

Virtual Group Knowledge on Topological Evidence Models

MSc Thesis (*Afstudeerscriptie*)

written by

Djanira dos Santos Gomes

(born October 15th, 1999 in Amsterdam, the Netherlands)

under the supervision of **Dr Alexandru Baltag** and **Dr Malvin Gattinger**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense: **Members of the Thesis Committee:**

March 27th, 2025

Dr Maria Aloni (chair)

Dr Alexandru Baltag (supervisor)

Dr Nick Bezhanishvili

Prof Dr Hans van Ditmarsch

Dr Malvin Gattinger (supervisor)

Dr Aybüke Özgün



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

In this thesis, we study notions of group knowledge and group belief in topological-evidential semantics. We define a multi-agent extension of the *topological evidence models* defined in [Özg17]. On these models, we present the evidence-based semantics of *virtual group knowledge*: this extends from individuals to groups the topological semantics of evidence-based belief and fallible knowledge, using the join topology (generated by the union of the individual topologies). As a result, group knowledge is non-monotonic with respect to inclusion. In contrast to distributed knowledge, this notion of virtual group knowledge is more realistic, since it pre-encodes the knowledge that a group can actually achieve after sharing all the individual evidence. We extend the language with a corresponding dynamic operator, which models the intended act of communication. We axiomatise the multi-agent language of evidence, as well as a knowledge-belief fragment of the language that restricts knowledge and belief to individual agents and the full group of agents. We bring these notions into practice by implementing symbolic and explicit approaches to Haskell-based model checking of the language of evidence. We compare the resulting model checkers with respect to performance.

Acknowledgements

I'm extremely grateful to Alexandru and Malvin for guiding me through the process of writing this (very long) thesis. Alexandru, thank you for the countless hours of explaining complicated proof structures to me (at a crazy speed), for the emails about new axioms in the middle of the night, for the IPAs in the pub, for pushing me to go to Bern, and for your endless motivation and enthusiasm. Malvin, thank you for supporting me for over two years –not only as my supervisor on multiple projects, but also as my mentor– and for always being there when I needed to vent, and for advising me on anything related to academia, lean, haskell, overleaf, git, mental health, and (more importantly) cute neighbours and other dilemmas. The two of you were the best combination of supervisors I could have wished for.

Wouter, thank you for your endless support and friendship, from dancing for me in front of the Oerknal counter in 2018(?) to cheering me to the finish line when I was writing this thesis. I hope we go to China together! Paul and Madi, thank you for being there for me in my final year(s) of the MoL, for a severely unhealthy amount of beers, cocktails, shots, and elderflower gin and tonics, but also for so many bouldering sessions together (which compensate the bad habits a little bit). Joris, thank you for going through probably the most similar thesis trajet that's possible at the MoL with me, for spending hours trying to prove the soundness of an axiom that might not have been sound after all (I guess we will never know), for struggling through the past six months together, for secret hitler, and many beers. Ramón and Dominique (Pantxo's Putivuelta), thank you for the early MoL years and for all the dancing we did throughout the years! I'm so glad that we're friends and I hope we will keep finding each other anywhere on earth.

Rosie, Marijn, and Yali, I love you guys and I got so close to each of you during the past year, even after having been friends with you for so long. Thank you for getting me through the last weeks before the deadline and for celebrating all those Friday nights with me afterwards. Oerknal, thank you for being my second home since 2017 and for always welcoming me with love after I left (in particular, thank you Tim, for discussing weather apps and contradicting beliefs with me during work hours!). Thank you Amsterdam, for being my first home always. I'm already counting down the days until I come back.

Dankjewel mamma, for everything.

Contents

1	Introduction	3
2	Background	6
2.1	Single-Agent Topological Models	6
2.1.1	Topological Evidence Models	6
2.1.2	Topological Accounts of Knowledge and Belief	8
2.1.3	Single-Agent Topological Semantics	9
2.1.4	Alexandroff Models: Relational Representation	10
2.2	Multi-Agent Topological Models	12
2.3	Group Knowledge According to Ramírez and Fernández	13
2.3.1	Distributed Knowledge	13
2.3.2	Failure of Group-Monotonicity in the Join Topology	14
2.3.3	Ramírez' Approach: Revising Individual Knowledge	15
2.3.4	Fernández' Approach: Revising Distributed Knowledge	16
2.4	Model Checking	16
2.4.1	Symbolic Model Checking	17
2.4.2	Existing Model Checkers	18
3	Virtual Group Knowledge	20
3.1	Virtual Group Knowledge: Motivation	20
3.2	Logics on Multi-Agent Topo-E-Models	22
3.2.1	Logic of Individual Evidence	23
3.2.2	Logics of Group Evidence	24
3.2.3	Logics of Group Knowledge and Belief	25
3.3	Dynamics: Sharing the Evidence	28
3.3.1	Dynamic Extensions	28
3.3.2	Group Knowledge Pre-encodes Knowledge after Sharing	29
3.3.3	Example: Missing Cake	31
3.4	Examples	33
3.5	Alexandroff Multi-Agent Models: Relational Representation	37
3.5.1	Relational Evidence Models	37
3.5.2	Alexandroff Topo-E-Models as Relational Evidence Models	38
3.5.3	Relational Evidence Models as Alexandroff Topo-E-Models	40

4	Axiomatisation	42
4.1	Validities	42
4.1.1	Single Groups	43
4.1.2	Groups and Subgroups	43
4.2	Proof Systems of the Static and Dynamic Languages	47
4.2.1	Axiomatisations of $\Box[\forall]_I$ and $\Box[\forall]_{iA}$	47
4.2.2	Axiomatisation of $KB_{i,A}$	48
4.2.3	Candidate Axiomatisation of KB_I	50
4.2.4	Axiomatisations of Dynamic Extensions	51
5	Completeness	54
5.1	Completeness of $\Box[\forall]_I$	56
5.1.1	Pseudo-Models for $\Box[\forall]_I$	56
5.1.2	Soundness and Completeness of $\Box[\forall]_I$ w.r.t. Pseudo-Models	58
5.1.3	From Pseudo-Models to Models	61
5.2	Completeness of $KB_{i,A}$	68
5.2.1	Pseudo-Models for $KB_{i,A}$	69
5.2.2	Soundness and Completeness of $KB_{i,A}$ w.r.t. Pseudo-Models	71
5.2.3	Back and Forth between Pseudo-Models for $KB_{i,A}$ and for $\Box[\forall]_I$	71
5.3	Completeness of Dynamic Extensions	80
6	Symbolic Model Checking	83
6.1	Symbolic Topo-Structures	83
6.2	Translations	88
6.3	Missing Cake Revisited	92
6.4	Implementation	92
6.4.1	Boolean Translation	93
6.4.2	Translations between Models	95
6.5	Benchmarks	95
7	Conclusions and Future Work	98
	Bibliography	100
	Appendices	104
A	Proofs	105
A.1	Proofs from Chapter 2	105
A.2	Proofs from Chapter 3	105
A.3	Proofs from Chapter 4	106
A.4	Proofs from Chapter 5	107

Chapter 1

Introduction

In this thesis, we study notions of group knowledge and group belief in topological-evidential semantics. We study and implement the topological semantics of a pragmatic approach to group knowledge, interpreted on a multi-agent extension of the *topological evidence models* defined in [Özg17].

In the digital age, distributed systems are ubiquitous. Whether we consider a network of coordinating computers, or the independent components of a single computer, we are surrounded by systems of independent agents that cooperate to achieve a shared objective. Moreover, this setting is a routine aspect of daily life in social contexts. Hence, the study of these systems is of both practical and philosophical significance.

In the context of distributed systems, it is common to reason about notions of knowledge and belief associated with groups [HM90]. The best-known notions of group knowledge are *distributed knowledge* and *common knowledge*. Distributed knowledge, which is a key concept in this thesis, is inherently linked to communication: it can be described as what a group of agents would come to know after sharing or combining their individual information or knowledge (see e.g. [HM92; Ram15; Fer18]).

A natural framework for reasoning about knowledge is Epistemic Logic: an umbrella term for modal logics that formalise notions of knowledge and belief for rational agents. Traditionally, these logics are interpreted on *Kripke models*, according to Hintikka’s semantics [Hin62]. Recently, topological interpretations of epistemic logics have gained popularity (see e.g. [Özg17; BP11; BRS12; BFS16; SSV17] for various approaches). An advantage of these interpretations over the traditional semantics is the incorporation of a richer notion of knowledge: it is not primitive, but based on *evidence*. Examples of evidence-based models for epistemic logic are *evidence models* [BP11], *justification models* [BFS16], and *topological evidence models* [Özg17].

Topological evidence models, or *topo-e-models*, are the subject of this thesis. Topo-e-models explicitly represent a topology of evidence, generated from basic evidence. The topological semantics of knowledge and belief involves an evaluation of the *interior semantics* of McKinsey and Tarski [MT44] on a restriction of the topology to the *justifications*: pieces of evidence that are consistent with all other evidence. Belief in a proposition amounts to a justification of that proposition, and knowledge is interpreted as *correctly justified belief*.

In the context of Dynamic Epistemic Logic and belief revision, various dynamic extensions to these logics have been proposed. The dynamic modalities include, among others, evidence addition and public announcements [Özg17; BP11; BRS12]. These dynamics highlight the *fallibility* of knowledge as correctly justified belief: an agent may lose her knowledge of a proposition in light of new evidence. This is a fundamental characteristic of knowledge evaluated on evidence-based models.

Given the existing dynamics concerning evidence addition and public announcements, a natural continuation of this research is to extend the framework to the multi-agent case and to incorporate a notion of group knowledge. The straightforward multi-agent extension evaluates knowledge and belief of a group on the *join topology*, which corresponds to pooling together individual evidence. However, under this interpretation, the main characteristic of distributed knowledge (as described in e.g. [HM90; FHV92]) fails: on topo-e-models, group knowledge is not monotonic with respect to group inclusion. In fact, a group may even know less than its members [Ram15; Fer18; Özg17]. In order to resolve this, two modifications of this semantics have been proposed. Both approaches ensure the validity of the *Group-Monotonicity* property, and adapt either the interpretation of individual knowledge [Ram15], or of distributed knowledge [Fer18].

Nevertheless, it can be argued that the models in [Ram15] and [Fer18] do not realistically capture group knowledge interpreted as *epistemic potential* [Özg17]. In the context of distributed systems, we would like to view group knowledge as an indication of what a group of agents can achieve through communication. A fallible notion of knowledge then justifies, and even demands the failure of the property, which strengthens the case for a straightforward interpretation on the join topology. This leads us to the first objective of this thesis:

The main objective of this thesis is to study, justify, and axiomatize a sound and complete logic of knowledge and belief, which interprets group knowledge and group belief on the join topology.

We list the main contributions of this thesis with respect to the first objective.

1. Topological notions *virtual group knowledge* and *virtual group belief*, which interpret knowledge and belief for groups on the join topology.
2. An axiomatisation $KB_{i,A}$ of the logic of virtual group knowledge and virtual group belief, which interprets knowledge and belief for individual agents and for the full group. The axiomatisation is presented in Definition 4.11, and the proof of soundness and completeness is presented in Theorem 4.12.
3. An axiomatisation $\Box[\forall]_I$ of the logic of group evidence, which interprets evidence for all subgroups of a given group of agents, and in which virtual group knowledge and virtual group belief can be expressed as abbreviations. The axiomatisation is presented in Definition 4.8, and the proof of soundness and completeness is presented in Theorem 4.9.
4. A sound candidate axiomatisation KB_I of the logic of virtual group knowledge and virtual group belief, which interprets knowledge and belief for all subgroups of a given group of agents. The axiomatisation is presented in Definition 4.13, and the proof of soundness is presented in Proposition 4.14. Completeness is left as an open question.

Axiomatizing the logic of group knowledge and group belief is challenging: the axiomatisation $KB_{i,A}$ of the logic of virtual group knowledge and virtual group belief required the introduction of unexpected axioms, which replace Group-Monotonicity and partially compensate its loss. As a consequence, the proof of completeness for $KB_{i,A}$ is rather complex and non-trivial, while finding a complete axiomatization of the larger logic KB_I remains an open question.

Motivated by the practical applications of group knowledge in distributed systems, our second objective is to implement a *model checker* for the proposed logic. The method of *model checking* is traditionally used for automatic verification of the correctness of (e.g. software) programs [Cla+18]:

model checking tools evaluate the behaviour of a given program by checking the validity of the intended properties on its structure. By representing the desired property as a logical formula, and the evaluated program as a model, the problem is reduced to logical satisfiability.

The majority of existing model checkers take as input Kripke models, and formulas represented in temporal logics. A bottleneck caused by these models is that the size of the state space is exponential in its vocabulary. In the case of topo-e-models, the addition of evidence sets adds to this problem.

The promising method of *symbolic model checking* has been widely implemented for temporal logics [Cla+18], and, more recently, extended to epistemic logics [Gat18]. A concise representation of the given model avoids explicitly spelling out the state space. Additionally, the evaluated formula is translated to a boolean function and represented as a *Binary Decision Diagram* [Bry86]. These adaptations result in a more efficient algorithm compared to the explicit approach.

Thus far, there exist explicit model checkers for non-epistemic spatial logics (see e.g. [Cia+17; Bez+22; Cia+18; ÁS23]). However, these model checkers are tailored to the Spatial Logic of Closure Spaces and implemented with the particular purpose of analysing images. Symbolic model checkers have not been implemented for any spatial epistemic logic. Nevertheless, there exists a symbolic model checker for epistemic logic on Kripke models [Gat18]. It was implemented in the functional programming language Haskell, on account of the mathematical nature of its syntax. Furthermore, the technique employed by [Gat18] suits the structure of topo-e-models. This leads us to the second objective of this thesis:

The second objective of this thesis is to implement a symbolic model checker for the proposed logic in Haskell.

We list the main contributions of this thesis with respect to the second objective.

1. An approach to symbolic model checking for the topological semantics of $\Box[\forall]_I$, based on symbolic model checking for Dynamic Epistemic Logic [Gat18]. We present and discuss the approach in Chapter 6.
2. A symbolic model checker implemented in Haskell, which follows this approach. We highlight the principal function of the implementation of the symbolic model checker in Section 6.4, and we evaluate its performance in Section 6.5.

This thesis is structured as follows. In Chapter 2, we introduce the relevant notions with respect to model checking, topological evidence models, and group knowledge. Chapter 3 contains the main theoretical contributions of this thesis: we motivate and define *virtual group knowledge* and *virtual group belief* on a multi-agent extension of topo-e-models; we define the language $\Box[\forall]_I$ and its fragments $KB_{i,A}$ and KB_I , for which we also introduce extensions with a dynamic operator for sharing. We argue against Group-Monotonicity through concrete examples. Finally, we present an alternative, relational representation of the topological models. Chapter 4 is dedicated to axiomatisations of the various logics. The completeness proofs for $\Box[\forall]_I$ and $KB_{i,A}$ are presented in Chapter 5. Chapter 6 contains the main applied contribution of this thesis: we extend the symbolic model-checking technique to the logic of $\Box[\forall]_I$ and we compare our implementations of a symbolic and explicit model checker for $\Box[\forall]_I$ with respect to performance. We conclude this thesis by summarising our findings and suggesting possible directions for future work.

Chapter 2

Background

In this chapter we introduce the relevant notions for this thesis. The majority of the background is dedicated to topological models: Section 2.1 introduces the relevant topological notions and models. In particular, we formally introduce the *topological evidence models*, which we will later extend to multiple agents. In Section 2.2 we discuss recent developments in multi-agent extensions of topological evidence models. We also introduce notions of *group knowledge* and discuss the problem associated with defining *distributed group knowledge* on topological evidence models. We conclude this chapter with a brief discussion of the theory and recent developments in model checking of modal logics in Section 2.4.

2.1 Single-Agent Topological Models

This section serves as an introduction to single-agent topological spaces and topological evidence models. The presentation of this section is based on Section 2 in [BB07a] and Chapters 2, 3, and 5 in [Özg17]. We start with the topological preliminaries, before introducing the topo-e-models on which our multi-agent model is based. We then introduce the topological accounts of knowledge and belief, as well as a language of individual evidence, in which these accounts can be expressed as abbreviations. Finally, we introduce an equivalent relational representation of a subclass of topo-e-models, called *Alexandroff* topo-e-models.

2.1.1 Topological Evidence Models

Throughout the thesis, we interpret modal logic on a class of *topological models*. The underlying structure of such a model is a *topological space*: a state space with a collection of sets that is closed under finite intersections and arbitrary unions:

Definition 2.1 (Topological Space). A *topological space* is a pair $\mathcal{X} = (X, \tau)$, where X (the *space*) is a nonempty set and τ (the *topology*) is a collection of subsets of X satisfying the following three conditions:

1. $\emptyset, X \in \tau$;
2. If $U, V \in \tau$, then $U \cap V \in \tau$;
3. If $\{U_i\}_{i \in I} \in \tau$, then $\bigcup_{i \in I} U_i \in \tau$.

Elements of τ are *open sets* (or *opens*) and their complements are *closed*. A set that is both closed and open is *clopen*. An open set containing $x \in X$ is an *open neighbourhood* of x . Given a space X , every $A \subseteq X$ has an *interior* and a *closure*. The corresponding operators satisfy a set of properties called the *Kuratowski axioms* (Proposition 2.3).

Definition 2.2 (Interior and Closure). Given a topological space (X, τ) , the *interior* of a set $A \subseteq X$, denoted by $\text{Int}_\tau(A)$, is the union of all opens in A : $\text{Int}_\tau(A) = \bigcup \{U \in \tau \mid U \subseteq A\}$. The *closure* of A , denoted by $\text{Cl}_\tau(A)$, is the dual of the interior: $\text{Cl}_\tau(A) = \bigcap \{C \in \tau \mid A \subseteq C\}$.

The interior of A is the largest open set contained in A and $\text{Cl}_\tau(A)$ is the least closed set containing A . Given $x \in X$, we have $x \in \text{Cl}_\tau(A)$ if $U \cap A \neq \emptyset$ for all nonempty open neighbourhoods U of x .

Proposition 2.3. *The interior and closure operators are known to satisfy the Kuratowski axioms [Eng89]: for any $A, B \subseteq X$, we have*

- | | |
|---|--|
| (I1) $\text{Int}_\tau(X) = X$ | (C1) $\text{Cl}_\tau(\emptyset) = \emptyset$ |
| (I2) $\text{Int}_\tau(A) \subseteq A$ | (C2) $A \subseteq \text{Cl}_\tau(A)$ |
| (I3) $\text{Int}_\tau(A \cap B) = \text{Int}_\tau(A) \cap \text{Int}_\tau(B)$ | (C3) $\text{Cl}_\tau(A \cup B) = \text{Cl}_\tau(A) \cup \text{Cl}_\tau(B)$ |
| (I4) $\text{Int}_\tau(\text{Int}_\tau(A)) = \text{Int}_\tau(A)$. | (C4) $\text{Cl}_\tau(\text{Cl}_\tau(A)) = \text{Cl}_\tau(A)$. |

In fact, the Kuratowski axioms correspond directly to the axioms of the logical system S4. Section 2.1.4 elaborates on this correspondence between S4 frames and topological spaces.

Definition 2.4 and 2.5 state that a topology can be generated from any collection of sets:

Definition 2.4 (Topological Basis and Subbasis). A family $\mathcal{B} \subseteq \tau$ is a *basis* for a topological space (X, τ) if every non-empty open subset of X can be written as a union of elements of \mathcal{B} . A family \mathcal{C} is a *subbasis* for (X, τ) if the set of finite intersections of members of \mathcal{C} forms a basis for (X, τ) .

\mathcal{B} is the set of *basic opens*. We can evaluate the interior of a set by referring only to a basis \mathcal{B} for a topological space (X, τ) : for any $A \subseteq X$, $x \in \text{Int}_\tau(A)$ if and only if there is an open set $U \in \mathcal{B}$ such that $x \in U \subseteq A$.

Definition 2.5 (Generated Topology). Given any family Σ of subsets of X , there exists a unique, smallest topology $\tau(\Sigma)$ with $\Sigma \subseteq \tau(\Sigma)$. The topology $\tau(\Sigma)$ is said to be *generated* by Σ and consists of \emptyset, X , all finite intersections of members of Σ , and all arbitrary unions of these finite intersections.

A topological space paired with a valuation on the state space is called a *topological model*. Topological models can be interpreted epistemically: states in the state space are possible worlds, open sets represent verifiable or observable properties, and closed sets are falsifiable properties. The verifiable properties are also called *evidence* [Özg17], which the agent uses to form her beliefs. Propositions are represented by subsets of the state space, and can be *supported* by a piece of evidence.

Topological evidence models, an extension of the *evidence models* defined by [BP11], are a variation of topological models designed specifically with this epistemic interpretation in mind. Evidence models describe only the *basic* or *direct* pieces of evidence that the agent possesses; topological evidence models extend these models with the topology generated by the basic evidence. If we consider the basic evidence to be directly observed, then the topology represents the *indirect* evidence that we obtain by reasoning.

Definition 2.6 (Topological Evidence Model). A *topological evidence model* (or *topo-e-model*) is a tuple $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ where

1. X is a nonempty set of possible worlds (or *states*);
2. $\mathcal{E}^0 \subseteq \mathcal{P}(X)$, called *basic evidence* (or *pieces of evidence*), satisfies $X \in \mathcal{E}^0$ and $\emptyset \notin \mathcal{E}^0$;
3. $\pi : X \rightarrow \mathcal{P}(V)$ is a valuation function;
4. and the *evidential topology* $\tau = \tau_{\mathcal{E}^0}$ is the unique topology generated (see Definition 2.5) by the subbasis \mathcal{E}^0 .

Equivalently, τ is generated by the topological basis \mathcal{E} , which is obtained by closing \mathcal{E}^0 under finite intersections. The term *basic evidence* is slightly confusing, as this collection of evidence \mathcal{E} technically constitutes the *subbasis* of the generated topology. On evidence models, the basis of the topology is referred to as *combined evidence*.

Definition 2.6 ensures that tautologies are always observed ($X \in \mathcal{E}^0$) and that contradictions are not ($\emptyset \notin \mathcal{E}^0$). Note that pieces of evidence can be false (evidence does not necessarily include the actual world) or mutually inconsistent.

On a topo-e-model, a proposition P and a piece of evidence e are both represented by a set of states. A piece of evidence e *supports* P if all states of e are states of P .

Definition 2.7 (Evidential Support). Given a topological evidence model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ and a proposition $P \subseteq X$, an element $U \in \tau$ *supports* (or *entails*) P if $U \subseteq P$.

Evidence can be false; it is only *factive* if it contains the actual state.

Definition 2.8 (Factivity). Given a topological evidence model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$, a piece of direct or combined evidence e is *factive* at $x \in X$ if it is true at x , that is, if $x \in e$. A body of evidence F is *factive* at x if each piece of evidence $e \in F$ is true at x , that is, if $x \in \bigcap F$.

2.1.2 Topological Accounts of Knowledge and Belief

Belief. If the evidence supporting a proposition P is convincing enough, the agent comes to *believe* P [Bal+16]. We consider a piece of direct or indirect evidence convincing if it is *undefeated* by all other evidence, that is, the agent possesses no other piece of evidence supporting its negation. The corresponding topological property is called *density*:

Definition 2.9 (Density). Given a topological space (X, τ) , a set $A \subseteq X$ is called *dense* in X if $Cl_\tau(A) = X$. Equivalently, A is dense if $A \cap U \neq \emptyset$ for all $U \in \tau \setminus \emptyset$.

A dense piece of evidence supporting P is called a *justification* for P :

Definition 2.10 (Justification). Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ and a proposition $P \subseteq X$, $U \in \tau$ is a *justification* for P if U is a dense open subset of P : $U \subseteq P$ and $Cl_\tau(U) = X$. Equivalently, $Cl_\tau(Int_\tau(P)) = X$.

With these notions, we can define belief on topo-e-models.

Definition 2.11 (Justified Belief). Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ and a state $x \in X$, the agent *believes* a proposition $P \subseteq X$ at x if she has a justification for P :

$$\exists U \in \tau \text{ s.t. } U \subseteq P \text{ and } Cl_\tau(U) = X.$$

Equivalently, $Cl_\tau(Int_\tau(P)) = X$. Note that the evaluated state x does not influence whether the agent believes P : if she believes P , then the belief holds at every state in the model. It can therefore be incorrect at the evaluated state. This leads us to the topological account of knowledge.

Knowledge. If the justification for P is factive at the evaluated state, the agent has *fallible knowledge*¹ of P . In other words, knowledge is *correctly justified belief*².

Definition 2.12 (Fallible Knowledge). Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ and a state $x \in X$, the agent *knows* a proposition $P \subseteq X$ at x if she has a correct justification for P :

$$\exists U \in \tau \text{ s.t. } x \in U \subseteq P \text{ and } Cl_\tau(U) = X.$$

Equivalently, $x \in Cl_\tau(Int_\tau(P)) = X$. We say that we read knowledge as the *dense interior* on τ . In fact, the restriction of τ to its dense open sets, denoted τ^* , is again a topology, and the interior with respect to τ^* coincides with the locally dense interior with respect to τ [Fer18]. Thus, we can equivalently say that knowledge is read as the interior on τ^* .

2.1.3 Single-Agent Topological Semantics

We define a modal language of individual evidence, in which knowledge and belief as in Definitions 2.11 and 2.12 can be expressed as abbreviations:

Definition 2.13 (Syntax of $\mathcal{L}_{\Box[\forall]}$). Given a countable vocabulary V , the language $\mathcal{L}_{\Box[\forall]}(V)$ of evidence, knowledge, and belief is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid [\forall]\varphi$$

where $p \in V$ is any propositional variable.

We interpret the *evidence modality* \Box on topological models according to the *interior semantics* of McKinsey and Tarski [MT44]. The evidence modality is interpreted as the interior operator: in contrast to the standard epistemic interpretation, \Box does not represent knowledge, but factive evidence. The *global modality* $[\forall]$ represents absolutely certain, *infallible* knowledge.

We define the semantics in *truth set*³ notation:

¹Notions of knowledge that is not absolutely certain, i.e. not *infallible*, are called *fallible* or *defeasible*.

²Note the difference between knowledge as *correctly* justified belief and knowledge as *true* justified belief [Özg17]. Justified true belief only requires for the belief to be true, while the justification might be wrong. This definition of knowledge has been criticized since Gettier's well-known counterexamples were published (see e.g. [Get63; Sta06; BRS12]). In the case of correctly justified belief, however, belief in φ is equated to having a justification for φ . As a result, the belief is correct if and only if the justification for φ is correct.

³Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$, let $\llbracket \varphi \rrbracket^{\mathfrak{M}} = \{x \in X \mid (\mathfrak{M}, x) \models \varphi\}$ denote the *truth set* of φ with respect to τ . When there is no ambiguity, we omit the superscript \mathfrak{M} .

Definition 2.14 (Topological Semantics for $\mathcal{L}_{\Box[V]}$). Given a countable vocabulary V and a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$, we interpret formulas φ in the language $\mathcal{L}_{\Box[V]}$ as follows:

$$\begin{aligned} \llbracket \top \rrbracket &= X \\ \llbracket p \rrbracket &= \{x \in X \mid p \in \pi(x)\} \\ \llbracket \neg \varphi \rrbracket &= X \setminus \llbracket \varphi \rrbracket \\ \llbracket \varphi \wedge \psi \rrbracket &= \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \\ \llbracket \Box \varphi \rrbracket &= \text{Int}_\tau(\llbracket \varphi \rrbracket) \\ \llbracket [\forall] \varphi \rrbracket &= \{x \in X \mid \llbracket \varphi \rrbracket = X\} \end{aligned}$$

where $p \in V$ is any propositional variable.

We define truth and validity of a formula $\varphi \in \mathcal{L}_{\Box}(V)$ in the topological semantics in the same way as in the relational semantics (see e.g. [BRV01] for an introduction into relational semantics for modal logic).

Abbreviating the dual of \Box as $\Diamond \varphi := \neg \Box \neg \varphi$, it can be checked ([Özg17]) that

$$\llbracket \Diamond \varphi \rrbracket = \text{Cl}_\tau(\llbracket \varphi \rrbracket) = \{x \in X \mid \forall U \in \tau_i(x \in U \Rightarrow U \cap \llbracket \varphi \rrbracket \neq \emptyset)\}.$$

Fact 2.15. The topological interpretations for belief (B) and knowledge (K) from Definitions 2.11 and 2.12 can be expressed as the following abbreviations in the language $\mathcal{L}_{\Box[V]}$:

$$\begin{aligned} B\varphi &:= [\forall] \Diamond \Box \varphi \\ K\varphi &:= \Box \varphi \wedge B\varphi. \end{aligned}$$

Given a basis \mathcal{B} of τ , we can equivalently interpret \Box in terms of combined evidence, instead of the interior:

$$\begin{aligned} \llbracket \Box \varphi \rrbracket &= \{x \in X \mid \exists U \in \tau(x \in U \subseteq \llbracket \varphi \rrbracket)\} \\ &= \{x \in X \mid \exists U \in \mathcal{B}(x \in U \subseteq \llbracket \varphi \rrbracket)\}. \end{aligned}$$

2.1.4 Alexandroff Models: Relational Representation

The properties of the Kuratowski axioms (Proposition 2.3) suggest a correspondence between topological structures and S4 relational frames (i.e. frames that are reflexive and transitive). In this subsection, we make the correspondence explicit. The content of this subsection follows Chapter 3.1.2 of [Özg17] and Section 2.4.1 of [BB07a]. We first introduce the notion of an Alexandroff space.

Definition 2.16 (Alexandroff Space). A topological space (X, τ) is an *Alexandroff space* if τ is closed under arbitrary intersections, i.e., $\bigcap \mathcal{A} \in \tau$ for any $\mathcal{A} \subseteq \tau$. A topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$ is Alexandroff if the underlying structure (X, τ) is an Alexandroff space.

Equivalently, a space (X, τ) is Alexandroff if every $x \in X$ has a least open neighbourhood. In particular, finite spaces are Alexandroff: they are automatically closed under arbitrary intersections.

The maps defined in Definition 2.18 and Definition 2.21 give us a one-to-one correspondence between S4 frames and Alexandroff spaces. In relational models, opens are represented by *up-sets*:

Definition 2.17 (Up-Set). Given an S4 relational frame (X, R) and some $x \in X$, the *up-set* of x is defined as $\uparrow x := R(x) = \{y \in X \mid xRy\}$. This is the smallest set containing all points that can be reached from x via R . The up-set of a set $S \subseteq X$ is given by $\uparrow S := \{y \in X \mid \exists x \in S(xRy)\}$.

Definition 2.18 (Up-set Topology). Given a finite S4 relational frame (X, R) , we construct an Alexandroff space (X, τ_R) by defining τ_R as the set of all up-sets of X , i.e. by defining $\tau_R := \{\uparrow S \mid S \subseteq X\}$. We call τ_R the *up-set topology*.

It can be checked that τ_R is a topology (see [Özg17; BB07a]). On the topological space (X, τ_R) , the least open neighbourhood of each $x \in X$ is given by the up-set of $\{x\}$. The closed sets on (X, τ_R) are given by the down-sets⁴ of (X, R) . The closure of a set $A \subseteq X$ is given by $Cl_{\tau_R}(A) = \downarrow A$.

To obtain a topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$, we can set $\mathcal{E}_i^0 := \{\uparrow x \mid x \in X\}$.

Proposition 2.19. *Let (X, R) be a finite S4 relational frame and construct an Alexandroff space (X, τ_R) following Definition 2.18. The topology τ on X generated by the subbasis $\{\uparrow x \mid x \in X\}$ equals the up-set topology $\tau_R = \{\uparrow S \mid S \subseteq X\}$.*

Proof. The proof is straightforward. For the left-to-right inclusion, it suffices to show that the intersection of any collection of up-sets is an up-set, and that the union of any collection of up-sets is an up-set; for the converse direction, we need to show that every up-set can be considered as a union of intersections of subbasic up-sets. In fact, a stronger claim holds, which states that any up-set of X can be considered as a union of subbasic up-sets. Details can be found in Appendix A.1. \square

For the converse direction, we define the following relation on the states of a topological space:

Definition 2.20 (Specialisation Pre-Order). The *specialisation pre-order* \sqsubseteq_τ on a topological space (X, τ) is defined as follows:

$$x \sqsubseteq_\tau y \text{ iff } x \in Cl_\tau(\{y\}) \text{ iff } (\forall U \in \tau)(x \in U \text{ implies } y \in U).$$

Definition 2.21. Given an Alexandroff space (X, τ) , we construct an S4 relational frame (X, \sqsubseteq_τ) by defining the relation on X to be the *specialisation pre-order* \sqsubseteq_τ .

It can be checked that the corresponding S4 relational frame (X, \sqsubseteq_τ) is reflexive and transitive.

Since $R = \sqsubseteq_{\tau_R}$ and $\tau = \tau_{\sqsubseteq_\tau}$ if and only if (X, τ) is Alexandroff, the correspondence between Alexandroff spaces and S4 frames is one-to-one. Because every finite topological space is Alexandroff, this also implies a one-to-one correspondence between finite topological spaces and finite S4 frames.

Let V denote the vocabulary. Setting $B(\mathcal{M}) = (X, \tau_R, \pi)$ for any reflexive and transitive relational model $\mathcal{M} = (X, R, \pi)$, and setting $A(\mathfrak{M}) = (X, \mathcal{E}^0, \tau, \pi)$ for any Alexandroff topo-e-model $\mathfrak{M} = (X, \mathcal{E}^0, \tau, \pi)$, we obtain the following proposition:

Proposition 2.22 ([Özg17, page 21]). *For all $\varphi \in \mathcal{L}_\square(V)$, for any reflexive and transitive Kripke model $\mathcal{M} = (X, R, \pi)$ and $x \in X$, and for any Alexandroff model $\mathcal{X} = (X', \tau, \pi')$ and $x' \in X'$,*

$$\begin{aligned} (\mathcal{M}, x) \models \varphi & \text{ iff } (B(\mathcal{M}), x) \models \varphi \\ (\mathcal{X}, x') \models \varphi & \text{ iff } (A(\mathcal{X}), x') \models \varphi. \end{aligned}$$

⁴The down-set of x is given by $\downarrow x = R^{-1}(x) = \{y \in X \mid yRx\}$.

2.2 Multi-Agent Topological Models

As the application of topological semantics to epistemic logic continues to gain recognition, interest in multi-agent extensions has followed suit. In this section we present recent developments. The first two approaches are outside the scope of this thesis; we mention them to be thorough, but we refer to the respective papers for details.

The first approach researches multi-agent epistemic logic and group knowledge through *product models* [Ben+06; BS04] (see also [BB07a] for a discussion of product models). These are products of topological spaces, representing individual topologies as dimensions on a plane. The approach specifically models the combination of topologies that are not defined on the same state space.

Multi-agent extensions have been researched extensively on a class of topological models called *subset spaces*. On the multi-agent subset models in [Özg17] and [DKÖ19] evidence is described by *neighbourhood functions*, known from neighbourhood semantics. On these models, evidence is not only relative to the agent, but also to the state. See also [WÅ13], who proposed a similar logic, or [Hei10].

As for topological evidence models, only two extensions have been proposed⁵ [Ram15; Fer18]⁶. In this section, we discuss the approach in [Fer18]; we discuss the approach in [Ram15] in Section 2.3, as it was designed to model a particular interpretation of distributed knowledge.

The main challenge associated with extending topo-e-models to multiple agents is how to deal with shifting perspectives when evaluating nested modalities. For example, on a *bi-topological space*⁷ (X, τ_1, τ_2, π) , every belief held by any agent is common knowledge⁸ [Ram15]. The solution, first defined by Ramírez [Ram15], is to *partition* the state space into equivalence classes for each agent. For two agents, [Fer18] defines the model as follows:

Definition 2.23 (Two-Agent Topological Models (Fernández)). Given a countable vocabulary V , a two-agent topological model is a tuple $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \pi)$ satisfying the following conditions:

1. X is a set of possible worlds;
2. τ_1, τ_2 are topologies on X ;
3. Π_1, Π_2 are partitions on X ;
4. $\Pi_1 \subseteq \tau_1, \Pi_2 \subseteq \tau_2$;
5. $\pi : V \rightarrow \mathcal{P}(X)$ is a valuation function.

Given a state $x \in X$ and an agent $i \in A$, the *information cell* $\Pi_i(x)$ is the equivalence class that the agent considers epistemically possible at state x . The equivalence class is *infallibly* known by i .

Definition 2.23 requires information cells to be open, that is, they are considered evidence. Multi-agent topological models make a distinction between *soft* and *hard* evidence: hard evidence is defined by the partition and characterises infallible knowledge; the soft evidence consists of the remaining

⁵Another approach outside the scope of our thesis, but related to evidence models, can be found in [Zot24]. It combines topo-e-models with threshold models for diffusion to define multi-agent models, referred to as *evidence diffusion models*. On these models, a threshold-based notion of group belief is defined.

⁶See also [BBF22] for the related paper that summarizes the main results from [Fer18].

⁷A *bi-topological space* is a model of the form (X, τ_1, τ_2, π) consisting of a shared state space and individual topologies.

⁸To see this, recall that if agent 1 believes some formula φ at *any* state in the model, then she believes φ at *every* state in the model. As a result, agent 2 has infallible knowledge that 1 believes φ , since both agents live in the same state space. It is, in fact, common knowledge that 1 believes φ .

open sets and characterises fallible knowledge. For each agent, knowledge is evaluated relative to their partitions: given a state x , we evaluate the modalities for agent i with respect to the *subspace topology* $\tau_i|_{\Pi_i(x)}$ that is induced by their information cell at x :

Definition 2.24 (Subspace Topology [Fer18]). Given a topological space (X, τ) and a nonempty subset $P \subseteq X$, the *subspace topology* $(P, \tau|_P)$ of (X, τ) is the topology induced by P , where $\tau|_P = \{U \cap P \mid U \in \tau\}$.

2.3 Group Knowledge According to Ramírez and Fernández

The term *Group knowledge* encompasses all notions of epistemic knowledge that are defined over groups of agents. The best-known kinds of group knowledge are *distributed knowledge*⁹ and *common knowledge*. These concepts admit multiple interpretations, both philosophically and technically, and are therefore a subject of debate in Epistemic Logic.

The most natural logics incorporating group knowledge are those that interpret the notion only on the full group of agents. Group knowledge was, in fact, first defined only for the full group [HM90]: the pioneers of group knowledge considered this the most relevant group on which to examine notions of group knowledge. This perspective is carried forward by Ramírez [Ram15] and Fernández [Fer18], whose approaches we examine in detail in this section, and, for example, [FHV92]. Group knowledge with respect to all subgroups, which we also consider in this thesis, is studied in [Fag+04; ÅW17; BS20], among other papers.

Roughly, we define common knowledge as what every member of the group knows (and every member knows that every member knows, etc.) *without* communicating. Distributed knowledge, on the other hand, is often interpreted as what the individual members of the group would know *after* communicating: by sharing, or combining, their individual information. As we do not address the problem of defining common knowledge on multi-agent topo-e-models in this thesis, we restrict our attention to distributed knowledge.

2.3.1 Distributed Knowledge

Distributed knowledge was first defined by Halpern and Moses [HM90], who described it as the "weakest state of knowledge": it can be attained for a group without any of its members possessing it. This is illustrated by the following, well-known example: if Alice only knows p , and Bob only knows $p \rightarrow q$, then Alice and Bob are said to have distributed knowledge of q .

Halpern and Moses also introduced the most common philosophical interpretation of the concept (see e.g. [Fer18; BS04; HM92; BH14; ÅW17; Fag+04]), as anything that is known by a fictitious third agent (sometimes called the *wise man*), whom every agent has told whatever they know, or who knows exactly what the agents in the model know. A more pragmatic interpretation describes it as what the group could come to know through communication (see e.g. [Özg17; Fer18; HM92]).

These two interpretations do not coincide on relational models [ÅW17], which we discuss in more detail in Section 3.1. Similarly, on topological models, the preferred definition of distributed knowledge will depend on the preferred interpretation of the concept.

⁹Sometimes also called *implicit knowledge*.

As for formal definitions, a widely accepted characteristic of distributed knowledge is that a group should not know less than its members. This condition corresponds to the following axiom (see e.g. [HM92; FHV92; Fer18; Ram15]), where D_I denotes distributed knowledge over the group I :

$$(KD) \quad K_i \varphi \rightarrow D_I \varphi.$$

In [HM92], this axiom is named (A9). We refer to it by (KD), following [FHV92]. When we incorporate subgroups of a full group A , we obtain the variant that describes monotonicity of knowledge with respect to groups, or *Group-Monotonicity*: $K_J \rightarrow K_I$ for all $I, J \subseteq A$ such that $J \subseteq I$ [Fag+04].

In the context of infallible knowledge (decided by an S5 relation) on relational models, the interpretation of distributed knowledge is straightforward: the indistinguishability relation for the group equals the intersection of the individual relations¹⁰ [DHK07]. This semantics satisfies (KD) and generally¹¹ coincides with the practical, communication-based interpretation of group knowledge.

On relational models, the act of communicating is unambiguously reduced to sharing knowledge. As for topo-e-models, the question now arises: should we evaluate the knowledge of the group of agents by aggregating all individual evidence, or by aggregating what the agents actually *know*? Both, according to Ramírez. Knowledge only, according to Fernández.

2.3.2 Failure of Group-Monotonicity in the Join Topology

Since the agents in topo-e-models cannot distinguish knowledge from belief, a natural approach to group knowledge is to aggregate the individual evidence of the respective agents. The resulting topology is called the *join topology*¹² [Ram15]:

Definition 2.25 (Two-Agent Join Topology $\tau_1 \vee \tau_2$). Given a group of agents $\{1, 2\}$ and a two-agent topological model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \pi)$, the *join topology* for the group $\{1, 2\}$ is the least topology generated by $\tau_1 \cup \tau_2$. Its basis is $\tau_1 \vee \tau_2 := \{U_1 \cap U_2 \mid U_1 \in \tau_1, U_2 \in \tau_2\}$.

This is the unique topology generated by the union of individual subbases. It contains all basic and combined evidence from the individual topologies. Analogous to reading individual knowledge as the locally dense interior on the individual topology, we read group knowledge of a group I as the locally dense interior on the join topology τ_I for the group¹³. Equivalently, denoting by $(\tau_1 \vee \tau_2)^*$ the restriction of $\tau_1 \vee \tau_2$ to the dense open sets, we can interpret knowledge as the interior on $(\tau_1 \vee \tau_2)^*$.

Under this interpretation, the evidence of groups is *distributed*: all soft and hard evidence possessed by any individual i , is also possessed by all groups containing i ¹⁴. Group knowledge, however, is not: the (KD) axiom fails [Ram15; Fer18; Özg17]. This is a consequence of *density* (see Definition 2.9): knowledge is dependent on belief, and belief requires a justification consistent with all available evidence. Clearly, justifications are not stable under evidence addition. As a consequence, justifications for individual agents are not generally justifications for the group.

¹⁰Note: since agents are represented by relations, distributed knowledge quite literally corresponds to a fictitious agent.

¹¹With the exception pointed out by [ÅW17], see Section 3.1 for a discussion.

¹²See also [LA75] for a discussion of the lattice of topologies over a space. This paper refers to the join topology by the *least upper bound*.

¹³Interestingly, [BS04] take this approach to defining distributed knowledge on product models, which we briefly mentioned in Section 2.2.

¹⁴In Section 4.1, we prove the generalisation of the corresponding validities to groups of agents: $\Box_i \varphi \rightarrow \Box_I \varphi$ and $\forall i \varphi \rightarrow \forall_I \varphi$ for $i \in I \subseteq A$.

In order to salvage the axiom, Ramírez [Ram15] and Fernández [Fer18] offer adaptations to the natural approach¹⁵ which, intuitively, make fallible knowledge slightly less fallible. In the remainder of this section, we discuss the respective approaches.

2.3.3 Ramírez' Approach: Revising Individual Knowledge

Ramírez preserves the original nature of distributed knowledge in terms of the join topology. His adaptation involves an additional constraint on individual knowledge.

Definition 2.26 (Two-Agent Topological Models (Ramírez)). Given a countable vocabulary V , a two-agent topological model is a tuple $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \tau, \pi)$ such that

1. $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \pi)$ satisfies the conditions from Definition 2.23;
2. τ is the *topology of learnable evidence* on X ;
3. $\tau_1, \tau_2 \subseteq \tau$.

The contents of the individual topologies τ_1 and τ_2 are now referred to as *current* evidence, whereas the topology τ consists of evidence that could potentially be learned in the future. In Definition 2.27, $\tau|_C$ denotes the subspace topology obtained by restricting τ to a nonempty subspace $C \subseteq X$ (Definition 2.24).

Definition 2.27 (Distributed Knowledge (Ramírez)). Given a countable vocabulary V , a two-agent topological model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \tau, \pi)$ according to Definition 2.26, an agent $i \in \{1, 2\}$, a state $x \in X$, and a formula φ in the language

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_1\varphi \mid K_2\varphi \mid D\varphi \mid C\varphi$$

we interpret the individual knowledge operator K_i and the distributed knowledge operator D as follows:

$$\begin{aligned} \llbracket K_i\varphi \rrbracket &:= \{x \in X \mid \exists U \in \tau_i|_{\Pi_i(x)} \text{ s.t. } x \in U \subseteq \llbracket \varphi \rrbracket \wedge Cl_{\tau|_{\Pi_i(x)}}(U) = \Pi_i(x)\} \\ \llbracket D\varphi \rrbracket &:= \{x \in X \mid \exists U \in \tau_1|_{\Pi_1(x)}, \exists V \in \tau_2|_{\Pi_2(x)} \\ &\quad \text{s.t. } x \in (U \cap V) \subseteq \llbracket \varphi \rrbracket \\ &\quad \wedge Cl_{\tau|_{\Pi_1(x) \cap \Pi_2(x)}}(U \cap V) = \Pi_1(x) \cap \Pi_2(x)\}. \end{aligned}$$

Thus, we read individual knowledge of φ as the (locally) dense interior in the topology of learnable evidence, instead of the individual topology. Conceptually, an agent only *really* knows φ if there does not exist any potentially learnable evidence – in particular from the topology of the other agent – that would defeat their knowledge. We evaluate distributed knowledge with respect to the same topology τ , such that justifications for individual agents are guaranteed to be justifications for the group.

Fernández mentions several issues with Ramírez' approach. First of all, the multi-agent semantics is not a straightforward generalisation of the single-agent case: it is unclear how to define the topology τ on a single-agent model. Furthermore, the dependency of the individual knowledge of one agent on the evidence possessed by a different agent poses two problems: first, it is not realistic. Second, it

¹⁵Although Fernández acknowledged the option to accept the failure of the axiom as a consequence of the defeasibility of knowledge, he did not find any relevant remaining validities and concluded that the natural approach did not warrant further investigation.

is unclear how to interpret private knowledge in the absence of a notion of distributed knowledge. Finally, Fernández argues that this approach is closer to *implicit evidence*¹⁶ than to *implicit knowledge*.

