_0)=1} g_0(\varphi)$$

where

$$g_0(\varphi) \stackrel{\text{def}}{=} \bigwedge_{p \in \text{Atm}(\varphi): g_0(p)=1} p \wedge \bigwedge_{p \in \text{Atm}(\varphi): g_0(p)=0} \neg p.$$

Moreover, for every $k > 1$, let us define:

$$\beta_{i,(f_0,\dots,f_k)} \stackrel{\text{def}}{=} \bigvee_{(g_0,\dots,g_{k-1}) \in Z_{k-1}: f_k(i)(g_0,\dots,g_{k-1})=1} (g_0(\varphi) \wedge \bigwedge_{j \in \text{Agt}} \Delta_j \beta_{j,(g_0,\dots,g_{k-1})}).$$

Note that $\beta_{i,(f_0,f_1)}$ and every $\beta_{i,(f_0,\dots,f_k)}$ are (finitary) formulas of \mathcal{L}_0 .

We define the multi-agent belief base $B^f = (B_1^f, \dots, B_n^f, V^f)$ as follows:

$$B_i^f = \bigcup_{k \geq 1} \{\beta_{i,(f_0,\dots,f_k)}\} \cup \bigcup_{k \geq 0, q \in \text{Atm} \setminus \text{Atm}(\varphi)} \{\text{MB}^k \neg q\},$$

$$V^f = \{p \in \text{Atm} : f_0(p) = 1\}.$$

Clearly, B^f belongs to W' . Moreover, it is the case that the multi-agent belief base B_i^f and the belief structure f define the same belief hierarchy. Specifically, we have $\tau(B^f) = f$. \square

By means of Proposition 28, we can show that the relation $\{(w, f) \in W \times W' : \tau(w) = f\}$ is a bisimulation between the Kripke models M and M' . Therefore, suppose $(M, v) \models \varphi$. It follows that $(M', \tau(v)) \models \varphi$. Viceversa, suppose $(M', f) \models \varphi$. Then, by Proposition 28, there exists $v \in W$ such that $\tau(v) = f$. Since v and f are bisimilar, we have $(M, v) \models \varphi$. \square

A.5 Proof of Lemma 1

Proof. In order to prove Lemma 1, we first prove the following intermediate result.

Lemma 2. *Let $V \subseteq \text{Atm}$ and let $\chi = \lambda.\pi$ be a quantified boolean formula. Moreover, let $B' = (B'_1, V')$, $B'' = (B''_1, V'') \in \mathbf{B}$ such that $(V \cap \text{Atm}(\pi)) = (V' \cap \text{Atm}(\pi)) = (V'' \cap \text{Atm}(\pi))$, $B_1^{(\chi, V)} \subseteq B'_1$, $B_1^{(\chi, V)} \subseteq B''_1$, $(B', \mathbf{B}_\top) \models \bigcirc_1 \lambda[0]$ and $(B'', \mathbf{B}_\top) \models \bigcirc_1 \lambda[0]$. Then, $(B', \mathbf{B}_\top) \models \text{tr}(\chi)$ iff $(B'', \mathbf{B}_\top) \models \text{tr}(\chi)$.*

Proof. Let $B' = (B'_1, V')$ and $B'' = (B''_1, V'')$. We assume $(V \cap \text{Atm}(\pi)) = (V' \cap \text{Atm}(\pi)) = (V'' \cap \text{Atm}(\pi))$, $B_1^{(\chi, V)} \subseteq B'_1$, $B_1^{(\chi, V)} \subseteq B''_1$, $(B', \mathbf{B}_\top) \models \bigcirc_1 \lambda[0]$ and $(B'', \mathbf{B}_\top) \models \bigcirc_1 \lambda[0]$.

The proof of the lemma is by induction on the length of the formula χ .

Base case Let λ be the empty sequence. Clearly, $B' \models \pi$ iff $B'' \models \pi$, since $(V' \cap \text{Atm}(\pi)) = (V'' \cap \text{Atm}(\pi))$. Therefore, $(B', \mathbf{B}_\top) \models \pi$ iff $(B'', \mathbf{B}_\top) \models \pi$.

Inductive case We prove that the statement is true for sequence λ of length $k + 1$, if we suppose it is true for sequence λ of length k . Suppose λ has length $k + 1$. Therefore, we can assume that $\lambda = Qp\lambda'$. We assume $Q = \forall$. The proof for $Q = \exists$ is analogous.