2.3.4 Fernández' Approach: Revising Distributed Knowledge

In an effort to solve the aforementioned issues associated with Ramírez' approach, Fernández retains the original interpretation of individual knowledge and instead redefines the semantics of distributed knowledge. Knowledge is interpreted as follows:

Definition 2.28 (Distributed Knowledge (Fernández)). Given a countable vocabulary V , a two-agent topological model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, \pi)$ according to Definition 2.23, an agent $i \in \{1, 2\}$, a state $x \in X$, and a formula φ in the language

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_1\varphi \mid K_2\varphi \mid D\varphi$$

we interpret the individual knowledge operator K_i and the distributed knowledge operator D as follows:

$$\begin{aligned} \llbracket K_i\varphi \rrbracket &:= \{x \in X \mid \exists U \in \tau_i|_{\Pi_i(x)} \text{ s.t. } x \in U \subseteq \llbracket \varphi \rrbracket \wedge Cl_{\tau_i|_{\Pi_i(x)}}(U) = \Pi_i(x)\} \\ \llbracket D\varphi \rrbracket &:= \{x \in X \mid \exists U_1 \in \tau_1|_{\Pi_1(x)}, \exists U_2 \in \tau_2|_{\Pi_2(x)} \\ &\quad \text{s.t. } x \in (U_1 \cap U_2) \subseteq \llbracket \varphi \rrbracket \\ &\quad \wedge Cl_{\tau_i|_{\Pi_i(x)}}(U_i) = \Pi_i(x)\}. \end{aligned}$$

We read individual knowledge for agent i as the (locally) dense interior in the topology τ_i or, equivalently, as the interior in the topology $(\tau_i)^*$. Distributed knowledge, on the other hand, is read as the interior in the topology $(\tau_1^*) \vee (\tau_2^*)$: the join of the individual topologies of dense open sets.

In other words, the agents only share what they know. The reasoning behind this approach is that, in contrast to a full evidence merge, it agrees with the *wise man* interpretation of distributed knowledge: the wise man knows exactly what the individual members of the group *know*.

To see that this, indeed, saves the (KD) axiom, suppose agent $i \in \{1, 2\}$ knows φ at x . Then for some open U_i which is locally dense in τ_i , we have $x \in U_i \subseteq \llbracket \varphi \rrbracket$. For the other agent j , take $U_j := \Pi_j(x)$, which is a locally dense open for j at x . It follows that $(U_1 \cap U_2) \subseteq \llbracket \varphi \rrbracket$, as required.

2.4 Model Checking

This section provides a short introduction into model checking. After a general introduction into the technique, we discuss the more efficient approach of *symbolic* model checking. We conclude the section with an overview of existing model checkers (Section 2.4.2), emphasising the absence of a symbolic model checker for topological evidence models.

The method of *model checking* is traditionally used for automatic verification of the correctness of (e.g. software) programs [Cla+18]. In practice, software programs are often evaluated through trial and error. The need for a method to globally check correctness of a program resulted in the development of model-checking tools. These tools evaluate the behaviour of a given program by checking the validity of the intended properties on its structure.

¹⁶And indeed, we mentioned that evidence on the join topology is distributed. Since the topology of learnable evidence is a superset of the join topology, this applies to the topology of learnable evidence as well.

Viewed as a transition structure, a program can be represented by a mathematical model \mathcal{M} (the *system description*). The intended behaviour of the model is expressed by a logical formula φ (the *system specification*) whose validity (or truth) we check on the model. If $\mathcal{M} \models \varphi$, (or $(\mathcal{M}, x) \models \varphi$) then the model behaves as desired. Otherwise, a concrete counterexample is returned: this is a state of the program that does not satisfy φ . Since the majority of existing model-checking tools are designed to evaluate transition structures, the evaluated properties often have a temporal aspect. Therefore, generally, the system description is a Kripke model, and the system specification is represented in temporal logic [Cla+18].

2.4.1 Symbolic Model Checking

A strength of model-checking tools is falsifiability: one non-satisfying state suffices to conclude that the model does not exhibit the intended behaviour. The verification of correctness, on the other hand, requires for the property to be checked on all states of a given model. On explicitly represented models, this causes a bottleneck: the state space of a model is exponential in the size of its vocabulary. For example, a program that describes all possible configurations of n propositional variables will have 2^n states. The method of *symbolic model checking* was developed in response to this problem. Using a concise, abstract representation of both the model and the property, we avoid specifying an exponential number of states (see e.g. [Gat18; Bry86]). The idea is to reduce the problem to boolean satisfiability. Given a model \mathcal{M} and a formula φ , we represent \mathcal{M} compactly and translate φ to a boolean function, that is, a formula in the boolean language.

Definition 2.29 (Syntax of $\mathcal{L}_{\text{Bool}}$). Given a countable vocabulary V , the boolean language $\mathcal{L}_{\text{Bool}}(V)$ is defined recursively as

$$\varphi ::= p \mid \neg \varphi \mid \varphi \wedge \varphi$$

where $p \in V$.

This allows a representation of φ as a *Binary Decision Diagram* (BDD). BDDs are a characteristic of symbolic model checking.

Definition 2.30 (Binary Decision Diagram [Bry86]). A *binary decision diagram* is a rooted, directed acyclic graph that represents a boolean function. It has two *terminal* nodes, which are labeled \top and \perp , respectively. Each *nonterminal* node is indexed and has two children. The indices of these nodes represent propositional variables and each assignment over these variables defines a unique path through the graph, such that the truth value of the variable decides to which child the path continues. The terminal node at the end of the path decides the truth value of the function on the given assignment.

A binary decision diagram is *ordered* if for each nonterminal node v and each child c of v we have that $\text{index}(v) < \text{index}(c)$. A binary decision diagram is *reduced* if it contains no nonterminal nodes v, v' such that the subgraphs rooted at v and v' are isomorphic.

With the abbreviation *BDD* we mean an ordered and reduced binary decision diagram. For any formula $\varphi \in \mathcal{L}_B$, the BDD of φ (denoted by $\text{Bdd}(\varphi)$) is the BDD of the boolean function given by φ .

Theorem 2.31 states that, given a boolean function φ , $\text{Bdd}(\varphi)$ is unique.

Theorem 2.31 ([Bry86]). *Given a total order on the propositional variables there is exactly one reduced and ordered binary decision diagram for each boolean function.*

As a consequence, two formulas are equivalent if and only if their BDDs are identical. We can combine existing BDDs for φ, ψ to compute BDDs for $\varphi \wedge \psi$, $\varphi \rightarrow \psi$, and all other boolean connectives. Once we have the BDD of a formula, we can easily check whether it is a tautology or contradiction: the terminal node of the BDD of the tautology equals \top , while the terminal node of the BDD of the contradiction equals \perp .

Thus, even though it can be computationally expensive to construct a BDD, the processes of using and combining existing BDDs are efficient ([Bry86]). This allows us to check, efficiently, on a given assignment to a set of propositional variables, whether it satisfies φ .

BDDs make use of boolean quantification, which is defined as follows:

Definition 2.32 (Boolean Quantification). For any boolean formula φ , the *boolean quantifier* $\forall p\varphi$ is an abbreviation for the formula $[p \mapsto \top]\varphi \wedge [p \mapsto \perp]\varphi$, where $[p \mapsto \psi]\varphi$ denotes substitution of ψ for all occurrences of p in φ . For sets of propositions S , the formula $\forall S\varphi$ is defined as $\forall p_1 \dots \forall p_n \varphi$ for $S = \{p_1, \dots, p_n\}$.

Under the modal interpretation, the global modality is computationally expensive: in the context of explicit model checking, the verification of the statement $(\mathcal{M}, x) \models \forall \varphi$ on some pointed relational model (\mathfrak{M}, x) requires the evaluation of φ at every state of the model. Compare this to the verification of $(\mathcal{M}, x) \models \varphi$, which requires only the evaluation at x . Similarly, the formula abbreviated by $\forall p\varphi$ in Definition 2.32 appears to be duplicated compared to φ . This is not the case: in fact, on BDDs, the addition of a boolean quantifier reduces computational cost: given a vocabulary V and a formula φ , the BDD is a decision diagram that defines which combinations of assignments to V satisfy φ . For any propositional variable $p \in V$, the formula $\forall p\varphi$ signifies that p need not be considered in the evaluation of φ , since φ is true for all assignments to p . As a result, in contrast to the BDD for φ , the BDD for $\forall p\varphi$ does not mention p . Therefore, the BDD for $\forall p\varphi$ is less complicated.

2.4.2 Existing Model Checkers

Two challenges are identified by [Cla+18]: solving the problem of exponential blowup, and developing model-checking methods beyond Kripke models and temporal logic. This thesis aims to contribute to both challenges, by proposing a symbolic approach to model checking of epistemic logic on topo-e-models. We present a brief overview of existing model checkers for epistemic and *spatial logics*¹⁷. In particular, we note the absence of a symbolic approach to model checking of spatial logics.

The most relevant symbolic model checker for this thesis was developed by [Gat18]: this is the first general approach to symbolic model checking of Dynamic Epistemic Logics (DEL) on Kripke models, as opposed to previous implementations, which each focused on a specific logic puzzle.

Multiple model checkers have been developed for spatial logics, none of which concern epistemic logic and all of which are explicit. For basic modal logic, there exists a model checking library for topological spaces, written in the Haskell programming language [ÁS23]. Aside from this library, the majority model checkers of spatial logics is intended for the Spatial Logic of Closure Spaces: an extension of modal logic interpreted on a generalisation of topological spaces (*closure spaces*). This

¹⁷Spatial logics are spatial interpretations of modal logics [APB07]; these include modal logics interpreted on topological spaces.

logic is concerned with describing paths between points, and contains operators such as the *until* operator¹⁸. For examples of model checkers on closure spaces, see e.g. [Cia+17; Bez+22; Cia+18] for several variants, developed with the purpose of analysing images¹⁹.

¹⁸See [BB07b] for a topological interpretation of this operator.

¹⁹[Bez+22] implemented a model checker specifically for the purpose of medical imaging, and [Cia+18] implemented a model checker to analyse vehicle locations on a street map.

Chapter 3

Virtual Group Knowledge

In this section, we present our main contributions in this thesis: a multi-agent generalization of topo-e-models and an accompanying language that incorporates notions similar to, but distinct from, distributed knowledge and distributed belief: *virtual group knowledge* and *virtual group belief*. We start by presenting our motivations for these definitions and for explicitly avoiding the terms *distributed knowledge* and *distributed belief*, in Section 3.1. In Section 3.2, we introduce multi-agent models, along with corresponding logical languages of evidence, and of knowledge and belief. In particular, the logics KB_I and $KB_{i,A}$ of group knowledge and belief are introduced in Section 3.2.3. Next, we introduce a dynamic modality $[share_I]$, which implements the act of sharing evidence within a group (Section 3.3). Examples that advocate against the validity of the (KD) axiom are presented in Section 3.4. We conclude the chapter by providing an alternative, relational representation of Alexandroff multi-agent topo-e-models. This representation is heavily referred to in Chapters 5 and 6.

3.1 Virtual Group Knowledge: Motivation

The main objection against evaluating group knowledge directly on the join topology, shared by Ramírez [Ram15] and Fernández [Fer18], and acknowledged by Özgün [Özg17], is the failure of the (KD) axiom. However, we argue that this non-monotonicity with respect to groups is rather a realistic feature of fallible knowledge than an issue.

In this section, we articulate the motivation for virtual group knowledge on topo-e-models. First, we argue for the relevance of a communication-based interpretation of group knowledge from both a practical and a theoretical viewpoint. Next, we argue from this perspective on group knowledge, that a fallible notion of knowledge justifies a non-monotonic notion of group knowledge.

Group Knowledge from a Communication-based Perspective. Our justification to study a communication-based interpretation of group knowledge stems mainly from practical applications: in the context of distributed systems (see e.g. [HM90]), the concepts of knowledge and communication are undeniably intertwined. We reason about individual knowledge and group knowledge on such systems to evaluate the effectivity of a given distributed protocol. These protocols control coordinated actions by independent agents with a shared objective and therefore, communication is essential. This underlines the importance of a communication-based notion of group knowledge, especially if the agents cannot distinguish knowledge from belief.

However, there is also a more technical consideration: even on relational models of S5 knowledge - and this type of knowledge is, in fact, distinguishable from belief - distributed knowledge does not

always encode what agents would come to know through communication.

The two concepts diverge with respect to formulas involving knowledge modalities, which is illustrated by the following example [ÅW17]: consider the sentence $K_{\{1,2\}}(p \wedge \neg K_1 p)$, which would be true in a situation in which agent 2 knows p , but agent 1 does not. Even though $p \wedge \neg K_1 p$ is distributed knowledge, it cannot remain true after an information exchange between the two agents: agent 1 would then know p . In fact, there does not seem to be any clear formal correspondence between group announcements¹ and distributed knowledge. [ÅAG22] shows that the extension of the *group announcement logic* (GAL) with distributed knowledge does not result in any relevant interaction axioms between the two modalities: the combination of the independent axioms for epistemic logic with distributed knowledge and GAL already results in a complete system.

Because distributed knowledge fails to model the effect of group announcements, [ÅW17] proposed the *resolution* operator as an alternative to distributed knowledge on relational models: for any group I of agents, $R_I \varphi$ holds for all formulas φ that are true after all members of the group I have exchanged their individual information. A similar operator that models the sharing of knowledge among groups is presented in [BS20].

It is clear that virtual group knowledge is more closely related to the concept of resolution than to distributed knowledge in the traditional sense: it pre-encodes individual knowledge after the event of evidence sharing. We further strengthen this reasoning in Section 3.3, when we define a dynamic operator, inspired by the resolution operator, that models this event.

Fallible Knowledge in a Group Setting. From the communication-based perspective on group knowledge, it is questionable to impose Group-Monotonicity on a fallible notion of knowledge. The crux of the problem is a discrepancy between our formal notions of knowledge and the intuitions we have about the concept: we want knowledge to be, in some sense, infallible, because a characteristic of our intuitive interpretation of *knowing* a fact is to be certain of it². The monotonicity implied by the (KD) axiom is inherently linked to this infallible notion of knowledge: if one member of the group is certain of a fact, then it is reasonable to assume that the group is implicitly certain of it. Unfortunately, for many non-trivial facts, absolute certainty is not realistic. So should we admit defeat and conclude that knowledge is unattainable³?

Fallible knowledge saves the day. On single-agent topo-e-models, fallible knowledge allows an agent to know a proposition, even if further evidence might defeat that knowledge again. Naturally, knowledge can be defeated by false evidence. On the other hand, if knowledge is defeated by a piece of factive evidence, we call that evidence *misleading* [Özg17].

This is considered a feature, not a bug, for the single-agent case [Özg17; Fer18]. Özgün [Özg17] argues that there have been intuitive examples of which one would like to say that an agent knows a certain proposition, even though she might lose her knowledge after being misled by another true fact. If we interpret group knowledge as what a group of agents can realistically come to know through communication, then it is only reasonable to extend this feature to multiple agents: if Alice’s knowledge can be defeated by new evidence observed by Alice herself, then it should also be defeasible by evidence received from her friend Bob⁴. On that account, we should accept losing the (KD) axiom

¹Group announcements are simultaneous public announcements by a group of agents [ÅAG22].

²For example, consider the classic S5 notion of knowledge as "absolutely unrevisable belief" [BRS12] on Kripke models.

³This philosophical stance is called epistemological skepticism [Str02]. It is not very fruitful for epistemic logic.

⁴Of course, one could argue that Alice might choose to reject Bob’s evidence and only trust her own. But in that case it is unclear why we would reason about what they know as a group in the first place.

as a validity⁵ if we reason about a fallible notion of knowledge.

Clearly, in the context of fallible knowledge, virtual group knowledge is a notion orthogonal to distributed knowledge according to Fernández [Fer18]. On topo-e-models, virtual group knowledge constitutes the practical *epistemic potential* of a group [Özg17]. Distributed knowledge, on the other hand, is reduced to a purely theoretical construct. This raises the question whether a traditional notion of distributed knowledge has any philosophical significance on topo-e-models.

Given this misalignment between distributed knowledge and virtual group knowledge, and given the association of the term *distributed knowledge* with monotonicity of knowledge, we chose to name our notions *virtual group knowledge* and *virtual group belief*, rather than distributed knowledge and distributed belief. This avoids confusion and stresses the fact that virtual group knowledge should not be mistaken for distributed knowledge in the traditional sense.

3.2 Logics on Multi-Agent Topo-E-Models

In this section, we define our multi-agent topo-e-models, by combining the single-agent topo-e-model from Definition 2.6 with a generalization of the notion of two-agent models from Definition 2.23. We formally define virtual group knowledge and belief on these models. We then present the multi-agent analogue of the language of individual evidence, after which we define a language of group evidence, as well as a language of group knowledge and group belief, incorporating these notions.

Definition 3.1 (Multi-Agent Topo-E-Model). Given a countable vocabulary V , a *multi-agent topo-e-model* for a group $A := \{1, \dots, n\}$ agents is a tuple $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ such that

1. X is a set of states;
2. $\mathcal{E}_i^0 \subseteq \mathcal{P}(X)$ consists of the basic evidence held by agent i , with $\emptyset \notin \mathcal{E}_i^0$;
3. $\tau_i \subseteq \mathcal{P}(X)$ is the topology for agent i defined on X ;
4. Π_i is the partition of X for agent i ;
5. $\pi : X \rightarrow \mathcal{P}(V)$ is a valuation;
6. we require $\Pi_i \subseteq \tau_i$ for all $i \in A$.

The *information cell* $\Pi_i(x)$ is the unique partition cell in Π_i that contains x . For $i \in A$ and $\pi_i \in \Pi_i$, we say that U is *locally dense* in π_i if $U \cap \pi_i$ is dense in the subspace topology⁶ $(\pi_i, \tau_i|_{\pi_i})$ (see Definition 2.24).

In our topological examples throughout the thesis, we will take the partition Π_i for each agent $i \in A$ to be a subset of their directly observed (i.e. subbasic) evidence \mathcal{E}_i^0 , so that it is automatically included in the generated topology τ_i . Conceptually, this is a natural assumption: it implies that the hard evidence is not obtained from the soft evidence by reasoning, but obtained directly through some kind of observation or prior knowledge.

⁵Fernández' argued that the resulting logic would not be useful, given that no validities remain [Fer18]. We did, in fact, find other validities, which we discuss in Section 4.1.

⁶Note that $U \cap \pi_i$ is dense in the subspace topology $(\pi_i, \tau_i|_{\pi_i})$ if and only if $\pi_i \subseteq Cl_{\tau_i}(U)$. We will often use the latter, simplified representation.

Observation 3.2. Technically, we interpret individual soft evidence on multi-agent topo-e-models relative to the partition cell: given a model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, a state $x \in X$, and an agent $i \in A$, i reasons about her evidence within the subspace topology $(\Pi_i(x), \tau_i|_{\Pi_i(x)})$ (see Definition 2.24). Thus, i has factive evidence for a proposition P if $x \in e \subseteq P$ for some $e \in \tau_i|_{\Pi_i(x)}$, rather than $x \in e' \subseteq P$ for some $e' \in \tau_i$. However, since each information cell $\pi \in \Pi_i$ is required to be open (Definition 3.1), it can be checked that the two statements are equivalent⁷. Thus, for simplicity, we evaluate evidence with respect to τ_i , instead of $\tau_i|_{\Pi_i(x)}$.

Similarly, individual knowledge and belief of an agent $i \in A$ are interpreted relative to the partition of i on the model. Instead of density of evidence, we require *local* density. The correspondence between τ_i and $\tau_i|_{\Pi_i(x)}$ allows us to define knowledge and belief in terms of evidence from τ_i , rather than evidence from $\tau_i|_{\Pi_i(x)}$.

Definition 3.3 (Individual Knowledge and Belief on Multi-Agent Topo-E-Models). Given a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, a state $x \in X$, an agent $i \in A$, and a proposition $P \subseteq X$, the agent i

$$\begin{aligned} \text{believes } P \text{ at } x & \text{ if } \exists U \in \tau_i \text{ s.t. } U \subseteq P \text{ and } \Pi_i(x) \subseteq Cl_{\tau_i}(U), \text{ and} \\ \text{knows } P \text{ at } x & \text{ if } \exists U \in \tau_i \text{ s.t. } x \in U \subseteq P \text{ and } \Pi_i(x) \subseteq Cl_{\tau_i}(U). \end{aligned}$$

3.2.1 Logic of Individual Evidence

We now define the multi-agent analogue $\Box[\forall]_i$ of the language of individual evidence from Definition 2.13. It defines the individual modalities for soft and hard evidence for each agent on the model.

Definition 3.4 (Syntax of $\Box[\forall]_i$). Given a countable vocabulary V and a group of agents A , the language $\Box[\forall]_i(V)$ of evidence is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_i\varphi \mid [\forall]_i\varphi$$

where $p \in V$ and $i \in A$ is any agent.

The semantics of soft and hard evidence in $\Box[\forall]_i$ is defined relative to the partition, following [Ram15; Fer18].

Definition 3.5 (Topological Semantics of $\Box[\forall]_i$). Given a countable vocabulary V and a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, the semantics of $\Box[\forall]_i(V)$ is defined as follows: the interpretation of atomic sentences, negation, and conjunction follows the semantics on single-agent topo-e-models, as defined for the language $\Box[\forall]$ (Definition 2.14). For the modalities of soft and hard evidence, for all $i \in A$, define

$$\begin{aligned} \llbracket \Box_i\varphi \rrbracket &:= Int_{\tau_i}(\llbracket \varphi \rrbracket) \\ &= \{x \in X \mid \exists U \in \tau_i(x \in U \subseteq \llbracket \varphi \rrbracket)\} \quad (\text{Obs. 3.2}) \\ \llbracket [\forall]_i\varphi \rrbracket &:= \{x \in X \mid \Pi_i(x) \subseteq \llbracket \varphi \rrbracket\}. \end{aligned}$$

⁷For one direction, suppose $x \in e \subseteq P$ for some $e \in \tau_i$ and consider the restricted open $(e \cap \Pi_i(x)) \in \tau_i|_{\Pi_i(x)}$. Since $x \in \Pi_i(x)$, we clearly also have that $x \in (e \cap \Pi_i(x))$. Furthermore, $(e \cap \Pi_i(x)) \subseteq P$, so $(e \cap \Pi_i(x)) \in \tau_i|_{\Pi_i(x)}$ supports P , as required. For the other direction, suppose $x \in e \subseteq P$ for $e \in \tau_i|_{\Pi_i(x)}$. Then $e = U \cap \Pi_i(x)$ (Definition 2.24) for some $U \in \tau_i$. Since $U \in \tau_i$ and $\Pi_i(x) \in \tau_i$, the intersection $U \cap \Pi_i(x)$ is also in τ_i . Thus, $e \in \tau_i$, which proves our claim.

Abbreviating the dual of \Box_i as $\Diamond_i\varphi := \neg B_i\neg\varphi$, it can be checked that, similar to the single-agent case,

$$\begin{aligned}\llbracket \Diamond_i\varphi \rrbracket &= Cl_{\tau_i}(\llbracket \varphi \rrbracket) \\ &= \{x \in X \mid \forall U \in \tau_i(x \in U \Rightarrow U \cap \llbracket \varphi \rrbracket \neq \emptyset)\}.\end{aligned}$$

Fact 3.6. Analogous to the single-agent case (see Fact 2.15), the notions of topological knowledge and belief from Definition 3.3 can be expressed in $\Box[\forall]_i$ as abbreviations:

$$\begin{aligned}B_i\varphi &= [\forall]_i\Diamond_i\Box_i\varphi \\ K_i\varphi &= \Box_i\varphi \wedge B_i\varphi.\end{aligned}$$

3.2.2 Logics of Group Evidence

We extend the language and semantics of individual evidence to groups by introducing group notions for soft and hard evidence. Throughout this thesis, all notions regarding subgroups of the full group of agents are defined only on nonempty groups: it is not straightforward to decide what the empty group of agents should know or believe, and if the empty group should know anything in the first place.

We represent the knowledge obtained by a group of agents exchanging evidence as a fictitious agent, by evaluating group notions of evidence on a different topology; we use the join topology. We generalize the two-agent join topology from Definition 2.25 to an arbitrary number of agents.

Definition 3.7 (Join Topology τ_I). Given a multi-agent topo-e-model $(X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ for a group of agents A , and a nonempty subgroup $I := \{1, \dots, n\} \subseteq A$, the *join topology* for I , denoted by τ_I , is given by $\bigvee_{i \in I} \tau_i$ and is generated by the subbasis $\mathcal{E}_I^0 := \bigcup_{i \in I} \mathcal{E}_i^0$. Equivalently, it is generated by the basis $\{\bigcap_{i \in I} U_i \mid U_1 \in \tau_1, \dots, U_n \in \tau_n\}$.

Being generated by the subbases from the individual topologies, the join captures exactly the notion of a group of agents sharing their direct evidence. We want to emphasize the fact that the \vee notation represents the algebraic join operation, rather than logical disjunction.

The partition for a group I in a multi-agent topo-e-model is obtained by intersecting all partition cells of its individual members $i \in I$. We generalize the two-agent definition from [Ram15; Fer18]:

Definition 3.8 (Group Partition). Given a multi-agent topo-e-model $(X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ nonempty subgroup $I := \{1, \dots, n\} \subseteq A$ and partitions Π_1, \dots, Π_n , define the *group partition* for I as $\Pi_I := \{\bigcap_{i \in I} \pi_i \mid \pi_i \in \Pi_i\}$.

Given a state $x \in X$, the information cell with respect to the group I at x is $\Pi_I(x) := \bigcap_{i \in I} \Pi_i(x)$.

The information cell $\Pi_I(x)$ represents the hard evidence, i.e. infallible knowledge, that the group $I \subseteq A$ has at state x , and combines the hard evidence that each of the agents $i \in I$ individually have at that state.

We can now define the language $\Box[\forall]_I$ of group evidence, which extends the language from Definition 3.4 to include group notions of the modalities for soft and hard evidence. Additionally, we define a fragment $\Box[\forall]_{i,A}$, which restricts the full language of evidence to individual agents $i \in A$ and the full group A . To be concise, we will often use the notation $\alpha \in \{A\} \cup A$ when we consider notions of evidence, knowledge, and belief restricted to individual agents and the full group A . We pattern match on both cases: if the label of the group is $\alpha \in \{A\} \cup A$, it can denote either a singleton set $\{i\} \subseteq A$, or A itself.

Definition 3.9 (Syntax of $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$). Given a countable vocabulary V , the language $\Box[\forall]_I$ of evidence is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_I \varphi \mid [\forall]_I \varphi$$

where $p \in V$ and $I \subseteq A$ is any nonempty subgroup of agents. The fragment $\Box[\forall]_{i,A}$ of $\Box[\forall]_I$ is obtained by denoting groups $\{i\} \subseteq A$ by i , and restricting the evidence modalities to \Box_α and $[\forall]_\alpha$, for all $\alpha \in \{A\} \cup A$.

As a natural extension of the semantics for the language of individual evidence (Definition 3.4), we propose the following semantics for the language of group evidence $\Box[\forall]_I$ and its fragment $\Box[\forall]_{i,A}$.

Definition 3.10 (Topological Semantics of $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$). Given a countable vocabulary V , and a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, the semantics of $\Box[\forall]_I(V)$ is defined as follows: the interpretation of atomic sentences, negation, and conjunction follows the semantics defined for $\Box[\forall]_i$ (Definition 3.5). For the modalities of soft and hard evidence, for all $I \subseteq A$, define

$$\begin{aligned} \llbracket \Box_I \varphi \rrbracket &:= \text{Int}_{\tau_I}(\llbracket \varphi \rrbracket) \\ &= \{x \in X \mid \exists U \in \tau_I (x \in U \subseteq \llbracket \varphi \rrbracket)\} \\ \llbracket [\forall]_I \varphi \rrbracket &:= \{x \in X \mid \Pi_I(x) \subseteq \llbracket \varphi \rrbracket\}. \end{aligned}$$

The semantics for $\Box[\forall]_{i,A}(V)$ on multi-agent topo-e-models is obtained by restricting the semantics for $\Box[\forall]_I(V)$ to $\Box[\forall]_{i,A}(V)$.

In particular, observe that hard evidence of the group, also referred to as infallible group knowledge, satisfies the (KD) axiom⁸: we have for all nonempty $I \subseteq J \subseteq A$ that $[\forall]_I \varphi$ implies $[\forall]_J \varphi$. Furthermore, we can derive from Definition 3.10 that $\llbracket \Diamond_I \varphi \rrbracket = \text{Cl}_{\tau_I}(\llbracket \varphi \rrbracket)$, analogous to the interpretation of \Diamond_i .

This definition is a clear generalization of the semantics from Definition 3.5 for single-agent evidence and partitions on multi-agent models: the respective interpretations of $\Box_{\{i\}}$ and $[\forall]_{\{i\}}$ in Definition 3.10 are equivalent to the interpretations of \Box_i and $[\forall]_i$ according to the semantics for $\Box[\forall]_i$ (Definition 3.5). For simplicity, in the remainder of this thesis, we denote singleton sets $\{i\}$ by i .

In Section 4.2, we introduce a proof system $\Box[\forall]_I$ for $\Box[\forall]_I$, as well as a restricted proof system $\Box[\forall]_{i,A}$ for the fragment $\Box[\forall]_{i,A}$. We prove completeness of these proof systems in Section 5.1.

3.2.3 Logics of Group Knowledge and Belief

It remains to define the group notions for knowledge and belief, i.e. *virtual group knowledge* and *virtual group belief*. We naturally extend the single-agent interpretation to the multi-agent case: group knowledge of a nonempty group $I \subseteq A$ is read as a factive justification in the join topology τ_I . Formally, we define:

Definition 3.11 (Virtual Group Knowledge and Belief ($\Box[\forall]_I$)). Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a multi-agent topo-e-model with agents A . Let $I \subseteq A$ be nonempty and let $x \in X$. Then, for any

⁸In fact, [Fag+04] discusses an interpretation of distributed knowledge, which coincides with our semantics of group evidence, on *Aumann structures*; these structures can roughly be described as a variation of multi-agent topo-e-models, which defines only partitions, and no soft evidence. On these structures, the interpretation of distributed knowledge D_G of a group G in [Fag+04] is defined as $\llbracket D_G(\varphi) \rrbracket := \{x \in X \mid (\bigcap_{i \in G} \Pi_i(x)) \subseteq \llbracket \varphi \rrbracket\}$.

proposition $P \subseteq X$, the group I has

$$\begin{aligned} \text{virtual group belief in } P & \quad \text{if } \exists U \in \tau_I \text{ s.t.} \quad U \subseteq P \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U), \text{ and} \\ \text{virtual group knowledge of } P & \quad \text{if } \exists U \in \tau_I \text{ s.t. } x \in U \subseteq P \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U). \end{aligned}$$

Comparing Definition 3.11 with Definitions 2.11 and 2.12 (defined on single-agent topo-e-models) and Definition 3.3 (defined on multi-agent topo-e-models), the symmetry with the topological interpretation of individual knowledge and belief is evident: where individual knowledge is read as the dense interior in the *individual* topology, group knowledge is read as the locally dense interior in the *join* topology.

To be able to reason about evidence-based knowledge and belief without explicitly mentioning notions of evidence, we introduce a language KB_I , which defines knowledge and belief as primitive operators and does not include any evidence operators. We will find, however, that formulas in KB_I can be translated to $\Box[\forall]_I$ and, as a result, we can consider KB_I as a fragment of $\Box[\forall]_I$.

Additionally, we define a fragment $KB_{i,A}$, which restricts KB_I to individual agents and the full group. The fragment $KB_{i,A}$, is of primary significance: as we discussed in Section 2.3, the logic of knowledge and belief for individual agents and the full group is one of the most natural logics of group knowledge. Furthermore, this logic has a more complicated completeness proof than the logics of evidence; it is not straightforward, and constitutes the main theoretical result of this thesis. We could not prove completeness of the logic of KB_I .

To ensure that virtual group knowledge is not confused with distributed knowledge, we deliberately avoid using the D_I notation for our notion. Instead, we use the notation K_I , which gives us a straightforward notation for virtual group belief as B_I .

Definition 3.12 (Syntax of KB_I and $KB_{i,A}$). Given a countable vocabulary V , the language fragment $KB_I(V)$ of belief and knowledge is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_I\varphi \mid B_I\varphi$$

where $p \in V$ and $I \subseteq A$ is any nonempty subgroup of agents. The fragment $KB_{i,A}$ of KB_I is obtained by denoting groups $\{i\} \subseteq A$ by i , and restricting the modalities of knowledge and belief to K_α and B_α , for all $\alpha \in \{A\} \cup A$.

The semantics of the modalities K_I for group knowledge and B_I for group belief capture precisely the topological notions from Definition 3.11 of virtual group knowledge virtual group belief, respectively.

Definition 3.13 (Topological Semantics of KB_I and $KB_{i,A}$). Given a countable vocabulary V and a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, the semantics of $KB_I(V)$ is defined as follows: the interpretation of atomic sentences, negation, and conjunction follows the semantics defined for $\Box[\forall]_I$ (Definition 3.5). For the modalities of soft and hard evidence, for all $I \subseteq A$, define

$$\begin{aligned} \llbracket B_I\varphi \rrbracket &:= \{x \in X \mid \exists U \in \tau_I \text{ s.t.} \quad U \subseteq \llbracket \varphi \rrbracket \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U)\} \\ \llbracket K_I\varphi \rrbracket &:= \{x \in X \mid \exists U \in \tau_I \text{ s.t. } x \in U \subseteq \llbracket \varphi \rrbracket \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U)\}. \end{aligned}$$

The semantics for $KB_{i,A}(V)$ on multi-agent topo-e-models is obtained by restricting the semantics of $KB_I(V)$ to $KB_{i,A}(V)$.

Observation 3.14. Recalling that the join topology τ_I is generated by the basis $\{\bigcap_{i \leq n} U_i \mid U_1 \in \tau_1, \dots, U_n \in \tau_n\}$ (Definition 3.7), we can equivalently evaluate knowledge and belief in terms of the basis: taking $I := \{1, \dots, n\}$, we can replace $U \in \tau_I$ in Definition 3.13 with an intersection $\bigcap_{i \leq n} U_i$ of evidence $U_1 \in \tau_1, \dots, U_n \in \tau_n$ from the individual topologies.

The following propositions show that the topological virtual group knowledge and virtual group belief, as defined in Definition 3.11, can be expressed as abbreviations in the language $\Box[\forall]_I$ of group evidence.

Proposition 3.15. Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a multi-agent topo-e-model with agents A . Let $I \subseteq A$ be nonempty and let $x \in X$. Then, for any formula φ in the language $\Box[\forall]_I$, we have:

- (1) $x \in \llbracket [\forall]_I \Diamond_I \Box_I \varphi \rrbracket$ iff I has virtual group belief in φ
- (2) $x \in \llbracket \Box_I \varphi \rrbracket \cap \llbracket [\forall]_I \Diamond_I \Box_I \varphi \rrbracket$ iff I has virtual group knowledge of φ .

Proof. We first prove (1). For the left-to-right direction, we have

$$\begin{aligned}
 x \in \llbracket [\forall]_I \Diamond_I \Box_I \varphi \rrbracket &\Leftrightarrow \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket)) && \text{(Def. 3.10)} \\
 &\Rightarrow \exists U \in \tau_I \text{ s.t. } U \subseteq \llbracket \varphi \rrbracket \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U) && (Int_{\tau_I}(\llbracket \varphi \rrbracket) \in \tau_I, Int_{\tau_I}(\llbracket \varphi \rrbracket) \subseteq \llbracket \varphi \rrbracket) \\
 &\Leftrightarrow I \text{ has virtual group belief in } \varphi && \text{(Def. 3.11)}
 \end{aligned}$$

For the converse direction, suppose there exists $U \in \tau_I$ such that $U \subseteq \llbracket \varphi \rrbracket$ and $\Pi_I(x) \subseteq Cl_{\tau_I}(U)$. Being an open subset of $\llbracket \varphi \rrbracket$, we know that $U \subseteq Int_{\tau_I}(\llbracket \varphi \rrbracket)$. By Proposition 2.3, this implies that $Cl_{\tau_I}(U) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket))$ ⁹. But then, $\Pi_I(x) \subseteq Cl_{\tau_I}(U) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket))$, and therefore, $x \in \llbracket [\forall]_I \Diamond_I \Box_I \varphi \rrbracket$ (Definition 3.10).

For (2), we use the result from (1):

$$\begin{aligned}
 x \in \llbracket \Box_I \varphi \rrbracket \cap \llbracket [\forall]_I \Diamond_I \Box_I \varphi \rrbracket &\Leftrightarrow \exists V \in \tau_I \text{ s.t. } x \in V \subseteq \llbracket \varphi \rrbracket \text{ and} \\
 &\Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket)) && \text{(Def. 3.10)} \\
 &\Leftrightarrow x \in Int_{\tau_I}(\llbracket \varphi \rrbracket) \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket)) && \text{(Def. 2.2)} \\
 &\Leftrightarrow \exists U \in \tau_I \text{ s.t. } x \in U \subseteq \llbracket \varphi \rrbracket \text{ and } \Pi_I(x) \subseteq Cl_{\tau_I}(U) && (Int_{\tau_I}(\llbracket \varphi \rrbracket) \in \tau_I, (1)) \\
 &\Leftrightarrow I \text{ has virtual group knowledge of } \varphi. && \text{(Def. 3.11)}
 \end{aligned}$$

□

Corollary 3.16. We can consider the language KB_I as a fragment of $\Box[\forall]_I$.

Proof. The claim can be proved via straightforward induction on the complexity of $\varphi \in KB_I$, by showing that $\varphi \in KB_I$ is equivalent to some $\varphi' \in \Box[\forall]_I$. The base case of atomic propositions and the boolean cases of the induction step are immediate; the cases of knowledge and belief follow from Proposition 3.15. □

In Section 4.2, we axiomatise the fragment $KB_{i,A}$. Additionally, we present a candidate axiomatisation for KB_I , for which we could not prove completeness. The completeness proof for the axiomatisation $KB_{i,A}$ is presented in Section 5.2.

⁹To see this, let (X, τ) be any topological space, let $A, B \subseteq X$. We have $A \subseteq B$, if and only if $B = A \cup B$, if and only if $Cl_{\tau}(B) = Cl_{\tau}(A \cup B)$, if and only if $Cl_{\tau}(B) = Cl_{\tau}(A) \cup Cl_{\tau}(B)$, if and only if $Cl_{\tau}(A) \subseteq Cl_{\tau}(B)$.

3.3 Dynamics: Sharing the Evidence

In this section, we formalize our motivation for group knowledge as the result of communication (see Section 3.1), by defining a dynamic modality $[share_I]$ that models the sharing of soft and hard evidence within a group of agents. We first define $[share_I]$ and extend the languages $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$ with this modality (Section 3.3.1). Next, we prove the crucial statement that virtual group knowledge of a group pre-encodes the individual knowledge of its members that results from an event of evidence sharing (Section 3.3.2). We conclude the section with a toy example that illustrates the interplay between virtual group knowledge and the $[share_I]$ modality (Section 3.3.3). Throughout this section, fix a finite set A of agents and a countable vocabulary V .

3.3.1 Dynamic Extensions

The dynamic modality $[share_I]$ updates a multi-agent topo-e-model \mathfrak{M} to a multi-agent topo-e-model \mathfrak{M}^{share_I} , which reflects the result of all members of a subgroup $I \subseteq A$ sharing their individual evidence in \mathfrak{M} . We formally define the updated model \mathfrak{M}^{share_I} in Definition 3.17. Following [ÅW17], we consider the event of information sharing common knowledge among all agents in the model, while keeping the contents of the shared information private.

Definition 3.17. Given a model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ and two nonempty subgroups $I, J \subseteq A$, let $\mathfrak{M}^{share_I} := (X, (\mathcal{E}_i^{0,S})_{i \in A}, (\tau_i^S)_{i \in A}, (\Pi_i^S)_{i \in A}, \pi)$ where we have for all $i \in I$ and for all $j \notin I$ that

$$\begin{array}{ll} \mathcal{E}_i^{0,S} &= \mathcal{E}_I^0 & \mathcal{E}_j^{0,S} &= \mathcal{E}_j^0 \\ \tau_i^S &= \tau_I & \tau_j^S &= \tau_j \\ \Pi_i^S &= \Pi_I & \Pi_j^S &= \Pi_j. \end{array}$$

We extend the languages $\Box[\forall]_I$, $\Box[\forall]_{i,A}$, and $KB_{i,A}$ with the $[share_I]$ modality. Given the restrictions on group knowledge, the fragments $\Box[\forall]_{i,A}$ and $KB_{i,A}$ allow sharing only within the full group. We define the extended languages and the corresponding semantics as follows.

Definition 3.18 (Syntax of the Dynamic Languages). Given a countable vocabulary V , the dynamic language $\Box[\forall]_I(V) + [share_I]$ of group evidence is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_I \varphi \mid [\forall]_I \varphi \mid [share_I] \varphi$$

where $p \in V$, and $I \subseteq A$ is any nonempty subgroup of agents.

The fragment $\Box[\forall]_{i,A} + (V) + [share_A]$ of evidence is obtained by denoting groups $\{i\} \subseteq A$ by i , and restricting the evidence modalities to \Box_α and $[\forall]_\alpha$, for all $\alpha \in \{A\} \cup A$, and restricting the dynamic modality to $[share_A]$.

The fragment $KB_{i,A} + (V) + [share_A]$ of evidence is defined recursively as

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i \varphi \mid K_A \varphi \mid B_i \varphi \mid B_A \varphi \mid [share_A] \varphi$$

where $p \in V$, and $i \in A$ is any agent.

A pointed model (\mathfrak{M}, x) satisfies the formula $[share_I] \varphi$ if and only if the updated pointed model $(\mathfrak{M}^{share_I}, x)$ satisfies the formula φ . Formally, we define, for each of the languages:

Definition 3.19 (Topological Semantics of the Dynamic Languages). For formulas over the language $\Box[\forall]_I(V) + [share_I]$, the interpretation of atomic sentences, negation, conjunction, and the modalities B_I and $[\forall]_I$ on multi-agent topo-e-models follows the semantics for $\Box[\forall]_I$ (Definition 3.10). We interpret the $[share_I]$ modality as follows: given a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ and a nonempty subgroup $I \subseteq A$, let

$$\llbracket [share_I] \varphi \rrbracket^{\mathfrak{M}} := \llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}}.$$

The semantics for the fragment $\Box[\forall]_{i,A}(V) + [share_A]$ on multi-agent topo-e-models is obtained by restricting the semantics of $\Box[\forall]_I + [share_I]$ to formulas in $\Box[\forall]_{i,A} + [share_A]$.

For formulas over the language $KB_{i,A}(V) + [share_A]$, the interpretation of atomic sentences, negation, conjunction, and the modalities K_i, K_A, B_i , and B_A on multi-agent topo-e-models follows the semantics for the language $KB_{i,A}$ (Definition 3.13). The $[share_A]$ modality is interpreted as follows:

$$\llbracket [share_A] \varphi \rrbracket^{\mathfrak{M}} := \llbracket \varphi \rrbracket^{\mathfrak{M}^{share_A}}.$$

3.3.2 Group Knowledge Pre-encodes Knowledge after Sharing

After the sharing of evidence within a subgroup J , all members $j \in J$ individually know everything that was previously implicitly known by the group. As a result, each member $j \in J$ can share their newly obtained evidence with other subgroups in future updates. Essentially, after an event $share_J$, the implicit knowledge of all groups containing j gets supplemented with what j has learned. In Proposition 3.22, we prove a generalisation of this statement to all modalities of evidence, knowledge, and belief in our languages, and is proved by Proposition 3.22.

Proposition 3.22 makes use of Lemma 3.20, which states that sharing updates do not influence the valuation function on the model, and Lemma 3.21, which defines the updated evidence for groups that overlap with the sharing group. We prove these lemmas first.

Lemma 3.20. Fix a countable vocabulary V and any $p \in V$, fix a multi-agent topo-e-model

$$\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$$

and fix any nonempty subgroup of agents $J \subseteq A$. Consider the semantics from Definition 3.19 for the language $\Box[\forall]_I + [share_I]$. Then we have $\llbracket p \rrbracket^{\mathfrak{M}^{share_J}} = \llbracket p \rrbracket^{\mathfrak{M}}$.

Proof of claim. The $[share_J]$ modality does not change the truth value of propositional variables on \mathfrak{M} . Since both models have the same state space, we get from the semantics (Definition 3.19) that the truth set of p is identical in \mathfrak{M} and \mathfrak{M}^{share_J} : we have for all $x \in X$ that

$$(\mathfrak{M}, x) \models p \quad \text{iff} \quad (\mathfrak{M}^{share_J}, x) \models p \quad \text{iff} \quad (\mathfrak{M}, x) \models [share_J]p.$$

and therefore, $\llbracket p \rrbracket^{\mathfrak{M}^{share_J}} = \llbracket p \rrbracket^{\mathfrak{M}}$. □

Lemma 3.21. Fix a multi-agent topological model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, and let $I, J \subseteq A$ such that $I \cap J \neq \emptyset$. Consider the updated model $\mathfrak{M}^{share_J} := (X, (\mathcal{E}_i^{0,S})_{i \in A}, (\tau_i^S)_{i \in A}, (\Pi_i^S)_{i \in A}, \pi)$. Let τ_I^S and Π_I^S respectively denote the join topology and the group partition for I over \mathfrak{M}^{share_J} . Then we have that $\tau_I^S = \tau_{I \cup J}$ and $\Pi_I^S = \Pi_{I \cup J}$.

Proof. The proof follows directly from the definitions of τ_I and Π_i and can be found in Appendix A.2. \square

It remains to prove the correspondence between the $[share_I]$ modality and virtual group knowledge. We state each of the following claims with respect to the full language $\Box[\forall]_I + [share_I]$. Since the languages $\Box[\forall]_{i,A} + [share_A]$ and $KB_{i,A} + [share_A]$ are fragments, their semantics are given by restrictions of the semantics of the larger language, thus, the proofs extend to the fragments.

The cases for knowledge and belief constitute the most important part of the proof: our motivation for virtual group knowledge and virtual group belief is based, in particular, on these instances of the claim. Therefore, we explicitly prove the case for knowledge. We omit the proofs for the other modalities: each of these proofs follows a similar line of reasoning to the case that we show.

Proposition 3.22. *Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a multi-agent topo-e-model, let $x \in X$ be a state, and let $I, J \subseteq A$ be nonempty subgroups. Fix a countable vocabulary V and let $p \in V$ be any atomic proposition. Then, for any modality $M \in \{\Box, [\forall], B, K\}$, the following equivalences are true:*

- (1) $(\mathfrak{M}^{share_I}, x) \models M_I p \Leftrightarrow (\mathfrak{M}, x) \models M_{I \cup J} p$ if $I \cap J \neq \emptyset$
- (2) $(\mathfrak{M}^{share_I}, x) \models M_I p \Leftrightarrow (\mathfrak{M}, x) \models M_I p$ if $I \cap J = \emptyset$.

Proof. We show the proof of (1) for the knowledge modality: assume that $I \cap J \neq \emptyset$. We have the following chain of equivalences:

$$\begin{aligned}
(\mathfrak{M}^{share_I}, x) \models K_I p & \text{ iff } x \in \llbracket K_I p \rrbracket^{\mathfrak{M}^{share_I}} \\
& \text{ iff } x \in Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}})) \quad (\text{Def. 3.19}) \\
\text{iff} & \quad x \in Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}})) \quad (\text{Lem. 3.20}) \\
\text{iff} & \quad x \in Int_{\tau_{I \cup J}}(\llbracket p \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_{I \cup J}(x) \subseteq Cl_{\tau_{I \cup J}}(Int_{\tau_{I \cup J}}(\llbracket p \rrbracket^{\mathfrak{M}})) \quad (\text{Lem. 3.21}) \\
\text{iff} & \quad x \in \llbracket K_{I \cup J} p \rrbracket^{\mathfrak{M}} \quad (\text{Def. 3.19}) \\
\text{iff} & \quad (\mathfrak{M}, x) \models K_{I \cup J} p.
\end{aligned}$$

For (2), assume that $I \cap J = \emptyset$. Then we have for all $i \in I$ that $\tau_i^S = \tau_i$ and $\Pi_i^S = \Pi_i$ (Definition 3.17), and therefore, $\tau_I^S = \tau_I$ and $\Pi_I^S = \Pi_I$ (*). We get the following chain of equivalences:

$$\begin{aligned}
(\mathfrak{M}^{share_I}, x) \models K_I p & \text{ iff } x \in \llbracket K_I p \rrbracket^{\mathfrak{M}^{share_I}} \\
& \text{ iff } x \in Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}})) \quad (\text{Def. 3.19}) \\
& \text{ iff } x \in Int_{\tau_I}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket p \rrbracket^{\mathfrak{M}^{share_I}})) \quad (*) \\
& \text{ iff } x \in Int_{\tau_I}(\llbracket p \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket p \rrbracket^{\mathfrak{M}})) \quad (\text{Lem. 3.20}) \\
& \text{ iff } x \in \llbracket K_I p \rrbracket^{\mathfrak{M}} \quad (\text{Def. 3.19}) \\
& \text{ iff } (\mathfrak{M}, x) \models K_I p.
\end{aligned}$$

\square

We conclude with a corollary that states the special cases from Proposition 3.22 that support our justification for virtual group knowledge and virtual group belief: the interplay between the sharing update and our group notions of knowledge and belief.

Corollary 3.23. *Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a multi-agent topo-e-model, let $x \in X$ be a state, and let $I, J \subseteq A$ be nonempty subgroups. Fix a countable vocabulary V and let $p \in V$ be any atomic proposition. For knowledge and belief, the following equivalences are true:*

$$\begin{aligned} (1_K) \quad & (\mathfrak{M}^{share_I}, x) \models K_I p \Leftrightarrow (\mathfrak{M}, x) \models K_{I \cup J} p \quad \text{if } I \cap J \neq \emptyset \\ (1_B) \quad & (\mathfrak{M}^{share_I}, x) \models B_I p \Leftrightarrow (\mathfrak{M}, x) \models B_{I \cup J} p \quad \text{if } I \cap J \neq \emptyset. \end{aligned}$$

Proof. These are special cases in Proposition 3.22. □

Note that the result from Proposition 3.22 does not hold when we substitute p for an arbitrary formula, as the truth value of formulas containing modalities may change with a sharing update. In Section 3.3.3, we demonstrate by means of an example how, in particular, implicit knowledge is updated through sharing events.

Proposition 4.17 proves an adaptation of the statement to arbitrary formulas, which we need to prove soundness of the axiomatisations of the dynamic languages (introduced in Section 4.2).

3.3.3 Example: Missing Cake

In the following example, we put our model transformer into practice. It shows how the implicit knowledge of one subgroup of agents gets updated when a member of this group is involved in an act of sharing within a different subgroup.

Missing Cake: initial situation. Alice, Bob, and Charles are roommates. Bob baked a chocolate cake earlier. There was only one piece left, which is suddenly missing: Alice secretly ate it. The roommates are trying to find out who did it. The initial situation is as follows: Alice remembers eating the cake, so she knows that she's guilty. Bob knows that he did not do it, and he has evidence that Charles did it: he found the empty baking tin in Charles' room, surrounded by cake crumbs. In fact, Alice planted this evidence. Charles, who drank too much yesterday, does not remember whether he did it or not. He did call some of Bob's friends to find out if Bob has an alibi for last night, which he does: his mom claims that Bob came by. Charles now believes that it was not Bob. If Bob and Charles were to share their individual evidence, they would both come to believe that Charles ate the cake.

We formalise the initial scenario in Example 3.24, depicted in Figure 3.1, and show for a particular instance of Proposition 3.22 that it holds.

Example 3.24. Let $V := \{p_a, p_b, p_c\}$ be our vocabulary and let the set of agents be $A := \{a, b, c\}$, representing Alice, Bob, and Charles. Let $\mathfrak{M}_0 := (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a three-agent topo-e-model where $X = \{w_a, w_b, w_c\}$, and

$$\begin{aligned} \mathcal{E}_a^0 &= \{\{w_a\}, \{w_b, w_c\}, X\} & \Pi_a &= \{\{w_a\}, \{w_b, w_c\}\} \\ \mathcal{E}_b^0 &= \{\{w_b\}, \{w_a, w_c\}, \{w_c\}, X\} & \Pi_b &= \{\{w_b\}, \{w_a, w_c\}\} \\ \mathcal{E}_c^0 &= \{\{w_a, w_c\}, X\} & \Pi_c &= \{X\} \end{aligned}$$

and $\pi : X \rightarrow \mathcal{P}(V)$ is a valuation function such that for each $i \in A$ we have that $\llbracket p_i \rrbracket^{\mathfrak{M}_0} = \{w_i\}$.

For each $i \in A$, the propositional variable p_i denotes the proposition " i ate the cake" and w_i is the state in which i ate the cake: $(\mathfrak{M}_0, w_i) \models p_i \wedge \bigwedge_{j \in A \setminus \{i\}} \neg p_j$.

The group of Bob and Charles implicitly believes that Charles ate the cake: it can be checked on the join topology $\tau_b \vee \tau_c$ that we have $(\mathfrak{M}_0, w_a) \models B_{\{b,c\}} p_c$. Clearly, $(\mathfrak{M}_0, w_a) \not\models K_{\{b,c\}} p_a$.

However, we can check that $(\mathfrak{M}_0, w_a) \models K_{\{b,c\} \cup \{a,b\}} p_a$, and that, in accordance with Proposition 3.22, $(\mathfrak{M}_0^{[\text{share}_{\{a,b\}}]}, w_a) \models K_{\{b,c\}} p_a$ also holds.

For the first claim, observe that $(\mathfrak{M}_0, w_a) \models K_A p$. This follows directly from the fact that $\Pi_a(w_a) = \{w_a\}$, so it must be that $\Pi_A(w_a) = \{w_a\}$. We discuss the second claim, that $(\mathfrak{M}_0^{[\text{share}_{\{a,b\}}]}, w_a) \models K_{\{b,c\}} p_a$, in Example 3.25, where we update the scenario and define the model $\mathfrak{M}_0^{[\text{share}_{\{a,b\}}]}$.




Figure 3.1: The initial model \mathfrak{M}_0 . On the left side, from left to right, are the individual topologies for Alice, Bob, and Charles. On the right side: the join topology for Bob and Charles. The actual world is w_a . We draw hard evidence using solid lines, and soft evidence using dashed lines.

Despite the fact that the group of Bob and Charles implicitly believes that Charles is guilty, by Proposition 3.22, an act of evidence sharing between Alice and Bob should result in the group of Bob and Charles implicitly knowing that Alice is guilty. We check that $K_{\{b,c\}} p_a$ indeed holds on $\mathfrak{M}_0^{[\text{share}_{\{a,b\}}]}$, by updating our model with an act of sharing between Alice and Bob:

Missing Cake: update. Bob, who believes Charles ate the cake, knocks on Alice's door to discuss his suspicions and find out what she knows (and Charles is aware of this information being shared). Alice feels guilty and confesses everything to Bob. So now Bob and Alice both individually know infallibly who ate the cake. If Bob were to share his evidence with Charles, then Charles would also individually know: in other words, the group of Bob and Charles, which initially implicitly believed that Charles was guilty, now implicitly knows that Alice is the culprit.

We formalise the result of this update in Example 3.25, depicted in Figure 3.2.

Example 3.25. Let \mathfrak{M}_0 be the model defined in Example 3.24. Let

$$\mathfrak{M}_1 := \mathfrak{M}_0^{[\text{share}_{\{a,b\}}]} = (X, (\mathcal{E}_i^{0,S})_{i \in A}, (\tau_i^S)_{i \in A}, (\Pi_i^S)_{i \in A}, \pi)$$

be the model resulting from an act of sharing within the group $\{a, b\}$. The state space is still $X = \{w_a, w_b, w_c\}$. We have

$$\begin{aligned} \mathcal{E}_a^{0,S} &= \mathcal{E}_b^{0,S} = \mathcal{E}_{\{a,b\}}^0 = \{\{w_a\}, \{w_b\}, \{w_c\}, X\} & \Pi_a^S &= \Pi_b^S = \Pi_{\{a,b\}}^S = \{\{w_a\}, \{w_b\}, \{w_c\}\} \\ \mathcal{E}_c^{0,S} &= \mathcal{E}_c^0 = \{\{w_a, w_c\}, X\} & \Pi_c^S &= \Pi_c = \{X\}. \end{aligned}$$

First, note that a and b both individually have infallible knowledge of the actual state at w_a . Moreover, contrary to the situation in \mathfrak{M}_0 , they individually have infallible knowledge of the actual state at any evaluated state in the model.

To see that $(\mathfrak{M}_1, w_a) \models K_{\{b,c\}} p_a$, we consider the model $\mathfrak{M}_1^{[share_{\{b,c\}}]}$, that would result from b and c sharing their individual evidence in model \mathfrak{M}_1 . The topologies for b and c on the resulting model $\mathfrak{M}_1^{[share_{\{b,c\}}]}$ are both given by the join topology $\tau_b^S \vee \tau_c^S$, depicted in Figure 3.2. We find that b and, in particular, c individually infallibly know p_a .




Figure 3.2: The updated model \mathfrak{M}_1 . On the left side are the individual topologies: the topology shared by Alice and Bob, and Charles' unchanged topology. On the right side is the updated join topology for Bob and Charles. The actual world is w_a .

Thus, in the model \mathfrak{M}_1 resulting from Alice's confession, Bob and Charles indeed have infallible virtual group knowledge that Alice ate the cake.

3.4 Examples

The main axiom associated with distributed knowledge fails when we use our proposed semantics on our proposed models. In fact, weaker, related axioms also fail. In this section we discuss three realistic scenarios that illustrate the consequences of fallible knowledge. Our first example concerns two agents who believe the same proposition, but both drop this belief after exchanging evidence. Consider the following scenario:

Lazy Teacher: scenario 1. Suppose Alice and Bob both believe their teacher will cancel class tomorrow. Alice bases her belief on evidence from her *AllWeather* weather app, which predicts that it will be raining all day; and the lazy teacher is famous for disliking biking to school when it rains. Bob, on the other hand, bases his belief on evidence from the *BestWeather* weather app - and more often than not, weather apps are mutually contradictory. *BestWeather* predicts that tomorrow will be the most beautiful day of the week and it is, in fact, also common knowledge that the teacher cancels class to go to the beach whenever the sun is out. Additionally, based on prior evidence regarding the accuracy of *AllWeather*, Alice considers it more plausible that tomorrow will be a regular day in weather terms^a, than that the day will be beautiful. Conversely, based on prior evidence regarding *BestWeather*, Bob considers a regular day more plausible than Alice's scenario. Now, after exchanging their evidence, none of the three scenarios seem to be any more likely than the others: neither Alice nor Bob has reason to believe that class will be cancelled.

^aLet's define regular weather as the gray kind of weather: there is no sunshine, but it is not raining either.

Even though Alice and Bob share a belief, their justifications are mutually contradictory; since they do not value one piece of evidence over another, they drop their individual beliefs after exchanging evidence. The scenario is formalised in Example 3.26 and depicted in Figure 3.3. It illustrates the failure of the axiom $\bigwedge_{i \in I} B_i \varphi \rightarrow B_I \varphi$. It is closely related to, but slightly different from Example 4.1.1 in [Ram15].

Example 3.26. Let $V := \{p\}$ be our vocabulary. Let $\mathfrak{M} := (X, \mathcal{E}_a^0, \mathcal{E}_b^0, \tau_a, \tau_b, \Pi_a, \Pi_b, \pi)$ be a two-agent topo-e-model where $X = \{x, y, z\}$, and we have that $\mathcal{E}_a^0 = \{\{x\}, \{x, y\}, X\}$, $\mathcal{E}_b^0 = \{\{z\}, \{y, z\}, X\}$, and $\Pi_a = \Pi_b = \{X\}$. Define $\pi : X \rightarrow \mathcal{P}(V)$ such that $\llbracket p \rrbracket = \{x, z\}$.

Alice and Bob are represented by a and b , respectively. The states x, y , and z respectively represent the situations in which tomorrow will be rainy, regular, and sunny. The proposition p denotes the sentence "class will be cancelled tomorrow".

We evaluate at x . Then $(\mathfrak{M}, x) \models B_a p \wedge B_b p$, since both agents have locally dense evidence supporting p : for agent a , we have $\tau_a \ni \{x\} \subseteq \{x, z\} = \llbracket p \rrbracket$ with $Cl_{\tau_a}(\{x\}) = X = \Pi_a$. For agent b , the evidence for p is given by $\{z\}$.

Consider the join topology $\tau_A = \tau_a \vee \tau_b = \{\emptyset, \{x\}, \{y\}, \{z\}, \{x, y\}, \{y, z\}, X\}$: combined evidence $\{y\} \in \tau_A$ resulted from sharing the evidence. Since $\{y\} \cap \{x\} = \{y\} \cap \{z\} = \emptyset$, the evidence from the individual topologies is not dense with respect to τ_A . In fact, X is the only dense set with respect to τ_A and $X \not\subseteq \llbracket p \rrbracket$. Therefore, $(\mathfrak{M}, x) \not\models B_A p$.




Figure 3.3: From left to right: Alice's topology, the join topology, and Bob's topology. The actual world is x . For clarity, we omit the partition cells from the diagram, since the partition for each agent is given by $\{X\}$. The new evidence is coloured gray.

It is important to note that at state x , Alice knows p , but Bob does not. Bob's belief in p is correct and justified, but it is not *correctly* justified. This shows that topological knowledge is indeed distinct from the Gettierizable definition of knowledge as justified true belief.

This example raises the question whether knowledge, or even belief in factive propositions, might be stable under an evidence merge. However, this is not a given in the real world. Let's adapt Alice and Bob's scenario such that the individual agents do not only believe that class will be cancelled; they both individually *know* it.

Lazy Teacher: scenario 2. Suppose that Alice's *AllWeather* app and Bob's *BestWeather* app both give separate predictions with respect to sunshine and rain. In this situation, Alice and Bob both have a correctly justified belief that the teacher will be cancelling class: Alice's *AllWeather* app predicts rain, and no sunshine. Bob's *BestWeather* app, on the other hand, predicts sunshine and no rain. The following situations are possible:

1. Analogous to scenario 1, the teacher cancels class in case of rain (he dislikes biking through it), and in case of sun (he will be found at the beach);
2. Whenever both weather conditions are satisfied, he will definitely take advantage of the situation and cancel class;
3. When the weather is regular, that is, when the sun does not shine and it does not rain, class goes on as usual.

It happens to be the case that tomorrow will be both sunny and rainy. Alice and Bob have correct justifications: *AllWeather* correctly predicted rain, and *BestWeather* correctly predicted sun. However, *AllWeather*'s false prediction contradicts *BestWeather*'s correct prediction, and vice versa. When they compare the contradictory predictions on both weather apps, they are - again - forced to drop their beliefs: all weather conditions are now equally plausible!

So even if both agents have a correct justification, an exchange of evidence can cause them to have reasonable doubt. In the light of new evidence, the possibility that class gets cancelled becomes relevant. Example 3.27, depicted by Figure 3.4, formalises this failure of the $\bigwedge_{i \in I} K_i \varphi \rightarrow K_I \varphi$ axiom.

Example 3.27. Again, let p denote the sentence "class will be cancelled tomorrow". Let $V := \{p\}$ be our vocabulary. Let $\mathfrak{M} := (X, \mathcal{E}_a^0, \mathcal{E}_b^0, \tau_a, \tau_b, \Pi_a, \Pi_b, \pi)$ be a two-agent topo-e-model where $X = \{w_1, w_2, w_3, w_4\}$, and

$$\begin{aligned}\mathcal{E}_a^0 &= \{\{w_2, w_4\}, \{w_3, w_4\}, X\} \\ \mathcal{E}_b^0 &= \{\{w_1, w_2\}, \{w_1, w_3\}, X\} \\ \Pi_a &= \Pi_b = \{X\}\end{aligned}$$

and $\pi : X \rightarrow \mathcal{P}(V)$ is a valuation function such that $\llbracket p \rrbracket = \{w_1, w_2, w_4\}$.

The evidence for rain is given by $\{w_2, w_4\}$; evidence for sunny weather is given by $\{w_1, w_2\}$.

We evaluate at w_2 : observe that $(\mathfrak{M}, w_2) \models K_a p \wedge K_b p$. For agent a , the correct justification is given by the open $\{w_2, w_4\}$; for agent b , it is the open $\{w_1, w_2\}$. However, $\{w_3\} \in \tau_A$. Since $\{w_3\} = X \setminus \llbracket p \rrbracket$, it cannot intersect with any evidence for p . Therefore, no evidence for p can be dense with respect to the join topology: $(\mathfrak{M}, w_2) \not\models K_{\{a,b\}} p$.

In this example, Alice and Bob make a rational decision to drop their beliefs: together they have combined evidence supporting the proposition that class will not be cancelled. Furthermore, both agents are dropping a false belief: had Alice's app not falsely predicted that there would be no sun, or had Bob's app not falsely predicted that there would be no rain, then they would have had no reason to drop their beliefs.

In both of the previous examples, one agent lost their knowledge of a proposition. The following scenario illustrates how a group of agents can collectively believe the wrong proposition, and still learn the truth through communication.




Figure 3.4: From left to right: Alice's topology, Bob's topology, and the join topology. For clarity, we omit the partition cells from the diagram, since the partition for each agent is given by $\{X\}$. The actual world is w_2 . The new evidence is coloured gray.

Alarm. Charles and Daisy have different jobs at the same factory and they both individually know that the two of them are working late today. It is a company rule that the last person to leave the office should activate the alarm. Charles is packing up when he hears the warning beep: Daisy just set the alarm without checking his office. He leaves through the closest exit and, hurrying to be in time before the alarm activates, forgets to turn off the lights. Little did he know, that the beep he heard was the sound of Daisy turning off the machines. When Daisy leaves a few minutes later, she sees the light on in Charles' office and concludes that she is not, in fact, the last person to leave: Charles is. Now Charles and Daisy both individually know infallibly that they did not activate the alarm themselves, and they both have evidence convincing them that the alarm was set (by their co-worker). Later that night, Charles calls Daisy to remind her to check the offices before setting the alarm next time, and they both learn that the alarm has not been activated.

We formalize this scenario in Example 3.28 (depicted by Figure 3.5). It shows that the axiom $\bigwedge_{i \in I} B_i p \rightarrow \neg K_I \neg p$ fails. It is depicted in Figure 3.5 and it is a variation on Example 3.26, with different partitions.

Example 3.28. Let $V := \{p\}$ be our vocabulary. Let $\mathfrak{M} := (X, \mathcal{E}_c^0, \mathcal{E}_d^0, \tau_c, \tau_d, \Pi_c, \Pi_d, \pi)$ be a two-agent topo-e-model where $X = \{x, y, z\}$, and

$$\begin{aligned} \Pi_c &= \{\{x, y\}, \{z\}\} & \mathcal{E}_c^0 &= \{\{x\}, \{x, y\}, \{z\}, X\} \\ \Pi_d &= \{\{x\}, \{y, z\}\} & \mathcal{E}_d^0 &= \{\{x\}, \{y, z\}, \{z\}, X\} \end{aligned}$$

and $\pi : X \rightarrow \mathcal{P}(V)$ is a valuation function such that $\llbracket p \rrbracket = \{x, z\}$.

Charles and Daisy are represented by c and d , respectively. The variable p denotes the sentence "the alarm has been set". At x , the alarm has been set by Daisy; Daisy knows this is not the actual state. In state z , it was set by Charles; Charles knows that this is not the actual state. The actual state is y .

We evaluate at y , thus evaluating the beliefs of agent c with respect to information cell $\{x, y\} \in \Pi_c$ and those of agent d with respect to information cell $\{y, z\} \in \Pi_d$. We see that $(\mathfrak{M}, y) \models B_c p \wedge B_d p$. Analogous to Example 1, the evidence for agents c and d is given by $\{x\}$ and $\{z\}$, respectively.

In the join topology, we evaluate with respect to the group partition $\Pi_c(y) \cap \Pi_d(y) = \{y\}$. Since $(\mathfrak{M}, y) \models \neg p$, the group A has implicit infallible knowledge that $\neg p$ is true. In other words, $(\mathfrak{M}, y) \models K_{\{c, d\}} \neg p$.




Figure 3.5: From left to right: Charles' topology, the join topology for Charles and Daisy, and Daisy's topology. The actual world is y . We draw hard evidence using solid lines, and soft evidence using dashed lines. The new partition cell is coloured gray.

Each of these examples shows why a realistic depiction of fallible knowledge and belief revision, in the context of communication-based group knowledge, *demands* the failure of the (KD) axiom: agents should be able to reconsider their beliefs in light of new evidence.

3.5 Alexandroff Multi-Agent Models: Relational Representation

In Section 2.1.4 we established a correspondence between Alexandroff single-agent topological models and relational Kripke models. In this section, we naturally extend this correspondence to the multi-agent case. It is useful for multiple reasons to define a relational counterpart to multi-agent topo-e-models. First of all, although a relational model does not preserve the intuitive representation of evidence sets that the topological model has, the relational model is in a sense less complicated: for instance, the evidence set of an agent, given by a subset of the power set of all states, is reduced to one single pre-order over the set of states. This insight also aids our symbolic implementation, and its correctness proof, which we further elaborate on in Chapter 6. Secondly, this correspondence allows us to prove completeness of our logic via the relational models and apply established modal completeness results such as the construction of the canonical model (see e.g. [BRV01]) in Chapter 5.

In Section 3.5.1 we define the alternative relational representation of Alexandroff multi-agent topo-e-models, which we refer to as *relational evidence models*. In Sections 3.5.2 and 3.5.3, we present maps between the two types and prove that these maps preserve truth with respect to formulas over the language $\Box[\forall]_I$. We conclude that any logic over $\Box[\forall]_I$, which is sound and complete with respect to multi-agent topo-e-models, is also sound and complete with respect to relational evidence models, and vice versa.

3.5.1 Relational Evidence Models

For the remainder of the section, fix a finite set of agents A and a countable vocabulary V . We work with the following models: Alexandroff multi-agent topo-e-models (Definition 3.29) and relational evidence models (Definition 3.30).

Definition 3.29 (Alexandroff Multi-Agent Topo-E-Model). A multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ is *Alexandroff* if for all $i \in A$, τ_i is closed under arbitrary intersections, i.e., $\bigcap \mathcal{C} \in \tau_i$ for any $\mathcal{C} \subseteq \tau_i$.

Definition 3.30 (Relational Evidence Model). A *relational evidence model* is a structure $\mathbf{X} = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ where X is a set of states and for each agent $i \in A$, $\leq_i \subseteq X \times X$ is a preorder,

$\sim_i \subseteq X \times X$ is an equivalence relation, and we have $\leq_i \subseteq \sim_i$. The valuation function is given by $\pi : S \rightarrow \mathcal{P}(V)$.

Definition 3.31. We define group relations on the model as abbreviations: for $I \subseteq A$, define $\leq_I := \bigcap_{i \in I} \leq_i$ and $\sim_I := \bigcap_{i \in I} \sim_i$.

We interpret the \Box_I and $[\forall]_I$ operators in terms of the respective relations \leq_I and \sim_I , respectively.

Definition 3.32 (Relational Semantics of $\Box[\forall]_I$). Given a countable vocabulary V , the semantics of $\Box[\forall]_I$ on relational evidence models is defined as follows: given a relational evidence model $\mathbf{X} = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ and a state $x \in X$, let

$(\mathbf{X}, x) \models \top$	always holds
$(\mathbf{X}, x) \models p$	iff $p \in \pi(x)$
$(\mathbf{X}, x) \models \neg \varphi$	iff $(\mathbf{X}, x) \not\models \varphi$
$(\mathbf{X}, x) \models \varphi \wedge \psi$	iff $(\mathbf{X}, x) \models \varphi$ and $(\mathbf{X}, x) \models \psi$
$(\mathbf{X}, x) \models \Box_I \varphi$	iff for all $y \in X$ s.t. $x \leq_I y : (\mathbf{X}, y) \models \varphi$
$(\mathbf{X}, x) \models [\forall]_I \varphi$	iff for all $y \in X$ s.t. $x \sim_I y : (\mathbf{X}, y) \models \varphi$

where $p \in V$ is any proposition, $I \subseteq A$ is any nonempty subgroup of agents, and \leq_I and \sim_I are the abbreviations from Definition 3.31.

Expressing knowledge and belief as abbreviations in the language $\Box[\forall]_I$ (see Proposition 3.15), we obtain that belief in φ holds if and only if φ is true at all worlds within the information cell, that are *plausible enough*: unfolding the abbreviation, we have for all $s \in X$ that $(\mathbf{X}, s) \models B_I \varphi$ if and only if

$$\forall t \in X : s \sim_I t \Rightarrow (\exists u \in X (t \leq_I u \text{ and } \forall v \in X : u \leq_I v \Rightarrow (\mathbf{X}, v) \models \varphi)).$$

3.5.2 Alexandroff Topo-E-Models as Relational Evidence Models

Definition 3.33 describes a map from Alexandroff multi-agent topo-e-models to relational evidence models¹⁰. It is a generalisation of the map for single-agent models described in Definition 2.21.

Definition 3.33. Given an Alexandroff multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, we construct a relational evidence model $Rel(\mathfrak{M}) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ as follows: for each $i \in A$ and any $x, y \in X$, define

1. $x \leq_i y$ if and only if $\Pi_i(x) = \Pi_i(y)$ and $x \sqsubseteq_{\tau_i} y$, where \sqsubseteq_{τ_i} is the specialisation pre-order for τ_i (see Definition 2.20);
2. $x \sim_i y$ if and only if $\Pi_i(x) = \Pi_i(y)$.

¹⁰If the model resulting from the map from Definition 3.33 is finite (or *maximal-dense*, see Definition 5.40), then it is an *epistemic plausibility model*, as defined in [BS16]. These models do not generally coincide. An epistemic plausibility model is a relational evidence model with the following property: for each $i \in A$, the connected components of the relation \leq_i are *well-pre-orders*. A relation R is a well-pre-order if it is a pre-order and it has at least one R -maximal (see Definition 5.34) element. For finite models, this condition is clearly satisfied.

To see that not all infinite models satisfy the condition, even if they are equivalent to an Alexandroff topo-e-model, consider the topology (τ, \mathbb{N}) on the natural numbers, where τ is generated by the subbasis $\mathcal{E}_0 = \{[n, \infty) \cap \mathbb{N} \mid n \in \mathbb{N}\}$. The topology is Alexandroff: it is closed under arbitrary intersections. We also have an infinite chain $0 \sqsubseteq_{\tau} 1 \sqsubseteq_{\tau} 2 \dots$, with $n+1 \not\sqsubseteq_{\tau} n$ for all $n \in \mathbb{N}$. Now consider a multi-agent topo-e-model for agents A over \mathbb{N} such that $\Pi_i = \{\mathbb{N}\}$ and $\tau_i = \tau$ for some $i \in A$. In the equivalent relational evidence model, \leq_i forms one single connected component, structured as an infinitely increasing chain $0 \leq_i 1 \leq_i 2 \leq_i \dots$ where for each $n \in \mathbb{N}$, we have $n+1 \not\leq_i n$. Thus, \leq_i is not a well-pre-order.

Observation 3.34. With respect to the abbreviations \leq_I and \sim_I of the group relations, it can be checked by spelling out Definition 3.8 of Π_I that we have $x \leq_I y$ if and only if $\Pi_I(x) = \Pi_I(y)$ and $x \sqsubseteq_{\tau_i} y$; and we have $x \sim_I y$ if and only if $\Pi_I(x) = \Pi_I(y)$.

We show that the map is correct (Proposition 3.35) and preserves truth (Theorem 3.36).

Proposition 3.35. *For any Alexandroff multi-agent topo-e-model, there exists a relational evidence model.*

Proof. Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be an Alexandroff multi-agent topo-e-model. We show that $Rel(\mathfrak{M}) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$, constructed according to Definition 3.33, is a relational evidence model. We check the conditions from Definition 3.30. Let $i \in A$.

The relation \leq_i is a pre-order. For reflexivity, we have $x \leq_i x$ if and only if $\Pi_i(x) = \Pi_i(x)$; $x \sqsubseteq_{\tau_i} x$ follows from Definition 2.20 of the specialisation pre-order. To see that \leq_i is transitive, let $x \leq_i y \leq_i z$, i.e. let $\Pi_i(x) = \Pi_i(y) = \Pi_i(z)$ and $x \sqsubseteq_{\tau_i} y \sqsubseteq_{\tau_i} z$. Then $\Pi_i(x) = \Pi_i(z)$ and $x \sqsubseteq_{\tau_i} z$, so $x \leq_i z$.

The relation \sim_i is an equivalence relation. This follows directly from the properties of a partition.

Inclusion is satisfied, i.e. $\leq_i \subseteq \sim_i$. Suppose $x \leq_i y$. Then by Definition 3.33, $\Pi_i(x) = \Pi_i(y)$ and therefore, $x \sim_i y$, as required. \square

Theorem 3.36. *The map $\mathfrak{M} \mapsto Rel(\mathfrak{M})$ from Definition 3.33 preserves truth: fix a vocabulary V ; then for any pointed Alexandroff multi-agent topo-e-model (\mathfrak{M}, x) and every formula $\varphi \in \Box[\forall]_I(V)$, we have*

$$(\mathfrak{M}, x) \models \varphi \quad \text{iff} \quad (Rel(\mathfrak{M}), x) \models \varphi.$$

Proof. By induction on the complexity of φ . Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a topo-e-model and let $Rel(\mathfrak{M}) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ be its relational correspondent. The base case of atomic propositions, and the boolean cases of the induction step, are standard. We only show the proof of the modality \Box_I ; the proof of $[\forall]_I$ is similar and less complicated, as the definition of the \sim_I relations only concerns the partition, whereas the \leq_I relations additionally involve the specialisation pre-order.

For the case where $\varphi = \Box_I \psi$, suppose for the left-to-right direction that $(\mathfrak{M}, x) \models \Box_I \psi$. Then there is $U_1 \in \tau_1, \dots, U_n \in \tau_n$ such that $U_1 \cap \dots \cap U_n = U \in \tau_I$ and $x \in U \subseteq \llbracket \psi \rrbracket$ (Definition 3.10). Now suppose for contradiction that $(Rel(\mathfrak{M}), x) \not\models \Box_I \psi$, i.e. suppose there is $y \in X$ such that $x \leq_I y$ but $(Rel(\mathfrak{M}), y) \not\models \psi$ (by Definition 3.32). Then by the induction hypothesis, $(\mathfrak{M}, y) \not\models \psi$. By definition of $x \leq_I y$, we have for all $i \in I$ that $x \leq_i y$ and thus, $x \sqsubseteq_{\tau_i} y$ (by Definition 3.33). So by Definition 2.20 of the specialisation pre-order we have for all $i \in I$, for all $U \in \tau_i$, that $x \in U$ implies $y \in U$. In particular, this means that $y \in (U_1 \cap \dots \cap U_n) = U \subseteq \llbracket \psi \rrbracket$ and therefore $(\mathfrak{M}, y) \models \psi$, giving us the desired contradiction. We conclude that $(\mathfrak{M}, x) \models \Box_I \psi$.

For the converse direction, suppose that $(Rel(\mathfrak{M}), x) \models \Box_I \psi$. Then we have for all $y \in X$ such that $x \leq_I y$, $(Rel(\mathfrak{M}), y) \models \psi$ (Definition 3.32). Let such y be arbitrary. By the induction hypothesis, $(\mathfrak{M}, y) \models \psi$. Furthermore, by definition of \leq_I , we have that $x \leq_i y$ for all $i \in I$ and so, by definition of \leq_i (Definition 3.33), that $x \sqsubseteq_{\tau_i} y$ for all $i \in I$. Let such $i \in I$ be arbitrary. By Definition 2.20 of the specialisation pre-order, for all $U \in \tau_i$ and for all $y \in X$ such that $x \leq_i y$, $x \in U$ implies $y \in U$. Now the intersection of all open neighbourhoods of x in its information cell for i must be a subset of $\llbracket \psi \rrbracket$:

$$\begin{aligned} \bigcap \{U \in \tau_i \mid x \in U\} \cap \Pi_i(x) &= \{y \in \Pi_i(x) \mid \forall U \in \tau_i (x \in U \Rightarrow y \in U)\} \\ &= \{y \in X \mid x \leq_i y\} \\ &\subseteq \llbracket \psi \rrbracket. \end{aligned} \tag{Def. 3.33}$$

Furthermore, because τ_i is, by assumption, Alexandroff, $\bigcap\{U \in \tau_i \mid x \in U\} \in \tau_i$. For $i \in I$, let $U_i := \bigcap\{U \in \tau_i \mid x \in U\}$. By Definition 3.7 of the join topology, the set $\bigcap_{i \in I} U_i$ is open in τ_I . Furthermore, $(\bigcap_{i \in I} U_i) \subseteq \llbracket \psi \rrbracket$. Because $x \in U_i$ for all $i \in I$, we have $x \in (\bigcap_{i \in I} U_i)$. But this gives us that $(\text{Rel}(\mathfrak{M}), x) \models \Box_I \psi$, as required. \square

3.5.3 Relational Evidence Models as Alexandroff Topo-E-Models

Going from relational evidence models to Alexandroff multi-agent topo-e-models, Definition 3.37 describes a map, which constructs the up-set topology for each relation on the given model. It is a multi-agent generalisation of Definition 2.18. Recall from Proposition 2.19 that the topology τ_i , generated by \mathcal{E}_i^0 , equals the up-set topology.

Definition 3.37. Given a relational evidence model $\mathbf{X} = (S, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$, we construct an Alexandroff multi-agent topo-e-model $\text{Top}(\mathbf{X}) = (S, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ as follows: for each $i \in A$, define

1. $\mathcal{E}_i^0 := \{\uparrow_{\leq_i} x \mid x \in X\}$, where $\uparrow_{\leq_i} x$ is the up-set of the singleton set $\{x\}$ with respect to the relation \leq_i ;
2. Let τ_i be the topology generated by \mathcal{E}_i^0 ;
3. $\Pi_i := S / \sim_i$, i.e. let Π_i be given by the quotient space of S by \sim_i .

We show that the map is correct (Proposition 3.38) and preserves truth (Theorem 3.39).

Proposition 3.38. *For any relational evidence model, there exists an Alexandroff multi-agent topo-e-model.*

Proof. Let $\mathbf{X} = (S, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ be a relational evidence model. We show that $\text{Top}(\mathbf{X}) = (S, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, constructed according to Definition 3.37, is a multi-agent topo-e-model (Definition 3.1) and is Alexandroff (Definition 3.29). We check the following three conditions. Let $i \in A$.

First, $\emptyset \notin \mathcal{E}_i^0$, by construction; and it is a subbasis for τ_i , by construction of τ_i .

Second, Π_i is a partition. This follows directly from the fact that \sim_i is an equivalence relation. The fact that $\Pi_i \subseteq \tau_i$ follows directly from the property of \mathbf{X} that $\leq_i \subseteq \sim_i$ (Definition 3.30): each $\pi_i \in \Pi_i$ is an equivalence class with respect to \sim_i , and by $\leq_i \subseteq \sim_i$, it is automatically an upset with respect to \leq_i .

Finally, the topology τ_i is Alexandroff: a space is Alexandroff if and only if every element of the space has a least open neighbourhood [BB07a]; in this case the least open neighbourhood of every $x \in X$ is given by $\uparrow_{\leq_i} x$. \square

Theorem 3.39. *The map $\mathbf{X} \mapsto \text{Top}(\mathbf{X})$ from Definition 3.37 preserves truth: for any pointed relational evidence model (\mathbf{X}, x) and every formula $\varphi \in \Box[\forall]_I(V)$, we have*

$$(\mathbf{X}, x) \models \varphi \quad \text{iff} \quad (\text{Top}(\mathbf{X}), x) \models \varphi.$$

Proof. By induction on the complexity of φ . Let $\mathbf{X} = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ be a relational evidence model and let $\text{Top}(\mathbf{X}) = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be its topological correspondent. The base case of atomic propositions and the boolean cases of the induction step are standard. So we focus on the cases involving modalities \Box_I and $[\forall]_I$.

For the case where $\varphi = \Box_I \psi$, suppose for the left-to-right direction that $(\mathbf{X}, x) \models \Box_I \psi$. Then, for all $y \in X$ such that $x \leq_I y$, we have $(\mathbf{X}, y) \models \psi$ (Definition 3.32). By the induction hypothesis, $(\text{Top}(\mathbf{X}), y) \models \psi$. So $\uparrow_{\leq_I} x = \{y \in X \mid x \leq_I y\} \subseteq \llbracket \psi \rrbracket$. The following equivalences show that $\uparrow_{\leq_I} x$ is equivalent to $\bigcap_{i \in I} \uparrow_{\leq_i} x$:

$$\begin{aligned} \uparrow_{\leq_I} x &= \{y \in X \mid x \leq_i y \text{ for all } i \in I\} \\ &= \bigcap_{i \in I} \{y \in X \mid x \leq_i y\} \\ &= \bigcap_{i \in I} \uparrow_{\leq_i} x. \end{aligned}$$

But $\bigcap_{i \in I} \uparrow_{\leq_i} x$ is open in the join topology τ_I . Thus, with $x \in (\bigcap_{i \in I} \uparrow_{\leq_i} x) \subseteq \llbracket \psi \rrbracket$, we can conclude that $(\text{Top}(\mathbf{X}), x) \models \Box_I \psi$ (Definition 3.10).

For the converse direction, suppose that $(\text{Top}(\mathbf{X}), x) \models \Box_I \psi$. Then for each $i \in I$ there is $U_i \in \tau_i$ such that $\bigcap_{i \in I} U_i = U \in \tau_I$ and $x \in U \subseteq \llbracket \psi \rrbracket$ (Definition 3.10). Now let $y \in X$ be arbitrary and suppose $x \leq_I y$. It remains to show that $(\mathbf{X}, y) \models \psi$. By $x \leq_I y$, we have for all $i \in I$ that $x \leq_i y$. So let $i \in I$ be arbitrary. Since U_i is an up-set (Proposition 2.19), we know that $y \in U_i$. Since i was arbitrary, we have $y \in \bigcap_{i \in I} U_i = U$. But then it follows from the fact that $U \subseteq \llbracket \psi \rrbracket$, that $(\text{Top}(\mathbf{X}), y) \models \psi$. By the induction hypothesis, $(\mathbf{X}, y) \models \psi$. Therefore, $(\mathbf{X}, x) \models \Box_I \psi$ (Definition 3.32).

For the case where $\varphi = [\forall]_I \psi$, suppose for the left-to-right direction that $(\mathbf{X}, x) \models [\forall]_I \psi$. Then for all $y \sim_I x$ we have $(\mathbf{X}, y) \models \psi$ (Definition 3.32). Now let $y \in \Pi_I(x)$ be arbitrary. By Definition 3.37 of the map, $y \sim_I x$, so automatically, by $(\mathbf{X}, y) \models \psi$ and the induction hypothesis, $(\text{Top}(\mathbf{X}), y) \models \psi$. But then $\Pi_I(x) \subseteq \llbracket \psi \rrbracket$, which gives us $(\text{Top}(\mathbf{X}), x) \models [\forall]_I \psi$ (Definition 3.10).

For the converse direction, suppose that $(\text{Top}(\mathbf{X}), x) \models [\forall]_I \psi$. Then $\Pi_I(x) \subseteq \llbracket \psi \rrbracket$ (Definition 3.10). So let $y \in X$ be arbitrary and suppose $x \sim_I y$. It suffices to show that $(\mathbf{X}, y) \models \psi$. But this follows directly from y being in $\Pi_I(x)$ (Definition 3.37), which gives us that $(\text{Top}(\mathbf{X}), y) \models \psi$. By the induction hypothesis, $(\mathbf{X}, y) \models \psi$ and therefore, $(\mathbf{X}, x) \models [\forall]_I \psi$. \square

Corollary 3.40. *For all Alexandroff multi-agent topo-e-models there exists an equivalent relational evidence model, and vice versa. Hence, any logic that is sound and complete with respect to the former class is also sound and complete with respect to the latter.*

Proof. The claim is an immediate consequence of Theorem 3.36 and Theorem 3.39. \square

Chapter 4

Axiomatisation

In this chapter, we axiomatise the language $KB_{i,A}$, as well as $\Box[\forall]_I$ and its fragment $\Box[\forall]_{i,A}$. Additionally, we present proof systems for their dynamic extensions. In Section 4.1, we first prove a number of validities for the given languages on multi-agent topo-e-models. Next, for each language, we provide a proof system containing a subset of the validities proved in Section 4.1. The proof system for the principal fragment $KB_{i,A}$ is presented in Definition 4.11. Additionally, we present a candidate proof system for the language KB_I , for which we could not prove completeness. As for the other languages, for the time being we leave soundness and completeness of the proof systems as a claim: the actual proofs will be presented in Chapter 5, which is dedicated to proving completeness.

4.1 Validities

In this section, we list and prove several validities in the language $\Box[\forall]_I$ on multi-agent topo-e-models. This directly gives us a proof of the restrictions of these validities to $\Box[\forall]_{i,A}$ and $KB_{i,A}$. We will make reference to these proofs of soundness in Section 4.2, where we introduce sound proof systems for the language $\Box[\forall]_I$ and its fragments $\Box[\forall]_{i,A}$ and $KB_{i,A}$.

Throughout the section, fix a countable vocabulary V and a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ with a finite set of agents $A := \{1, \dots, n\}$. Finally, for any modality M , let $\langle M \rangle$ denote its dual, i.e. let $\langle M \rangle \varphi := \neg M \neg \varphi$.

In the following proofs, whenever no confusion can arise, we sometimes simply say that a set is *locally dense* without mentioning the information cell in which it is locally dense.

We start by observing the following fact, which states that we can view the join topology τ_I of a group I as a fictitious, individual agent, whose knowledge is given by virtual group knowledge of I . We will use it in multiple proofs throughout this section.

Fact 4.1. For all nonempty $I \subseteq A$, the multi-agent semantics for formulas of the form $\Box_I \varphi$, $[\forall]_I \varphi$, $K_I \varphi$, and $B_I \varphi$, evaluated on the subspace topology $(\Pi_I(x), \tau_I|_{\Pi_I(x)})$, coincides with the single-agent semantics when we consider I as a single agent. This has the following consequence: let $\varphi \in \mathcal{L}_{\Box[\forall]}$ and, given a group of agents I , let the multi-agent variant $\varphi_I \in \Box[\forall]_I$ of φ be obtained by replacing every occurrence of \Box , $[\forall]$, K , and B in φ with \Box_I , $[\forall]_I$, K_I , and B_I respectively. Then, if φ is valid on topo-e-models, so is φ_I .

For example, consider the modality for group belief. Given a formula of the form $B_I \varphi \in \Box[\forall]_I$ and a pointed multi-agent topo-e-model (\mathfrak{M}, x) , we evaluate whether $\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$ is locally dense in $\Pi_I(x)$.

To see the equivalence with single-agent semantics, let $X' := \Pi_I(x)$, let $\tau' := \tau_I|_{\Pi_I(x)}$ and let (X, τ') be the subspace topology; then $\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$ is locally dense in $\Pi_I(x)$ if and only if $\text{Int}_{\tau'}(\llbracket \varphi \rrbracket)$ is dense in (X', τ') .

In the remainder of this section, we work within the fixed language $\Box[\forall]_I$.

4.1.1 Single Groups

First, we prove the more obvious validities involving single groups $I \subseteq A$. On single-agent models, the logic S4 is sound for soft evidence and S5 is sound for hard evidence. Additionally, Stalnaker's axioms ([Sta06]) have been proved to be valid on single-agent topo-e-models [Özg17]. These proofs easily extend to the multi-agent case (Proposition 4.3).

Lemma 4.2. *For all nonempty $I \subseteq A$, (1) the S4 axioms and rules for \Box_I and (2) the S5 axioms and rules for $[\forall]_I$ are valid on all multi-agent topo-e-models.*

Proof. For (1), the claim follows immediately from the Kuratowski axioms on interior operators (Proposition 2.3) and the interpretation of the modality for group belief as the interior operator with respect to τ_I . For (2), the claim follows immediately from the fact that $[\forall]_I$ partitions the model into equivalence classes. □

Proposition 4.3. *For all nonempty $I \subseteq A$, Stalnaker's axioms for K_I and B_I are valid on multi-agent topo-e-models.*

Proof. The proof is straightforward and can be found in Appendix A.3. □

4.1.2 Groups and Subgroups

In this subsection, we prove the validity of a number of new formulas. With the exception of the (Inclusion) validity, which is defined for single groups I , all validities in this subsection relate supergroups to their subgroups. The proofs are presented in Proposition 4.5. Not all validities are equally intuitive. Therefore, we first discuss the properties defined by these validities. Table 4.1.1 presents a short description and an abbreviation for each of the validities that we prove later in this section. We discuss the intuitions behind these validities below.

The validities associated with evidence are straightforward: for both soft and hard evidence, we prove monotonicity on groups of agents with respect to inclusion. We call the respective axioms $(\Box - \text{Monotonicity})$ and $([\forall] - \text{Monotonicity})$. The symmetry between these axioms and the (KD) axiom for distributed knowledge is apparent: its generalisation to subgroups states that $K_J \rightarrow K_I$, for all nonempty subgroups $J \subseteq I \subseteq A$. This is no coincidence. As we briefly discussed in Section 2.3, evidence is *distributed*: both soft and hard evidence that is available to a subgroup, is also available to all of its supergroups. This property is formalised by the axioms $(\Box - \text{Monotonicity})$ and $([\forall] - \text{Monotonicity})$.