Let

$$\mathcal{R}_{[\lambda']} = \{(B''', B'''') \in \mathcal{R}_1 : (B'''' , \mathbf{B}_\top) \models \bigcirc_1 \lambda'[0]\}.$$

Clearly, for all $B''' \in \mathbf{B}_\top$, we have

$$(B''', \mathbf{B}_\top) \models [\lambda']\psi \text{ iff } (B'''' , \mathbf{B}_\top) \models \psi \text{ for all } B'''' \in \mathcal{R}_{[\lambda']}(B''').$$

Now, suppose $(B', \mathbf{B}_\top) \models \text{tr}(\forall p\lambda'.\pi)$. The latter is equivalent to $(B', \mathbf{B}_\top) \models [\lambda']\text{tr}(\lambda'.\pi)$. The latter is equivalent to

$$(B''', \mathbf{B}_\top) \models \text{tr}(\lambda'.\pi) \text{ for all } B''' \in \mathcal{R}_{[\lambda']}(B').$$

For all $B''' \in (\mathcal{R}_{[\lambda']}(B') \cup \mathcal{R}_{[\lambda']}(B''))$, we clearly have:

$$(B''', \mathbf{B}_\top) \models \bigcirc_1 \lambda'[0].$$

since $B_1^{(\lambda', V)} \subseteq B'_1$ and $B_1^{(\lambda', V)} \subseteq B''_1$. Moreover, by the initial assumption that $(V \cap \text{Atm}(\pi)) = (V' \cap \text{Atm}(\pi)) = (V'' \cap \text{Atm}(\pi))$, $B_1^{(\lambda', V)} \subseteq B'_1$ and $B_1^{(\lambda', V)} \subseteq B''_1$, for all $B''' \in (\mathcal{R}_{[\lambda']}(B') \cup \mathcal{R}_{[\lambda']}(B''))$ we have:

$$\begin{aligned} B_1^{(\lambda'.\pi, V \cup \{p\})} &\subseteq B'''_1; \\ B_1^{(\lambda'.\pi, V \setminus \{p\})} &\subseteq B'''_1; \\ ((V \cup \{p\}) \cap \text{Atm}(\pi)) &= (V''' \cap \text{Atm}(\pi)) \text{ or} \\ ((V \setminus \{p\}) \cap \text{Atm}(\pi)) &= (V''' \cap \text{Atm}(\pi)). \end{aligned}$$

Furthermore, by the fact that $(B', \mathbf{B}_\top) \models \bigcirc_1 p$ and $(B'', \mathbf{B}_\top) \models \bigcirc_1 p$, we have that:

$$\begin{aligned} \forall B'''' \in \mathcal{R}_{[\lambda']}(B''), \exists B'''' \in \mathcal{R}_{[\lambda']}(B') \\ \text{such that } (V'''' \cap \text{Atm}(\pi)) &= (V''' \cap \text{Atm}(\pi)). \end{aligned}$$

Therefore, by induction hypothesis and the fact that $(B''', \mathbf{B}_\top) \models \text{tr}(\lambda'.\pi)$ for all $B''' \in \mathcal{R}_{[\lambda']}(B')$, we have:

$$(B'''' , \mathbf{B}_\top) \models \text{tr}(\lambda'.\pi) \text{ for all } B'''' \in \mathcal{R}_{[\lambda']}(B'').$$

The latter is equivalent to $(B'', \mathbf{B}_\top) \models [\lambda']\text{tr}(\lambda'.\pi)$ which in turn is equivalent to $(B'', \mathbf{B}_\top) \models \text{tr}(Qp\lambda'.\pi)$. In an analogous way, we can prove that $(B'', \mathbf{B}_\top) \models \text{tr}(Qp\lambda'.\pi)$ implies $(B', \mathbf{B}_\top) \models \text{tr}(Qp\lambda'.\pi)$. \square

We can now go back to the statement of Lemma 1 and prove it. The proof of the lemma is by induction on the length of the formula λ .

Base case Suppose λ is the empty sequence and $V \models \pi$. By construction of $B^{(\lambda, V)}$, the latter is equivalent to $B^{(\lambda, V)} \models \pi$. The latter is equivalent to $(B^{(\lambda, V)}, \mathbf{B}_\top) \models \pi$.