The (Inclusion) axiom relates hard evidence to soft evidence: every piece of hard evidence, possessed by a group, is also included in the soft evidence of that group. This axiom corresponds to the property of topo-e-models that for each agent $i \in A$, Π_i is open in τ_i .

Abbreviation	Validity	Description
	For all $\emptyset \neq J \subseteq I \subseteq A$ and $\varphi \in \Box[\forall]_I$:	
(\Box -Monotonicity)	$\Box_J \varphi \rightarrow \Box_I \varphi$	<i>Supergroups possess all soft evidence possessed by their subgroups.</i>
($[\forall]$ -Monotonicity)	$[\forall]_J \varphi \rightarrow [\forall]_I \varphi$	<i>Supergroups possess all hard evidence possessed by their subgroups.</i>
(Inclusion)	$[\forall]_I \varphi \rightarrow \Box_I \varphi$	<i>Hard evidence implies soft evidence.</i>
(KPB)	$B_J \varphi \rightarrow K_I B_J \varphi$	<i>Group Knowledge of Positive Subgroup Beliefs.</i>
(KNB)	$\neg B_J \varphi \rightarrow K_I \neg B_J \varphi$	<i>Group Knowledge of Negative Subgroup Beliefs.</i>
(BDK)	$\bigwedge_{\emptyset \neq J' \subseteq I} K_{J'} \varphi_{J'} \rightarrow \langle B_I \rangle (\bigwedge_{\emptyset \neq J' \subseteq I} \varphi_{J'})$	<i>Consistency of Group Belief with Distributed Knowledge.</i>
(KBK)	$(K_\varphi \wedge B_I \varphi) \rightarrow K_I \varphi$	<i>Subgroup Knowledge and Group Belief imply Group Knowledge.</i>

Table 4.1.1: A list of validities with corresponding abbreviations and descriptions given a group of agents A . In the validity corresponding to (BDK), $\{\varphi_{J'} \mid \emptyset \neq J' \subseteq I\}$ are arbitrary formulas over the language $\Box[\forall]_I$; one for each nonempty subgroup of I .

The remaining validities are more involved and therefore deserve a more detailed explanation. Whenever we say that "the group knows φ " or "the group implicitly knows φ ", we refer to knowledge in the sense of virtual group knowledge.

The validities *Group Knowledge of Positive Subgroup Beliefs* (KPB) and *Group Knowledge of Negative Subgroup Beliefs* (KNB) go hand in hand. These formulas state that the group I implicitly knows the beliefs of its subgroups: for any proposition $\varphi \in \Box[\forall]_I$ and for any (nonempty) subgroup $J \subseteq I$, I knows whether J believes φ . The reader might notice the similarity with Stalnaker's Strong Introspection axioms (Strong PI) and (Strong NI). Compared to Stalnaker's axioms, the validities (KPB) and (KNB) represent an even stronger form of introspection, which generalises across supergroups; one might refer to it as positive and negative *Super-Introspection*.

The validity *Consistency of Group Belief with Distributed Knowledge* (BDK) states that I cannot believe the negation of any conjunction of the propositions known by its subgroups. In Section 3.4 we found that it is possible for I to believe or know the negation of the beliefs of its subgroups. Due to the validity of (BDK), this does not occur when the beliefs are *factive*. A different intuition of the axiom is that it describes a consistency between virtual group belief and traditional distributed knowledge: let $J \subseteq I \subseteq A$ be nonempty subgroups. Then, under the traditional interpretation of distributed knowledge, followed by Ramírez [Ram15] and Fernández [Fer18], the conjunction $\bigwedge_{\emptyset \neq J' \subseteq I} \varphi_{J'}$ would be distributed knowledge of I . The (BDK) axiom states that the group I cannot implicitly believe the negation of a proposition that is distributedly known by I .

The validity *Subgroup Knowledge and Group Belief imply Group Knowledge* (KBK) could alternatively be referred to as *weak (group) monotonicity*, as it is, in a sense, a weakening of the group variant $K_J \rightarrow K_I$

(for all $I, J \subseteq A$) of the (KD) axiom. It specifies sufficient conditions for a group I to know a proposition φ that is known by one of its subgroups J : given that J knows φ , the supergroup I will not necessarily know φ . However, if I already believes φ , then I *does* know φ .

Proofs. We prove each of the validities in Proposition 4.5. This proposition makes use of the following fact:

Fact 4.4. Given a nonempty subgroup $I \subseteq A$, we can alternatively define τ_I and Π_I in terms of its nonempty subgroups $J \subseteq I$, instead of all $i \in I$. We have

$$\begin{aligned}\tau_I &= \tau_{\{i|i \in I\}} = \tau_{\{\emptyset \neq J | J \subseteq I\}} \\ \Pi_I &= \Pi_{\{i|i \in I\}} = \Pi_{\{J | \emptyset \neq J \subseteq I\}}.\end{aligned}$$

This is a direct consequence of the definitions of the join topology and group partition (Definition 3.7 and Definition 3.8). The commutative and associative properties of intersections allow us to view I as any combination of subgroups $J \subseteq I$ whose union is I , considering each J as a single agent.

Proposition 4.5 proves each of the discussed validities.

Proposition 4.5. *The formulas from Table 4.1.1 are valid on \mathfrak{M} .*

Proof. In the following proofs, let $I \subseteq A$ be nonempty and let $\varphi \in \Box[\forall]_I$ be any formula.

\Box -Monotonicity. Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and suppose $(\mathfrak{M}, x) \models \Box_J \varphi$. Then there is $U \in \tau_J$ such that $x \in U \subseteq \llbracket \varphi \rrbracket$. Note that $(U \cap \Pi_I(x)) \subseteq \llbracket \varphi \rrbracket$. It is clear that $x \in (U \cap \Pi_I(x))$. To see that $(U \cap \Pi_I(x)) \in \tau_I$, apply Fact 4.4 to consider τ_I as $\tau_J \vee \tau_{I \setminus J}$. Observe that $U \cap \Pi_I(x) = (U \cap \Pi_J(x)) \cap \Pi_{I \setminus J}(x)$, where $(U \cap \Pi_J(x)) \in \tau_J$ and $\Pi_{I \setminus J}(x) \in \tau_{I \setminus J}$ (as information cells are open). In other words, $(U \cap \Pi_I(x))$ is the intersection of an open in τ_J and an open in $\tau_{I \setminus J}$. By definition, the set is open in τ_I . So x is included in a τ_I -open subset of $\llbracket \varphi \rrbracket$, therefore, $(\mathfrak{M}, x) \models \Box_I \varphi$.

$[\forall]$ -Monotonicity. Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and suppose $(\mathfrak{M}, x) \models [\forall]_J \varphi$. Then $\Pi_J(x) \subseteq \llbracket \varphi \rrbracket$. Since $J \subseteq I$, we have $\Pi_I(x) \subseteq \Pi_J(x) \subseteq \llbracket \varphi \rrbracket$, i.e. $(\mathfrak{M}, x) \models [\forall]_I \varphi$.

Inclusion. This follows from the fact that we require $\Pi_i \subseteq \tau_i$ for all $i \in A$: for all $I \subseteq A$, we define $\Pi_I(x) = \bigcap_{i \in I} \Pi_i(x)$, which gives us that $\Pi_I(x)$ is open in τ_I .

Group Knowledge of Positive Subgroup Beliefs (KPB). Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and suppose that $(\mathfrak{M}, x) \models B_J \varphi$. By the semantics of B_J , this gives us $\Pi_J(x) \subseteq \llbracket B_J \varphi \rrbracket$. Since $J \subseteq I$, we have that $\Pi_I(x) \subseteq \Pi_J(x)$. Therefore, $\Pi_I(x) \subseteq \llbracket B_J \varphi \rrbracket$. Observe that $\Pi_I(x)$ is a τ_I -open subset of $\llbracket B_J \varphi \rrbracket$ which is locally dense in $\Pi_I(x)$ and contains x , we have $(\mathfrak{M}, x) \models K_I B_I \varphi$.

Group Knowledge of Negative Subgroup Beliefs (KNB). Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and suppose that $(\mathfrak{M}, x) \models \neg B_J \varphi$. Then $\Pi_J(x) \subseteq \llbracket \neg B_J \varphi \rrbracket$. Recall that $\Pi_I(x) \subseteq \Pi_J(x)$, since we have $J \subseteq I$. Thus, $\Pi_I(x) \subseteq \llbracket \neg B_J \varphi \rrbracket$. Clearly, $\Pi_I(x)$ itself is locally dense in $\Pi_I(x)$, and we have $x \in \Pi_I(x)$ by definition. Therefore, $(\mathfrak{M}, x) \models K_I \neg B_I \varphi$.

Consistency of Group Belief with Distributed Knowledge (BDK). By contradiction. Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and let $\{\varphi_J \mid J \subseteq I\}$ be arbitrary formulas, one for each subgroup J . Suppose that $(\mathfrak{M}, x) \models \bigwedge_{\emptyset \neq J \subseteq I} K_J \varphi_J$. Then $x \in \bigcap_{J \subseteq I} \llbracket K_J \varphi_J \rrbracket$. So for each $J \subseteq I$, $\text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket)$ is locally dense in $\Pi_J(x)$ and contains x : we have $x \in \bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket)$.

Now consider the restriction of $\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket)$ to $\Pi_I(x)$ and call it U ; we know that $x \in U = (\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket)) \cap \Pi_I(x)$. We claim that U is also open in τ_I .

Claim 4.6. $U \in \tau_I$.

Proof of claim. Recall from Fact 4.4 that we can define τ_I as the join of all τ_J of nonempty $J \subseteq I$. So to prove that $U \in \tau_I$, it suffices to show that there exist opens U_J for each $J \subseteq I$ such that U can be written as the intersection $\bigcap_{J \subseteq I} U_J$. Define for all $J \subseteq I$:

$$U_J := \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \cap \Pi_J(x) \quad (*)$$

Each U_J is open in τ_J , as U_J is an intersection of the interior of a set with an information cell; both are open by definition.

Applying the definition of Π_I as $\Pi_{\{J \mid J \subseteq I\}}$ (**), we obtain the following equivalences:

$$\begin{aligned} U &= \left(\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \right) \cap \Pi_I(x) \\ &= \left(\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \right) \cap \left(\bigcap_{J \subseteq I} \Pi_J(x) \right) \quad (**) \\ &= \bigcap_{J \subseteq I} (\text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \cap \Pi_J(x)) \\ &= \bigcap_{J \subseteq I} U_J \quad (*) \end{aligned}$$

which concludes our claim. \square

Now, for contradiction, suppose that $(\mathfrak{M}, x) \not\models \langle B_I \rangle (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J)$. That is, $(\mathfrak{M}, x) \models B_I \neg (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J)$. So $\text{Int}_{\tau_I} \llbracket \neg (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J) \rrbracket$ is locally dense in $\Pi_I(x)$. By the Kuratowski axioms, $\llbracket \neg (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J) \rrbracket = X \setminus \bigcap_{J \subseteq I} \llbracket \varphi_J \rrbracket$ is locally dense in $\Pi_I(x)$. Since it intersects with all opens in τ_I , it intersects with U :

$$\left(\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \right) \cap \Pi_I(x) \cap \left(X \setminus \bigcap_{J \subseteq I} \llbracket \varphi_J \rrbracket \right) \neq \emptyset.$$

But recall that $\text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \subseteq \llbracket \varphi_J \rrbracket$ for all $J \subseteq I$ (Proposition 2.3). So in particular, we have

$$\bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \cap \left(X \setminus \bigcap_{J \subseteq I} \text{Int}_{\tau_J}(\llbracket \varphi_J \rrbracket) \right) \neq \emptyset$$

which is a clear contradiction; we conclude that $(\mathfrak{M}, x) \models \langle B_I \rangle (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J)$, after all.

Subgroup Knowledge and Group Belief imply Group Knowledge (KBK). Let $x \in X$, let $J \subseteq I \subseteq A$ be nonempty, and suppose that $(\mathfrak{M}, x) \models (K_J \varphi \wedge B_I \varphi)$. Since $(\mathfrak{M}, x) \models K_I \varphi$ if and only if $(\mathfrak{M}, x) \models \Box_I \varphi \wedge B_I \varphi$ and we assumed $(\mathfrak{M}, x) \models B_I \varphi$, it remains to show that $(\mathfrak{M}, x) \models \Box_I \varphi$. But this follows from the fact that $x \in \llbracket K_J \varphi \rrbracket$: this gives us that $x \in \text{Int}_{\tau_J} \llbracket \varphi \rrbracket$; then $\text{Int}_{\tau_J} \llbracket \varphi \rrbracket \cap \Pi_I(x)$ is open in τ_I : using Fact 4.4, we consider τ_I as the join topology of τ_J and $\tau_{I \setminus J}$. Then $\text{Int}_{\tau_J} \llbracket \varphi \rrbracket \cap \Pi_I(x)$ can be written as an intersection of the τ_J -open set $(\text{Int}_{\tau_J} \llbracket \varphi \rrbracket \cap \Pi_J(x))$ with the information cells $\bigcap_{i \in (I \setminus \{j\})} \Pi_i(x)$ of the other agents. Finally, $\text{Int}_{\tau_J} \llbracket \varphi \rrbracket \cap \Pi_I(x)$ is clearly a subset of $\llbracket \varphi \rrbracket$, and it contains x . Therefore, $(\mathfrak{M}, x) \models \Box_I \varphi$, as required. \square

The following validity was not included in any axiomatisation. Nevertheless, we mention it, because it states the following interesting fact on multi-agent topo-e-models, for any nonempty

subgroup $I \subseteq A$ of agents: if there is any agent $i \in I$ for which I does not have any evidence that i does not know φ , then I believes φ .

We use the following abbreviation: define the global analogue of \Box as $E := \exists \Box \varphi$, where \exists is the dual of $[\forall]$. The sentence $E\varphi$ is interpreted as the existence of combined evidence for φ , that is not necessarily factive at the actual world. We define the group variant of the modality E as follows: given any nonempty $I \subseteq A$, let

$$E_I \varphi := \exists_I \Box_I \varphi$$

where \exists_I denotes the dual of the modality of infallible group knowledge.

Proposition 4.7. *Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be a multi-agent topo-e-model. For all $i \in I \subseteq A$, we have*

$$\mathfrak{M} \models \langle E_I \rangle K_i \varphi \rightarrow B_I \varphi.$$

Proof. Suppose $(\mathfrak{M}, x) \models \langle E_I \rangle K_i \varphi$, i.e. let $(\mathfrak{M}, x) \models \neg E_I \neg K_i \varphi$. Then there exists no open $U \in \tau_I|_{\Pi_I(x)}$ such that $U \subseteq \llbracket \neg K_i \varphi \rrbracket$. That is, all $U \in \tau_I|_{\Pi_I(x)}$ have nonempty intersection with $X \setminus \llbracket \neg K_i \varphi \rrbracket = \llbracket K_i \varphi \rrbracket$. So the set $\llbracket K_i \varphi \rrbracket$ is locally dense in $\Pi_I(x)$. Now, recall that $\llbracket K_i \varphi \rrbracket = \text{Int}_{\tau_i}(\llbracket \varphi \rrbracket)$ is open in τ_i by definition, hence $\llbracket K_i \varphi \rrbracket \cap \Pi_I(x)$ is open in τ_I : it can be written as the intersection of $\llbracket K_i \varphi \rrbracket \cap \Pi_i(x)$ (open in τ_i), with $\bigcap_{j \in (A \setminus \{i\})} \Pi_j(x)$ (with each $\Pi_j(x)$ open in τ_j). Furthermore, since $\llbracket K_i \varphi \rrbracket$ and $\Pi_I(x)$ are both locally dense in $\Pi_I(x)$, so is their intersection¹. Now, since it is obvious that $(\llbracket K_i \varphi \rrbracket \cap \Pi_i(x)) \subseteq \llbracket \varphi \rrbracket$, it is a subset of $\llbracket \varphi \rrbracket$ that is locally dense in $\Pi_I(x)$. We have $(\mathfrak{M}, x) \models B_I \varphi$, as required. \square

4.2 Proof Systems of the Static and Dynamic Languages

In this section, we present the principal proof system in this thesis: $KB_{i,A}$. Additionally, we provide sound and complete proof systems for the static and dynamic variant of the larger language $\Box[\forall]_I$, and its fragment $\Box[\forall]_{i,A}$, using the validities from Section 4.1. For now, we state that the proof systems are sound and complete without providing a proof. The proofs can be found in Chapter 5; the completeness proofs are more involved, and are therefore detailed in a separate chapter.

Section 4.2.1 concerns the proof systems $\Box[\forall]_I$ for $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$ for $\Box[\forall]_{i,A}$; Section 4.2.2 concerns the proof system $KB_{i,A}$ for $KB_{i,A}$; and in Section 4.2.3 we present a candidate proof system KB_I for KB_I . We conclude with proof systems for the dynamic extensions of $\Box[\forall]_I$, $\Box[\forall]_{i,A}$, and $KB_{i,A}$ (Section 4.2.4). Throughout this section, fix a finite set of agents A and a finite vocabulary V .

4.2.1 Axiomatisations of $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$

We first introduce our proof system $\Box[\forall]_I$ for the full language $\Box[\forall]_I$ and the restricted proof system $\Box[\forall]_{i,A}$ for the fragment $\Box[\forall]_{i,A}$.

Definition 4.8 (Proof Systems $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$). The proof system $\Box[\forall]_I$ is listed in Table 4.2.1.

The proof system $\Box[\forall]_{i,A}$ for the fragment $\Box[\forall]_{i,A}$ of the language $\Box[\forall]_I$ is obtained by restricting $\Box[\forall]_I$ to the full group A and subgroups consisting of individual agents (that is, groups of the form $\{i\}$ for $i \in A$).

¹The intersection of any two dense τ -open sets $U, V \subseteq X$ in any topological space (X, τ) is dense: let $W \in \tau$. Then $W \cap U \neq \emptyset$ and furthermore, being the intersection of two opens in τ , $W \cap U$ is open in τ . Thus, it also has nonempty intersection with the dense set V . But then $W \cap (U \cap V)$ is nonempty, as required.

(CPL)	Axioms and rules of classical propositional logic
Epistemic-Doxastic Axioms and Rules:	
For all nonempty $I \subseteq A$:	
(S4 $_{\Box}$)	All S4 axioms and rules for \Box_I
(S5 $_{[\forall]}$)	All S5 axioms and rules for $[\forall]_I$
Group Knowledge Axioms:	
For all nonempty $J \subseteq I \subseteq A$:	
(\Box -Monotonicity)	$\Box_J \varphi \rightarrow \Box_I \varphi$
($[\forall]$ -Monotonicity)	$[\forall]_J \varphi \rightarrow [\forall]_I \varphi$
(Inclusion)	$[\forall]_I \varphi \rightarrow \Box_I \varphi$

Table 4.2.1: The proof system $\Box[\forall]_I$, where A represents the group of all agents.

We remind the reader of the descriptions corresponding to the axioms associated with groups and their subgroups:

- (\Box – Monotonicity) *Supergroups possess all soft evidence possessed by their subgroups.*
- ($[\forall]$ – Monotonicity) *Supergroups possess all hard evidence possessed by their subgroups.*
- (Inclusion) *Hard evidence implies soft evidence.*

We do not prove yet that $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$ are sound and weakly complete with respect to topo-e-models; the proofs of the following theorems follow in Section 5.1 on completeness.

Theorem 4.9. *The proof system $\Box[\forall]_I$ from Definition 4.8 is sound and weakly complete with respect to multi-agent topo-e-models. Furthermore, the logic of $\Box[\forall]_I$ is decidable.*

Proof. The proof can be found at the end of Section 5.1. □

Theorem 4.10. *The proof system $\Box[\forall]_{i,A}$ from Definition 4.8 is sound and weakly complete with respect to multi-agent topo-e-models. Furthermore, the logic of $\Box[\forall]_{i,A}$ is decidable.*

Proof. The proof can be found at the end of Section 5.1. □

4.2.2 Axiomatisation of $KB_{i,A}$

In this section, we present the principal axiomatisation in this thesis: the proof system $KB_{i,A}$ for $KB_{i,A}$. The group knowledge axioms are obtained by restricting the corresponding validities in $\Box[\forall]_I$ (with knowledge and belief given by abbreviations) to $KB_{i,A}$.

Definition 4.11 (Proof System $KB_{i,A}$). The proof system $KB_{i,A}$ is listed in Table 4.2.2.

The axioms for the full group A mirror those for single agents $i \in A$. This analogy between individual agents and the full group also appears in multiple statements and definitions in the completeness proof of the axiomatisation, presented in Section 5.2. We use the notation α , which can denote singleton sets $\{i\} \subseteq A$, as well as A itself.

The belief-analogue of (K-Distributivity), the (B-Distributivity) axiom, states that

$$B_\alpha(\varphi \rightarrow \psi) \rightarrow (B_\alpha \varphi \rightarrow B_\alpha \psi)$$

(CPL)	Axioms and rules of classical propositional logic
Stalnaker's Epistemic-Doxastic Axioms:	
For all $\alpha \in \{A\} \cup A$:	
(K-Distributivity)	$K_\alpha(\varphi \rightarrow \psi) \rightarrow (K_\alpha\varphi \rightarrow K_\alpha\psi)$
(T)	$K_\alpha\varphi \rightarrow \varphi$
(KK)	$K_\alpha\varphi \rightarrow K_\alpha K_\alpha\varphi$
(CB)	$B_\alpha\varphi \rightarrow \neg B_\alpha\neg\varphi$
(Strong PI)	$B_\alpha\varphi \rightarrow K_\alpha B_\alpha\varphi$
(Strong NI)	$\neg B_\alpha\varphi \rightarrow K_\alpha\neg B_\alpha\varphi$
(KB)	$K_\alpha\varphi \rightarrow B_\alpha\varphi$
(FB)	$B_\alpha\varphi \rightarrow B_\alpha K_\alpha\varphi$
Inference Rules:	
For all $\alpha \in \{A\} \cup A$:	
(Modus Ponens)	From φ and $\varphi \rightarrow \psi$ infer ψ .
(K-Necessitation)	From φ infer $K_\alpha\varphi$.
(B-Necessitation)	From φ infer $B_\alpha\varphi$.
Group Knowledge Axioms:	
For all $i \in A$:	
(KBK)	$(K_i\varphi \wedge B_A\varphi) \rightarrow K_A\varphi$
(KPB)	$B_i\varphi \rightarrow K_A B_i\varphi$
(BDK)	$\bigwedge_{i \in A} K_i\varphi_i \rightarrow \langle B_A \rangle (\bigwedge_{i \in A} \varphi_i)$ (where $\{\varphi_i \mid i \in A\}$ are arb. formulas)

Table 4.2.2: The proof system $\mathbf{KB}_{i,A}$, where A represents the group of all agents.

and is provable from Stalnaker's axioms [Sta06].

We remind the reader of the descriptions corresponding to the validities associated with groups and their subgroups. Note that the axioms of $\mathbf{KB}_{i,A}$ restrict these to $\mathbf{KB}_{i,A}$, such that the only group that we consider is A , and the only subgroups that we consider are A and singleton groups, consisting of individual agents $i \in A$.

(KPB) *Group Knowledge of Positive Subgroup (Individual) Beliefs.*

(KNB) *Group Knowledge of Negative Subgroup (Individual) Beliefs* (provable² from (KPB)).

(BDK) *Consistency of Group Belief with Distributed Knowledge.*

(KBK) *Subgroup (Individual) Knowledge and Group Belief imply Group Knowledge.*

Again, we do not prove yet that $\mathbf{KB}_{i,A}$ is sound and weakly complete with respect to topo-e-models; the proof of the following theorem, which is not trivial, follows in Section 5.2 on completeness.

Theorem 4.12. *The proof system $\mathbf{KB}_{i,A}$ from Definition 4.11 for $\mathbf{KB}_{i,A}$ is sound and weakly complete with respect to multi-agent topo-e-models. Furthermore, the logic of $\mathbf{KB}_{i,A}$ is decidable.*

Proof. The proof can be found at the end of Section 5.2. □

²Let $i \in A$ and $\varphi \in \mathbf{KB}_{i,A}$ be arbitrary. Then by (Strong NI), we have that $\vdash \neg B_i\varphi \rightarrow K_i\neg B_i\varphi$. Applying (KB), we get that $\vdash \neg B_i\varphi \rightarrow B_i\neg B_i\varphi$. We use (KPB) to obtain $\vdash \neg B_i\varphi \rightarrow K_A B_i\neg B_i\varphi$. It remains to show that $\vdash B_i\neg B_i\varphi \rightarrow \neg B_i\varphi$. This is straightforward: since $\vdash B_i\varphi \rightarrow B_i B_i\varphi$, we have that $\vdash B_i\neg B_i\varphi \wedge B_i\varphi \rightarrow B_i\neg B_i\varphi \wedge B_i B_i\varphi$, contradicting the (CB) axiom. Thus, $\vdash B_i\neg B_i\varphi \wedge B_i\varphi \rightarrow \perp$, i.e. $\vdash B_i\neg B_i\varphi \rightarrow \neg B_i\varphi$, which concludes our proof.

4.2.3 Candidate Axiomatisation of KB_I

For the language fragment KB_I , we present a sound candidate proof system KB_I . We leave completeness as an open question. In contrast to $KB_{i,A}$, the group knowledge axioms for KB_I are given by the validities for all subgroups. Soundness and decidability are proved in Proposition 4.14. We conjecture that KB_I is complete with respect to multi-agent topo-e-models.

Definition 4.13 (Candidate Proof System KB_I). The candidate proof system KB_I is listed in Table 4.2.3.

(CPL)	Axioms and rules of classical propositional logic
Stalnaker's Epistemic-Doxastic Axioms:	
For all nonempty $I \subseteq A$:	
(K)	$K_I(\varphi \rightarrow \psi) \rightarrow (K_I\varphi \rightarrow K_I\psi)$
(T)	$K_I\varphi \rightarrow \varphi$
(KK)	$K_I\varphi \rightarrow K_IK_I\varphi$
(CB)	$B_I\varphi \rightarrow \neg B_I\neg\varphi$
(Strong PI)	$B_I\varphi \rightarrow K_IB_I\varphi$
(Strong NI)	$\neg B_I\varphi \rightarrow K_I\neg B_I\varphi$
(KB)	$K_I\varphi \rightarrow B_I\varphi$
(FB)	$B_I\varphi \rightarrow B_IK_I\varphi$
Inference Rules:	
For all nonempty $I \subseteq A$:	
(Modus Ponens)	From φ and $\varphi \rightarrow \psi$ infer ψ .
(K-Necessitation)	From φ infer $K_I\varphi$.
(B-Necessitation)	From φ infer $B_I\varphi$.
Group Knowledge Axioms:	
For all nonempty $J \subseteq I \subseteq A$:	
(KBK)	$(K_J\varphi \wedge B_I\varphi) \rightarrow K_I\varphi$
(KPB)	$B_J\varphi \rightarrow K_IB_J\varphi$
(BDK)	$\bigwedge_{\emptyset \neq J \subseteq I} K_J\varphi \rightarrow \langle B_I \rangle (\bigwedge_{\emptyset \neq J \subseteq I} \varphi_J)$ (where $\{\varphi_J \mid \emptyset \neq J \subseteq I\}$ are arb. formulas)

Table 4.2.3: The candidate proof system KB_I , where $H \subseteq J \subseteq I$.

Proposition 4.14. *The axioms and rules in the proof system (Table 4.2.3) for KB_I are sound with respect to multi-agent topo-e-models. Furthermore, the logic axiomatised by KB_I is decidable.*

Proof. The soundness of the axioms and rules of classical propositional logic is a routine check, as well as the soundness of the inference rules (Modus Ponens, and Necessitation for both \Box_I and $[\forall]_I$ for all $I \subseteq A$); therefore, we omit these proofs. Soundness of Stalnaker's epistemic-doxastic axioms follows from Lemma 4.2. As for the group knowledge axioms, soundness was proved in Proposition 4.5.

Furthermore, the logic of KB_I is decidable: in Theorem 4.9, we prove decidability for the logic of the larger language $\Box[\forall]_I$. This automatically implies decidability for all fragments, and we established in Corollary 3.16 that we can consider KB_I as a fragment of $\Box[\forall]_I$. \square

4.2.4 Axiomatisations of Dynamic Extensions

Finally, we present proof systems for the dynamic extensions of the languages $\Box[\forall]_I$, $\Box[\forall]_{i,A}$, and $KB_{i,A}$. For each language, we extend the proof system with a set of reduction axioms. In Section 5.3 we will prove soundness and completeness for the resulting proof systems, by showing that each of the extended languages is equally expressive as the corresponding static language.

In the proof systems $\Box[\forall]_I + [\text{share}_I]$ and $\Box[\forall]_{i,A} + [\text{share}_A]$, the reduction axioms for the cases of soft and hard evidence are symmetrical to the reduction axioms for distributed knowledge defined in [BS20] for the *logic of resolution*.

Definition 4.15 (Proof Systems $\Box[\forall]_I + [\text{share}_I]$ and $\Box[\forall]_{i,A} + [\text{share}_A]$). The proof system $\Box[\forall]_I + [\text{share}_I]$ for the language $\Box[\forall]_I + [\text{share}_I]$ (Definition 3.18), is listed in Table 4.2.4.

The proof system $\Box[\forall]_{i,A} + [\text{share}_A]$ for the language fragment $\Box[\forall]_{i,A} + [\text{share}_A]$ (in Definition 3.18) is obtained by restricting the proof system $\Box[\forall]_I + [\text{share}_I]$ to $\Box[\forall]_{i,A}$.

$(\Box[\forall]_I)$	Axioms and rules of $\Box[\forall]_I$	
$([\text{share}_I])$	Axioms and rules of normal logic for $[\text{share}_I]$	for all nonempty $I \subseteq A$
Reduction Axioms:		
For all nonempty $I, J \subseteq A$:		
(1)	$[\text{share}_I]p \leftrightarrow p$	for atomic propositions p
(2)	$[\text{share}_I]\neg\varphi \leftrightarrow \neg[\text{share}_I]\varphi$	
(3)	$[\text{share}_I]\varphi \wedge \psi \leftrightarrow [\text{share}_I]\varphi \wedge [\text{share}_I]\psi$	
(4)	$[\text{share}_I]\Box_J\varphi \leftrightarrow \Box_{J \cup I}[\text{share}_I]\varphi$	if $I \cap J \neq \emptyset$
(5)	$[\text{share}_I]\Box_J\varphi \leftrightarrow \Box_J[\text{share}_I]\varphi$	if $I \cap J = \emptyset$
(6)	$[\text{share}_I][\forall]_J\varphi \leftrightarrow [\forall]_{J \cup I}[\text{share}_I]\varphi$	if $I \cap J \neq \emptyset$
(7)	$[\text{share}_I][\forall]_J\varphi \leftrightarrow [\forall]_J[\text{share}_I]\varphi$	if $I \cap J = \emptyset$

Table 4.2.4: The proof system $\Box[\forall]_I + [\text{share}_I]$, where A represents the group of all agents.

For the proof system $KB_{i,A} + [\text{share}_A]$, we define reduction axioms for the modalities of knowledge and belief.

Definition 4.16. The proof system $KB_{i,A} + [\text{share}_A]$ for the language $KB_{i,A} + [\text{share}_A]$ (in Definition 3.18) is listed in Table 4.2.5.

Note that there are no cases where $i \notin A$, since i has to be in A .

It remains to show that the axiomatisations of the dynamic languages are sound and weakly complete with respect to topo-e-models. Again, we do not prove soundness and completeness yet; we state the theorems, which we prove in Section 5.3. We do prove Proposition 4.17, which we use in the proofs of soundness. In its proof, we show only the case for the knowledge modality, since the knowledge modality is the principal modality in this thesis, and the proofs for the other modalities are similar.

Proposition 4.17. Fix a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, a state $x \in X$, an agent $i \in A$, and any two nonempty subgroups $I, J \subseteq A$. Fix a countable vocabulary V and $\varphi \in \Box[\forall]_I + [\text{share}_I](V)$. Then, for any modality $M \in \{\Box, [\forall], B, K\}$, the following equivalences are true:

- (1) $(\mathfrak{M}^{\text{share}_I}, x) \models M_I\varphi \Leftrightarrow (\mathfrak{M}, x) \models M_{I \cup J}[\text{share}_I]\varphi$ if $I \cap J \neq \emptyset$
- (2) $(\mathfrak{M}^{\text{share}_I}, x) \models M_I\varphi \Leftrightarrow (\mathfrak{M}, x) \models M_I[\text{share}_I]\varphi$ if $I \cap J = \emptyset$.

$(KB_{i,A})$	Axioms and rules of $KB_{i,A}$
$([share_A])$	Axioms and rules of normal logic for $[share_A]$
Reduction Axioms:	
For all nonempty $i \in A$:	
(1)	$[share_A]p \leftrightarrow p$ for atomic propositions p
(2)	$[share_A]\neg\varphi \leftrightarrow \neg[share_A]\varphi$
(3)	$[share_A]\varphi \wedge \psi \leftrightarrow [share_A]\varphi \wedge [share_A]\psi$
(4)	$[share_A]K_i\varphi \leftrightarrow K_A[share_A]\varphi$
(5)	$[share_A]K_A\varphi \leftrightarrow K_A[share_A]\varphi$
(6)	$[share_A]B_i\varphi \leftrightarrow B_A[share_A]\varphi$
(7)	$[share_A]B_A\varphi \leftrightarrow B_A[share_A]\varphi$

Table 4.2.5: The proof system $KB_{i,A} + [share_A]$, where A represents the group of all agents.

Proof. We show the proof for the knowledge modality, that is, we prove (1_K) and (2_K) .

For (1_K) , suppose $I \cap J \neq \emptyset$. We distinguish two cases: either φ is atomic, or it is not. For the atomic case, let $p \in V$ be any propositional variable. We have $(\mathfrak{M}^{share_I}, x) \models K_I p$ if and only if $(\mathfrak{M}, x) \models K_{I \cup J} p$ (Proposition 3.22), if and only if $(\mathfrak{M}, x) \models K_{I \cup J} [share_J] p$ (Definition 3.19).

If φ is not atomic, we have the following chain of equivalences:

$$\begin{aligned}
(\mathfrak{M}^{share_I}, x) \models K_I \varphi & \text{ iff } x \in \llbracket K_I \varphi \rrbracket^{\mathfrak{M}^{share_I}} \\
& \text{ iff } x \in Int_{\tau_I^S}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}})) \quad (\text{Def. 3.19}) \\
& \text{ iff } x \in Int_{\tau_I^S}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}})) \quad (\text{Def. 3.19}) \\
& \text{ iff } x \in Int_{\tau_{I \cup J}}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_{I \cup J}(x) \subseteq Cl_{\tau_{I \cup J}}(Int_{\tau_{I \cup J}}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}})) \quad (\text{Lem. 3.21}) \\
& \text{ iff } x \in \llbracket K_{I \cup J} [share_J] \varphi \rrbracket^{\mathfrak{M}} \quad (\text{Def. 3.19}) \\
& \text{ iff } (\mathfrak{M}, x) \models K_{I \cup J} [share_J] \varphi.
\end{aligned}$$

For (2_K) , suppose $I \cap J = \emptyset$. Again, we distinguish two cases: either φ is atomic, or it is not. For the atomic case, the proof is identical to the proof of the atomic case for (1_K) .

If φ is not atomic, we have the following chain of equivalences. We use the fact that $I \cap J = \emptyset$ implies that for all $i \in I$ we have $\tau_i^S = \tau_i$ and $\Pi_i^S = \Pi_i$ (Definition 3.17), from which it directly follows that $\tau_I^S = \tau_I$ and $\Pi_I^S = \Pi_I$ (*).

$$\begin{aligned}
(\mathfrak{M}^{share_I}, x) \models K_I \varphi & \text{ iff } x \in \llbracket K_I \varphi \rrbracket^{\mathfrak{M}^{share_I}} \\
& \text{ iff } x \in Int_{\tau_I^S}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I^S(x) \subseteq Cl_{\tau_I^S}(Int_{\tau_I^S}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}})) \quad (\text{Def. 3.19}) \\
& \text{ iff } x \in Int_{\tau_I}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}}) \\
& \quad \text{and } \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket \varphi \rrbracket^{\mathfrak{M}^{share_I}})) \quad (*) \\
& \text{ iff } x \in Int_{\tau_I}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}}) \\
& \quad \text{and } \Pi_I(x) \subseteq Cl_{\tau_I}(Int_{\tau_I}(\llbracket [share_J] \varphi \rrbracket^{\mathfrak{M}})) \quad (\text{Def. 3.19}) \\
& \text{ iff } x \in \llbracket K_I [share_J] \varphi \rrbracket^{\mathfrak{M}} \quad (\text{Def. 3.19}) \\
& \text{ iff } (\mathfrak{M}, x) \models K_I [share_J] \varphi.
\end{aligned}$$

□

The proof of the following theorem follows in Section 5.3, which is dedicated to proving completeness of the dynamic languages.

Theorem 4.18. *The following dynamic proof systems are sound and weakly complete with respect to multi-agent topo-e-models:*

1. $\Box[\forall]_I + [\mathit{share}_I]$, from Definition 4.15, for $\Box[\forall]_I + [\mathit{share}_I]$;
2. $\Box[\forall]_{i,A} + [\mathit{share}_A]$, from Definition 4.15, for the language fragment $\Box[\forall]_{i,A} + [\mathit{share}_A]$;
3. $\mathbf{KB}_{i,A} + [\mathit{share}_A]$ from Definition 4.16 for the fragment $\mathbf{KB}_{i,A} + [\mathit{share}_A]$.

Proof. The proof can be found at the end of Section 5.3.

□

Chapter 5

Completeness

In this chapter, we prove the claims that we stated in Section 4.2: the proof systems $\Box[\forall]_I$, $\Box[\forall]_{i,A}$ and, most importantly, $KB_{i,A}$, are complete, as well as their dynamic extensions. Before presenting the proofs, we provide a short overview of our approach to proving completeness for the static languages.

For each of the proof systems – $\Box[\forall]_I$, $\Box[\forall]_{i,A}$, and $KB_{i,A}$ – we prove completeness via pseudo-models and relational evidence models, instead of directly for topo-e-models. That is, for each proof system, we first prove the claim for intermediate structures, which are tailored to the respective languages. Next, we define correspondences between these structures and relational evidence models. We use the correspondence from Corollary 3.40 of the latter class of models to multi-agent topo-e-models, to obtain completeness with respect to the intended multi-agent topo-e-models. A complete overview of correspondences is depicted in Figure 5.1.

For $\Box[\forall]_I$ and its fragment $\Box[\forall]_{i,A}$, the pseudo-models are generalisations of relational evidence models. On relational evidence models, evidence relations are only defined for individuals: those for groups can be obtained through intersection. In contrast, on pseudo-models, relations for groups are explicitly represented. These evidence relations are *not* distributed, that is, given a group $I \subseteq A$, a relation describing (soft or hard) evidence of I is not necessarily the intersection of the corresponding individual relations associated with its members. If all relations on a given pseudo-model do have this property, then we call the pseudo-model *standard*, and it is simply an equivalent representation of a relational evidence model.

For $KB_{i,A}$, the pseudo-models are given by a differently interpreted relational structure which, instead of evidence relations, represents only relations directly corresponding to knowledge and belief. Again, relations for subgroups are explicitly defined.

We prove completeness for each proof system with respect to relational evidence models, by defining a truth-preserving map from the pseudo-models for that proof system to *standard* pseudo-models. The standard pseudo-model that we construct is referred to as the *associated model* (see Definition 5.24). This is a tree-like model, obtained by *unravelling* a pseudo-model for $\Box[\forall]_I$ (resp. $\Box[\forall]_{i,A}$) and ensuring the properties of a standard pseudo-model on the resulting tree.

It is important to note that the *associated model* is simply a *standard pseudo-model* for $\Box[\forall]_I$ that is *associated* with the (general) pseudo-model from which it is constructed; and that the terms *standard pseudo-model* for $\Box[\forall]_I$ and *relational evidence model* refer to the same class of models.

Completeness for $\Box[\forall]_I$ (resp. $\Box[\forall]_{i,A}$) is therefore proved by the chain of correspondences from pseudo-models for $\Box[\forall]_I$ (resp. $\Box[\forall]_{i,A}$), to associated models, to topo-e-models. Since the associated model is constructed from a pseudo-model containing evidence relations, the correspondence proof

with respect to pseudo-models for $KB_{i,A}$ requires an additional step: we first map pseudo-models for $KB_{i,A}$ to pseudo-models for $\Box[\forall]_{i,A}$. The main result of this chapter is the representation theorem that proves this correspondence: Theorem 5.44. Completeness for $KB_{i,A}$ is therefore proved by the chain of correspondences from pseudo-models for $KB_{i,A}$, to pseudo-models for $\Box[\forall]_{i,A}$, to associated models, to topo-e-models.

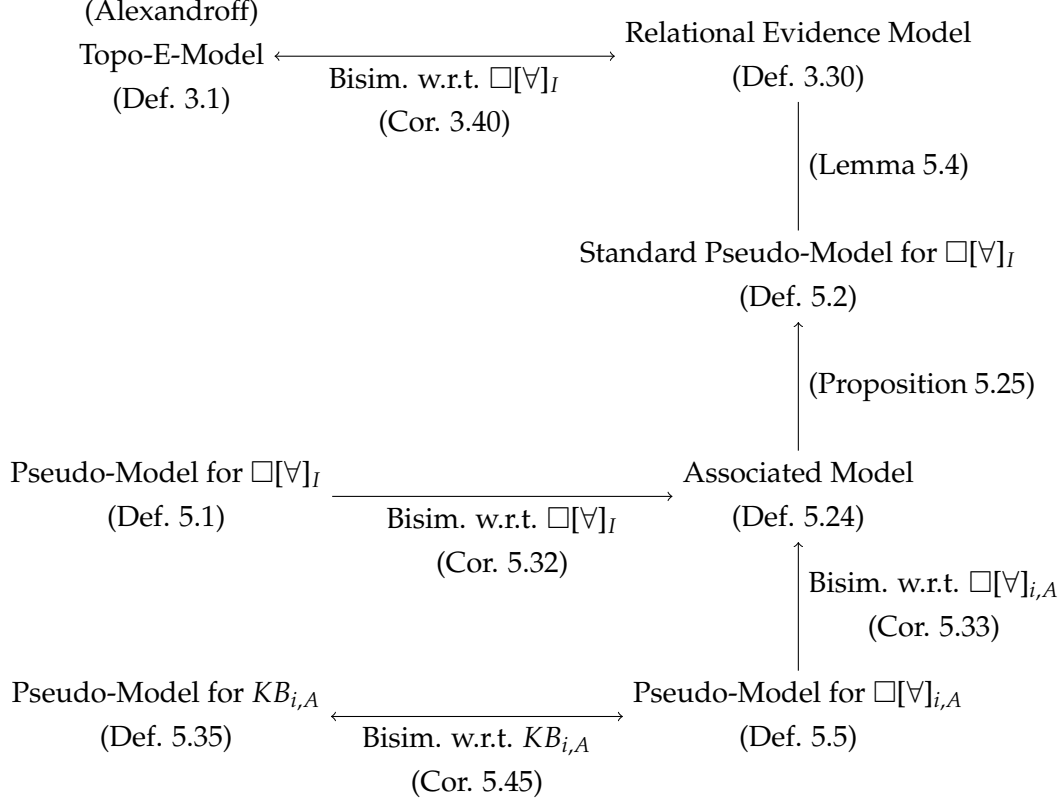


Figure 5.1: Flowchart of the correspondences proved in this chapter. An arrow from model X to model Y signifies a map from models of type X to models of type Y . The absence of arrow heads between relational evidence models and standard pseudo-models denotes that these terms refer to the same class of models. Associated models are standard pseudo-models; however, not every standard pseudo-model is an associated model.

Figure 5.1 illustrates that for each language that we consider, it suffices to show completeness with respect to the corresponding pseudo-models for that language. Completeness with respect to topo-e-models then follows from the depicted truth-preserving maps. In particular, any formula $\varphi \in KB_{i,A}$, that is satisfiable on a pseudo-model for $KB_{i,A}$, is satisfiable on a topo-e-model: $KB_{i,A}$ is a fragment of $\Box[\forall]_{i,A}$, which is a fragment of $\Box[\forall]_I$. Thus, it can be checked that every map in the chain of correspondences is truth-preserving with respect to $KB_{i,A}$.

In Table 5.0.1, we list an overview of completeness results throughout this chapter (in order of appearance), for the various proof systems, with respect to multi-agent topo-e-models. The main result is the proof of completeness for $KB_{i,A}$.

This chapter is structured as follows. The majority of Section 5.1 is concerned with proving completeness for $\Box[\forall]_I$; the claim follows virtually immediately for $\Box[\forall]_{i,A}$. The latter completeness result will be relevant in Section 5.2, where we prove completeness of $KB_{i,A}$. Finally, we conclude this

Proof System	Proof of Completeness
$\Box[\forall]_I$	Theorem 4.9
$\Box[\forall]_{i,A}$	Theorem 4.10
$KB_{i,A}$	Theorem 4.12
$\Box[\forall]_I + [share_I], \Box[\forall]_{i,A} + [share_A], \text{ and } KB_{i,A} + [share_A]$	Theorem 4.18

Table 5.0.1: The proof systems and the theorems presenting their respective proofs of completeness with respect to multi-agent topo-e-models.

chapter by presenting the completeness proofs for the dynamic extensions, in Section 5.3.

5.1 Completeness of $\Box[\forall]_I$

This section presents the completeness proof of $\Box[\forall]_I$ on multi-agent topo-e-models. The result of Corollary 3.40 allows us to prove the claim via relational evidence models. Instead of proving the claim directly, we first show completeness of $\Box[\forall]_I$ on a more general class of relational models, which we will call *pseudo-models for $\Box[\forall]_I$* (Section 5.1.1 and Section 5.1.2). Next, we define a truth-preserving map from pseudo-models to relational evidence models to derive the claim for the latter class of models (Section 5.2.3). To be precise, we prove a *p-morphism*.

As we discussed in Section 2.3.2, in contrast to group knowledge, soft and hard evidence of the group are *distributed*: evidence of the group is given by an aggregation of individual evidence possessed by its members. As a result, the completeness proof for $\Box[\forall]_I$ is similar to existing completeness proofs for logics that incorporate distributed knowledge. In particular, our proof closely resembles the proof in Appendix A of [BS20], which proves completeness of a logic incorporating, among other notions, distributed knowledge for all subgroups: Section 5.1.2 and Section 5.1.3 are based on Appendix A.1 and A.2 from [BS20], respectively.

Throughout the proof, fix a finite set of agents A and a finite vocabulary V .

5.1.1 Pseudo-Models for $\Box[\forall]_I$

This section introduces pseudo-models for $\Box[\forall]_I$ (Definition 5.1). By restricting the relations on these pseudo-models, we obtain pseudo-models for the fragment $\Box[\forall]_{i,A}$ (Definition 5.5).

Definition 5.1 (Pseudo-Model for $\Box[\forall]_I$). Given a countable vocabulary V , a *pseudo-model for $\Box[\forall]_I$* is a structure $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi)$ where S is a set of states. For each nonempty group of agents $I \subseteq A$, the relation $\leq_I \subseteq S \times S$ is a preorder and $\sim_I \subseteq S \times S$ is an equivalence relation; and $\pi : S \rightarrow \mathcal{P}(V)$ is a valuation function. The relations on this structure are required to satisfy the following two conditions:

1. **Anti-Monotonicity.** For all nonempty $I \subseteq A$, and $s, t \in X$:
 - If $s \leq_I t$ and $I \supseteq J \neq \emptyset$, then $s \leq_J t$;
 - If $s \sim_I t$ and $I \supseteq J \neq \emptyset$, then $s \sim_J t$.
2. **Inclusion.** For all nonempty $I \subseteq A$: $\leq_I \subseteq \sim_I$.

Definition 5.2 (Standard Pseudo-Model for $\Box[\forall]_I$). A pseudo-model is *standard* if it also satisfies the following condition¹:

3. **Intersection.** For all nonempty $I, J \subseteq A$:

- $\leq_{I \cup J}$ is the intersection of \leq_I and \leq_J ;
- $\sim_{I \cup J}$ is the intersection of \sim_I and \sim_J .

We define the following semantics of $\Box[\forall]_I$ on pseudo-models for $\Box[\forall]_I$.

Definition 5.3 (Pseudo-Model Semantics of $\Box[\forall]_I$). Given a countable vocabulary V , the semantics of $\Box[\forall]_I$ on pseudo-models for $\Box[\forall]_I$ is defined as follows: given a pseudo-model $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi)$ for $\Box[\forall]_I$ and a state $s \in S$, let

$$\begin{aligned}
(\mathbf{S}, s) &\models \top && \text{always holds} \\
(\mathbf{S}, s) &\models p && \text{iff } p \in \pi(s) \\
(\mathbf{S}, s) &\models \neg \varphi && \text{iff } (\mathbf{S}, s) \not\models \varphi \\
(\mathbf{S}, s) &\models \varphi \wedge \psi && \text{iff } (\mathbf{S}, s) \models \varphi \text{ and } (\mathbf{S}, s) \models \psi \\
(\mathbf{S}, s) &\models \Box_I \varphi && \text{iff for all } t \in S \text{ s.t. } s \leq_I t : (\mathbf{S}, t) \models \varphi \\
(\mathbf{S}, s) &\models [\forall]_I \varphi && \text{iff for all } t \in S \text{ s.t. } s \sim_I t : (\mathbf{S}, t) \models \varphi
\end{aligned}$$

where $p \in V$ is any proposition and $I \subseteq A$ is any nonempty subgroup of agents.

The reader might notice the symmetry with the semantics on relational evidence models (Definition 3.32). In fact, *standard* pseudo-models for $\Box[\forall]_I$ are relational evidence models:

Lemma 5.4. *Relational evidence models (Definition 3.30) are the same as standard pseudo-models.*

Proof. We can represent a relational evidence model \mathbf{X} as a standard pseudo-model \mathbf{S} for $\Box[\forall]_I$ by setting $\leq_I := \bigcap_{i \in I} \leq_i$ and setting $\sim_I := \bigcap_{i \in I} \sim_i$. Conversely, we represent a standard pseudo-model \mathbf{S} as a relational evidence model \mathbf{X} by setting $\leq_i := \leq_{\{i\}}$ and $\sim_i := \sim_{\{i\}}$.

The interpretation of any formula $\varphi \in \Box[\forall]_I$ on the relational evidence model \mathbf{X} (according to Definition 3.32) agrees with the interpretation of φ the standard pseudo-model \mathbf{S} (according to Definition 5.3), because the abbreviations $\leq_I := \bigcap_{i \in I} \leq_i$ and $\sim_I := \bigcap_{i \in I} \sim_i$ (Definition 3.31) on \mathbf{X} coincide with the directly defined group relations \leq_I and \sim_I on \mathbf{S} . \square

Thus, in order to prove completeness with respect to relational evidence models, it suffices to prove completeness with respect to standard pseudo-models for $\Box[\forall]_I$. We first prove completeness with respect to general pseudo-models for $\Box[\forall]_I$.

With general pseudo-models for $\Box[\forall]_I$ having the same type as standard pseudo-models for $\Box[\forall]_I$, we can inherit the semantics of $\Box[\forall]_I$ on this class of models from the semantics for the same language on relational evidence models.

Given this definition of pseudo-models for $\Box[\forall]_I$, we can construct pseudo-models for the fragment $\Box[\forall]_{i,A}$. We obtain these by restricting the relations and conditions in the definition of the pseudo-model for $\Box[\forall]_I$: we allow only relations with respect to the full group and individual agents.

¹For one direction, the intersection condition reduces to anti-monotonicity: let $s, t \in S$ and let $I, J \subseteq A$ be nonempty. Then, if $s \sim_{I \cup J} t$, we have by $I \subseteq I \cup J$ that $s \sim_I t$; analogously, with $J \subseteq I \cup J$, we have $s \sim_J t$.

Definition 5.5 (Pseudo-Model for $\Box[\forall]_{i,A}$). Given a countable vocabulary V , a *pseudo-model* for $\Box[\forall]_{i,A}$ is a structure $\mathbf{S} = (S, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \leq_A, \sim_A, \pi)$ where S is a set of states; $\leq_A \subseteq S \times S$; $\leq_i \subseteq S \times S$ (for each agent $i \in A$) are preorders; and $\sim_A \subseteq S \times S$ and $\sim_i \subseteq S \times S$ (for each agent $i \in A$) are equivalence relations. The valuation function is given by $\pi : S \rightarrow \mathcal{P}(V)$. These models are required to satisfy the same conditions as pseudo-models, restricted to \leq_i, \leq_A, \sim_i , and \sim_A :

1. **Anti-Monotonicity.** For all $i \in A$ and $s, t \in X$:

- If $s \leq_A t$ and $i \in A$, then $s \leq_i t$;
- If $s \sim_A t$ and $i \in A$, then $s \sim_i t$.

2. **Inclusion.** For all $i \in A$:

- $\leq_i \subseteq \sim_i$;
- $\leq_A \subseteq \sim_A$.

We restrict the semantics accordingly.

Definition 5.6 (Pseudo-Model Semantics of $\Box[\forall]_{i,A}$). The semantics for the fragment $\Box[\forall]_{i,A}$ of the language $\Box[\forall]_I$ on pseudo-models for $\Box[\forall]_{i,A}$ is obtained by restricting the semantics from Definition 5.3 on pseudo-models for $\Box[\forall]_I$ (or, equivalently, the semantics from Definition 3.32 on relational evidence models) to formulas over $\Box[\forall]_{i,A}$.

We can consider a pseudo-model for $\Box[\forall]_I$ as a pseudo-model for $\Box[\forall]_{i,A}$ by omitting all relations that are not allowed by the definition of a pseudo-model for $\Box[\forall]_{i,A}$. Proposition 5.7 shows that the two representations agree on the interpretation of all formulas over the restricted language $\Box[\forall]_{i,A}$.

Proposition 5.7. Fix a countable vocabulary V , a pseudo-model $\mathbf{S} = (S, (\leq_J)_{J \subseteq I}, (\sim_J)_{J \subseteq I}, \pi)$ for $\Box[\forall]_I$, and let $\mathbf{S}_f = (S, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \leq_A, \sim_A, \pi)$ be the corresponding pseudo-model for the fragment $\Box[\forall]_{i,A}$. We have for all formulas $\varphi \in \Box[\forall]_{i,A}(V)$ and all states $s \in S$ that $(\mathbf{S}, s) \models \varphi$ if and only if $(\mathbf{S}_f, s) \models \varphi$.

Proof. Let $\varphi \in \Box[\forall]_{i,A}(V)$. The proof is trivial when we make the following two observations: first, the semantics on the restricted pseudo-model constitutes a restriction of the semantics on the original pseudo-model. Second, φ is a formula over the restricted language, on which the modalities \Box and $[\forall]$ are defined only for individual agents and for A . Therefore, the truth value of φ does not depend in any way on relations \leq_J and \sim_J defined on subgroups J other than A or any $i \in A$. It depends only on (a subset of) the relations $(\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \leq_A$, and \sim_A , which are identical on \mathbf{S} and \mathbf{S}_f , and on which the semantics for $\Box[\forall]_I$ and the restricted semantics for $\Box[\forall]_{i,A}$ impose the same interpretation. \square

5.1.2 Soundness and Completeness of $\Box[\forall]_I$ w.r.t. Pseudo-Models

We first prove soundness and completeness of $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$ with respect to pseudo-models (for $\Box[\forall]_I$, and for $\Box[\forall]_{i,A}$, respectively). The structure of this subsection follows the structure of the proof in Appendix A.1 of [BS20]. Proposition 5.8 gives us soundness of $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$.

Proposition 5.8. The proof system $\Box[\forall]_I$ is sound with respect to pseudo-models for $\Box[\forall]_I$, and the proof system $\Box[\forall]_{i,A}$ is sound with respect to pseudo-models for $\Box[\forall]_{i,A}$.

Proof. The proof is a routine check and can be found in Appendix A.4. \square

It remains to prove completeness of $\Box[\forall]_I$ with respect to pseudo-models for $\Box[\forall]_I$ (resp. completeness of $\Box[\forall]_{i,A}$ for the language $\Box[\forall]_{i,A}$). We show the proof for $\Box[\forall]_I$ and explain how it can be adapted to $\Box[\forall]_{i,A}$.

Throughout the section, fix a finite vocabulary V and a finite set A of agents. We show that any consistent formula $\varphi_0 \in \Box[\forall]_I(V)$ is satisfiable in a finite pseudo-model, which additionally gives us the finite model property for the logic of $\Box[\forall]_I$. Using the *filtration* method (see e.g. [BRV01]), we construct a finite model by identifying each set of states in the canonical model that agrees on a finite set of formulas: the *closure* of φ_0 . We define the closure of a formula with respect to $\Box[\forall]_I$ as follows:

Definition 5.9 (Closure ($\Box[\forall]_I$)). Given a formula $\varphi_0 \in \Box[\forall]_I(V)$, the closure $\Phi = \Phi(\varphi_0)$ of φ is the smallest set of formulas over $\Box[\forall]_I(V)$ satisfying, for all nonempty $J \subsetneq I \subseteq A$, and for all formulas $\psi, \theta \in \Box[\forall]_I(V)$:

1. $\varphi_0 \in \Phi$;
2. If $\psi \in \Phi$ and θ is a subformula of ψ , then $\theta \in \Phi$;
3. Φ is closed under single negations² \sim : if $\psi \in \Phi$, then $(\sim \psi) \in \Phi$.
4. If $[\forall]_J \psi \in \Phi$, then $[\forall]_I \psi \in \Phi$;
5. If $[\forall]_I \psi \in \Phi$, then $\Box_I [\forall]_I \psi \in \Phi$;
6. If $\neg [\forall]_I \psi \in \Phi$, then $\Box_I \neg [\forall]_I \psi \in \Phi$;
7. If $\Box_J \psi \in \Phi$, then $\Box_I \psi \in \Phi$;
8. If $[\forall]_I \psi \in \Phi$, then $\Box_I \psi \in \Phi$.

Lemma 5.10. Every formula $\varphi_0 \in \Box[\forall]_I(V)$ has a finite closure $\Phi(\varphi_0)$.

Proof. The proof is straightforward and can be found in Appendix A.4. □

Filtrated pseudo-model for $\Box[\forall]_I$. Now fix a consistent formula φ_0 and let $\Phi = \Phi(\varphi_0)$ be its closure. By Lemma 5.10, Φ is finite. It remains to show that φ is satisfiable in a finite pseudo-model for $\Box[\forall]_I$. We define the *filtrated pseudo-model* for $\Box[\forall]_I$, which can be thought of as a finite filtration of the usual notion of a canonical model, with respect to Φ (see [BRV01] for details). We define the filtrated pseudo-model as follows:

Definition 5.11. Fix a finite set of agents A and a consistent formula $\varphi_0 \in \Box[\forall]_I(V)$, fix a maximally consistent theory³ $T_0 \subseteq \Phi$ with $\varphi_0 \in T_0$. The *filtrated pseudo-model* for $\Box[\forall]_I$ for formula φ_0 is the finite structure $\mathbf{S}^C = (S^C, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi)$, where S^C is defined as

$$S^C := \{T \subseteq \Phi \mid T \subseteq \Box[\forall]_I(V) \text{ and } T \text{ is a maximally consistent subset of } \Phi\}$$

and for all nonempty $I \subseteq A$, the relations \leq_I and \sim_I on S^C are given by

$$\begin{aligned} T \sim_I W & \quad \text{iff} \quad \forall \emptyset \neq J \subseteq I \quad ([\forall]_J \varphi \in T \Leftrightarrow [\forall]_J \varphi \in W) \\ T \leq_I W & \quad \text{iff} \quad \forall \emptyset \neq J \subseteq I \quad (\Box_J \varphi \in T \Rightarrow \Box_J \varphi \in W) \end{aligned}$$

²The single negation $\sim \varphi$ is defined as: $\sim \varphi := \theta$ if φ is of the form $\neg \theta$; and $\sim \varphi := \neg \varphi$ if φ is not of the form $\neg \theta$.

³This theory exists by the Lindenbaum Lemma (see e.g. [BRV01]) and consistency of φ_0 .

for $T, W \in S^C$. Finally, we define for all $T \in S^C$

$$\pi(T) := \{p \in V \mid p \in T\}.$$

The relations on the filtrated pseudo-model for $\Box[\forall]_I$ are constructed with the purpose of facilitating a finite closure Φ : first, the relations \sim_I and \leq_I on the model take subgroups into account, such that anti-monotonicity is ensured. Furthermore, the definition of \sim_I ensures that all \sim_I relations are equivalence relations and, similarly, the definition of the relations \leq_I ensures that these are pre-orders.

Moreover, since we ensured that Φ is finite, the model S^C is finite: its size is $|S^C| \leq |2^\Phi|$, as the collection of maximally consistent subsets of Φ is a subset of the powerset of Φ . It can be checked that S^C is indeed a pseudo-model:

Lemma 5.12. *The filtrated pseudo-model S^C for $\Box[\forall]_I$ is a pseudo-model.*

Proof. The proof is straightforward and can be found in Appendix A.4. □

To prove completeness with respect to S^C , we use the Truth Lemma.

Lemma 5.13 (Truth Lemma). *Given a filtrated pseudo-model S^C for $\Box[\forall]_I$ over a closure Φ , we have for all $\varphi \in \Phi$:*

$$T \models_{S^C} \varphi \text{ iff } \varphi \in T, \text{ for every } T \in S^C.$$

Proof. The Truth Lemma is a standard lemma in canonical-model constructions (see e.g. [BRV01]) and its proof is straightforward. Details can be found in Appendix A.4. The cases for soft and hard evidence are similar to the case for distributed knowledge in the proof of Lemma 1.2 in Appendix A.1 in [BS20]. □

Corollary 5.14. *The proof system $\Box[\forall]_I$ (Definition 4.8) is sound and weakly complete with respect to pseudo-models for $\Box[\forall]_I$. Moreover, the logic of $\Box[\forall]_I$ is decidable.*

Proof. Soundness was established in Proposition 5.8. For completeness, let $\varphi_0 \in \Box[\forall]_I$ be any consistent formula and construct the filtrated pseudo-model S^C for $\Box[\forall]_I$, for φ_0 . By the Lindenbaum Lemma, there exists some maximally consistent theory T_0 in S^C with $\varphi_0 \in T_0$. By the Truth Lemma 5.13, T_0 satisfies φ_0 in S^C . Since S^C is finite, this gives us weak completeness with respect to finite pseudo-models for $\Box[\forall]_I$ (and hence also with respect to all pseudo-models).

Since $\Box[\forall]_I$ is weakly complete with respect to finite pseudo-models for the language, the logic $\Box[\forall]_I$ has the finite pseudo-model property. Therefore, it is decidable: to decide $\varphi_0 \in \Box[\forall]_I$, let $\Phi := \Phi(\varphi_0)$ be its closure and generate all pseudo-models (up to isomorphism) that are at most of the size $2^{|\Phi|}$. Then model-check φ_0 on these models: if φ_0 is satisfied at any state in any of the models, then it is satisfiable (on pseudo-models for $\Box[\forall]_I$); otherwise, it is unsatisfiable. □

Corollary 5.15. *The proof system $\Box[\forall]_{i,A}$ (Definition 4.8) is sound and weakly complete with respect to pseudo-models for $\Box[\forall]_{i,A}$. Moreover, the logic of $\Box[\forall]_{i,A}$ is decidable.*

Proof. The proof is obtained simply by restricting to $\Box[\forall]_{i,A}$ all the constructions in the proof for $\Box[\forall]_I$: to construct the filtrated pseudo-model S^C for $\Box[\forall]_{i,A}$, for any consistent formula $\varphi_0 \in \Box[\forall]_{i,A}$, restrict

the formulas in the closure $\Phi(\varphi_0)$ to $\Box[\forall]_{i,A}$; and define the filtrated pseudo-model \mathbf{S}^C as a pseudo-model for $\Box[\forall]_{i,A}$ (that is, restrict the relations from the definition of the canonical pseudo-model for $\Box[\forall]_I$ to those labelled by A and $\{i\}$ for all $i \in A$). The rest of the proof goes through exactly as in the proof of Corollary 5.14. \square

5.1.3 From Pseudo-Models to Models

To prove completeness with respect to *standard* pseudo-models, it remains to show how to go from a general pseudo-model to a standard pseudo-model satisfying the same formulas, for both $\Box[\forall]_I$ and its fragment $\Box[\forall]_{i,A}$. We define the correspondence with respect to pseudo-models for $\Box[\forall]_I$, after which we show how to adapt the proof for $\Box[\forall]_{i,A}$.

Given a pseudo-model $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi_S)$ for $\Box[\forall]_I$, we use *model unravelling* to construct an *associated* model $\mathbf{X} = (H, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi_X)$. The associated model will be a relational evidence model structured as a *tree*, which consists of infinitely many copies of the states in the original pseudo-model. On this structure we impose the desired properties. The challenge of this proof is to ensure that the relations on the unravelled tree satisfy the intersection condition of a pseudo-model for $\Box[\forall]_I$ (Definition 5.2), such that it is indeed standard.

The structure of this proof closely follows the structure of the proof in Appendix A.2 of [BS20], which uses the same technique to construct an associated model, from a relational model similar to our pseudo-model. For an introduction into model unravelling for completeness proofs, we refer to [BRV01]. Throughout this section, fix a pseudo-model $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi_S)$ for $\Box[\forall]_I$, and a state $s_0 \in S$. The state space of the associated model will consist of all s_0 -originated *histories*:

Definition 5.16 (Histories). The set H consisting of all histories in the pseudo-model \mathbf{S} is defined as the set of all finite sequences $h = (s_0, R_{G_1}, \dots, R_{G_n}, s_n)$ satisfying the following conditions:

1. The sequence h has length $n \geq 0$ and we have $s_i \in S$ for all $i \leq n$ (with s_0 being the fixed state in the model);
2. The subgroups $G_1, \dots, G_n \subseteq A$ are nonempty;
3. For each $k \in \{1, \dots, n\}$, we have one of the following two cases:
 - (a) R_{I^k} refers to \leq_{I^k} , and we have $s_{k-1} \leq_{I^k} s_k$
 - (b) R_{I^k} refers to \sim_{I^k} , and we have $s_{k-1} \sim_{I^k} s_k$.

Given a history $h = (s_0, R_{G_1}, \dots, R_{G_n}, s_n) \in H$, we denote by $\text{last}(h) := s_n$ the last state in the history.

Next, we construct the relations \leq_I and \sim_I for all nonempty $I \in A$ in intermediate steps (described in Definitions 5.17, 5.18, and 5.19), in a manner that, in particular, ensures that $\leq_I = \bigcap_{i \in I} \leq_i$ and $\sim_I = \bigcap_{i \in I} \sim_i$. We show in Proposition 5.25 that the relations resulting from Definition 5.19 satisfy the intersection condition on pseudo-models for $\Box[\forall]_I$.

Definition 5.17. Let H denote the set of (s_0 -generated) histories of \mathbf{S} . By \xrightarrow{P}_I and \xrightarrow{E}_I we denote the natural *one-step relations* on histories, labelled by P (pre-order) or E (equivalence relation), and by a

nonempty subgroup $I \subseteq A$. We let

$$\begin{aligned} h \xrightarrow{P}_I h' & \text{ iff } h' = (h, \leq_I, s') \text{ with } \text{last}(h) \leq_I s' = \text{last}(h') \\ h \xrightarrow{E}_I h' & \text{ iff } h' = (h, \sim_I, s') \text{ with } \text{last}(h) \sim_I s' = \text{last}(h'). \end{aligned}$$

Define the *immediate successor relation* \rightarrow on histories as the union of all one-step relations:

$$h \rightarrow h' \text{ iff } h (\xrightarrow{P}_I \cup \xrightarrow{E}_I) h' \text{ for some } I \subseteq A.$$

Note that H has the structure of a tree rooted at s_0 (that is, the history given by the sequence (s_0)): the immediate successor relation on H has the tree property, i.e. it connects every two nodes h, h' of the tree by a unique non-redundant path [BRV01].

For the tree to be a relational evidence model, the relations \leq_I and \sim_I for all nonempty $I \subseteq A$ need to satisfy the conditions from Definition 5.1. The first condition is anti-monotonicity. For each nonempty $J \subseteq A$, we close the existing \xrightarrow{P}_J and \xrightarrow{E}_J relations from Definition 5.17 under anti-monotonicity, by extending them with all the corresponding one-step relations labelled by supergroups of J . We refer to the resulting extended one-step relations by $\xrightarrow{\leq}_J$ and $\xrightarrow{\sim}_J$ for all $J \subseteq A$:

Definition 5.18. Let \xrightarrow{P}_I and \xrightarrow{E}_I be as defined in Definition 5.17. For all nonempty $J \subseteq A$, define

$$\begin{aligned} h \xrightarrow{\leq}_J h' & \text{ iff } h \xrightarrow{P}_I h' \text{ for some } I \supseteq J \\ h \xrightarrow{\sim}_J h' & \text{ iff } h \xrightarrow{E}_I h' \text{ for some } I \supseteq J. \end{aligned}$$

We now define the final relations \leq_I and \sim_I , which satisfy the conditions of a relational evidence model. In particular, we obtain individual relations $\leq_i := \leq_{\{i\}}$ and $\sim_i := \sim_{\{i\}}$.

Definition 5.19 (Relations on the Associated Model). Let $I \subseteq A$ be nonempty and let $\xrightarrow{\leq}_I$ and $\xrightarrow{\sim}_I$ be as defined in Definition 5.18. We define

$$\begin{aligned} \leq_I &:= \left(\xrightarrow{\leq}_I \right)^* \\ \sim_I &:= \left(\xrightarrow{\leq}_I \cup \xleftarrow{\leq}_I \cup \xrightarrow{\sim}_I \cup \xleftarrow{\sim}_I \right)^* \end{aligned}$$

where R^* denotes the reflexive-transitive closure of R , and $\xleftarrow{\leq}_I$ and $\xleftarrow{\sim}_I$ denote the converses of $\xrightarrow{\leq}_I$ and $\xrightarrow{\sim}_I$, respectively.

The following lemmas state a number of properties of the relations from Definition 5.19, which we will impose on our associated model. These will be useful when we show that the relations satisfy the conditions in Definition 5.1 and Definition 5.2 of a standard pseudo-model for $\Box[\forall]_I$ (Lemma 5.22) and, subsequently, when we prove a bisimulation between the associated model and the original pseudo-model for $\Box[\forall]_I$ (Proposition 5.31).

Lemma 5.20. For all nonempty $I \subseteq A$, and histories $h, h' \in H$, the following are equivalent:

1. $h \leq_I h'$;
2. the non-redundant path from h to h' consists only of steps of the form $h_{n-1} \xrightarrow{P}_{G_n} h_n$, with $I \subseteq G_n$.

Proof. Let $I \subseteq A$ be nonempty, and let $h, h' \in H$. For the left-to-right direction, suppose $h \leq_I h'$. Then, by definition of \leq_I (Definition 5.19), we have $h \left(\xrightarrow{I} \right)^* h'$, that is, from h we can reach h' via a finite non-redundant path under the relation \xrightarrow{I} . More importantly, by the properties of a tree-like model, this non-redundant path is unique. The claim now follows immediately from the definition of \xrightarrow{I} (Definition 5.18): each step $h_{n-1} \xrightarrow{I} h_n$ on the path implies that for some $G_n \supseteq I$ we have $h_{n-1} \xrightarrow{P_{G_n}} h_n$.

For the converse direction, the claim is immediate: assuming that the non-redundant path from h to h' consists only of steps of the form $h_{n-1} \xrightarrow{P_{G_n}} h_n$, with $I \subseteq G_n$, we have for every step $h_{n-1} \xrightarrow{P_{G_n}} h_n$ on the path that $h_{n-1} \xrightarrow{I} h_n$ (Definition 5.18), and thereby, $h \left(\xrightarrow{I} \right)^* h'$ (Definition 5.19). \square

Lemma 5.21 is the analogue of the previous lemma, for the equivalence relations \sim_I .

Lemma 5.21. *The following are equivalent, for all $I \subseteq A$ and histories $h, h' \in H$:*

1. $h \sim_I h'$;
2. *each of the steps on the non-redundant path from h to h' is of one of the following forms:*

$$\begin{array}{ll}
 (a) \ h_{n-1} \xrightarrow{P_{G_n}} h_n & (c) \ h_{n-1} \xrightarrow{E_{G_n}} h_n \\
 (b) \ h_{n-1} \xleftarrow{P_{G_n}} h_n & (d) \ h_{n-1} \xleftarrow{E_{G_n}} h_n
 \end{array}$$

with $I \subseteq G_n$.

Proof. Let $I \subseteq A$ be nonempty, and let $h, h' \in H$. For the left-to-right direction, suppose $h \sim_I h'$. Then, by definition of \sim_I (Definition 5.19), we have $h \left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim_I} \cup \xleftarrow{\sim_I} \right)^* h'$, that is, from h we can reach h' via a finite non-redundant path under the relation $\left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim_I} \cup \xleftarrow{\sim_I} \right)$. More importantly, by the properties of a tree-like model, this non-redundant path is unique. Consider an arbitrary step $h_{n-1} \left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim_I} \cup \xleftarrow{\sim_I} \right) h_n$ on this path. We have one of the following four cases:

$$\begin{array}{ll}
 (a) \ h_{n-1} \xrightarrow{I} h_n; & (c) \ h_{n-1} \xrightarrow{\sim_I} h_n; \\
 (b) \ h_{n-1} \xleftarrow{I} h_n; & (d) \ h_{n-1} \xleftarrow{\sim_I} h_n.
 \end{array}$$

The claim then follows from unfolding the respective definitions of these relations (Definition 5.18).

For the converse direction, the claim is immediate: assuming that the non-redundant path from h to h' consists only of steps of the form (a)-(d) as listed in Lemma 5.21, with $I \subseteq G_n$ for each step from h_{n-1} to h_n , we can apply the corresponding definitions from Definition 5.18 to each step, to obtain that $h \left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim_I} \cup \xleftarrow{\sim_I} \right)^* h'$, i.e. $h \sim_I h'$ (Definition 5.19). \square

Lemma 5.22 shows that the relations from Definition 5.19 satisfy the requirements of a *standard pseudo-model* for $\Box[\forall]_I$:

Lemma 5.22. *Let $I \subseteq A$ be nonempty. The relations \leq_I and \sim_I from Definition 5.19 satisfy the relational conditions from Definition 5.1: $\leq_I \subseteq \sim_I$ (the inclusion condition); \leq_I is a pre-order; and \sim_I is an equivalence relation. Furthermore, for all nonempty $J \subseteq A$, \leq_I and \leq_J satisfy the anti-monotonicity and intersection conditions, as well as \sim_I and \sim_J .*

Proof. The inclusion condition is satisfied by construction of \sim_I : let $h, h' \in H$ and suppose $h \leq_I h'$. Then from h , we can reach h' via a unique non-redundant path under the relation $\xrightarrow{\leq}_I$. Since the relation $\xrightarrow{\leq}_I$ is a subset of the relation $\left(\xrightarrow{\leq}_I \cup \xleftarrow{\leq}_I \cup \xrightarrow{\sim}_I \cup \xleftarrow{\sim}_I \right)$, h and h' are automatically connected by the same path, under the relation $\left(\xrightarrow{\leq}_I \cup \xleftarrow{\leq}_I \cup \xrightarrow{\sim}_I \cup \xleftarrow{\sim}_I \right)$. By definition of \sim_I , we have $h \sim_I h'$.

The relation \leq_I is a pre-order by construction: it is the reflexive-transitive closure of $\xrightarrow{\leq}_I$.

Reflexivity and transitivity of \sim_I are immediate by Definition 5.19, since \sim_I is the reflexive-transitive closure of a union of relations. For symmetry, let $h, h' \in H$ and suppose $h \sim_I h'$. Then each of the steps on the non-redundant path from h to h' is of one of the forms listed in Lemma 5.21. Observe that the converse of each of these steps is also listed, which means that each of the steps on the non-redundant path from h' to h is also of one of the listed forms, i.e. we have $h' \sim_I h$.

We prove the anti-monotonicity claim only for the \sim relations, since the proof for \leq is similar and less complicated. To see that the \sim relations satisfy the anti-monotonicity condition, let $I, J \subseteq A$ be two nonempty subgroups and let $h, h' \in H$. Suppose that $J \subseteq I$ and $h \sim_I h'$. We claim that $h \sim_J h'$. By $h \sim_I h'$, we know that each of the steps on the non-redundant path from h to h' is of one of the forms listed in Lemma 5.21. Consider an arbitrary step on this path, from a history h_{n-1} to another history h_n . We distinguish the four⁴ cases from Lemma 5.21, with G_n being an arbitrary superset of I :

1. $h_{n-1} \xrightarrow{P}_{G_n} h_n$. With $J \subseteq I$, clearly, $J \subseteq G_n$. By construction of $\xrightarrow{\leq}_J$ (Definition 5.18), we get $h_{n-1} \xrightarrow{\leq}_J h_n$.
2. $h_{n-1} \xleftarrow{P}_{G_n} h_n$. This is equivalent to having $h_n \xrightarrow{P}_{G_n} h_{n-1}$ and thus, by item (a), we have $h_n \xrightarrow{\leq}_J h_{n-1}$, i.e. $h_{n-1} \xleftarrow{\leq}_J h_n$.
3. $h_{n-1} \xrightarrow{E}_{G_n} h_n$. Similar to case (a): clearly, $J \subseteq G_n$. By construction of $\xrightarrow{\sim}_J$, we get $h_{n-1} \xrightarrow{\sim}_J h_n$.
4. $h_{n-1} \xleftarrow{E}_{G_n} h_n$. Similar to (b). We get $h_{n-1} \xleftarrow{\sim}_J h_n$.

Thus, each step on the path is of the form $h_{n-1} \left(\xrightarrow{\leq}_J \cup \xleftarrow{\leq}_J \cup \xrightarrow{\sim}_J \cup \xleftarrow{\sim}_J \right) h_n$. By definition of \leq_J (Definition 5.19), we get that $h \leq_J h'$, as required.

Similarly, we prove the intersection condition only for \sim : the proof for \leq_I is similar and less complicated. Let $I, J \subseteq A$ be nonempty. We show that for any $h, h' \in H$, we have $h \sim_{I \cup J} h'$ if and only if $h \sim_I h'$ and $h \sim_J h'$. Observe that if we assume $\sim_I = \bigcap_{i \in I} \sim_i$, then the result follows directly: by $\sim_{I \cup J} = \bigcap_{i \in I \cup J} \sim_i$, and $\sim_I = \bigcap_{i \in I} \sim_i$ and $\sim_J = \bigcap_{j \in J} \sim_j$, we get that

$$\sim_{I \cup J} = \bigcap_{i \in I \cup J} \sim_i = \left(\bigcap_{i \in I} \sim_i \right) \cap \left(\bigcap_{j \in J} \sim_j \right) = \sim_I \cap \sim_J.$$

It remains to prove the claim. We state it for both \sim_I and \leq_I .

Claim 5.23. *For all nonempty $I \subseteq A$, we have that $\sim_I = \bigcap_{i \in I} \sim_i$ and $\leq_I = \bigcap_{i \in I} \leq_i$.*

Proof of claim. We prove the claim only for \sim_I : the proof for \leq_I is similar and less complicated.

For the left-to-right direction, the claim reduces to anti-monotonicity, which we already proved. For the converse direction, let $h, h' \in H$ and suppose that $h \sim_i h'$ for all $i \in I$. Let $i \in I$ be arbitrary. By definition of \sim_i , each of the steps on the non-redundant path from h to h' is of one of the forms listed in Lemma 5.21. Consider an arbitrary step on this path, from a history h_{n-1} to a history h_n . Since the proofs for the different cases from Lemma 5.21 are symmetrical, we only show the proof for case (a)⁵.

⁴For the proof of anti-monotonicity for \leq , case (a) is the only possible case for any step on the path (Lemma 5.20).

⁵For the proof of intersection for \leq_I , (a) is the only possible case for any step on the path (Lemma 5.20).

Suppose that (a) the step is of the form $h_{n-1} \xrightarrow{P_{G_n}} h_n$ for some $G_n \supseteq \{i\}$. Recall that i was arbitrary, and that this path is unique. It follows that $G_n \supseteq \{i'\}$ for all $i' \in I$. But then $G_n \supseteq I$. Thus, by definition of \xrightarrow{I} (Definition 5.18), we have that $h_{n-1} \xrightarrow{I} h_n$.

Combining this with the proofs of the other cases, we get that h_{n-1} and h_n must be related by one of the one-step relations \xrightarrow{I} , \xleftarrow{I} , $\xrightarrow{\sim}_I$, or $\xleftarrow{\sim}_I$ for I . In other words, $h_{n-1} \left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim}_I \cup \xleftarrow{\sim}_I \right) h_n$. Since this was an arbitrary step on the unique non-redundant path from h to h' , we can conclude that $h \left(\xrightarrow{I} \cup \xleftarrow{I} \cup \xrightarrow{\sim}_I \cup \xleftarrow{\sim}_I \right)^* h'$, i.e. $h \sim_I h'$, as required. \square

In conclusion, all relational conditions of a pseudo-model for $\Box[\forall]_I$ are satisfied by the given relations. \square

As our relations \leq_I and \sim_I on H have the desired properties, we can now define our associated model \mathbf{X} for \mathbf{S} . In Definition 5.24, the model is represented as a relational evidence model, that is, we explicitly define only the individual relations.

Definition 5.24 (Associated Model). The *associated model* for \mathbf{S} is a structure $\mathbf{X} = (H, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi_{\mathbf{X}})$, where

1. H is the set of all histories on S ;
2. For all $i \in A$, $\leq_i = \leq_{\{i\}}$ and $\sim_i = \sim_{\{i\}}$, with $\leq_{\{i\}}$ and $\sim_{\{i\}}$ as defined in Definition 5.19;
3. The valuation $\pi_{\mathbf{X}} : H \rightarrow \mathcal{P}(V)$ on histories is defined as $\pi_{\mathbf{X}}(h) := \pi_{\mathbf{S}}(\text{last}(h))$.

To see that \mathbf{X} is a relational evidence model (as defined in Definition 3.30), recall that for each $i \in A$, \leq_i is a preorder, \sim_i is an equivalence relation, and that $\leq_i \subseteq \sim_i$ (as we proved in Lemma 5.22).

Proposition 5.25. We can consider \mathbf{X} as a standard pseudo-model (for $\Box[\forall]_I$) $\mathbf{X} = (H, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi_{\mathbf{X}})$ by explicitly representing the group relations \leq_I and \sim_I for all nonempty $I \subseteq A$, as defined in Definition 5.19.

Proof. We proved in Lemma 5.22 that the group relations \leq_I and \sim_I for nonempty groups $I \subseteq A$ on \mathbf{X} satisfy all conditions of a pseudo-model for $\Box[\forall]_I$, and in particular, that the intersection condition for standard pseudo-models for $\Box[\forall]_I$ (Definition 5.2) is satisfied. \square

Since the pseudo-model-based semantics from Definition 5.3 agrees with the model-based semantics from Definition 3.32 for \mathbf{X} , we can compare \mathbf{S} and \mathbf{X} directly as pseudo-models, that is, by explicitly representing the group relations.

For pseudo-models for the fragment $\Box[\forall]_{i,A}$, the associated model is constructed in the same way:

Fact 5.26. Given a pseudo-model \mathbf{S}_f for the fragment $\Box[\forall]_{i,A}$, the *associated model* for \mathbf{S}_f is a structure $\mathbf{X}_f = (H, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi_{\mathbf{X}_f})$, which is obtained by restricting the construction of the associated model from Definition 5.24 to relations for individual agents and for the full group, i.e. the relations labelled by A itself or by groups of the form $\{i\} \subseteq A$. The resulting structure \mathbf{X}_f is a relational evidence model.

Mapping general pseudo-models to relational evidence models. To extend our completeness proof from Theorem 4.9 to relational evidence models, it suffices to show that any formula φ that is satisfiable on our arbitrary pseudo-model \mathbf{S} for $\Box[\forall]_I$, is also satisfiable on its associated model \mathbf{X} .

We show that the map $last(\cdot)$ from histories to states is a p-morphism, i.e. a functional bisimulation. To prove this, we need Lemma 5.27 and Lemma 5.28. These lemmas state properties of the one-step relations \xrightarrow{I} and $\xrightarrow{\sim}_I$, which we inductively extend to properties of \leq_I and \sim_I in Lemma 5.29 and Lemma 5.30.

Lemma 5.27. *For all nonempty $I \subseteq A$, if $h \xrightarrow{I} h'$, then $last(h) \leq_I last(h')$.*

Proof. Suppose $h \xrightarrow{I} h'$. By Definition 5.18 of \xrightarrow{I} , there is $G \supseteq I$ such that $h \xrightarrow{P}_G h'$. By Definition 5.17 of \xrightarrow{P}_G , we have $h' = (h, \leq_G s')$ with $last(h) \leq_G s' = last(h')$. By the anti-monotonicity condition on pseudo-models (Definition 5.1), we get that $last(h) \leq_I last(h')$. \square

Lemma 5.28. *For all nonempty $I \subseteq A$, if $h \xrightarrow{\sim}_I h'$, then $last(h) \sim_I last(h')$.*

Proof. Analogous to the proof of Lemma 5.27: if $h \xrightarrow{\sim}_I h'$, then for some $G \supseteq I$ we have $h' = (h, \sim_G s')$ with $last(h) \sim_G s' = last(h')$, which by the anti-monotonicity condition on pseudo-models gives us that $last(h) \sim_I last(h')$. \square

For our final step, we extend the properties from Lemmas 5.27 and 5.28 to groups:

Lemma 5.29. *For all nonempty $I \subseteq A$, if $h \leq_I h'$, then $last(h) \leq_I last(h')$.*

Proof. By induction on the length n of the non-redundant path from h to h' . For the base case, where $h \leq_I h'$ with $n = 0$, we have $h = h'$. So the claim that $last(h) \leq_I last(h')$ follows immediately from reflexivity of \leq_I .

For the inductive step, suppose the claim holds for paths of length n , and suppose the non-redundant path from h to h' has length $n + 1$. By Lemma 5.20, the last step of the non-redundant path from h to h' must be of the form $h_n \xrightarrow{P}_{I^{n+1}} h_{n+1} = h'$, with $I^{n+1} \supseteq I$. So by definition of \xrightarrow{I} , we have $h_n \xrightarrow{I} h_{n+1}$. Using Lemma 5.27 we obtain that $last(h_n) \leq_I last(h_{n+1})$. By transitivity of \leq_I , it now suffices to show that $last(h) \leq_I last(h_n)$ (since that would give us that $last(h) \leq_I last(h_{n+1})$).

Since the path from h to h_n has length n , we can apply the induction hypothesis to the fact that $h \leq_I h_n$ (which follows from our assumption that $h \leq_I h_{n+1}$). This gives us that $last(h) \leq_I last(h_n)$. \square

Lemma 5.30. *For all nonempty $I \subseteq A$, if $h \sim_I h'$, then $last(h) \sim_I last(h')$.*

Proof. By induction on the length n of the non-redundant path from h to h' . For the base case, where $h \sim_I h'$ with $n = 0$ we have $h = h'$. So the claim that $last(h) \sim_I last(h')$ follows immediately from reflexivity of \sim_I .

For the inductive step, suppose the claim holds for paths of length n , and suppose the non-redundant path from h to h' has length $n + 1$. By Lemma 5.21, the last step of the non-redundant path from h to h' must be of one of the forms

1. $h_n \xrightarrow{P}_{I^{n+1}} h_{n+1}$
2. $h_n \xleftarrow{P}_{I^{n+1}} h_{n+1}$
3. $h_n \xrightarrow{E}_{I^{n+1}} h_{n+1}$
4. $h_n \xleftarrow{E}_{I^{n+1}} h_{n+1}$

with $h_{n+1} = h'$ and $I^{n+1} \supseteq I$. So applying the definitions of \xrightarrow{I} and $\xrightarrow{\sim}_I$ to these cases, one of the following is the case:

1. $h_n \xrightarrow{\leq}_{I^{n+1}} h_{n+1}$
2. $h_n \xleftarrow{\leq}_{I^{n+1}} h_{n+1}$
3. $h_n \xrightarrow{\sim}_{I^{n+1}} h_{n+1}$
4. $h_n \xleftarrow{\sim}_{I^{n+1}} h_{n+1}$.

First observe that the path from h to h_n has length n and we can therefore apply the induction hypothesis to the fact that $h \sim_I h_n$ (which follows from the assumption that $h \sim_I h_{n+1}$ and from the definition of \sim_I). This gives us that $\text{last}(h) \sim_I \text{last}(h_n)$. It remains to show that $\text{last}(h_n) \sim_I \text{last}(h_{n+1})$, which by transitivity of \sim_I will give us that $\text{last}(h) \sim_I \text{last}(h_{n+1}) = \text{last}(h')$, as required.

We use Lemma 5.27 for cases (1) and (2), and Lemma 5.28 for cases (3) and (4), to obtain that either $\text{last}(h_n) \leq_I \text{last}(h_{n+1})$, or $\text{last}(h_n) \sim_I \text{last}(h_{n+1})$, or one of their converses is true. In the cases of $\text{last}(h_n) \sim_I \text{last}(h_{n+1})$ and $\text{last}(h_{n+1}) \sim_I \text{last}(h_n)$ we are done, so suppose that $\text{last}(h_n) \leq_I \text{last}(h_{n+1})$ or $\text{last}(h_{n+1}) \leq_I \text{last}(h_n)$ is the case. But then we have by the inclusion condition on pseudo-models that $\text{last}(h_n) \sim_I \text{last}(h_{n+1})$, so we can conclude that $\text{last}(h) \sim_I \text{last}(h')$. \square

The following proposition states that \mathbf{S} and \mathbf{X} (represented as a pseudo-model) are bisimilar, in particular, that the function $\text{last}(\cdot)$ from associated models to pseudo-models is a p -morphism⁶.

Proposition 5.31. *Let \mathbf{S} be a pseudo-model and let its associated model be given by \mathbf{X} . Then the map $\text{last} : H \rightarrow S$, mapping every history $h \in H$ to its last element $\text{last}(h)$, defines a p -morphism from \mathbf{X} to \mathbf{S} (with \mathbf{X} and \mathbf{S} seen as Kripke models with basic relations \sim_I for all $I \subseteq A$).*

Proof. The function $\text{last}(\cdot)$ is well-defined: since every history $h \in H$ is by definition a nonempty sequence, it contains at least one state. Since it is also finite, it must have a last state: $\text{last}(h)$ exists. To see that $\text{last}(\cdot)$ is a p -morphism, we check the following three conditions:

Atomic preservation for basic atoms: $p \in V$, i.e. $p \in \pi_{\mathbf{X}}(h)$ if and only if $p \in \pi_{\mathbf{S}}(\text{last}(h))$. This is immediate by definition of the valuation function $\pi_{\mathbf{X}}$ for associated models.

Forth condition: let $I \subseteq A$ be nonempty. For \leq_I , assume $h \leq_I h'$; then $\text{last}(h) \leq_I \text{last}(h')$ is immediate from Lemma 5.27. For \sim_I , assume $h \sim_I h'$ then $\text{last}(h) \sim_I \text{last}(h')$ is immediate from Lemma 5.28.

Back condition: let $I \subseteq A$ be nonempty. For \leq_I , assume $\text{last}(h) \leq_I s'$. We need to prove that there is $h' \in H$ such that $h \leq_I h'$ and $\text{last}(h') = s'$. From $\text{last}(h) \leq_I s'$, we know that (h, \leq_I, s') is a history in H . So we can take $h' := (h, \leq_I, s')$. Similarly, for \sim_I , assume $\text{last}(h) \sim_I s'$. Again, we can take $h' := (h, \leq_I, s')$ to prove that there is $h' \in H$ such that $h \sim_I h'$ and $\text{last}(h') = s'$. \square

Corollary 5.32. *The same formulas in $\square[\forall]_I$ are satisfiable in the associated model \mathbf{X} , as in its p -morphic image contained in the pseudo-model \mathbf{S} for $\square[\forall]_I$. More precisely, for every history $h \in H$ and every formula $\varphi \in \square[\forall]_I$, we have:*

$$h \models_{\mathbf{X}} \varphi \quad \text{iff} \quad \text{last}(h) \models_{\mathbf{S}} \varphi.$$

Proof. By Proposition 5.31, the map $\text{last}(\cdot) : H \rightarrow S$ is a bisimulation between \mathbf{S} and its image in \mathbf{X} , seen as Kripke models for the language with modalities \square_I and $[\forall]_I$ for all nonempty $I \subseteq A$. Since $\square[\forall]_I$ is the basic modal language for this vocabulary, formulas in $\square[\forall]_I$ are preserved by $\text{last}(\cdot)$ (by the standard results on preservation of modal formulas under bisimulations, cf. [BRV01]). \square

We naturally extend Corollary 5.32 to the fragment of the language:

⁶A functional bisimulation, see [BRV01].