Inductive case We prove that the statement is true for sequence λ of length $k + 1$, if we suppose it is true for sequence λ of length k . Suppose λ has length $k + 1$. Therefore, χ can be written as $Qp\lambda'.\pi$. We assume $Q = \exists$, as the proof for $Q = \forall$ is analogous.

(\Rightarrow) Suppose $V \models \exists p\lambda'.\pi$. The latter is equivalent to saying that $V \cup \{p\} \models \lambda'.\pi$ or $V \setminus \{p\} \models \lambda'.\pi$.

By induction hypothesis, the latter implies that

$$\begin{aligned} (B^{(\lambda'.\pi, V \cup \{p\})}, \mathbf{B}_\top) &\models tr(\lambda'.\pi), \text{ or} \\ (B^{(\lambda'.\pi, V \setminus \{p\})}, \mathbf{B}_\top) &\models tr(\lambda'.\pi). \end{aligned}$$

It is easy to check that

$$\begin{aligned} B^{(\chi, V)} \mathcal{R}_1 B^{(\lambda'.\pi, V \cup \{p\})}, \text{ and} \\ B^{(\chi, V)} \mathcal{R}_1 B^{(\lambda'.\pi, V \setminus \{p\})}. \end{aligned}$$

Moreover,

$$\begin{aligned} (B^{(\lambda'.\pi, V \cup \{p\})}, \mathbf{B}_\top) &\models \bigcirc_1 \lambda'[0], \text{ and} \\ (B^{(\lambda'.\pi, V \setminus \{p\})}, \mathbf{B}_\top) &\models \bigcirc_1 \lambda'[0]. \end{aligned}$$

Therefore,

$$(B^{(\chi, V)}, \mathbf{B}_\top) \models \langle\langle \lambda' \rangle\rangle tr(\lambda'.\pi).$$

The latter is equivalent to $(B^{(\chi, V)}, \mathbf{B}_\top) \models tr(\chi)$.

(\Leftarrow) Suppose $(B^{(\chi, V)}, \mathbf{B}_\top) \models tr(\exists p\lambda'.\pi)$. Hence, $(B^{(\chi, V)}, \mathbf{B}_\top) \models \langle\langle \lambda' \rangle\rangle tr(\lambda'.\pi)$. The latter implies that

$$\exists B' \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)}) \text{ such that } (B', \mathbf{B}_\top) \models tr(\lambda'.\pi)$$

where $\mathcal{R}_{[\lambda']}$ is defined as in the proof of Lemma 2.

For all $B' \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)})$, we clearly have:

$$(B', \mathbf{B}_\top) \models \bigcirc_1 \lambda'[0].$$

Moreover, for all $B' \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)})$, we have:

$$\begin{aligned} B_1^{(\lambda'.\pi, V \cup \{p\})} &\subseteq B'_1; \\ B_1^{(\lambda'.\pi, V \setminus \{p\})} &\subseteq B'_1; \\ ((V \cup \{p\}) \cap \text{Atm}(\pi)) &= (V' \cap \text{Atm}(\pi)) \text{ or} \\ ((V \setminus \{p\}) \cap \text{Atm}(\pi)) &= (V' \cap \text{Atm}(\pi)). \end{aligned}$$

Finally, we have:

$$B^{(\lambda'.\pi, V \cup \{p\})} \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)}) \text{ and } B^{(\lambda'.\pi, V \setminus \{p\})} \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)}).$$

Thus, by Lemma 2 and the fact that there exists $B' \in \mathcal{R}_{[\lambda']} (B^{(\chi, V)})$ such that $(B', \mathbf{B}_\top) \models tr(\lambda'.\pi)$, we have:

$$\begin{aligned} (B^{(\lambda'.\pi, V \cup \{p\})}, \mathbf{B}_\top) &\models tr(\lambda'.\pi), \text{ or} \\ (B^{(\lambda'.\pi, V \setminus \{p\})}, \mathbf{B}_\top) &\models tr(\lambda'.\pi). \end{aligned}$$

By induction hypothesis, the latter implies that $V \cup \{p\} \models \lambda'.\pi$ or $V \setminus \{p\} \models \lambda'.\pi$. The latter is equivalent to $V \models \exists p\lambda'.\pi$. \square