Corollary 5.33. *Let \mathbf{S}_f be a pseudo-model for $\Box[\forall]_{i,A}$. The same formulas in $\Box[\forall]_{i,A}$ are satisfiable in the associated model \mathbf{X}_f , as in its p -morphic image in \mathbf{S}_f .*

Proof. The proof is obtained by restricting all the constructions in the proof of Corollary 5.32 to $\Box[\forall]_{i,A}$. This gives us a bisimulation between \mathbf{S}_f and \mathbf{X}_f . An argument following the same line of reasoning as Corollary 5.32 then concludes our proof. \square

To finish the proof of soundness, completeness, and decidability for the proof systems $\Box[\forall]_I$ and $\Box[\forall]_{i,A}$ with respect to the intended models, we finally prove Theorems 4.9 and 4.10, stated in Section 4.2. We first show that $\Box[\forall]_I$ from Definition 4.8 is sound and weakly complete with respect to multi-agent topo-e-models, and that the logic of $\Box[\forall]_I$ is decidable; the proof for $\Box[\forall]_{i,A}$ is similar.

Proof of Theorem 4.9. Soundness of the axioms and rules of classical propositional logic is a routine check, as well as the soundness of the inference rules (Modus Ponens, and Necessitation for both \Box_I and $[\forall]_I$ for all $I \subseteq A$); therefore, we omit these proofs. Soundness of the epistemic-doxastic axioms (the S4 axioms for each \Box_I , and the S5 axioms for each $[\forall]_I$) was proved in Lemma 4.2. As for the group knowledge axioms, soundness was proved in Proposition 4.5.

As for completeness, let $\varphi \in \Box[\forall]_I(V)$ be any consistent formula. By Corollary 5.14, there exists a pseudo-model $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi_S)$ for $\Box[\forall]_I$ and some state $s_0 \in S$, such that $(\mathbf{S}, s) \models \varphi$. Consider the associated model $\mathbf{X} = (H, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi_X)$ for \mathbf{S} , where H is given by the set of s_0 -generated histories in the pseudo-model \mathbf{S} .

By Corollary 5.32, the same formulas in $\Box[\forall]_I(V)$ are satisfied in the associated model \mathbf{X} as in its p -morphic image in \mathbf{S} . Note that s_0 is contained in the p -morphic image of \mathbf{X} in \mathbf{S} , since the sequence $h := (s_0)$ is an s_0 -generated history in H with $\text{last}(h) = s_0$. Therefore, φ is satisfied on \mathbf{X} .

This gives us weak completeness of $\Box[\forall]_I$ with respect to relational evidence models. By Corollary 3.40, we obtain weak completeness with respect to multi-agent topo-e-models. Decidability of the logic of $\Box[\forall]_I$ follows from the fact that it has the finite pseudo-model property (see Corollary 5.14). \square

We automatically obtain the same result for the proof system $\Box[\forall]_{i,A}$ for the fragment $\Box[\forall]_{i,A}$ with respect to topo-e-models.

Proof of Theorem 4.10. Soundness of $\Box[\forall]_{i,A}$ follows directly from soundness for the proof system $\Box[\forall]_I$ (Theorem 4.9), given that the axioms and rules of $\Box[\forall]_{i,A}$ are contained in $\Box[\forall]_I$.

The completeness proof for $\Box[\forall]_{i,A}$ follows the same line of reasoning as the proof of Theorem 4.9 for the larger language: let $\varphi \in \Box[\forall]_{i,A}(V)$ be any consistent formula. By Corollary 5.15, there exists a pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$ that satisfies φ at some state s_0 . By Corollary 5.33, there exists an associated model \mathbf{X} for \mathbf{S} , such that its state space H is defined by the s_0 -originated histories of \mathbf{S} , and therefore satisfies φ . The associated model is a relational evidence model, which gives us weak completeness for $\Box[\forall]_{i,A}$ with respect to relational evidence models. By Corollary 3.40, we obtain weak completeness for topo-e-models. Decidability of the logic of $\Box[\forall]_{i,A}$ follows from the fact that it has the finite pseudo-model property (see Corollary 5.15). \square

5.2 Completeness of $KB_{i,A}$

In this section we present the main completeness result in this thesis: a proof of completeness of the proof system $KB_{i,A}$, of knowledge and belief for individual agents and virtual group knowledge and

belief for the full group. The structure of our proof bears a superficial resemblance to the structure of the completeness proof presented in Section 5.1. As we discussed, the construction for $KB_{i,A}$ requires an extra step. We first define the relevant pseudo-models (Section 5.2.1): structures that explicitly represent the relations corresponding to knowledge and belief. We prove completeness with respect to these structures via the standard canonical-model construction (Proposition 5.38); we refer to [BRV01] for a detailed discussion of this construction.

The crucial step is the *representation theorem*⁷ (Theorem 5.44), which allows us to extend the completeness proof to the intended models. It states that each pseudo-model \mathbf{M} for $KB_{i,A}$ can alternatively be represented as a pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$, which agrees with \mathbf{M} on the interpretation of formulas over the language $KB_{i,A}$. This proof involves recovering relations for knowledge and belief on pseudo-models for $\Box[\forall]_{i,A}$ of evidence and, conversely, recovering evidence relations on a pseudo-model for $KB_{i,A}$ of knowledge and belief. The former construction is straightforward, since the relations for knowledge and belief are uniquely determined by the evidence relations. The latter is more complicated.

The correspondence between pseudo-models for $KB_{i,A}$ and for $\Box[\forall]_{i,A}$ allows us to extend the completeness proof of $KB_{i,A}$ to topo-e-models, via the previously proved correspondences from Section 5.1.

This section is structured as follows: we introduce the pseudo-models for $KB_{i,A}$ in Section 5.2.1; soundness and completeness with respect to these structures is proved in Section 5.2.2. Section 5.2.3 consists of two parts: we first define the map from pseudo-models for $\Box[\forall]_{i,A}$ to pseudo-models for $KB_{i,A}$. Next, we define the more complicated converse map, from which we derive completeness of the proof system $KB_{i,A}$ with respect to topo-e-models. Throughout this section, fix a finite group of agents A and a finite vocabulary V .

5.2.1 Pseudo-Models for $KB_{i,A}$

We first define the relevant pseudo-models, for which we initially prove completeness of the proof system $KB_{i,A}$. These models contain knowledge and belief relations for all individual agents $i \in A$, as well as the group relations for knowledge and belief for the group A .

For several relations R on our models, we will use the notion of R -maximal worlds. We define R -maximality as follows:

Definition 5.34 (R -maximal worlds). Given a set of states S and a relation R on S , define R -maximal worlds of S as $Max_R(S) := \{s \in S \mid \forall w \in S (sRw \Rightarrow wRs)\}$.

We use the notation α (first defined in Section 4.2.4) for conciseness. If the label of a given relation is $\alpha \in \{A\} \cup A$, it denotes either a singleton set $\{i\} \subseteq A$, or A itself.

Definition 5.35 (Pseudo-Model for $KB_{i,A}$). A *pseudo-model* for $KB_{i,A}$ is a structure $\mathbf{M} = (S, (\leq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \leq_A, \rightarrow_A, \pi)$, where S is a set of states and π is a valuation. A pseudo-model for $KB_{i,A}$ is required to satisfy the following conditions:

⁷*Representation theorems* exist in various fields of research. In abstract terms, a representation theorem is a structure-preserving mapping between two classes of objects or models. Given such a correspondence between the two classes, we can refer to any object in one class as a *representation* of some object in the other class. See e.g. [Özç19] for a general introduction into representation theorems.

1. **Stalnaker's conditions.** The knowledge and belief modalities \sqsubseteq_i and \rightarrow_i , for $i \in A$, and \sqsubseteq_A and \rightarrow_A , for the full group A , each satisfy the relational correspondents of Stalnaker's axioms (see [Sta06]). That is, for all $\alpha \in \{A\} \cup A$ we have:

- The \sqsubseteq_α (knowledge) relation is **S4**, i.e. \sqsubseteq_α is a preorder;
- The \rightarrow_α (belief) relation is **KD45**, i.e. \rightarrow_α is serial, transitive, and Euclidean;
- **Inclusion.** $\rightarrow_\alpha \subseteq \sqsubseteq_\alpha$;
- **Strong Transitivity.** For all $s, t, u \in S$, if $s \sqsubseteq_\alpha t$ and $t \rightarrow_\alpha u$, then $s \rightarrow_\alpha u$;
- **Strong Euclideanity.** For all $s, t, u \in S$, if $s \sqsubseteq_\alpha t$ and $s \rightarrow_\alpha u$, then $t \rightarrow_\alpha u$;
- **Full Belief.** For all $s, t, u \in S$, if $s \rightarrow_\alpha t$ and $t \sqsubseteq_\alpha u$, then $s \rightarrow_\alpha u$.

2. **KBK-Condition.** For all $i \in A$, $\sqsubseteq_A \subseteq (\sqsubseteq_i \cup \rightarrow_A)$;

3. **Group Knowledge of Individual Beliefs.** For all $s, t, u \in S$, if $s \sqsubseteq_A t$, then we have for all $i \in A$, that $s \rightarrow_i u$ if and only if $t \rightarrow_i u$;

4. **BDK-Condition.** For all $s \in S$ there exists $w \in S$ such that $s (\rightarrow_A \cap \bigcap_{i \in A} \sqsubseteq_i) w$.

It can be checked that Stalnaker's conditions imply that the knowledge relation on pseudo-models for $KB_{i,A}$ is weakly directed⁸ (see [Bal+13]), that is, the knowledge relation is S4.2.

The conditions correspond to the axioms in the proof system $KB_{i,A}$ from Definition 4.11. The KBK-condition corresponds to the (KBK) axiom: *Individual Knowledge and Group Belief imply Group Knowledge*. The group knowledge of individual beliefs condition corresponds to the (KPB) axiom (as well as the provable (KNB) validity): *Group Knowledge of Positive (and Negative) Individual Beliefs*. Finally, the BDK-condition corresponds to the (BDK) axiom: *Consistency of Group Belief with Distributed Knowledge*. This correspondence is made explicit in our soundness proof for $KB_{i,A}$ with respect to these models (Proposition 5.38) and in the proof that the canonical structure constructed in Definition A.5 is a pseudo-model for $KB_{i,A}$ (Proposition A.6).

This implies the following property for the belief relations on pseudo-model for $KB_{i,A}$: $\alpha \in \{A\} \cup A$ believes φ if and only if φ is true in the \sqsubseteq_α -maximal worlds within the current information cell (Lemma 5.36).

Lemma 5.36. *On a pseudo-model for $KB_{i,A}$ $\mathbf{M} = (S, (\sqsubseteq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \sqsubseteq_A, \rightarrow_A, \pi)$ we have for all $\alpha \in \{A\} \cup A$ and for all $s, w \in S$ that*

$$s \rightarrow_\alpha w \text{ iff } s \sqsubseteq_\alpha w \in \text{Max}_{\sqsubseteq_\alpha}(S).$$

Proof. Let $i \in A$. We show the proof for the individual relation \rightarrow_i . The proof for \rightarrow_A is symmetrical, as it refers only to Stalnaker's conditions on \rightarrow_A , which are analogous to Stalnaker's conditions on \rightarrow_i . For the left-to-right direction, fix $s, w \in S$ and suppose $s \rightarrow_i w$. Then the inclusion condition gives us that $s \sqsubseteq_i w$. To show that $w \in \text{Max}_{\sqsubseteq_i}(S)$, suppose that $w \sqsubseteq_i w'$. It suffices to show that $w' \sqsubseteq_i w$. Observe that $s \sqsubseteq_i w \sqsubseteq_i w'$ gives us $s \sqsubseteq_i w'$ (by transitivity of \sqsubseteq_i); now we have $s \rightarrow_i w$ and $s \sqsubseteq_i w'$ which, by strong Euclideanity, gives us $w' \rightarrow_i w$. But then, again by the inclusion condition, we obtain $w' \sqsubseteq_i w$, as required.

⁸A relation R on a relational frame $\mathcal{M} = (X, R)$ is *weakly directed* (also called *directed* or *confluent*) if we have for all $x, y, z \in X$ with xRy and xRz , that there exists $u \in X$ such that yRu and zRu .

For the right-to-left direction, fix $s, w \in S$ and suppose $s \sqsubseteq_i w \in \text{Max}_{\sqsubseteq_i}(S)$. We show $s \rightarrow_i w$. By seriality of \rightarrow_i , there exists $w' \in S$ such that $w \rightarrow_i w'$. Using the inclusion condition, we obtain $w \sqsubseteq_i w'$. But then, since $w \in \text{Max}_{\sqsubseteq_i}(S)$, we also have $w' \sqsubseteq_i w$. Now we have $s \sqsubseteq_i w \rightarrow_i w'$ which, by strong transitivity, gives us $s \rightarrow_i w'$. Finally, $s \rightarrow_i w' \sqsubseteq_i w$ gives us (by full belief) that $s \rightarrow_i w$, as required. \square

We interpret formulas over $KB_{i,A}$ on pseudo-models for $KB_{i,A}$ as follows:

Definition 5.37 (Pseudo-Model Semantics of $KB_{i,A}$). The topological semantics of $KB_{i,A}(V)$ on pseudo-model $\mathbf{M} = (S, (\sqsubseteq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \sqsubseteq_A, \rightarrow_A, \pi)$ for $KB_{i,A}$ is defined recursively as

$(\mathbf{M}, s) \models \top$	always holds
$(\mathbf{M}, s) \models p$	iff $p \in \pi(s)$
$(\mathbf{M}, s) \models \neg \varphi$	iff $(\mathbf{M}, s) \not\models \varphi$
$(\mathbf{M}, s) \models \varphi \wedge \psi$	iff $(\mathbf{M}, s) \models \varphi$ and $(\mathbf{M}, s) \models \psi$
$(\mathbf{M}, s) \models K_i \varphi$	iff for all $t \in S$ s.t. $s \sqsubseteq_i t : (\mathbf{M}, t) \models \varphi$
$(\mathbf{M}, s) \models B_i \varphi$	iff for all $t \in S$ s.t. $s \rightarrow_i t : (\mathbf{M}, t) \models \varphi$
$(\mathbf{M}, s) \models K_A \varphi$	iff for all $t \in S$ s.t. $s \sqsubseteq_A t : (\mathbf{M}, t) \models \varphi$
$(\mathbf{M}, s) \models B_A \varphi$	iff for all $t \in S$ s.t. $s \rightarrow_A t : (\mathbf{M}, t) \models \varphi$

where $s \in S$ is any state and $p \in V$ is any propositional variable.

5.2.2 Soundness and Completeness of $KB_{i,A}$ w.r.t. Pseudo-Models

In this section, we prove that the proof system $KB_{i,A}$ is sound and complete with respect to pseudo-models for $KB_{i,A}$. Proposition 5.38 takes care of the soundness proof.

Proposition 5.38. *The proof system $KB_{i,A}$ for $KB_{i,A}$ (Definition 4.11) is sound with respect to relational pseudo-models for $KB_{i,A}$.*

Proof. The proof is a routine check of the correspondences between the axioms of $KB_{i,A}$ and the properties of a pseudo-models for $KB_{i,A}$. It can be found in Appendix A.4. \square

Corollary 5.39 summarises our results:

Corollary 5.39. *The proof system $KB_{i,A}$ is sound and weakly complete with respect to pseudo-models for $KB_{i,A}$.*

Proof. We prove completeness with respect to pseudo-models for $KB_{i,A}$ by showing that every consistent formula $\varphi \in KB_{i,A}(V)$ is satisfiable in the canonical pseudo-model for $KB_{i,A}$ for $KB_{i,A}(V)$. The canonical pseudo-model is defined according to the standard "canonical model" construction (see e.g. [BRV01]); details of the proof can be found in Appendix A.4. \square

5.2.3 Back and Forth between Pseudo-Models for $KB_{i,A}$ and for $\Box[\forall]_I$

It now remains to represent the pseudo-models from Definition 5.35 as pseudo-models for $\Box[\forall]_{i,A}$, such that we can apply the results from Section 5.1.3: by unravelling of the structure represented as a pseudo-model for $\Box[\forall]_{i,A}$, we obtain a standard pseudo-model, i.e. a relational evidence model. The correspondence from Corollary 3.40 then gives us the desired result: completeness of $KB_{i,A}$ with respect to multi-agent topo-e-models.

We show both directions of the correspondence. The proof is structured as follows:

- (1) From pseudo-models for $\Box[\forall]_{i,A}$ to pseudo-models for $KB_{i,A}$. This is the straightforward direction of the proof. We recover the (uniquely determined) relations for knowledge and belief on a pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$, and show that the semantics in terms of these relations (from Definition 5.37), applied to \mathbf{S} , agrees with the original semantics on \mathbf{S} (from (Definition 5.6)), on all formulas over $KB_{i,A}$.
- (2) The crucial step: from pseudo-models for $KB_{i,A}$ to pseudo-models for $\Box[\forall]_{i,A}$. We present an approach to recover the evidence relations on a pseudo-model \mathbf{M} for $KB_{i,A}$ using the existing relations for knowledge and belief.
- (3) Representing \mathbf{M} as a pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$, consisting of the newly defined evidence relations, we use the approach from (1) to recover the corresponding knowledge and belief relations from these evidence relations.
- (4) Finally, we show that the newly recovered relations for knowledge and belief on \mathbf{S} coincide with the original relations for knowledge and belief on \mathbf{M} .
- (5) Using the result from (1), we conclude that on the two representations of \mathbf{M} , as a pseudo-model for $KB_{i,A}$ and as a pseudo-model for $\Box[\forall]_{i,A}$, the interpretations of knowledge and belief coincide.
- (6) We derive that the semantics for the two representations agree on all formulas in the language $KB_{i,A}$.

We first prove step (1) in Proposition 5.42. Steps (2)-(4) are proved in Theorem 5.44. Finally, the conclusion from step (6) is presented in Corollary 5.45. We conclude this section with Theorem 4.12, which states the desired result: $KB_{i,A}$ is sound and weakly complete with respect to multi-agent topo-e-models.

From pseudo-models for $\Box[\forall]_{i,A}$ to pseudo-models for $KB_{i,A}$. On pseudo-models for $\Box[\forall]_{i,A}$, knowledge and belief are interpreted indirectly in terms of relational evidence modalities (Definition 5.6). On pseudo-models for $KB_{i,A}$, on the other hand, knowledge and belief are interpreted directly in terms of the relations \sqsubseteq and \rightarrow , respectively (Definition 5.37). These relations can directly be recovered from the evidence relations by unfolding the interpretations of knowledge and belief as abbreviations. Proposition 5.42 shows how we recover these relations.

The map requires the \leq relations on the pseudo-model to have a particular property, which we refer to as *max-density*:

Definition 5.40 (Max-dense). Given a set of states S and a pre-order R on S , we say that R is *max-dense* if, for all $s \in S$, there exists $t \in \text{Max}_R(S)$, such that sRt . Equivalently, the pre-order R is max-dense if the set $\text{Max}_R(S)$ is dense in the up-set topology (see Definition 2.18) with respect to R .

Observation 5.41. Relations on finite models are automatically max-dense, by the absence of infinite R -chains. Thus, we can consider max-density as a generalisation of finiteness.

Proposition 5.42. Let $\mathbf{S} = (S, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \leq_A, \sim_A, \pi)$ be a pseudo-model for $\Box[\forall]_{i,A}$ such that for each $\alpha \in \{A\} \cup A$, the relation \leq_α is max-dense. Let $\alpha \in \{A\} \cup A$. If we set

$$\begin{aligned} s \rightarrow_\alpha^{\mathbf{S}} w & \quad \text{iff} \quad s \sim_\alpha w \in \text{Max}_{\leq_\alpha}(S) \\ s \leq_\alpha^{\mathbf{S}} w & \quad \text{iff} \quad s \leq_\alpha w \text{ or } s \rightarrow_\alpha^{\mathbf{S}} w \end{aligned}$$

then the following statements hold for \mathbf{S} :

1. the structure $\mathbf{M_S} = (S, (\leq_i^S)_{i \in A}, (\rightarrow_i^S)_{i \in A}, \leq_A^S, \rightarrow_A^S, \pi)$ is a pseudo-model for $KB_{i,A}$;
2. we have for all $s \in S$ and for all formulas φ over the language $KB_{i,A}$:

- (a) $(\mathbf{S}, s) \models B_\alpha \varphi$ iff for all $t \in S$ s.t. $s \rightarrow_\alpha^S t : (\mathbf{M_S}, t) \models \varphi$
- (b) $(\mathbf{S}, s) \models K_\alpha \varphi$ iff for all $t \in S$ s.t. $s \leq_\alpha^S t : (\mathbf{M_S}, t) \models \varphi$.

Proof. The interpretations of the modalities and the definitions of the corresponding relations for the full group A are analogous to those for individual agents, therefore we only show the cases for individual agents in both proofs.

For (1), we show that $\mathbf{M_S}$ satisfies the conditions of a pseudo-model for $KB_{i,A}$ (Definition 5.35). Fix an agent $i \in A$.

- Stalnaker's conditions. First, \leq_i^S is a pre-order. For reflexivity, observe that $s \leq_i^S s$ follows from the fact that $s \leq_i s$ (by definition of \leq_i in Definition 5.1). For transitivity, suppose $s \leq_i^S w \leq_i^S v$. Applying the definition of \leq_i^S , we have one of the following four cases:
 - (a) $s \leq_i w \leq_i v$. Then $s \leq_i v$ (by definition of \leq_i in Definition 5.1), so $s \leq_i^S v$.
 - (b) $s \leq_i w \rightarrow_i^S v$. So $w \sim_i v$ with $v \in \text{Max}_{\leq_i}(S)$. By the inclusion condition on \mathbf{S} , we have $s \sim_i w$, so with \sim_i being an equivalence relation $s \sim_i v$. But then $s \rightarrow_i^S v$, so $s \leq_i^S v$.
 - (c) $s \rightarrow_i^S w \rightarrow_i^S v$. Then $s \sim_i w$ and $w \sim_i v$, with $v \in \text{Max}_{\leq_i}(S)$, so $s \rightarrow_i^S v$, and therefore, $s \leq_i^S v$.
 - (d) $s \rightarrow_i^S w \leq_i v$. Then $s \sim_i w$ with $w \in \text{Max}_{\leq_i}(S)$, so with $w \leq_i v$, it must be that $v \in \text{Max}_{\leq_i}(S)$. With $w \leq_i v$, we have by the inclusion condition on \mathbf{S} that $w \sim_i v$, so by \sim_i being an equivalence relation, we have $s \sim_i v$. So $s \leq_i^S v$.

Next, we show that \rightarrow_i^S is serial, transitive, and Euclidean. For seriality, let $s \in S$. Note that \leq_i is max-dense. Thus, there is $t \in \text{Max}_{\leq_i}(S)$ such that $s \leq_i t$. By the inclusion condition on \mathbf{S} , we obtain $s \sim_i t$ which, by the fact that $t \in \text{Max}_{\leq_i}(S)$, gives us that $s \rightarrow_i^S t$. For transitivity, see item (c) above. For Euclideanity, let $s \rightarrow_i^S w$ and $s \rightarrow_i^S v$. Then $s \sim_i w$ and $s \sim_i v$, with both $w \in \text{Max}_{\leq_i}(S)$ and $v \in \text{Max}_{\leq_i}(S)$. As \sim_i is an equivalence relation, we have $w \sim_i v$, giving us $w \rightarrow_i^S v$.

Inclusion. Suppose $s \rightarrow_i^S w$. Then, by definition, $s \leq_i^S w$.

Strong transitivity. See case (b) above.

Strong Euclideanity. Suppose $s \leq_i^S w$ and $s \rightarrow_i^S v$. We claim that $w \rightarrow_i^S v$. Given the assumption that $s \leq_i^S w$, there are two possible cases: either (1) $s \leq_i w$, or (2) $s \rightarrow_i^S w$. In either case, we have $s \sim_i w$: in the case of (1), it follows from inclusion on \mathbf{S} ; in the case of (2), it follows from the definition of \rightarrow_i^S . Since \sim_i is an equivalence relation, we have $w \sim_i v$. By definition of \rightarrow_i^S , we get $w \rightarrow_i^S v$, as required.

Full belief. Suppose $s \rightarrow_i^S w$ and $w \leq_i^S v$. We claim that $s \rightarrow_i^S v$. By definition of \rightarrow_i^S , we have $s \sim_i w$ with $w \in \text{Max}_{\leq_i}(S)$. By $w \leq_i^S v$, it must be that $v \in \text{Max}_{\leq_i}(S)$. We have $w \sim_i v$ by definition of \rightarrow_i^S , which gives us that $w \rightarrow_i^S v$, as required.

- KBK-condition. We show that $\leq_A^S \subseteq (\leq_i^S \cup \rightarrow_A^S)$. Let $s \leq_A^S w$; we show that $s(\leq_i^S \cup \rightarrow_A^S)w$. The assumption gives us two possible cases: either (1) $s \leq_A w$, or (2) $s \rightarrow_A^S w$. In the case of (1), anti-monotonicity of \leq gives us $s \leq_i w$ so, by definition, $s \leq_i^S w$, and therefore, $s(\leq_i^S \cup \rightarrow_A^S)w$. In the case of (2), the claim is immediate from $s \rightarrow_A^S w$.
- Group knowledge of individual beliefs. Suppose $s \leq_A^S t$. We show that $s \rightarrow_i^S u$ if and only if $t \rightarrow_i^S u$. We show one direction; the converse direction is symmetrical. Suppose $s \rightarrow_i^S u$. Then $s \sim_i u$ with $u \in \text{Max}_{\leq_i}(S)$. By the inclusion condition on \mathbf{S} , the assumption $s \leq_A^S t$ gives us that $s \sim_i t$. Since \sim_i is an equivalence relation, $t \sim_i u$. With $u \in \text{Max}_{\leq_i}(S)$, we have $t \rightarrow_i^S u$, as required.
- BDK-condition. Let $s \in S$. We show that there exists $w \in S$ such that $s(\rightarrow_A^S \cap \leq_i^S)w$ (for our fixed, arbitrary $i \in A$). We use max-density of \leq_A : there exists $w \in \text{Max}_{\leq_A}(S)$ such that $s \leq_A w$. By the inclusion condition on \mathbf{S} , we have $s \sim_A w$. Thus, $s \rightarrow_A^S w$. To see that we also have $s \leq_i^S w$, note that with $s \leq_A w$, anti-monotonicity of \leq gives us that $s \leq_i w$. Thereby, $s \leq_i^S w$, as required.

For the proofs of (2), let $s \in S$ and let $\varphi \in KB_{i,A}$. Fix an agent $i \in A$.

- (a) $(\mathbf{S}, s) \models B_i \varphi$ iff for all $t \in S$ s.t. $s \rightarrow_i^S t : (\mathbf{M}_S, t) \models \varphi$.

Unfolding the semantic definition of B_i on pseudo-models for $\Box[\forall]_{i,A}$ (Definition 5.6) in terms of the evidence relations, we obtain the following interpretation for B_i , which we will use:

$$(\mathbf{S}, s) \models B_i \varphi \quad \text{iff} \quad \forall t \in S : s \sim_i t \Rightarrow (\exists u \in S(t \leq_i u \text{ and } \forall v \in S : u \leq_i v \Rightarrow (\mathbf{S}, v) \models \varphi)).$$

For the left-to-right direction, suppose $(\mathbf{S}, s) \models B_i \varphi$. We need to show that for all $t \in S$ with $s \rightarrow_i^S t$, we have $(\mathbf{M}_S, t) \models \varphi$. So let $t \in S$ such that $s \rightarrow_i^S t$. By definition of \rightarrow_i^S , we have $s \sim_i t$ with $t \in \text{Max}_{\leq_i}(S)$. So by the unfolding of $B_i \varphi$, there exists $u \in S$ such that $t \leq_i u$ and for all $v \in S, u \leq_i v$ implies $(\mathbf{S}, v) \models \varphi$. Furthermore, by \leq_i -maximality of t , we have that $t \leq_i u$ implies $u \leq_i t$. So by $u \leq_i t$, we have $(\mathbf{S}, t) \models \varphi$, i.e., $(\mathbf{M}_S, t) \models \varphi$, as required.

For the converse direction, suppose that $s \rightarrow_i^S t$ implies that $(\mathbf{M}_S, t) \models \varphi$ (i.e. $(\mathbf{S}, t) \models \varphi$), for all $t \in S$. Let $w \in S$ with $s \sim_i w$. We want to find $u \in S$ such that $w \leq_i u$ and for all $v \in S, u \leq_i v$ implies $(\mathbf{S}, v) \models \varphi$.

By max-density of \leq_i , there exists $u \in \text{Max}_{\leq_i}(S)$ such that $w \leq_i u$. By the inclusion condition on \mathbf{S} , we have $w \rightarrow_i^S u$. If we prove that for all $v \in S, u \leq_i v$ implies $(\mathbf{S}, v) \models \varphi$, then we are done. So let $v \in S$ and suppose $u \leq_i v$. As a property of \leq_i -maximality, it must be that $v \in \text{Max}_{\leq_i}(S)$. Now we claim that $s \rightarrow_i^S v$: we have a chain $s \sim_i w \leq_i u \leq_i v$, so by the inclusion condition on \mathbf{S} and by the properties of \sim_i , we have $s \sim_i v$. Therefore, $s \rightarrow_i^S v$. But then, by assumption, $(\mathbf{S}, v) \models \varphi$, as required.

- (b) $(\mathbf{S}, s) \models K_i \varphi$ iff for all $t \in S$ s.t. $s \leq_i^S t : (\mathbf{M}_S, t) \models \varphi$.

For the left-to-right direction, suppose $(\mathbf{S}, s) \models K_i \varphi$. Unfolding the interpretation of K_i , we have that $(\mathbf{S}, s) \models \Box_i \varphi \wedge B_i \varphi$. To prove the claim, let $t \in S$ and suppose that $s \leq_i^S t$. We have two possible cases: either (1) $s \leq_i t$ or (2) $s \rightarrow_i^S t$. In either case we have $(\mathbf{S}, t) \models \varphi$:

- (1) If $s \leq_i t$, then by $(\mathbf{S}, s) \models \Box_i \varphi$, we get that $(\mathbf{S}, t) \models \varphi$, i.e. $(\mathbf{M}_S, t) \models \varphi$.
- (2) If $s \rightarrow_i^S t$, then with $(\mathbf{S}, s) \models B_i \varphi$, (a) gives us that $(\mathbf{M}_S, t) \models \varphi$.

For the converse direction, suppose that for all $t \in S$, $s \leq_i^S t$ implies that $(\mathbf{M}_S, t) \models \varphi$. It suffices to show that $(\mathbf{S}, s) \models \Box_i \varphi \wedge B_i \varphi$. To see that $(\mathbf{S}, s) \models \Box_i \varphi$, let $t \in S$ such that $s \leq_i t$. We show that $(\mathbf{S}, t) \models \varphi$. By definition of \leq_i^S , $s \leq_i t$ gives us that $s \leq_i^S t$. By assumption, this implies that $(\mathbf{M}_S, t) \models \varphi$, as required.

Next, to see that $(\mathbf{S}, s) \models B_i \varphi$, recall from (a) that it suffices to show that for all $t \in S$, $s \rightarrow_i^S t$ implies $(\mathbf{M}_S, t) \models \varphi$. So let $t \in S$ and suppose $s \rightarrow_i^S t$. Then, by definition of \leq_i^S , we have $s \leq_i^S t$. The claim then follows directly from our assumption that $s \leq_i^S t$ implies that $(\mathbf{M}_S, t) \models \varphi$.

□

Thus, assuming a pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$ with max-dense \leq relations, we can recover the relations corresponding to knowledge and belief to obtain a pseudo-model \mathbf{M}_S for $KB_{i,A}$. The resulting model agrees with \mathbf{S} on the interpretation of knowledge and belief.

From pseudo-models for $KB_{i,A}$ to pseudo-models for $\Box[\forall]_{i,A}$. For the converse direction, we recover evidence relations on pseudo-models for $KB_{i,A}$. The representation theorem 5.44 constructs the desired relations. It uses the following lemma, which gives us equivalent definitions of the \sim relations that we will define on \mathbf{M} .

Lemma 5.43. *Let $s, w, t \in S$ and let $\alpha \in \{A\} \cup A$. Let $\mathbf{M} := (S, (\leq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \leq_A, \rightarrow_A, \pi)$ be a pseudo-model for $KB_{i,A}$. Then the following are equivalent on \mathbf{M} :*

$$\begin{aligned} \exists t(s \leq_\alpha t \text{ and } w \leq_\alpha t) & \quad \text{iff} \quad \exists t(s \rightarrow_\alpha t \text{ and } w \rightarrow_\alpha t) \\ & \quad \text{iff} \quad \forall t(s \rightarrow_\alpha t \text{ iff } w \rightarrow_\alpha t). \end{aligned}$$

Proof. We prove the following chain of implications: given (1) $\exists t(s \leq_\alpha t, w \leq_\alpha t)$, (2) $\exists t(s \rightarrow_\alpha t, w \rightarrow_\alpha t)$, and (3) $\forall t(s \rightarrow_\alpha t \text{ iff } w \rightarrow_\alpha t)$, we show that (1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (1).

(1) implies (2). Suppose (1). By seriality of \rightarrow_i , there exists $t_1 \in S$ such that $s \rightarrow_i t_1$. By the Strong Euclidean condition on KB-pseudo-models, we have $t \rightarrow_i t_1$. Now, applying the Strong Transitivity condition to $w \leq_i t \rightarrow_i t_1$, we have $w \rightarrow_i t_1$. Thus, there is t_1 such that $s \rightarrow_i t_1, w \rightarrow_i t_1$.

(2) implies (3). Suppose (2). Let $u \in S$ be arbitrary and assume, without loss of generality, that $s \rightarrow_i t$. Then $w \rightarrow_i t$ follows immediately from Euclidean and transitivity of \rightarrow_i .

(3) implies (1). Suppose (3). Then, by seriality of \rightarrow_i , there is t such that $s \rightarrow_i t$ if and only if $w \rightarrow_i t$. The claim then follows directly from the Inclusion condition on KB-pseudo-models. □

We can now prove the main completeness result: Theorem 5.44. It is important to note that, in contrast to the relations in the construction of Proposition 5.42, the evidence relations defined in Theorem 5.44 are *not* uniquely determined⁹.

Theorem 5.44 (Representing Pseudo-Models for $KB_{i,A}$ as Pseudo-Models for $\Box[\forall]_{i,A}$). *Let $\mathbf{M} := (S, (\leq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \leq_A, \rightarrow_A, \pi)$ be a pseudo-model for $KB_{i,A}$. We introduce the following relations on \mathbf{M} , for*

⁹In particular, an alternative, weaker condition for the relation \leq_A^M on a pseudo-model \mathbf{M} (for $KB_{i,A}$) replaces condition (2) in Theorem 5.44 with the following condition: (2') if $s \in \text{Max}_{\leq_A}(S)$, then $w (\bigcap_{i \in A} \leq_i) s$. The resulting alternative definition of relations also results in a max-dense pseudo-model for $\Box[\forall]_{i,A}$. Nevertheless, we chose condition (2), as it simplifies the proof of Theorem 5.44.

all $i \in A$:

$$\begin{aligned}
s \sim_i^{\mathbf{M}} w &\text{ iff } \exists t (s \sqsubseteq_i t \text{ and } w \sqsubseteq_i t) \\
s \sim_A^{\mathbf{M}} w &\text{ iff } \exists t (s \sqsubseteq_A t \text{ and } w \sqsubseteq_A t) \\
s \leq_i^{\mathbf{M}} w &\text{ iff } s \sqsubseteq_i w \\
s \leq_A^{\mathbf{M}} w &\text{ iff } \begin{cases} (1) & s (\sqsubseteq_A \cap \bigcap_{i \in A} \sqsubseteq_i) w; \\ (2) & \text{if } s \in \text{Max}_{\sqsubseteq_A}(S), \text{ then } w = s. \end{cases}
\end{aligned}$$

On \mathbf{M} , have for all $\alpha \in \{A\} \cup A$ that

- a. $\leq_\alpha^{\mathbf{M}} \subseteq \sqsubseteq_\alpha \subseteq \sim_\alpha^{\mathbf{M}}$;
- b. $\leq_\alpha^{\mathbf{M}}$ is a preorder and $\sim_\alpha^{\mathbf{M}}$ is an equivalence relation;
- c. $\text{Max}_{\leq_\alpha^{\mathbf{M}}}(S) = \text{Max}_{\sqsubseteq_\alpha}(S)$;
- d. $s \rightarrow_\alpha w$ if and only if $s \sim_\alpha^{\mathbf{M}} w \in \text{Max}_{\leq_\alpha^{\mathbf{M}}}(S)$;
- e. $s \sqsubseteq_\alpha w$ if and only if $(s \leq_\alpha^{\mathbf{M}} w \text{ or } s \rightarrow_\alpha w)$.

The following statements hold for the structure $\mathbf{S}_{\mathbf{M}} := (S, (\leq_i^{\mathbf{M}})_{i \in A}, (\sim_i^{\mathbf{M}})_{i \in A}, \leq_A^{\mathbf{M}}, \sim_A^{\mathbf{M}}, \pi)$:

- (I) $\mathbf{S}_{\mathbf{M}}$ is a pseudo-model for $\Box[\forall]_{i,A}$, such that for each $\alpha \in \{A\} \cup A$, the relation \leq_α is max-dense.
- (II) The pseudo-model \mathbf{M} for $KB_{i,A}$ and the pseudo-model $\mathbf{S}_{\mathbf{M}}$ for $\Box[\forall]_{i,A}$ agree on the interpretation of the modalities K_i, K_A, B_i , and B_A .

Proof. We first prove statements (a)-(e) for \mathbf{M} .

- (a) Let $i \in A$. We prove $\leq_i^{\mathbf{M}} \subseteq \sqsubseteq_i \subseteq \sim_i^{\mathbf{M}}$. For $\leq_i \subseteq \sqsubseteq_i$, let $s \leq_i^{\mathbf{M}} w$. Then $s \sqsubseteq_i w$ by definition. Now for $\sqsubseteq_i \subseteq \sim_i^{\mathbf{M}}$, let $s \sqsubseteq_i w$. For $s \sim_i^{\mathbf{M}} w$, by Lemma 5.43 it suffices to show that there exists $t \in S$ such that $s \sqsubseteq_i t, w \sqsubseteq_i t$. By seriality of \rightarrow_i , there is $t \in S$ such that $w \rightarrow_i t$. By the inclusion condition on pseudo-models for $KB_{i,A}$ we obtain $w \sqsubseteq_i t$, so that we have $s \sqsubseteq_i w \sqsubseteq_i t$. By transitivity of \sqsubseteq_i we get $s \sqsubseteq_i t$. But then we have $t \in S$ such that $w \sqsubseteq_i t$ and $s \sqsubseteq_i t$, as required.

For the full group, we prove $\leq_A^{\mathbf{M}} \subseteq \sqsubseteq_A \subseteq \sim_A^{\mathbf{M}}$. For $\leq_A^{\mathbf{M}} \subseteq \sqsubseteq_A$, let $s \leq_A^{\mathbf{M}} w$. Then we obtain $s \sqsubseteq_A w$ from (1) of the definition of $\leq_A^{\mathbf{M}}$. Next, for $\sqsubseteq_A \subseteq \sim_A^{\mathbf{M}}$, let $s \sqsubseteq_A w$. By Lemma 5.43, it suffices to show that there exists $t \in S$ such that $s \rightarrow_A t, w \rightarrow_A t$. By seriality of \rightarrow_A , there is $t \in S$ such that $w \rightarrow_A t$. Having $s \rightarrow_A t$ and $s \sqsubseteq_A w$, strong Euclideanity on pseudo-models for $KB_{i,A}$ gives us that $w \rightarrow_A t$, which concludes our proof.

- (b) Let $i \in A$. We show that $\leq_i^{\mathbf{M}}$ is a preorder. Since \sqsubseteq_i is a preorder, we have that $\leq_i^{\mathbf{M}}$ is a pre-order by definition.

For the full group, we show that $\leq_A^{\mathbf{M}}$ is a preorder. For reflexivity of $\leq_A^{\mathbf{M}}$, we show that conditions (1) and (2) of the definition of $\leq_A^{\mathbf{M}}$ hold for s with respect to s itself. For (1), observe that we have $s(\sqsubseteq_A \cap \bigcap_{i \in A} \sqsubseteq_i)s$: since \sqsubseteq_A and \sqsubseteq_i for $i \in A$ are preorders, we have $s \sqsubseteq_A s$ and $s \sqsubseteq_i s$ for all $i \in A$. (2) follows immediately from the fact that $s = s$. For transitivity of $\leq_A^{\mathbf{M}}$, suppose that $s \leq_A^{\mathbf{M}} w \leq_A^{\mathbf{M}} t$. We show that $s \leq_A^{\mathbf{M}} t$. Condition (1) follows directly from \sqsubseteq_A and all \sqsubseteq_i being transitive: we have $s \sqsubseteq_A w \sqsubseteq_A t$ and $s \sqsubseteq_i w \sqsubseteq_i t$ by assumption, which gives us $s \sqsubseteq_A t$

and $s \sqsubseteq_i t$ for $i \in A$, by transitivity of the respective relations. Thus, $s(\sqsubseteq_A \cap \bigcap_{i \in A} \sqsubseteq_i)t$. For condition (2), suppose $s \in \text{Max}_{\sqsubseteq_A}(S)$. Since we assumed $s \leq_A^M w$, we know that $w = s$. But then $w = s \in \text{Max}_{\sqsubseteq_A}(S)$. With $w \leq_A^M t$, we obtain that $t = s$, as required. So $s \leq_A^M t$, giving us transitivity for \leq_A^M .

Let $i \in A$. We show that \sim_i^M is an equivalence relation. For reflexivity, by definition of \sim_i^M , to obtain $s \sim_i^M s$ it suffices to show that there exists t such that $s \sqsubseteq_i t$. By reflexivity of \sqsubseteq_i , we have $s \sqsubseteq_i s$ and we are done. For transitivity of \sim_i^M , let $s \sim_i^M w \sim_i^M t$. So by definition of \sim_i^M , there are $u_1, u_2 \in S$ such that $s \sqsubseteq_i u_1$ and $w \sqsubseteq_i u_1$, and $w \sqsubseteq_i u_2$ and $t \sqsubseteq_i u_2$. To show $s \sim_i^M t$, we need u_3 such that $s \sim_i^M u_3$ and $t \sim_i^M u_3$. This state exists, because \sqsubseteq_i is weakly directed: by $w \sqsubseteq_i u_1$, and $w \sqsubseteq_i u_2$, there must be u_3 such that $u_1 \sqsubseteq_i u_3$ and $u_2 \sqsubseteq_i u_3$. Now we apply transitivity of \sqsubseteq_A to the fact that $s \sqsubseteq_i u_1 \sqsubseteq_i u_3$ and $t \sqsubseteq_i u_2 \sqsubseteq_i u_3$, and we are done. Finally, for symmetry of \sim_i^M , let $s \sim_i^M w$. By the symmetric nature of the definition of \sim_i^M , we automatically obtain $w \sim_i^M s$.

For the full group, \sim_A^M is an equivalence relation: the proofs for \sim_A^M are analogous to those for \sim_i^M , replacing each occurrence of \sqsubseteq_i with \sqsubseteq_A .

- (c) Let $i \in A$. Then $\text{Max}_{\leq_i^M}(S) = \text{Max}_{\sqsubseteq_i}(S)$ follows immediately from the definition of \leq_i^M on pseudo-models for $KB_{i,A}$: we have $s \leq_i^M t$ if and only if $s \sqsubseteq_i t$, for all $s, t \in S$.

For the full group, we show that $\text{Max}_{\leq_A^M}(S) = \text{Max}_{\sqsubseteq_A}(S)$: for the left-to-right direction, let $s \in \text{Max}_{\leq_A^M}(S)$ and suppose for contradiction that $s \notin \text{Max}_{\sqsubseteq_A}(S)$. By the BDK-condition on pseudo-models for $KB_{i,A}$, there exists $t \in S$ such that $s \rightarrow t$ and $s(\bigcap_{i \in A} \sqsubseteq_i)t$. By the inclusion condition, $s \rightarrow t$ gives us that $s \sqsubseteq t$. So with $s(\bigcap_{i \in A} \sqsubseteq_i)t$, condition (1) for $s \leq_A^M t$ is satisfied. Furthermore, since we assumed that $s \notin \text{Max}_{\sqsubseteq_A}(S)$, condition (2) holds trivially. Thus, we have $s \leq_A^M t$. Now, since we assumed that $s \in \text{Max}_{\leq_A^M}(S)$, we get $t \leq_A^M s$. By definition of \leq_A^M , we have $t \sqsubseteq s$. However, recall that we also have $s \rightarrow_A t$. Strong transitivity on pseudo-models for $KB_{i,A}$ then gives us that $t \rightarrow_A t$, so by Lemma 5.36, $t \in \text{Max}_{\sqsubseteq_A}(S)$. With $t \sqsubseteq_A s$, it must be that $s \in \text{Max}_{\sqsubseteq_A}(S)$: we have reached a contradiction, and we conclude that $s \in \text{Max}_{\sqsubseteq_A}(S)$ after all.

For the converse direction, let $s \in \text{Max}_{\sqsubseteq_A}(S)$ and let $t \in S$ such that $s \leq_A^M t$. We show that $t \leq_A^M s$. Since $s \in \text{Max}_{\sqsubseteq_A}(S)$, we get by $s \leq_A^M t$, and by definition of \leq_A^M , that $t = s$. Thus, it remains to show that $t \leq_A^M t$. Condition (1) of the definition of \leq_A^M follows from reflexivity of \sqsubseteq_A and \sqsubseteq_i for all $i \in A$; condition (2) is trivially true, since $t = t$. Therefore, $t \leq_A^M t = s$ and so, $s \in \text{Max}_{\leq_A^M}(S)$.

- (d) Let $i \in A$; we show that $s \rightarrow_i w$ if and only if $s \sim_i^M w$ and $w \in \text{Max}_{\leq_i^M}(S)$. For the left-to-right direction, let $s \rightarrow_i w$. For $s \sim_i^M w$, we know by the assumption that $s \rightarrow_i w$ and by Lemma 5.36, that $w \in \text{Max}_{\sqsubseteq_i}(S)$. So (by reflexivity of \sqsubseteq_i and, again by Lemma 5.36), we get $w \rightarrow_i w$. But then, with $s \rightarrow_i w$ and $w \rightarrow_i w$, we have $s \sim_i^M w$. It remains to show that $w \in \text{Max}_{\leq_i^M}(S)$: recall that $w \in \text{Max}_{\sqsubseteq_i}(S)$, so by our proof of (c) we know that $w \in \text{Max}_{\leq_i^M}(S)$.

For the converse direction, let $s \sim_i^M w$ with $w \in \text{Max}_{\leq_i^M}(S)$. For $s \rightarrow_i w$, by Lemma 5.36 it suffices to show that $s \sqsubseteq_i w$ and $w \in \text{Max}_{\sqsubseteq_i}(S)$. By our proof of (c) we know that $w \in \text{Max}_{\sqsubseteq_i}(S)$, so it remains to show that $s \sqsubseteq_i w$. By definition of \sim_i^M , there exists $t \in S$ such that $s \sqsubseteq_i t$ and $w \sqsubseteq_i t$. But then we must have $t \sqsubseteq_i w$. So transitivity of \sqsubseteq_i gives us $s \sqsubseteq_i w$, as required.

For the full group, the proof that $s \rightarrow_A w$ if and only if $s \sim_A^M w \in \text{Max}_{\leq_A^M}(S)$ is analogous to the proof for individual agents $i \in A$, by replacing each individual relation with its counterpart for the full group.

- (e) Let $i \in A$; we show that $s \sqsubseteq_i w$ if and only if $(s \leq_i^M w \text{ or } s \rightarrow_i w)$. For the left-to-right direction, $s \sqsubseteq_i w$ implies (by definition of \leq_i^M) that $s \leq_i^M w$, as required. For the converse direction, we make a case distinction. If $s \leq_i^M w$, then we obtain $s \sqsubseteq_i w$ by definition of \leq_i^M ; on the other hand, if $s \rightarrow_i w$, then we have by Lemma 5.36 that $s \sqsubseteq_i w$, as required.

For the full group, we show that $s \sqsubseteq_A w$ if and only if $(s \leq_A^M w \text{ or } s \rightarrow_A w)$: for the left-to-right direction, let $s \sqsubseteq_A w$. By the KBK-condition on pseudo-models for $KB_{i,A}$, we have for all $i \in A$ that $s(\sqsubseteq_i \cup \rightarrow_A)w$. If $s \rightarrow_A w$, we are done. So suppose not. Then $s \sqsubseteq_i w$ for all $i \in A$. We claim that this implies $s \leq_A^M w$. For condition (1) of the definition of \leq_A^M on pseudo-models for $KB_{i,A}$, observe that we have $s \sqsubseteq_A w$ by assumption, and $s(\bigcap_{i \in A} \sqsubseteq_i)w$ by the KBK-condition on for $KB_{i,A}$ (and by the assumption that $s \not\rightarrow_A w$). Thus, $s(\sqsubseteq_A \cap \bigcap_{i \in A} \sqsubseteq_i)w$. For condition (2), suppose that $s \in \text{Max}_{\sqsubseteq_A}(S)$. Then by $s \sqsubseteq_A w$ we obtain that $w \in \text{Max}_{\sqsubseteq_A}(S)$. With $s \sqsubseteq_A w$, this means that $s \rightarrow_A w$ (Lemma 5.36), which we assumed was not the case: a contradiction. Therefore, s cannot be contained in $\text{Max}_{\sqsubseteq_A}(S)$ and condition (2) is vacuously true. We conclude that $s \leq_A^M w$, as required.

It remains to prove statements (I) and (II).

- (I) First, to see that \mathbf{S}_M is a pseudo-model for $\Box[\forall]_{i,A}$, we check the following conditions (Definition 5.5) on the model \mathbf{M}^{10} .

Relations \leq_A^M and \leq_i^M for $i \in A$ are preorders, and \sim_A^M and \sim_i^M for $i \in A$ are equivalence relations. This is stated and proved in (b).

The relations \leq^M are anti-monotone: suppose that $s \leq_A^M t$ and let $i \in A$. By definition of \leq_A^M , $s \leq_A^M t$ implies $s \sqsubseteq_i t$. The definition of \leq_i^M then gives us $s \leq_i^M t$, as required.

The relations \sim^M are anti-monotone: suppose that $s \sim_A^M t$ and let $i \in A$. We show that $s \sim_i^M t$. By definition of \sim_i^M , there exists $u \in S$ such that $s \sqsubseteq u$ and $t \sqsubseteq u$. Now by seriality of \rightarrow_i , there exists $v_1 \in S$ such that $s \rightarrow_i v_1$. With $s \sqsubseteq u$ and $s \rightarrow_i v_1$, the condition of group knowledge of individual beliefs on pseudo-models for $KB_{i,A}$ then gives us that $u \rightarrow_i v_1$. By the same condition, applied to $t \sqsubseteq u$ and $u \rightarrow_i v_1$, we have $t \rightarrow_i v_1$. But then there exists v_1 such that both $s \rightarrow_i v_1$ and $u \rightarrow_i v_1$, which gives us $s \sim_i^M u$ by the definition of \sim_i^M .

Inclusion: we have $\leq_i^M \subseteq \sim_i^M$ for all $i \in A$, and $\leq_A^M \subseteq \sim_A^M$. This is stated and proved in (a).

It remains to show that the \leq^M relations are max-dense. For individual agents, let $i \in A$ and let $s \in S$. We need to find some $t \in \text{Max}_{\leq_i^M}(S)$, such that $s \leq_i^M t$. By seriality of \rightarrow_i , there exists $t \in S$ such that $s \rightarrow_i t$. Now, by (d), $t \in \text{Max}_{\leq_i^M}(S)$. The inclusion condition on \mathbf{M} gives us, by $s \rightarrow_i t$, that $s \sqsubseteq_i t$. By definition of \leq_i^M , we have $s \leq_i^M t$, as required.

For the full group, we find some $t \in \text{Max}_{\leq_A^M}(S)$, such that $s \leq_A^M t$. We consider two cases:

- (a) $s \in \text{Max}_{\leq_A^M}(S)$. We take $t := s$ and show that conditions (1) and (2) are satisfied for $s \leq_A^M s$: condition (1) follows from reflexivity of \sqsubseteq_A , as well as all \sqsubseteq_i relations. Condition (2) is trivially true, by the fact that $s = s$. Therefore, $s \leq_A^M s$, as required.
- (b) $s \notin \text{Max}_{\leq_A^M}(S)$. By the BDK-condition on \mathbf{M} , there exists $w \in S$ such that $s \rightarrow_A w$ and for all $i \in A$, $s \sqsubseteq_i w$. We take $t := w$. By (d), we have $w \in \text{Max}_{\leq_A^M}(S)$. It remains to show that

¹⁰The only difference between \mathbf{S}_M and \mathbf{M} is that \mathbf{S}_M does not contain the relations for knowledge and belief from \mathbf{M} , and we need these relations to prove our claims.

conditions (1) and (2) are satisfied for $s \leq_A^M w$. For (1), note that $s \rightarrow_A w$ implies $s \sqsubseteq_A w$ (by the inclusion condition on \mathbf{M}), and recall that for all $i \in A$, $s \sqsubseteq_i w$. Now, (2) is vacuously satisfied, as $s \notin \text{Max}_{\sqsubseteq_A}(S)$. Therefore, $s \leq_A^M w$, as required.

In conclusion, \mathbf{S}_M is a pseudo-model for $\Box[\forall]_{i,A}$, such that all \leq^M relations are max-dense.

- (II) To show that the pseudo-model \mathbf{M} for $KB_{i,A}$ and the pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$ agree on the interpretation of the modalities K_i , B_i , K_A , and B_A , we prove that the primitive knowledge and belief relations on \mathbf{M} coincide with the recovered knowledge and belief relations on \mathbf{S} .

First, recall the result from Proposition 5.42, which states that for all $\alpha \in \{A\} \cup A$, we can recover relations \rightarrow_α^S and \sqsubseteq_α^S for knowledge and belief on the pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$, given that the \leq relations for all agents and the full group are max-dense.

Now consider our pseudo-model \mathbf{S}_M , which was obtained from \mathbf{M} by recovering the evidence relations \sim_α^M and \leq_α^M for all $\alpha \in \{A\} \cup A$. By (I), it satisfies the conditions from Proposition 5.42. So suppose we apply Proposition 5.42 to recover the uniquely determined knowledge and belief relations \rightarrow_α^S and \sqsubseteq_α^S in terms of the (recovered, and not uniquely determined) evidence relations \sim_α^M and \leq_α^M .

Then, by combining the results from Proposition 5.42 and from (d) and (e) of this proposition, we have that $\rightarrow_\alpha^S = \rightarrow_\alpha$ and $\sqsubseteq_\alpha^S = \sqsubseteq_\alpha$, where \rightarrow_α and \sqsubseteq_α represent the primitive relations for knowledge and belief on the pseudo-model \mathbf{M} .

Thus, by Proposition 5.42, we have for all $\alpha \in \{A\} \cup A$ and $s \in S$, that

$$\begin{aligned}
(\mathbf{M}, x) \models B_\alpha \varphi & \quad \text{iff} \quad \text{for all } t \in S \text{ s.t. } s \rightarrow_\alpha t : (\mathbf{M}, t) \models \varphi & \quad (\text{Def. 5.37}) \\
& \quad \text{iff} \quad \text{for all } t \in S \text{ s.t. } s \rightarrow_\alpha^S t : (\mathbf{M}, t) \models \varphi & \quad (\text{Prop. 5.42, (d), (e)}) \\
& \quad \text{iff} \quad (\mathbf{S}_M, x) \models B_\alpha \varphi & \quad (\text{Prop. 5.42}) \\
\\
(\mathbf{M}, x) \models K_\alpha \varphi & \quad \text{iff} \quad \text{for all } t \in S \text{ s.t. } s \sqsubseteq_\alpha t : (\mathbf{M}, t) \models \varphi & \quad (\text{Def. 5.37}) \\
& \quad \text{iff} \quad \text{for all } t \in S \text{ s.t. } s \sqsubseteq_\alpha^S t : (\mathbf{M}, t) \models \varphi & \quad (\text{Prop. 5.42, (d), (e)}) \\
& \quad \text{iff} \quad (\mathbf{S}_M, x) \models K_\alpha \varphi & \quad (\text{Prop. 5.42})
\end{aligned}$$

which concludes our proof. □

Thus, we can recover evidence relations on the pseudo-model \mathbf{M} for $KB_{i,A}$, that result in a pseudo-model \mathbf{S}_M for $\Box[\forall]_{i,A}$, that agrees with \mathbf{M} on the interpretation of knowledge and belief. It remains to extend this claim to all formulas in the language $KB_{i,A}$:

Corollary 5.45. *The same formulas in $KB_{i,A}$ are satisfiable in the pseudo-model \mathbf{M} for $KB_{i,A}$, as in the pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$.*

Proof. By induction on the complexity of φ . We compare the interpretation of formulas $\varphi \in KB_{i,A}$ on the pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$, according to Definition 5.6, with the interpretation on the pseudo-model \mathbf{M} for $KB_{i,A}$ according to Definition 5.37.

For the atomic case, and for the boolean cases of the inductive step, we observe that the interpretations agree on the semantics of atomic propositions and boolean combinations. For formulas of the form $K_i\psi$, $B_i\psi$, $K_A\psi$, and $B_A\psi$, the claim follows from Proposition 5.42 and Theorem 5.44 (II). □

To finish the proof of soundness, completeness, and decidability for the proof system $KB_{i,A}$ with respect to the intended models, we finally prove Theorem 4.12, stated in Section 4.2: we show that $KB_{i,A}$ from Definition 4.11 is sound and weakly complete with respect to multi-agent topo-e-models, and that the logic of $KB_{i,A}$ is decidable.

Proof of Theorem 4.12. Soundness of the axioms and rules of classical propositional logic is a routine check, as well as the soundness of the inference rules (Modus Ponens, and Necessitation for both \Box_α and $[\forall]_\alpha$ for all $\alpha \in \{A\} \cup A$); therefore, we omit these proofs. Soundness of Stalnaker's epistemic-doxastic axioms follows from Proposition 4.3. As for the group knowledge axioms, soundness follows directly from Proposition 4.5.

As for completeness, let $\varphi \in KB_{i,A}(V)$ be any consistent formula. By Corollary 5.39, φ is satisfiable on a pseudo-model \mathbf{M} for $KB_{i,A}$. By Corollary 5.45, there exists an equivalent pseudo-model \mathbf{S} for $\Box[\forall]_{i,A}$. Thus, φ is satisfied on \mathbf{S} . By Corollary 5.33, φ is satisfiable on the associated model for \mathbf{S} . Since this is a relational evidence model, we obtain weak completeness for $KB_{i,A}$ with respect to relational evidence models. Finally, by Corollary 3.40, we obtain weak completeness with respect to topo-e-models. Decidability of $KB_{i,A}$ follows from decidability of the larger language $\Box[\forall]_I$ (see Corollary 5.14). \square

5.3 Completeness of Dynamic Extensions

In this section we prove completeness of the proof systems for the dynamic extensions of our languages (presented in 4.2.4), with respect to multi-agent topo-e-models. Our approach follows Appendix B in [BS20].

If the dynamic extension of a complete logic is provably equally expressive as the static logic, then proving completeness of the extended proof system amounts to proving completeness of the static proof system. We proved weak completeness with respect to multi-agent topo-e-models for the static proof systems $\Box[\forall]_I$, $\Box[\forall]_{i,A}$, and $KB_{i,A}$. Therefore, it suffices to show that each of the static languages is provably equally expressive as its dynamic extension. We will use the following two Lemmas.

Lemma 5.46. *Let φ be any formula over the static language $\Box[\forall]_I$. Then, for every nonempty $I \subseteq A$, there exists some formula φ_s over the language $\Box[\forall]_I$, such that*

$$\vdash [\text{share}_I]\varphi \leftrightarrow \varphi_s$$

is provable in $\Box[\forall]_I$.

Proof. By induction on the complexity of the static formula φ . Let $I \subseteq A$ be nonempty. For the atomic case, where $\varphi := p$, reduction axiom (1) from Definition 4.15 gives us that $\vdash [\text{share}_I]p \leftrightarrow p$, so we take $\varphi_s := p$.

For $\varphi := \neg\psi$, we apply the induction hypothesis to ψ to obtain that there exists $\psi_s \in \Box[\forall]_I$ such that $\vdash [\text{share}_I]\psi \leftrightarrow \psi_s$. By reduction axiom (2) for negation, we have $\vdash [\text{share}_I]\neg\psi \leftrightarrow \neg[\text{share}_I]\psi$ which, with the induction hypothesis, gives us that $\vdash [\text{share}_I]\neg\psi \leftrightarrow \neg\psi_s$. Thus, we take $\varphi_{sI} := \neg\psi_s$.

For $\varphi := \psi \wedge \chi$, the proof is similar.

For $\varphi := \Box_J \psi$, with $\emptyset \neq J \subseteq A$, first suppose that $I \cap J \neq \emptyset$. We use reduction axiom (4): we have that $\vdash [\text{share}_I] \Box_J \psi \leftrightarrow \Box_{J \cup I} [\text{share}_I] \psi$. By the induction hypothesis, there exists $\psi_s \in \Box[\forall]_I$ such that $\vdash [\text{share}_I] \psi \leftrightarrow \psi_s$. Thus, $\vdash [\text{share}_I] \Box_J \psi \leftrightarrow \Box_{J \cup I} \psi_s$, therefore, we take $\varphi_s := \Box_{J \cup I} \psi_s$.

Now, for the same case, suppose that $I \cap J = \emptyset$. Then we use reduction axiom (5): $\vdash [\text{share}_I] \Box_J \psi \leftrightarrow \Box_J [\text{share}_I] \psi$. By applying the induction hypothesis, we obtain that $\vdash [\text{share}_I] \Box_J \psi \leftrightarrow \Box_J \psi_s$, for some $\psi_s \in \Box[\forall]_I$. Thus, we take $\varphi_s := \Box_J \psi_s$.

For $\varphi := [\forall]_J \psi$, with $\emptyset \neq J \subseteq A$, the proof is symmetrical to the proof for $\varphi := \Box_J \psi$. \square

Lemma 5.47. *For every formula φ over the dynamic language $\Box[\forall]_I + [\text{share}_I]$, there exists some formula $\varphi' \in \Box[\forall]_I$ such that*

$$\vdash \varphi \leftrightarrow \varphi'$$

is provable in $\Box[\forall]_I$.

Proof. By induction on the complexity of the dynamic formula φ . Let $I \subseteq A$ be nonempty. For the atomic case, where $\varphi := p$, we have that $\varphi \in \Box[\forall]_I$, so we take $\varphi' := \varphi$.

For $\varphi := \neg \psi$, apply the induction hypothesis to ψ to obtain $\psi' \in \Box[\forall]_I$ such that $\vdash \psi \leftrightarrow \psi'$. But then $\vdash \neg \psi \leftrightarrow \neg \psi'$, so we take $\varphi' := \neg \psi'$.

For $\varphi := \psi \wedge \chi$, the proof is similar.

For $\varphi := \Box_I \psi$, apply the induction hypothesis to ψ to obtain $\psi' \in \Box[\forall]_I$ such that $\vdash \psi \leftrightarrow \psi'$. Then $\vdash \Box_I \psi \leftrightarrow \Box_I \psi'$ by Necessitation and Distribution for \Box_I , so we take $\varphi' := \Box_I \psi'$.

For $\varphi := [\forall]_I \psi$, the proof is symmetrical to the proof for $\varphi := \Box_I \psi$.

For $\varphi := [\text{share}_I] \psi$, apply the induction hypothesis to ψ to obtain $\psi' \in \Box[\forall]_I$ such that $\vdash \psi \leftrightarrow \psi'$. By Necessitation and Distribution for $[\text{share}_I]$, we have $\vdash [\text{share}_I] \psi \leftrightarrow [\text{share}_I] \psi'$. Now, by Lemma 5.46, we have $\vdash [\text{share}_I] \psi' \leftrightarrow \psi_s$. Combined with the induction hypothesis, this gives us $\vdash [\text{share}_I] \psi \leftrightarrow \psi_s$. So we take $\varphi' := \psi_s$. \square

Having proved Lemma 5.46 and Lemma 5.47, we can finally prove Theorem 4.18, stated in Section 4.2: we show that each of the presented dynamic proof systems is sound and weakly complete with respect to multi-agent topo-e-models.

Proof of Theorem 4.18. We prove the claims for each of the proof systems:

1. For soundness of $\Box[\forall]_I + [\text{share}_I]$, the proof system $\Box[\forall]_I$ is sound with respect to multi-agent topo-e-models by Theorem 4.9. It remains to show that the reduction axioms from Definition 4.15 are sound with respect to multi-agent topo-e-models. The proof of axiom (1) follows directly from Proposition 3.22. We omit the cases for axioms (2) and (3), as they are a routine check. Finally, the proofs for axioms (4)-(7) follow directly from Proposition 4.17.

For completeness, we use the fact that the dynamic language $\Box[\forall]_I + [\text{share}_I]$ is provably equally expressive as the static language $\Box[\forall]_I$: let $\varphi \in \Box[\forall]_I + [\text{share}_I]$ be a consistent formula. By Lemma 5.47, there exists $\varphi' \in \Box[\forall]_I$ such that $\vdash \varphi \leftrightarrow \varphi'$ is a theorem in the logic of $\Box[\forall]_I + [\text{share}_I]$. By Theorem 4.9, there exists a pointed multi-agent topo-e-model (\mathfrak{M}, x) such that $(\mathfrak{M}, x) \models \varphi'$. Now, since we proved soundness of $\Box[\forall]_I + [\text{share}_I]$, $\vdash \varphi \leftrightarrow \varphi'$ is sound on (\mathfrak{M}, x) and so, $(\mathfrak{M}, x) \models \varphi$.

2. For soundness of $\Box[\forall]_{i,A} + [\textit{share}_A]$, it suffices to observe that the language $\Box[\forall]_{i,A} + [\textit{share}_A]$ is a fragment of the full language $\Box[\forall]_I + [\textit{share}_I]$, and that $\Box[\forall]_I + [\textit{share}_I]$ is sound with respect to multi-agent topo-e-models by Theorem 4.9.

For completeness, the proof for $\Box[\forall]_{i,A} + [\textit{share}_A]$ is similar to the proof for $\Box[\forall]_I + [\textit{share}_I]$; co-expressivity can be proved by restricting the proofs of Lemma 5.46 and Lemma 5.47 to $\Box[\forall]_{i,A}$ (Lemma 5.46) and $\Box[\forall]_I + [\textit{share}_I]$ (Lemma 5.47). For completeness, the proof goes through as in Theorem 4.18, using Theorem 4.10 for soundness and completeness of $\Box[\forall]_{i,A}$.

3. For soundness of $KB_{i,A} + [\textit{share}_A]$, the proof follows the same line of reasoning as the proof for $\Box[\forall]_I + [\textit{share}_I]$. The proof system $KB_{i,A}$ is sound with respect to multi-agent topo-e-models by Theorem 4.12. It remains to show that the additional reduction axioms from Definition 4.16 are sound. We omit the proofs for axioms (1)-(3) from Definition 4.15, as they are a routine check. As for the axioms (1)-(4) for knowledge and belief, the claim follows directly from Proposition 4.17.

For completeness, the proof for $KB_{i,A} + [\textit{share}_A]$ is similar to the proof of $\Box[\forall]_I + [\textit{share}_I]$: co-expressivity can be proved by restricting Lemma 5.46 and Lemma 5.47 to $KB_{i,A}$. The cases for soft and hard evidence are replaced with the cases for knowledge and belief, by using the reduction axioms from Definition 4.16. The proofs are similar to the cases for evidence, though simpler: there is no case distinction whether two groups overlap, since the only allowed groups are individual agents and the full group A . For completeness, the proof goes through as in Theorem 4.18, using Theorem 4.12 for soundness and completeness of $KB_{i,A}$.

□

Chapter 6

Symbolic Model Checking

This chapter concerns symbolic model checking of the language $\Box[\forall]_I$ and is heavily inspired by the approach taken by [Gat18] to model checking of Dynamic Epistemic Logic. In Section 6.1, we introduce the *symbolic topo-structures*: the symbolic counterparts of topo-e-models on which the model checking is performed. The equivalence proof, which states that these translations preserve truth, is presented in Section 6.2. Section 6.3 revisits the Missing Cake scenario, which we represent in a symbolic topo-structure. Finally, Section 6.4 and Section 6.5 are dedicated to the implementation and benchmarking of a Haskell-based model checker for symbolic topo-structures.

In line with the implementation, which allows model checking only on finite models, we only consider finite models throughout the chapter¹. As a consequence, each of the considered models is Alexandroff and admits a translation, as established in Section 3.5, to an equivalent relational model. We make use of this in our equivalence proof. Throughout this chapter, fix a finite set V of propositions and a finite set A of agents.

6.1 Symbolic Topo-Structures

We start by introducing a number of preliminary definitions. Throughout this chapter, we use the abbreviation $\text{outof}(X, Y)$, which is interpreted as "from the propositions in Y , exactly those in X are true" [Gat18]:

$$\text{outof}(X, Y) := \bigwedge X \wedge \bigwedge \{\neg p \mid p \in Y \setminus X\}.$$

Model checking involves a model and a formula whose validity or truth we check on the model. Recall from Section 2.4 that for the symbolic approach, we represent both the model and the formula succinctly. Definition 6.1 presents a general symbolic representation of a state space [Gat18].

Definition 6.1 (Symbolic Encoding). Given a set of worlds X and an injective valuation function $\pi : X \rightarrow \mathcal{P}(V)$, a boolean formula $\theta \in \mathcal{L}_{\text{Bool}}(V)$ is a *symbolic encoding* of X if for all $s \subseteq V$,

$$s \models \theta \Leftrightarrow \exists x \in X : s = \pi(x).$$

Fact 6.2. Given V , X , and π , all formulas equivalent to $\theta := \bigvee_{x \in X} \text{outof}(\pi(x), V)$ are symbolic encodings of X .

¹As van Benthem and Pacuit remarked, "the lure of the infinite must be left for another occasion" ([BP11], p. 3).

Each state $x \in X$ is represented symbolically by the set of propositions that are true at x ; the symbolic encoding of the state space is therefore a formula equivalent to a disjunction of these symbolic representations, one for each state of the model.

Definition 6.3 presents the symbolic counterpart of our multi-agent topo-e-models. For clarity, we will sometimes refer to multi-agent topo-e-models as *explicit* topo-e-models. The generated topologies τ_i for each agent $i \in A$ are not explicitly represented. This aids us in representing the symbolic structure more efficiently than the explicit model.

Definition 6.3 (Symbolic Topo-Structure). A *symbolic topo-structure* is a tuple $\mathcal{F} = (V, \theta, E, O)$, satisfying the following conditions:

1. The *vocabulary* V is a finite set of propositional variables;
2. the *state law* $\theta \in \mathcal{L}_B(V)$ is a boolean formula over V ;
3. A *state* $s \subseteq V$ in a symbolic topo-structure is a set of propositions such that $s \models \theta$;
4. The *evidence* $E = (E_i)_{i \in A}$ is a collection of sets such that for each $i \in A$, $E_i \subseteq V$;
5. The *observables* $O = (O_i)_{i \in A}$ are a collection of sets such that for each $i \in A$, $O_i \subseteq V$.
6. For each $i \in A$, the propositions in O_i decide a partition over the states of \mathcal{F} . Formally, for all states s of \mathcal{F} and for all $i \in A$, the following conditions are satisfied:
 - there exists $q \in O_i$ such that $q \in s$;
 - for all $q' \in O_i$ such that $q \neq q'$, we have $q' \notin s$.

The state law θ is a symbolic encoding (see Definition 6.1) of the states of the structure: it describes which valuations over V represent states. Whereas explicit topo-e-models represent soft and hard evidence by sets of states, symbolic topo-structures represent these notions by propositional variables. For each agent $i \in A$, E_i consists of the evidence available to i , and O_i describes the information cells contained in the partition of i . More precisely, there exists a bijection between the observables O_i and the partition Π_i . Since E_i and O_i are subsets of V , soft and hard evidence are encoded in θ .

For each $i \in A$, the set O_i describes an equivalence relation \sim_i , where $s \sim_i t$ if and only if $s \cap O_i = t \cap O_i$. Similarly, the set E_i describes a pre-order \leq_i , where $s \leq_i t$ if and only if $s \cap E_i \subseteq t \cap E_i$. These properties play an important role in the semantics of $\Box[\forall]_I$ on symbolic topo-structures (Definition 6.5).

The propositions in both E_i and O_i correspond to sets of states in the explicit model. Nevertheless, the (philosophical) interpretation of the symbolic evidence differs from the interpretation of the observables. If $e_i \in E_i$ for agent i , then i does not know whether e_i is *factive* at the evaluated state. On the other hand, if $q_i \in O_i$, then she can distinguish whether q_i is factive at any state². To avoid confusion, by the noun *observable* we refer only to hard evidence, such that the interpretation of observables on symbolic topo-structures corresponds to the interpretation of the observable variables in the symbolic representation of relational models as defined in [Gat18].

Soft and hard evidence for groups is defined in terms of the join topology and group partition:

²Under Definition 6.3, the observables of \mathcal{F} have the same asymmetric nature as the evidence of \mathcal{F} : the fact that q is observed does not imply that $\neg q$ is observed. This is enforced by condition (6). This approach to modelling knowledge is referred to as *knowing that* [Gat18]. However, the observation of hard evidence can equivalently be modelled symmetrically, as a formalisation of *knowing whether*: under the semantics that we define in Definition 6.5, the observation of a proposition q is equivalent to the observation of its negation. Thus, instead of explicitly representing each equivalence class with an observable, one could choose a more economic representation in line with [Gat18]. We chose Definition 6.3, because it

Definition 6.4. For nonempty $I \subseteq A$, we let $E_I := \bigcup_{i \in I} E_i$ and $O_I := \bigcup_{i \in I} O_i$.

We explain these definitions in more detail after defining the symbolic semantics.

Definition 6.5 (Symbolic Semantics of $\Box[\forall]_I$). Given a nonempty subgroup $I \subseteq A$, and $p \in V$, the semantics for $\Box[\forall]_I(V)$ on symbolic topo-structures is defined as follows:

$$\begin{aligned}
(\mathcal{F}, s) &\models \top && \text{always holds} \\
(\mathcal{F}, s) &\models p && \text{iff } p \in s \\
(\mathcal{F}, s) &\models \neg\varphi && \text{iff } (\mathcal{F}, s) \not\models \varphi \\
(\mathcal{F}, s) &\models \varphi \wedge \psi && \text{iff } (\mathcal{F}, s) \models \varphi \text{ and } (\mathcal{F}, s) \models \psi \\
(\mathcal{F}, s) &\models \Box_I && \text{iff for all states } t \text{ of } \mathcal{F}, \text{ if } s \cap O_I = t \cap O_I \text{ and } s \cap E_I \subseteq t \cap E_I, \text{ then } (\mathcal{F}, t) \models \varphi \\
(\mathcal{F}, s) &\models [\forall]_I \varphi && \text{iff for all states } t \text{ of } \mathcal{F}, \text{ if } s \cap O_I = t \cap O_I, \text{ then } (\mathcal{F}, t) \models \varphi.
\end{aligned}$$

Knowledge and belief are interpreted by the usual abbreviations (see e.g. Definition 3.10). The cases concerning evidence deserve a more detailed explanation; we discuss them below.

Hard evidence. The semantics of the modality of infallible group knowledge $[\forall]_I \varphi$ is, in fact, identical to the individual semantics of the S5 knowledge modality $K_i \varphi$ presented in Definition 2.2.3 of [Gat18]. The modality of infallible group knowledge describes an equivalence relation on explicit models, i.e. it is also an S5 modality. Our symbolic representation of equivalence classes is similar to the representation defined in [Gat18]. As a result, the symbolic semantics agree.

Soft evidence. Whereas the modality of infallible group knowledge describes an equivalence relation, the modality of group evidence represents a preorder. This is reflected in the use of an inclusion $s \cap E_I \subseteq t \cap E_I$ (as opposed to an equality in the case of $s \cap O_I = t \cap O_I$) in the semantics of $\Box_i \varphi$. A state s satisfies a formula of the form $\Box_i \varphi$, if the set of states satisfying propositions $s \cap E_i$ is a subset of the truth set of φ (where $s \cap E_i$ is the evidence available to i and true at s). This is true if and only if all states within the same information cell $\Pi_i(s)$ that satisfy $s \cap E_i$ (and possibly other pieces of evidence), satisfy φ .

Group notions. To understand the definitions of the group notions $O_I := \bigcup_{i \in I} O_i$ and $E_I := \bigcup_{i \in I} E_i$ for nonempty $I \subseteq A$, observe the following: regarding the partition, in relational terms, we want to have for any two states s and t of the model that $s \sim_I t$, if and only if $s \sim_i t$ for all $i \in I$, that is, if and only if $s \cap O_i = t \cap O_i$ for all $i \in I$. This is equivalent to requiring that $s \cap \bigcup_{i \in I} O_i = t \cap \bigcup_{i \in I} O_i$. Thus, the group partition for I is simply given by $\bigcup_{i \in I} O_i$. For soft evidence, we apply the same line of reasoning to obtain that $E_I = \bigcup_{i \in I} E_i$. Note also that, as E_i is the subbasis of the topology τ_i , E_I corresponds to the subbasis of the join topology, as required.

In order to facilitate symbolic model checking, we present the boolean formulas as BDDs. Moreover, we provide a boolean translation of the full language. The cases for soft and hard evidence (both individual and group notions) use boolean quantification, which we defined in Definition 2.32.

Definition 6.6 (Local Boolean Translation). Given a symbolic topo-structure \mathcal{F} , we define the *local*

simplifies the equivalence proofs from Section 6.2. Nevertheless, the proofs also apply to the economic representation, as an economically represented structure can easily be transformed into a symbolic topo-structure according to Definition 6.3, by the addition of propositional variables.

boolean translation $\|\cdot\|_{\mathcal{F}} : \varphi \in \Box[\forall]_I(V) \rightarrow \mathcal{L}_{\text{Bool}}(V)$ on \mathcal{F} as follows:

$$\begin{aligned}
\|\top\|_{\mathcal{F}} &:= \top \\
\|p\|_{\mathcal{F}} &:= p \\
\|\neg\varphi\|_{\mathcal{F}} &:= \neg\|\varphi\|_{\mathcal{F}} \\
\|\varphi \wedge \psi\|_{\mathcal{F}} &:= \|\varphi\|_{\mathcal{F}} \wedge \|\psi\|_{\mathcal{F}} \\
\|\Box_I\varphi\|_{\mathcal{F}} &:= \forall E'_I (\bigwedge_{e_i \in E_I} (e'_i \leftrightarrow e_i) \rightarrow \forall (V \setminus O_I) (\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\varphi\|_{\mathcal{F}})) \\
\|[\forall]_I\varphi\|_{\mathcal{F}} &:= \forall (V \setminus O_I) (\theta \rightarrow \|\varphi\|_{\mathcal{F}})
\end{aligned}$$

where $p \in V$ is any propositional variable.

Similar to the semantics in Definition 6.5, the boolean translation of $[\forall]_I\varphi$ coincides with the boolean translation of $K_i\varphi$ provided in Definition 2.2.6 of [Gat18]. Given some state s , we check the formula φ on all states within the same information cell for I . We access these states by ranging over all possible valuations over V that satisfy the state law, without changing the valuation of the observables for I .

For the case of $\Box_I\varphi$, the translation is slightly more involved: given a structure \mathcal{F} , fix some state s of \mathcal{F} and any nonempty subgroup $I \subseteq A$. Then, in order to access states that satisfy a superset of the evidence $s \cap E_I$, we need to fix those pieces of evidence that are *factive at s* and access all states within the same information cell at which those pieces of evidence are factive. Thus we cannot simply range over the complement of $(E_I \cup O_I)$; we would need to range over the complement of $(s \cap E_I) \cup O_I$. Since the boolean translation does not have access to s , we first fix those $e_i \in E_I$ that are factive at s . Because E_I is included in the vocabulary $(V \setminus O_I)$, we make a copy e'_i for each $e_i \in (s \cap E_I)$. The copies allow us to access evidence factive at s while ranging over $(V \setminus O_I)$. We then check for each state t , that satisfies both the state law and all $e_i \in E_I$ for which we fixed e'_i , whether it satisfies φ .

The following Theorem proves that on any symbolic topo-structure, the local boolean translation of any formula φ is locally equivalent to φ .

Theorem 6.7. *For any formula $\varphi \in \Box[\forall]_I(V)$ and any pointed model (\mathcal{F}, s) we have that $(\mathcal{F}, s) \models \varphi$ if and only if $s \models \|\varphi\|_{\mathcal{F}}$.*

Proof. By induction on the complexity of φ . Let $I \subseteq A$ be any nonempty subgroup of A . The base case for atomic propositions is immediate. In the induction step, negation and conjunction are standard.

For the case of $\varphi = \Box_I\psi$ we have the following equivalences:

$$\begin{aligned}
&(\mathcal{F}, s) \models \Box_I\psi \\
&\iff \text{for all states } t \text{ of } \mathcal{F} \text{ s.t. } s \cap O_I = t \cap O_I \text{ and } s \cap E_I \subseteq t \cap E_I : (\mathcal{F}, t) \models \psi & (\text{Def. 6.5}) \\
&\iff \forall t \subseteq V \text{ s.t. } t \models \theta \text{ and } s \cap O_I = t \cap O_I \text{ and } s \cap E_I \subseteq t \cap E_I : (\mathcal{F}, t) \models \psi & (\text{Def. 6.3}) \\
&\iff \forall t \subseteq V \text{ s.t. } s \cap O_I = t \cap O_I \text{ and } s \cap E_I \subseteq t \cap E_I \text{ and } t \models \theta : t \models \|\psi\|_{\mathcal{F}} & (\text{IH}) \\
&\iff \forall t \subseteq V \text{ s.t. } s \cap O_I = t \cap O_I \text{ and } s \cap E_I \subseteq t \cap E_I : t \models \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad ((s \cap E_I)' = t_e \cap E'_I \text{ and } s \cap O_I = t_e \cap O_I \text{ and } s \cap E_I \subseteq t_e \cap E_I) \Rightarrow t_e \models \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad (s \cap O_I = t_e \cap O_I \text{ and } (s \cap E_I)' = t_e \cap E'_I) \Rightarrow (s \cap E_I \subseteq t_e \cap E_I \Rightarrow t_e \models \theta \rightarrow \|\psi\|_{\mathcal{F}})
\end{aligned}$$

$$\begin{aligned}
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad (s \cap O_I = t_e \cap O_I \text{ and } (s \cap E_I)' = t_e \cap E'_I) \Rightarrow ((s \cap E_I)' \subseteq (t_e \cap E_I)' \Rightarrow t_e \models \theta \rightarrow \|\psi\|_{\mathcal{F}}) \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad (s \cap O_I = t_e \cap O_I \text{ and } (s \cap E_I)' = t_e \cap E'_I) \Rightarrow (t_e \cap E'_I \subseteq (t_e \cap E_I)' \Rightarrow t_e \models \theta \rightarrow \|\psi\|_{\mathcal{F}}) \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad (s \cap O_I = t_e \cap O_I \text{ and } (s \cap E_I)' = t_e \cap E'_I) \Rightarrow t_e \models \bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad ((s \cup (s \cap E_I)') \cap O_I = t_e \cap O_I \text{ and } (s \cap E_I)' = t_e \cap E'_I) \Rightarrow t_e \models \bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff \forall t_e \subseteq (V \cup E'_I) : \\
&\quad ((s \cup (s \cap E_I)') \cap O_I = t_e \cap O_I \text{ and } (s \cup (s \cap E_I)') \cap E'_I = t_e \cap E'_I) \\
&\quad \Rightarrow t_e \models \bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff s \cup (s \cap E_I)' \models \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \\
&\iff \forall e' \subseteq E'_I : e' = (s \cap E_I)' \Rightarrow s \cup e' \models \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \\
&\iff \forall e' \subseteq E'_I : e'_i \in e' \text{ iff } e_i \in s \Rightarrow s \cup e' \models \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \\
&\iff \forall e' \subseteq E'_I : s \cup e' \models \bigwedge_{e_i \in E_I} (e'_i \leftrightarrow e_i) \Rightarrow s \cup e' \models \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \\
&\iff \forall e' \subseteq E'_I : s \cup e' \models \left(\bigwedge_{e_i \in E_I} (e'_i \leftrightarrow e_i) \rightarrow \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \right) \\
&\iff s \models \forall E'_I \left(\bigwedge_{e_i \in E_I} (e'_i \leftrightarrow e_i) \rightarrow \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\psi\|_{\mathcal{F}} \right) \right) \\
&\iff s \models \|\Box_I \psi\|_{\mathcal{F}}. \tag{Def. 6.6}
\end{aligned}$$

For the case of $\varphi = [\forall]_I \psi$, recall that the semantics for $[\forall]_I \psi$ on symbolic topo-structures coincides with the semantics for $K_i \psi$ on knowledge structures and that the boolean translation of $[\forall]_I \psi$ is identical to the boolean translation of $K_i \psi$ provided in Definition 2.2.6 of [Gat18]. As a consequence, the proof for $[\forall]_I \psi$ is identical to the proof for $K_i \psi$ in Theorem 2.2.8 of [Gat18]:

$$\begin{aligned}
&(\mathcal{F}, s) \models [\forall]_I \psi \\
&\iff \text{for all states } t \text{ of } \mathcal{F} \text{ s.t. } s \cap O_I = t \cap O_I : (\mathcal{F}, t) \models \psi \tag{Def. 6.5} \\
&\iff \forall t \subseteq V \text{ s.t. } t \models \theta \text{ and } s \cap O_I = t \cap O_I : (\mathcal{F}, t) \models \psi \tag{Def. 6.3} \\
&\iff \forall t \subseteq V \text{ s.t. } s \cap O_I = t \cap O_I \text{ and } t \models \theta : t \models \|\psi\|_{\mathcal{F}} \tag{IH} \\
&\iff \forall t \subseteq V \text{ s.t. } s \cap O_I = t \cap O_I : t \models \theta \rightarrow \|\psi\|_{\mathcal{F}} \\
&\iff s \models \forall (V \setminus O_I) (\theta \rightarrow \|\psi\|_{\mathcal{F}}) \\
&\iff s \models \|\Box_I \psi\|_{\mathcal{F}}. \tag{Def. 6.6}
\end{aligned}$$

6.2 Translations

The following definitions and theorems show that for each explicit topo-e-model, as defined in Definition 3.1, there exists an equivalent symbolic topo-structure as defined in Definition 6.3, and vice versa. The translations are given by Definition 6.9 and Definition 6.10. Via relational evidence models, we prove that these translations preserve truth.

Following the approach in Section 2.4 of [Gat18], we first show a general approach to proving that a given finite relational evidence model \mathbf{X} and finite symbolic topo-structure \mathcal{F} satisfy the same formulas (Lemma 6.8). Next, in Theorem 6.11, we show that Lemma 6.8 applies with respect to a relationally represented topo-e-model and its symbolic translation. The proof for the converse translation is presented in Theorem 6.12.

Lemma 6.8 defines a truth-preserving surjective map between a relational evidence model that uses (a subset of) a vocabulary V and a symbolic topo-structure that has a subset of $\mathcal{P}(V)$ as its set of worlds. The reader might notice the resemblance with a surjective morphism³.

Lemma 6.8. *Suppose we have a finite relational evidence model $\text{Rel}(\mathfrak{M}) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$ corresponding to a topo-e-model \mathfrak{M} , with a set of agents A , a set of primitive propositions $V_{\text{init}} \subseteq V$, and a finite symbolic topo-structure $\mathcal{F} = (V, \theta, (E_i)_{i \in A}, (O_i)_{i \in A})$. Furthermore, suppose we have a function $g : X \rightarrow \mathcal{P}(V)$ satisfying the following conditions:*

- (1.1) *For all $x_1, x_2 \in X$, and all $i \in A$, we have that $g(x_1) \cap O_i = g(x_2) \cap O_i$ if and only if $x_1 \sim_i x_2$;*
- (1.2) *For all $x_1, x_2 \in X$, and all $i \in A$, we have that $g(x_1) \cap O_i = g(x_2) \cap O_i$ and $g(x_1) \cap E_i \subseteq g(x_2) \cap E_i$ if and only if $x_1 \leq_i x_2$;*
- (2) *For all $x \in X$ and $p \in V_{\text{init}}$, we have that $p \in g(x)$ if and only if $p \in \pi(x)$;*
- (3) *For every $s \subseteq V$, s is a state of \mathcal{F} if and only if $s = g(x)$ for some $x \in X$.*

Then, for every $\Box[\forall]_I(V_{\text{init}})$ -formula φ we have $(\mathcal{F}, g(x)) \models \varphi$ if and only if $(\text{Rel}(\mathfrak{M}), x) \models \varphi$.

Proof. By induction on the complexity of φ . First, note that surjectivity of $g(\cdot)$ is ensured by condition (3). This allows us to refer to states of \mathcal{F} as the image $g(y)$ of some $y \in X$.

For the atomic case, the claim follows immediately from condition (2).

For $\varphi := \neg\psi$, apply the induction step to ψ to obtain that $(\mathcal{F}, g(x)) \models \psi$ if and only if $(\text{Rel}(\mathfrak{M}), x) \models \psi$. Then the claim immediately follows for $\neg\psi$.

The case for $\varphi := \psi \wedge \chi$ is similar.

We show the case for $\varphi := \Box_I \psi$ for any nonempty $I \subseteq A$, and omit the case for $\varphi := [\forall]_I \psi$: it is

³Since the symbolic topo-structure is not a relational model, the map that we define is technically not a p-morphism, so we cannot simply refer to standard results on the preservation of modal formulas under bisimulations [BRV01]; therefore, we prove by induction that the map preserves truth.

similar and less complicated. For $\varphi := \Box_I \psi$, we have

$$\begin{aligned}
(\mathcal{F}, g(x)) \models \Box_I \psi & \text{ iff } \text{for all states } g(y) \text{ of } \mathcal{F}, \\
& (g(x) \cap O_I = g(y) \cap O_I \text{ and } g(x) \cap E_I \subseteq g(y) \cap E_I) \\
& \text{implies } (\mathcal{F}, g(y)) \models \psi \quad (\text{Def. 6.5}) \\
& \text{iff } \text{for all } y \in X, (g(x) \cap (\bigcup_{i \in I} O_i) = g(y) \cap (\bigcup_{i \in I} O_i) \text{ and} \\
& g(x) \cap (\bigcup_{i \in I} E_i) \subseteq g(y) \cap (\bigcup_{i \in I} E_i)) \\
& \text{implies } (\mathcal{F}, g(y)) \models \psi \quad (\text{C. (3), Def. 6.4}) \\
& \text{iff } \text{for all } y \in X \text{ and for all } i \in I, \\
& (g(x) \cap O_i = g(y) \cap O_i \text{ and } g(x) \cap E_i \subseteq g(y) \cap E_i) \\
& \text{implies } (\mathcal{F}, g(y)) \models \psi \\
& \text{iff } \text{for all } y \in X \text{ and for all } i \in I, x \leq_i y \text{ implies } (\mathcal{F}, g(y)) \models \psi \quad (\text{C. (1.2)}) \\
& \text{iff } \text{for all } y \in X \text{ and for all } i \in I, x \leq_i y \text{ implies } (Rel(\mathfrak{M}), y) \models \psi \quad (\text{IH}) \\
& \text{iff } \text{for all } y \in X, x \leq_I y \text{ implies } (Rel(\mathfrak{M}), y) \models \psi \quad (\text{Def. 3.31}) \\
& \text{iff } (Rel(\mathfrak{M}), y) \models \Box_I \psi. \quad (\text{Def. 3.32})
\end{aligned}$$

□

Note that condition (2) in Lemma 6.8 is not equivalent to stating that $g(x) = \pi(x)$, as $g(x)$ might also contain elements from $V \setminus V_{init}$.

Definition 6.9 and Definition 6.10 describe how to translate an explicit topo-e-model to a symbolic topo-structure and vice versa.

Definition 6.9. Given a finite explicit topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$, where the valuation $\pi : X \rightarrow \mathcal{P}(V_{init})$ ranges over a primitive vocabulary V_{init} , we construct a symbolic topo-structure $\mathcal{F}(\mathfrak{M}) = (V, \theta, E, O)$ as follows.

1. Define $V := V_{init} \cup V^E \cup (V_i^P)_{i \in A}$, where
 - (a) $V^E := \{e_1, \dots, e_n\}$ are fresh variables with respect to V_{init} , corresponding to $\mathcal{E}^0 = \{e_1, \dots, e_n\}$ (where $\mathcal{E}^0 = \bigcup_{i \in A} \mathcal{E}_i^0$);
 - (b) For each $i \in A$, $V_i^P := \{q\pi_i \mid \pi_i \in \Pi_i\}$ such that $\bigcup_{i \in A} V_i^P$ are fresh variables with respect to V_{init} and V^E ;
2. Define $E := (E_i)_{i \in A}$, where for each $i \in A$, $E_i := \{e_i \in V \mid e_i \in \mathcal{E}_i^0\}$;
3. Define $O := (O_i)_{i \in A}$, where for each $i \in A$, $O_i := \{q\pi_i \in V \mid \pi_i \in \Pi_i\}$.
4. Define

$$\theta := \bigvee_{x \in X} \text{outof} \left(\left(\pi(x) \cup \{e \in \mathcal{E}^0 \mid x \in e\} \cup \left\{ q\pi \in \bigcup_{i \in A} O_i \mid x \in \pi \right\} \right), V \right).$$

Let $g : X \rightarrow \mathcal{P}(V)$ be the translation from worlds in \mathfrak{M} to states in \mathcal{F} , where $g(x) := \pi(x) \cup \{e \in \mathcal{E}^0 \mid x \in e\} \cup \{q\pi \in \bigcup_{i \in A} O_i \mid x \in \pi\}$.

We note that the translation from explicit topo-e-models to symbolic structures can be defined more economically by following the approach in [Gat18]: for simplicity, we chose to represent the pieces

soft and hard evidence with one fresh variable each. This means adding $n + m$ fresh propositions to represent the $n := |\bigcup_{i \in A} E_i|$ and $m := |\bigcup_{i \in A} O_i|$ pieces of evidence. Instead, one could add only $l := \lceil \log_2(n + m) \rceil$ many propositions and represent each piece of evidence with a subset of l .

We would like to emphasise that in general, a directly defined symbolic representation of a topo-e-model will have a smaller vocabulary than the automated translation of an explicit representation, and it will therefore be more concise. Thus, ideally, a symbolic model checking algorithm for the language $\Box[\forall]_I$ on topo-e-models will take its input represented directly as a symbolic topo-structure.

For the converse direction, the translation is as follows:

Definition 6.10. Given a finite symbolic topo-structure $\mathcal{F} = (V, \theta, (E_i)_{i \in A}, (O_i)_{i \in A})$, we construct an explicit topo-e-model $\mathfrak{M}(\mathcal{F}) = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ as follows.

1. Define $X := \{s \subseteq V \mid s \models \theta\}$;
2. For each $i \in A$:
 - (a) define $\mathcal{E}_i^0 := \{S_e \subseteq X \mid e \in E_i \text{ and } S_e = \{s \in X \mid e \in s\}\}$;
 - (b) let τ_i be the topology generated by \mathcal{E}_i^0 ;
 - (c) define $\Pi_i := \{\pi_q \subseteq X \mid q\pi_i \in O_i \text{ and } \pi_q = \{s \in X \mid q\pi_i \in s\}\}$
3. Define $\pi : X \rightarrow \mathcal{P}(V)$ as follows: for each $s \in X$, let $\pi(s) := s$.

Theorem 6.11. The function $\mathcal{F}(\cdot)$ from Definition 6.9 preserves truth: for any finite pointed explicit topo-e-model (\mathfrak{M}, x) with a set of primitive propositions V_{init} , and for every formula $\varphi \in \Box[\forall]_I(V_{init})$, we have $(\mathfrak{M}, x) \models \varphi$ if and only if $(\mathcal{F}(\mathfrak{M}), g(x)) \models \varphi$.

Proof. Let $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ and $\mathcal{F}(\mathfrak{M}) = (V, \theta, (E_i)_{i \in A}, (O_i)_{i \in A})$ be the explicit topo-e-model and its symbolic counterpart constructed following Definition 6.9, respectively; let $x \in X$. We first translate the pointed explicit model (\mathfrak{M}, x) to an equivalent relational evidence model $(Rel(\mathfrak{M}), x)$, where $Rel(\mathfrak{M}) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$, following Definition 3.33. Let $g(\cdot)$ be defined as at the end of Definition 6.9. We check the conditions from Lemma 6.8 with respect to $\mathcal{F}(\mathfrak{M})$ and $Rel(\mathfrak{M})$.

For (1.1), let $x_1, x_2 \in X$ and let $i \in A$. We have

$$\begin{aligned}
g(x_1) \cap O_i = g(x_2) \cap O_i & \quad \text{iff} \quad (q_\pi \in g(x_1) \text{ iff } q_\pi \in g(x_2)) \\
& \quad \text{for all } q_\pi \in V \text{ corresp. to } \pi \in \Pi_i \quad (\text{step 3 in Def. 6.9}) \\
& \quad \text{iff} \quad (x_1 \in \pi \text{ iff } x_2 \in \pi) \\
& \quad \text{for all } q_\pi \in V \text{ corresp. to } \pi \in \Pi_i \quad (\text{def. of } g(\cdot) \text{ in Def. 6.9}) \\
& \quad \text{iff} \quad \Pi_i(x_1) = \Pi_i(x_2) \quad (\text{Def. 3.1}) \\
& \quad \text{iff} \quad x_1 \sim_i x_2. \quad (\text{Def. 3.33})
\end{aligned}$$

For (1.2), let $x_1, x_2 \in X$ and let $i \in A$. By our proof of (1.1), we have $g(x_1) \cap O_i = g(x_2) \cap O_i$ if and only if $\Pi_i(x_1) = \Pi_i(x_2)$. It remains to show that we have $g(x_1) \cap E_i \subseteq g(x_2) \cap E_i$ if and only if $x_1 \sqsubseteq_{\tau_i} x_2$, since having $\Pi_i(x_1) = \Pi_i(x_2)$ and $x_1 \sqsubseteq_{\tau_i} x_2$ is equivalent to having $x_1 \leq_i x_2$ (Definition 3.33). Now, we have

$$\begin{aligned}
g(x_1) \cap E_i \subseteq g(x_2) \cap E_i & \quad \text{iff} \quad (e \in g(x_1) \text{ implies } e \in g(x_2)) \text{ for all } e \in \mathcal{E}_i^0 \quad (\text{step 2 in Def. 6.9}) \\
& \quad \text{iff} \quad (x_1 \in e \text{ implies } x_2 \in e) \text{ for all } e \in \mathcal{E}_i^0 \quad (\text{def. of } g(\cdot) \text{ in Def. 6.9}) \\
& \quad \text{iff} \quad (x_1 \in U \text{ implies } x_2 \in U) \text{ for all } U \in \tau_i \quad (\text{by construction of } \tau_i) \\
& \quad \text{iff} \quad x_1 \sqsubseteq_{\tau_i} x_2. \quad (\text{Def. 2.20})
\end{aligned}$$

For (2), let $x \in X$ and $p \in V_{init}$. It follows immediately from our definition of $g(\cdot)$ (Definition 6.9) that $p \in g(x)$ if and only if $p \in \pi(x)$, using the fact that p is not a fresh variable: p cannot be contained in $\{e \in \mathcal{E}^0 \mid x \in e\} \cup \{q_\pi \in \bigcup_{i \in A} O_i \mid x \in \pi\}$, so it must be that $p \in g(x)$ if and only if $p \in \pi(x)$.

Finally, for (3), let $s \subseteq V$. Then s is a state of $\mathcal{F}(\mathfrak{M})$ if and only if $s \subseteq V$ and $s \models \theta$ (Definition 6.3), if and only if there is $x \in X$ such that $s = \pi(x) \cup \{e \in \mathcal{E}^0 \mid x \in e\} \cup \{q_\pi \in \bigcup_{i \in A} O_i \mid x \in \pi\}$ (by definition of θ in Definition 6.9), if and only if $s = g(x)$ for some $x \in X$ (by our definition of $g(\cdot)$, Definition 6.9).

By Lemma 6.8 we obtain for every $\Box[\forall]_I(V_{init})$ -formula φ , that $(\mathcal{F}(\mathfrak{M}), g(x)) \models \varphi$ if and only if $(\text{Rel}(\mathfrak{M}), x) \models \varphi$. We apply Corollary 3.40 to conclude that $(\mathcal{F}(\mathfrak{M}), g(x)) \models \varphi$ if and only if $(\mathfrak{M}, s) \models \varphi$. \square

Theorem 6.12. *The function $\mathfrak{M}(\cdot)$ from Definition 6.10 preserves truth: for any finite pointed symbolic topo-structure (\mathcal{F}, s) over vocabulary V , and every formula $\varphi \in \Box[\forall]_I(V)$, we have $(\mathcal{F}, s) \models \varphi$ if and only if $(\mathfrak{M}(\mathcal{F}), s) \models \varphi$.*

Proof. Let $\mathcal{F} = (V, \theta, (E_i)_{i \in A}, (O_i)_{i \in A})$ be the symbolic structure, let s be a state of \mathcal{F} , and let $\mathfrak{M}(\mathcal{F}) = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ be the explicit counterpart of \mathcal{F} , constructed following Definition 6.10. We first translate the pointed explicit model $(\mathfrak{M}(\mathcal{F}), s)$ to an equivalent relational evidence model $(\text{Rel}(\mathfrak{M}(\mathcal{F})), s)$, where $\text{Rel}(\mathfrak{M}(\mathcal{F})) = (X, (\leq_i)_{i \in A}, (\sim_i)_{i \in A}, \pi)$, following Definition 3.33. Take $g(\cdot)$ to be the identity function. We check the conditions from Lemma 6.8 with respect to \mathcal{F} and $\text{Rel}(\mathfrak{M}(\mathcal{F}))$.

For (1.1), let $x_1, x_2 \in X$ and let $i \in A$. Taking $g(\cdot)$ to be the identity function, we have

$$\begin{aligned} g(x_1) \cap O_i = g(x_2) \cap O_i & \text{ iff } x_1 \cap O_i = x_2 \cap O_i & (\text{ass. } g(x) = x) \\ & \text{ iff } (q \in x_1 \text{ iff } q \in x_2) \text{ for all } q \in O_i \\ & \text{ iff } (x_1 \in \pi_q \text{ iff } x_2 \in \pi_q) \text{ for all } \pi_q \in \Pi_i & (\text{step 2c in Def. 6.10}) \\ & \text{ iff } \Pi_i(x_1) = \Pi_i(x_2) & (\text{Def. 3.1}) \\ & \text{ iff } x_1 \sim_i x_2. & (\text{Def. 3.33}) \end{aligned}$$

For (1.2), let $x_1, x_2 \in X$ and let $i \in A$. Take $g(\cdot)$ to be the identity function. By our proof of (1.1), we have $g(x_1) \cap O_i = g(x_2) \cap O_i$ if and only if $\Pi_i(x_1) = \Pi_i(x_2)$. It remains to show that we have $g(x_1) \cap E_i \subseteq g(x_2) \cap E_i$ if and only if $x_1 \sqsubseteq_{\tau_i} x_2$, since having $\Pi_i(x_1) = \Pi_i(x_2)$ and $x_1 \sqsubseteq_{\tau_i} x_2$ is equivalent to having $x_1 \leq_i x_2$ (Definition 3.33). Now, having $g(\cdot)$ as the identity function, we have

$$\begin{aligned} g(x_1) \cap E_i \subseteq g(x_2) \cap E_i & \text{ iff } x_1 \cap E_i \subseteq x_2 \cap E_i & (\text{ass. } g(x) = x) \\ & \text{ iff } (e \in x_1 \text{ implies } e \in x_2) \text{ for all } e \in E_i \\ & \text{ iff } (x_1 \in e \text{ implies } x_2 \in e) \text{ for all } e \in \mathcal{E}_i^0 & (\text{step 2 in Def. 6.10}) \\ & \text{ iff } x_1 \in U \text{ implies } x_2 \in U \text{ for all } U \in \tau_i & (\text{by construction of } \tau_i) \\ & \text{ iff } x_1 \sqsubseteq_{\tau_i} x_2 & (\text{Def. 2.20}) \end{aligned}$$

For (2), let $x \in X$ and $p \in V_{init}$. With $g(\cdot)$ as the identity function, we have $p \in g(x)$ if and only if $p \in x$, if and only if $p \in \pi(x)$ (step 5 in Definition 6.10). Finally, for (3), let $s \subseteq V$. Then s is a state of \mathcal{F} if and only if $s \subseteq V$ and $s \models \theta$ (Definition 6.3), if and only if $s \in X$ (step 1 in Definition 6.10).

By Lemma 6.8 we obtain for every $\Box[\forall]_I(V_{init})$ -formula φ , that $(\mathcal{F}(\mathfrak{M}), g(s)) \models \varphi$ if and only if $(\text{Rel}(\mathfrak{M}(\mathcal{F})), s) \models \varphi$. Using the result in Corollary 3.40, we conclude that $(\mathcal{F}(\mathfrak{M}), g(s)) \models \varphi$ if and only if $(\mathfrak{M}, s) \models \varphi$. \square

6.3 Missing Cake Revisited

We illustrate how symbolic topo-structures relate to their explicit counterparts by revisiting the Missing Cake scenario from Section 3.3.3. We define a symbolic topo-structure equivalent to the topo-e-model from Example 3.24, which represents the initial scenario, as follows:

Example 6.13. We represent the topo-e-model \mathfrak{M}_0 from Example 3.24 symbolically as follows: define $\mathcal{F} := (V, \theta, E, O)$, where $V := \{p_a, p_b, p_c, p_{ac}, p_{na}, p_{nb}, p_t\}$ is the vocabulary, the evidence is

$$\begin{aligned} E_a &:= \emptyset & O_a &:= \{p_a, p_{na}\} \\ E_b &:= \{p_c\} & O_b &:= \{p_b, p_{nb}\} \\ E_c &:= \{p_{ac}\} & O_c &:= \{p_t\} \end{aligned}$$

and the state law is the boolean translation of

$$\theta := \text{outof}(\{p_a, p_{ac}, p_{nb}, p_t\}, V) \vee \text{outof}(\{p_b, p_{na}, p_t\}, V) \vee \text{outof}(\{p_c, p_{ac}, p_t, p_{nb}, p_{na}\}, V).$$

The actual state is $\{p_a, p_{ac}, p_{nb}, p_t\}$.

We add a number of variables to the original vocabulary $\{p_a, p_b, p_c\}$. To represent Charles' evidence for $p_a \vee p_c$, we include p_{ac} . To represent his partitionm we add p_t , which is true at all states. To represent the partitions for a and b , we add variables p_{na} ("a is not guilty") and p_{nb} ("b is not guilty").

Compared to the explicit representation, it is clear that the symbolic topo-structure is more concise than the explicit model: consider Bob's evidence. Bob's basic evidence $\mathcal{E}_b^0 = \{\{w_b\}, \{w_a, w_c\}, \{w_c\}, X\}$ and partition $\Pi_b = \{\{w_b\}, \{w_a, w_c\}\}$, are reduced to $E_b = \{p_c\}$ and $O_b = \{p_b, p_{nb}\}$, and the generated topology can be derived from these sets.

6.4 Implementation

In this section, we discuss the symbolic model checker for topo-e-models, which we implemented following the theory presented in this chapter. The program is written in Haskell and model-checks formulas over the larger language $\Box[\forall]_I$; as a consequence, it is also a model checker for the fragment KB_I and its subfragment $KB_{i,A}$. We have not implemented any dynamics: the $[share_I]$ modality does not add any expressivity to the language (see Section 5.3). Therefore, formulas in $\Box[\forall]_I + [share_I]$ can be evaluated via translation to $\Box[\forall]_I$, using the reduction axioms from Definition 4.15.

Besides the symbolic model checker, we implemented two explicit model checkers – one for topological representations, and one for relational representations – and several translations. The implementation of the explicit model checkers is based on a model checker by [ÁS23], while the symbolic implementation is heavily inspired by [Gat18]. Since the implementation choices for the explicit model checkers were relatively straightforward, here we focus on the symbolic model checker. The source code can be found at <https://github.com/DdosSantosGomes/Symbolic-Topo-E-Models>. In particular, the code discussed in Section 6.4.1 is in <https://github.com/DdosSantosGomes/Symbolic-Topo-E-Models/blob/main/lib/Symbolic/Semantics.hs>.

6.4.1 Boolean Translation

The model checker evaluates the boolean equivalents (Definition 6.6) of formulas in the language $\Box[\forall]_I$ on *symbolic topo-structures* (Definition 6.3). In this section, we highlight the implementation of the boolean translation. The boolean translation and the translation into a BDD happen simultaneously and are implemented by the function `bdd0f`. This function takes as input a symbolic structure and a formula in $\Box[\forall]_I$, whose types we discuss below, and returns a BDD.

Formulas have type `Form`. The type is defined recursively and includes connectives that could also be done as abbreviations, such as the knowledge and belief operators. For boolean connectives, this increases efficiency [Gat18]. For the more complex operators, it makes the model checker more accessible and it increases readability of formulas. All modalities are defined over `Groups`, which are sets of agents. For example, $K_I\varphi$ is of the form $(K \text{ Group } \text{Form})$.

Symbolic structures follow Definition 6.3 and have the following type:

```
data SymTopoEModel = SymTEM
  { vocab      :: [Prp]
  , theta     :: Bdd
  , evidence  :: M.Map Agent [Prp]
  , obs       :: M.Map Agent [Prp]
  } deriving (Eq, Show)
```

A state s is represented by the list of propositional variables true at s . A pointed structure is a tuple of type `PtSymTopoEModel`, consisting of a `SymTopoEModel` and a state, which has type `[Prp]`.

The recursively defined `bdd0f` function takes the types presented above as input and returns a `Bdd`, i.e. a Binary Decision Diagram as defined in Definition 2.30. We use existing functions, imported from the implementation by [Gat18], that implement boolean combinations of existing BDDs (as described in Section 2.4). For example, the function `neg :: Bdd -> Bdd` takes a BDD representing φ as input and returns the (unique) BDD representing $\neg\varphi$. The cases for boolean connectives are straightforward:

```
bdd0f :: SymTopoEModel -> Form -> Bdd
bdd0f _ Top      = top
bdd0f _ Bot      = bot
bdd0f _ (PrpF (P n)) = var n
bdd0f stm (Neg f)  = neg $ bdd0f stm f
bdd0f stm (Conj fs) = conSet $ map (bdd0f stm) fs
bdd0f stm (Disj fs) = disSet $ map (bdd0f stm) fs
bdd0f stm (Impl f g) = imp (bdd0f stm f) (bdd0f stm g)
```

Before presenting the translation of the case for formulas of type $\Box_I\varphi$, we discuss the less involved case of $[\forall]_I\varphi$, represented by `(Forall ags f)`. We remind the reader of the boolean formula from Definition 6.6 corresponding to the resulting BDD: $\forall(V \setminus O_I)(\theta \rightarrow \|\varphi\|_{\mathcal{F}})$.

```
bdd0f stm (Forall ags f) = forallSet otherps $ imp (theta stm) (bdd0f stm f) where
  otherps = map fromEnum $ vocab stm \\\ ev0r0bs0fGroup ags (obs stm)
```

The vocabulary V is given by `(vocab stm)`, and the observables O_I of the group are given by `(ev0r0bs0fGroup ags (obs stm))`. We obtain the set $V \setminus O_I$, which we call `otherps`, by subtracting the observables from the vocabulary. We range over the resulting vocabulary with `(forallSet otherps)`, while evaluating the implication $\theta \rightarrow \|\varphi\|_{\mathcal{F}}$, represented by the BDD `(imp (theta stm) (bdd0f stm f))`, on the accessible states. The recursion occurs in `(bdd0f stm f)`, which is the BDD of the formula φ , represented by `f`.

The case of type $\Box_I \varphi$ is represented by $(\text{Box ags } f)$. We remind the reader of the boolean formula from Definition 6.6 corresponding to the resulting BDD:

$$\forall E'_I \left(\bigwedge_{e_i \in E_I} (e'_i \leftrightarrow e_i) \rightarrow \forall (V \setminus O_I) \left(\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\varphi\|_{\mathcal{F}} \right) \right).$$

The translation is as follows.

```
bddOf stm (Box ags f) = forallSet evPrime $ imp evAtState evImpliesf
  where
    ev          = map fromEnum $ evObsOfGroup ags (evidence stm)
    evPrime     = map fromEnum $ take (length ev) [freshp (vocab stm) ..]
    primeMap    = Data.IntMap.fromList $ zip ev evPrime
    evAtState   = conSet [equ (var $ primeMap Data.IntMap.! e) (var e) | e <- ev]
    stateSatEv  = conSet [imp (var $ primeMap Data.IntMap.! e) (var e) | e <- ev]
    otherps     = map fromEnum $ vocab stm \ \ evObsOfGroup ags (obs stm)
    evImpliesf  = forallSet otherps $ imp (con stateSatEv (theta stm)) (bddOf stm f)
```

We discuss each of the abbreviations below.

ev, evPrime, primeMap. The evidence of the group E_I is given by `ev`. The duplicated set E'_I , given by `evPrime`, consists of the copies e'_i for all $e_i \in E_I$. These copies are fresh variables, i.e. variables that are not included in the vocabulary of the given symbolic structure. The function `freshp` takes as input a vocabulary consisting of variables of the form $P \ i$, and returns the first fresh variable. That is, given the highest integer i for which $P \ i$ is in the vocabulary, it returns $P \ (i+1)$. For `evPrime`, we take the required amount $(\text{length } \text{ev})$ of fresh variables from the infinitely increasing list of variables $[\text{freshp } (\text{vocab } \text{stm}) \dots]$. The correspondence $e'_i \leftrightarrow e_i$ for all $e_i \in E_I$ is stored in `primeMap`.

evAtState. In the BDD `evAtState` we store the evidence of the group that is factive at the evaluated state: this BDD is a conjunction of BDDs, each representing the formula $e'_i \leftrightarrow e_i$, for all evidence from E_I . The resulting BDD accepts the unique assignment over `evPrime` that sets to true precisely those e'_i for which the original e_i is set to true at the evaluated state.

stateSatEv. The similarly constructed BDD `stateSatEv` stores the implications $e'_i \rightarrow e_i$ for all $e_i \in E_I$. This BDD accepts precisely those valuations, i.e. other states, that set to true a superset of the set of factive evidence of the group, stored in `evAtState`.

otherps, evImpliesf. The set $V \setminus O_I$ is given by `otherps`, analogous to the previous case. Recall that the evidence E_I is included in this set, but that we stored copies of the factive evidence. With the BDD `evImpliesf`, we range over the information cell of the group using $(\text{forallSet } \text{otherps})$, and we evaluate the formula $\bigwedge_{e_i \in E_I} (e'_i \rightarrow e_i) \wedge \theta \rightarrow \|\varphi\|_{\mathcal{F}}$ on the accepted states. This formula is represented by the implication $(\text{imp } (\text{con } \text{stateSatEv } (\text{theta } \text{stm})) (\text{bddOf } \text{stm } f))$, which accepts the following valuations (i.e. states): given a valuation, if the valuation satisfies a superset of the factive evidence, as well as the state law, i.e. if the antecedent, $(\text{con } \text{stateSatEv } (\text{theta } \text{stm}))$, is true, then it satisfies φ : the consequent, $(\text{bddOf } \text{stm } f)$, is true. The recursion occurs in $(\text{bddOf } \text{stm } f)$.

We then define the BDD for $(\text{Box ags } f)$ as $(\text{forallSet } \text{evPrime } \$ \text{imp } \text{evAtState } \text{evImpliesf})$: it ranges over all possible assignments over `evPrime`, to fix the assignment that stores the factive evidence (`evAtState`); given this valuation, it checks `evImpliesf`.

The remaining operators are given by abbreviations:

```
bddOf stm (Dia ags f) = neg $ bddOf stm (Box ags $ Neg f)
bddOf stm (B ags f) = forallSet otherps $ imp (theta stm) (bddOf stm $ Dia ags $ Box ags f)
  where
    otherps = map fromEnum $ vocab stm \ \ evObsOfGroup ags (obs stm)
```

```
bddOf stm (K ags f) = con (bddOf stm $ Box ags f) (bddOf stm $ B ags f)
```

Given an instance f of `Form` and a pointed symbolic topo-structure (stm, x) of type `PtSymTopoEModel`, we use the function `evalViaBdd`, which has type `PtSymTopoEModel -> Form -> Bool`, to check whether $(stm, x) \models f$.

```
evalViaBdd :: PtSymTopoEModel -> Form -> Bool
evalViaBdd (stm, x) f = let
    bdd = bddOf stm f
    b = restrictSet bdd props
    props = [(n, P n 'elem' x) | (P n) <- vocab stm]
in
    case (b == top, b == bot) of
        (True, _) -> True
        (_, True) -> False
        _ -> error \$ "evalViaBdd failed: Composite BDD leftover!"
```

The BDD of f is computed by the `bddOf` function. The function `restrictSet` (imported from the implementation by [Gat18]) *restricts* this BDD by substituting the propositions in the BDD with their assignments at the evaluated state x . If this results in the BDD of \top , then x satisfies f ; if we obtain the BDD of \perp , then x does not satisfy f ; and otherwise, the resulting BDD contains propositional variables whose truth value was not decided by x . This should not happen when the program is used correctly.

The formula f is valid on the model stm if the state law $(\text{theta } stm)$ logically implies f : the BDD of the implication $(\text{imp } (\text{theta } stm) (\text{bddOf } stm f))$ is then identical to the BDD of \top .

```
validViaBdd :: SymTopoEModel -> Form -> Bool
validViaBdd stm f = top == imp (theta stm) (bddOf stm f)
```

6.4.2 Translations between Models

The implementation includes back-and-forth translations between (explicit) relational models and explicit topo-e-models, as well as translations between explicit topo-e-models and symbolic structures. The translations between relational models and explicit topo-e-models are justified by the fact that the input is finite and therefore, all topo-e-models are Alexandroff; we implement the translations from Definition 3.33 and Definition 3.37. The translations between symbolic structures and explicit topo-e-models are based on Definition 6.9 and Definition 6.10.

6.5 Benchmarks

We compare the two model checkers and two translations with respect to performance, on an implementation of a variant of the Missing Cake scenario from Example 3.24, that has a state space which is exponential in the number of agents. The source code can be found at <https://github.com/DdosSantosGomes/Symbolic-Topo-E-Models/blob/main/bench/cake.hs>.

Missing Cake: parents. Today is Daisy’s birthday. Her mother baked Daisy’s favourite chocolate cake for her to eat for breakfast. Yesterday night, someone broke in and ate the cake! Daisy finds the broken window and wakes up her mother, crying. When asked who she thinks broke in, Daisy reluctantly says: "My classmates, maybe. They always bully me at school. They know it is my birthday today and that you always bake a cake." In fact, the culprit is the next-door neighbour: he smelled the freshly-baked cake, waited until nightfall, and broke in. Daisy’s mom immediately calls a parents’ meeting to find out whether any of the classmates are guilty. None of the parents are certain that their child is not guilty, but they each have (factive!) fallible evidence that their own child was in their room all night. For Daisy’s mom, the also fallible (and also factive) evidence absolving her daughter is what Daisy said this morning. Which beliefs are formed by the group of parents?

We check a formula that is closely related to the validity (BDK) of *Consistency of Group Belief with Distributed Knowledge* (see Table 4.1.1) and states that "Each nonempty subgroup of parents has fallible group knowledge of the fact that none of their own children are guilty; and the full group of parents does not believe that any (nonempty) group of children is guilty." It is true in the above scenario.

The considered model is parametrised by the number of involved children: the model for n children has a state space of size 2^n , representing all possible (empty or nonempty) subgroups of children that could have eaten the cake. For each child $i \in \{1, \dots, n\}$, the proposition (P_i) is true if and only if the child is guilty. Each parent has soft evidence absolving their own child from the crime, but none of the parents have hard evidence. In the actual world, none of the children ate the cake.

We directly implemented an explicit topological model and an equivalent symbolic structure. We used four methods to evaluate the formula: directly evaluating on (1) the explicit model and (2) the symbolic structure; (3) translating the explicit model to a symbolic structure, on which we evaluate; and (4) translating the symbolic structure to an explicit model, on which we evaluate.

The results are presented in Figure 6.1. Despite underperforming on smaller models, the symbolic approach is significantly faster than the explicit approach on larger models. Furthermore, we observe that the least efficient approach is to translate symbolic structures into explicit representations and to evaluate on the resulting models. Surprisingly, model checking via the converse translation is roughly as fast as model checking on the directly defined symbolic structure.




Figure 6.1: Benchmarking results. We only show the results of runs that terminated within 1000 seconds: the maximal number of children in the symbolic and explicit-to-symbolic cases is 11; in the explicit case, it is 10; and in the symbolic-to-explicit case, it is only 9.

Chapter 7

Conclusions and Future Work

The contributions of this thesis to the field of Formal Epistemology are both theoretical and practical. Our main theoretical contribution is a pragmatic, non-monotonic, evidence-based interpretation of knowledge and belief with respect to groups. Compared to previous attempts at topological accounts of group knowledge, corresponding to a *traditional* interpretation of distributed knowledge, our notion is better suited to topo-e-models: since virtual group knowledge aligns with the fallibility of knowledge on these models, it is a straightforward, as well as realistic, generalisation of individual knowledge. The argument that virtual group knowledge is realistic was strengthened by the dynamic modality that models how group knowledge can be *resolved* through communication. As a result, virtual group knowledge takes us a step nearer toward applying topological semantics to epistemic languages in practical settings, such as distributed computing and the modelling of real-life situations.

We defined a logic of evidence for groups over the larger language $\Box[\forall]_I$, and showed that it is sound and complete on topo-e-models. The principal axiomatisation presented in this thesis, however, is $KB_{i,A}$, for the subfragment $KB_{i,A}$ of knowledge and belief for individual agents and the full group. This logic allows us to reason topologically about knowledge and belief, without having to explicitly mention the underlying evidence.

The evident suggestion for future research in this respect is to axiomatise the in-between fragment KB_I , in order to obtain a logic of knowledge and belief for all subgroups. We conjecture that the proposed candidate proof system from Table 4.2.3 is complete, and that the proof of its completeness will be similar to the proof for the language $KB_{i,A}$ from Section 5.2. One unresolved factor in the adaptation of the proof to KB_I is the crucial representation theorem (Theorem 5.44), which has to be adapted to pseudo-models for KB_I ; these structures will represent relations for knowledge and belief for all subgroups. As we mentioned, the evidence relations constructed from existing relations for knowledge and belief are not uniquely determined. In particular, the currently used definition of the relations for soft evidence in Theorem 5.44 presents a challenge for the proof of consequence (c) in the same theorem; in the case of pseudo-models for KB_I , it states that for all groups I , the set of \leq_I -maximal states, where \leq_I is a recovered evidence relation, is identical to the set of \trianglelefteq_I -maximal states.

Our practical contribution consists of the two implemented model checkers for the language $\Box[\forall]_I$ on topo-e-models. Model checkers have various benefits: first, having an automated model checker aids our understanding of group knowledge and of the validities associated with these notions. Second, it allows the automated checking of larger models. The latter benefit holds, in particular, for symbolic implementations.

The symbolic implementation brings together two directions of research into model checking: model checking for spatial logics and model checking for epistemic logics. The symbolic model checker for topo-e-models is, to the best of our knowledge, the first symbolic model checker for spatial logics. As for epistemic logic, it expands on the functionalities of the existing symbolic model checker for epistemic logic from [Gat18]. In particular, if we disregard hard evidence, our implementation can be used as a model checker tailored to *plausibility models*: relational models such that the relation on the underlying frame is a pre-order. This could provide a beneficial extension to the model checker from [Gat18], which is optimised with respect to S5 frames. It separately implements a different approach for K frames, including pre-orders, but it is not optimal.

The benchmarking results suggest that currently, the differences in performance between the explicit and symbolic model checker are relatively small, compared to the radical improvements achieved by [Gat18]. Nevertheless, the improvement in performance of the symbolic approach with respect to the explicit approach is promising. An evident direction of future research into symbolic model checking on topo-e-models is therefore to further optimise the implementation. In particular, as we mentioned, the number of variables used to represent soft and hard evidence symbolically can be optimised.

Bibliography

- [ÅAG22] Thomas Ågotnes, Natasha Alechina, and Rustam Galimullin. “Logics with Group Announcements and Distributed Knowledge: Completeness and Expressive Power”. In: *J. Logic Lang. Inf.* 31.2 (2022), pages 141–166. <https://doi.org/10.1007/s10849-022-09355-0> (cited on page 21).
- [APB07] Marco Aiello, Ian Pratt-Hartmann, and Johan van Benthem. “What is Spatial Logic?” In: *Handbook of Spatial Logics*. Edited by Marco Aiello, Ian Pratt-Hartmann, and Johan van Benthem. Springer-Verlag, 2007. Chapter 1, pages 1–12. ISBN: 978-1-4020-5586-7. <https://doi.org/10.1007/978-1-4020-5587-4> (cited on page 18).
- [ÁS23] Álvarez Lombardi, David and Paulius Skaisgiris. *Topomodels in Haskell*. <https://gitlab.com/dqalvarez/topomodels>. 2023 (cited on pages 5, 18, 92).
- [ÅW17] Thomas Ågotnes and Yi N Wáng. “Resolving distributed knowledge”. In: *Artificial Intelligence* 252 (2017), pages 1–21. <https://doi.org/10.1016/j.artint.2017.07.002> (cited on pages 13, 14, 21, 28).
- [Bal+13] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. “The Topology of Belief, Belief Revision and Defeasible Knowledge”. In: *Lecture notes in computer science*. Springer, 2013, pages 27–40. https://doi.org/10.1007/978-3-642-40948-6_3 (cited on page 70).
- [Bal+16] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. “Justified belief and the topology of evidence”. In: *Logic, Language, Information, and Computation: 23rd International Workshop, WoLLIC 2016*. Springer, 2016, pages 83–103 (cited on page 8).
- [BB07a] Johan van Benthem and Guram Bezhanishvili. “Modal Logics of Space”. In: *Handbook of Spatial Logics*. Edited by Marco Aiello, Ian Pratt-Hartmann, and Johan van Benthem. Springer-Verlag, 2007. Chapter 5, pages 217–298. ISBN: 978-1-4020-5586-7. <https://doi.org/10.1007/978-1-4020-5587-4> (cited on pages 6, 10–12, 40).
- [BB07b] Johan van Benthem and Guram Bezhanishvili. “Spatial Logic + Temporal Logic = ?” In: *Handbook of Spatial Logics*. Edited by Marco Aiello, Ian Pratt-Hartmann, and Johan van Benthem. Springer-Verlag, 2007. Chapter 9, pages 497–564. ISBN: 978-1-4020-5586-7. <https://doi.org/10.1007/978-1-4020-5587-4> (cited on page 19).
- [BBF22] Alexandru Baltag, Nick Bezhanishvili, and Saúl Fernández González. “Topological Evidence Logics: Multi-agent Setting”. In: *Language, Logic, and Computation*. Edited by Aybüke Özgün and Yulia Zinova. 2022, pages 237–257. https://doi.org/10.1007/978-3-030-98479-3_12 (cited on page 12).

- [Ben+06] Johan van Benthem, Guram Bezhanishvili, Balder ten Cate, and Darko Sarenac. “Multimodal Logics of Products of Topologies”. In: *Studia Logica* 84.3 (2006), pages 369–392. <https://doi.org/10.1007/s11225-006-9013-x> (cited on page 12).
- [Bez+22] Nick Bezhanishvili, Vincenzo Ciancia, David Gabelaia, Gianluca Grilletti, Diego Latella, and Mieke Massink. “Geometric Model Checking of Continuous Space”. In: *Logical Methods in Computer Science* 18.4 (2022). <https://doi.org/10.48550/arXiv.2105.06194> (cited on pages 5, 19).
- [BFS16] Alexandru Baltag, Virginie Fiutek, and Sonja Smets. “Beliefs and Evidence in Justification Models”. In: *Advances in Modal Logic*. Edited by Lev Beklemishev, Stéphane Demri, and András Máté. 2016, pages 156–176. <https://dare.uva.nl/search?identifier=8fc600a9-dc18-491c-b6a8-fa670c1e3133> (cited on page 3).
- [BH14] Nick Bezhanishvili and Wiebe van der Hoek. “Structures for Epistemic Logic”. In: *Johan van Benthem on Logic and Information Dynamics. Outstanding Contributions to Logic*. Edited by Alexandru Baltag and Sonja Smets. Volume 5. Cham, 2014, pages 339–380. https://doi.org/10.1007/978-3-319-06025-5_12 (cited on page 13).
- [BP11] Johan van Benthem and Eric Pacuit. “Dynamic Logics of Evidence-Based Beliefs”. In: *Studia Logica* 99.1-3 (2011), pages 61–92. <https://doi.org/10.1007/s11225-011-9347-x> (cited on pages 3, 7, 83).
- [BRS12] Alexandru Baltag, Bryan Renne, and Sonja Smets. “The Logic of Justified Belief Change, Soft Evidence and Defeasible Knowledge”. In: *Logic, Language, Information and Computation*. Edited by Luke Ong and Ruy de Queiroz. Volume 7456. 2012, pages 168–190. https://doi.org/10.1007/978-3-642-32621-9_13 (cited on pages 3, 9, 21).
- [BRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science 53. Cambridge University Press, 2001. ISBN: 978-0-521-52714-9. <https://doi.org/10.1017/CB09781107050884> (cited on pages 10, 37, 59–62, 67, 69, 71, 88, 112, 115).
- [Bry86] Randal Bryant. “Graph-Based Algorithms for Boolean Function Manipulation”. In: *IEEE Transactions on Computers* C-35.8 (1986), pages 677–691. <https://doi.org/10.1109/TC.1986.1676819> (cited on pages 5, 17, 18).
- [BS04] Johan van Benthem and Darko Sarenac. *The Geometry of Knowledge*. <https://eprints.illc.uva.nl/id/eprint/131/>. 2004 (cited on pages 12–14).
- [BS16] Alexandru Baltag and Sonja Smets. “A Qualitative Theory of Dynamic Interactive Belief Revision”. In: *Readings in Formal Epistemology*. 2016, pages 813–858. https://doi.org/10.1007/978-3-319-20451-2_39 (cited on page 38).
- [BS20] Alexandru Baltag and Sonja Smets. “Learning What Others Know. EPiC Series in Computing”. In: *LPAR23 proceedings of the International Conference on Logic for Programming AI and Reasoning*. Edited by L. Kovacs and Albert E. Volume 73. 2020, pages 90–110. <https://doi.org/10.48550/arXiv.2109.07255> (cited on pages 13, 21, 51, 56, 58, 60, 61, 80, 107).
- [Cia+17] Vincenzo Ciancia, Diego Latella, Michele Loreti, and Mieke Massink. “Model Checking Spatial Logics For Closure Spaces”. In: *Logical Methods in Computer Science* 12.4 (2017). [https://doi.org/10.2168/LMCS-12\(4:2\)2016](https://doi.org/10.2168/LMCS-12(4:2)2016) (cited on pages 5, 19).

- [Cia+18] Vincenzo Ciancia, Stephen Gilmore, Gianluca Grilletti, Diego Latella, Michele Loreti, and Mieke Massink. “Spatio-temporal model checking of vehicular movement in public transport systems”. In: *International Journal on Software Tools for Technology Transfer* 20.3 (2018), pages 289–311. <https://doi.org/10.1007/s10009-018-0483-8> (cited on pages 5, 19).
- [Cla+18] Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors. *Handbook of Model Checking*. Springer International Publishing, 2018. ISBN: 978-3-319-10574-1. <https://doi.org/10.1007/978-3-319-10575-8> (cited on pages 4, 5, 16–18).
- [DHK07] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer-Verlag, 2007. ISBN: 9781402058394 (cited on page 14).
- [DKÖ19] Hans van Ditmarsch, Sophia Knight, and Aybüke Özgün. “Announcement as effort on topological spaces”. In: *Synthese* (2019). <https://doi.org/10.1007/s11229-017-1592-8> (cited on page 12).
- [Eng89] Ryszard Engelking. *General Topology*. Volume 6. Sigma Series in Pure Mathematics. Heldermann, 1989. ISBN: 9783885380061 (cited on page 7).
- [Fag+04] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. The MIT Press, 2004. ISBN: 9780262562003 (cited on pages 13, 14, 25).
- [Fer18] Saül Fernández González. “Generic Models for Topological Evidence Logics”. Master’s thesis. University of Amsterdam, 2018. <https://eprints.illc.uva.nl/id/eprint/1641/1/MoL-2018-27.text.pdf> (cited on pages 3, 4, 9, 12–15, 20–24, 44).
- [FHV92] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi. “What can machines know? On the properties of knowledge in distributed systems”. In: *J. ACM* 39.2 (1992), pages 328–376. <https://doi.org/10.1145/128749.150945> (cited on pages 4, 13, 14).
- [Gat18] Malvin Gattinger. “New Directions in Model Checking Dynamic Epistemic Logic”. PhD thesis. University of Amsterdam, 2018. ISBN: 978-94-028-1025-7. <https://malv.in/phdthesis> (cited on pages 5, 17, 18, 83–89, 92, 93, 95, 99).
- [Get63] Edmund L. Gettier. “Is Justified True Belief Knowledge?” In: *Analysis* 23.6 (1963), pages 121–123. <https://doi.org/10.1093/analys/23.6.121> (cited on page 9).
- [Hei10] Bernhard Heinemann. “Logics for multi-subset spaces”. In: *Journal of Applied Non-Classical Logics* 20.3 (2010), pages 219–240. <https://doi.org/10.3166/janc1.20.219-240> (cited on page 12).
- [Hin62] Jaakko Hintikka. *Knowledge and Belief*. Ithaca, N.Y.: Cornell University Press, 1962. ISBN: 0801401879 (cited on page 3).
- [HM90] Joseph Y. Halpern and Yoram Moses. “Knowledge and common knowledge in a distributed environment”. In: *J. ACM* 37.3 (1990), pages 549–587. <https://doi.org/10.1145/79147.79161> (cited on pages 3, 4, 13, 20).
- [HM92] Joseph Y. Halpern and Yoram Moses. “A guide to completeness and complexity for modal logics of knowledge and belief”. In: *Artificial Intelligence* 54.3 (1992), pages 319–379. [https://doi.org/10.1016/0004-3702\(92\)90049-4](https://doi.org/10.1016/0004-3702(92)90049-4) (cited on pages 3, 13, 14).

- [LA75] Roland E. Larson and Susan J. Andima. “The Lattice of Topologies: a Survey.” In: *The Rocky Mountain Journal of Mathematics* 5.2 (1975), pages 177–198. <https://doi.org/10.1216/RMJ-1975-5-2-177> (cited on page 14).
- [MT44] John C. C. McKinsey and Alfred Tarski. “The algebra of topology”. In: *The Annals of Mathematics* 45.1 (1944), pages 141–191. <https://doi.org/10.2307/1969080> (cited on pages 3, 9).
- [Özç19] Özgür L. Özçep. *Representation Theorems in Computer Science*. Springer International Publishing, 2019. ISBN: 978-3-030-25785-9. <https://doi.org/10.1007/978-3-030-25785-9> (cited on page 69).
- [Özg17] Aybüke Özgün. “Evidence in Epistemic Logic: a Topological Perspective”. PhD thesis. University of Amsterdam, 2017. <https://eprints.illc.uva.nl/id/eprint/2147/1/DS-2017-07.text.pdf> (cited on pages a, 3, 4, 6, 7, 9–14, 20–22, 43).
- [Ram15] Aldo Iván Ramírez Abarca. “Topological Models for Group Knowledge and Belief”. Master’s thesis. University of Amsterdam, 2015. <https://eprints.illc.uva.nl/id/eprint/2250/1/MoL-2015-32.text.pdf> (cited on pages 3, 4, 12–15, 20, 23, 24, 34, 44).
- [SSV17] Chenwei Shi, Sonja Smets, and Fernando R. Velázquez-Quesada. “Argument-based Belief in Topological Structures”. In: *Electronic Proceedings in Theoretical Computer Science* 251 (2017), pages 489–503. <https://doi.org/10.4204/EPTCS.251.36> (cited on page 3).
- [Sta06] Robert Stalnaker. “On Logics of Knowledge and Belief”. In: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 128.1 (2006), pages 169–199. <https://doi.org/10.1007/s11098-005-4062-y> (cited on pages 9, 43, 49, 70).
- [Str02] Barry Stroud. “Scepticism and the Possibility of Knowledge”. In: *Understanding Human Knowledge: Philosophical Essays*. 2002. <https://doi.org/10.1093/0199252130.003.0001> (cited on page 21).
- [WÅ13] Yi N. Wáng and Thomas Ågotnes. “Multi-agent subset space logic”. In: *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*. IJCAI ’13. 2013, pages 1155–1161. <https://www.ijcai.org/Proceedings/13/Papers/174.pdf> (cited on page 12).
- [Zot24] Teodor-Ştefan Zotescu. “Multi-agent Topological Models for Evidence Diffusion”. Master’s thesis. University of Amsterdam, 2024. <https://dspace.uba.uva.nl/server/api/core/bitstreams/a00dc057-7a90-4071-a4e1-a78f60b0e99b/content> (cited on page 12).

Appendices

Appendix A

Proofs

A.1 Proofs from Chapter 2

Proof of Proposition 2.19. For $\tau \subseteq \tau_R$, it suffices to show that the intersection of any collection of up-sets is an up-set, and that the union of any collection of up-sets is an up-set. Let $(U_i \subseteq X)_{i \in A}$ be such a collection and first consider $\bigcap_{i \in A} U_i$. To see that $\bigcap_{i \in A} U_i$ is an up-set, let $x \in \bigcap_{i \in A} U_i$ and let $y \in X$ such that xRy . Then $y \in \bigcap_{i \in A} U_i$, since $x \in U_i$ for all $i \in A$ and each U_i is an up-set; thus, $y \in U_i$ for all $i \in A$. So $y \in \bigcap_{i \in A} U_i$. Now consider $\bigcup_{i \in A} U_i$. Again, let $x \in \bigcup_{i \in A} U_i$ and let $y \in X$ such that xRy . Then $y \in \bigcup_{i \in A} U_i$, since there is $i \in A$ such that $x \in U_i$ and, with U_i being an up-set, $y \in U_i$. Thus, $y \in \bigcup_{i \in A} U_i$. In conclusion, since τ is generated from a subbasis of up-sets, and since both intersection and union of any collection of up-sets result in an up-set, τ itself must consist of up-sets.

For $\tau_R \subseteq \tau$, we prove the stronger claim that any upset of X is a union of elements of $\{\uparrow x \mid x \in X\}$ (and therefore, closure under intersections is unnecessary given this subbasis). Let $S \subseteq X$ be arbitrary. It suffices to show that $(\uparrow S) \in \tau$. We claim that $\uparrow S = \bigcup_{s \in S} \uparrow s$. For the left-to-right inclusion, let $t \in \uparrow S$. Then there is $s \in S$ such that sRt . In other words, $t \in \uparrow s$. But then $t \in \bigcup_{s \in S} \uparrow s$, as required. For the converse direction, let $t \in \bigcup_{s \in S} \uparrow s$. Then there is $s \in S$ for which $t \in \uparrow s$. In other words, there is $s \in S$ for which sRt . But then by definition of an up-set, $t \in \uparrow S$, as required. \square

A.2 Proofs from Chapter 3

Proof of Lemma 3.21. Partition I into the set $N := I \cap J$ and its complement $M := I \setminus N$. We make the following two observations:

1. For all $i \in N$, we have $i \in J$. By Definition 3.17, we get that for all $i \in N$, $\tau'_i = \tau_j$. Therefore, $\tau'_N = \tau_j$. Analogously, we have that $\Pi'_N = \Pi_j$.
2. For all $i \in M$, we have that $i \notin J$, i.e. $\tau'_i = \tau_i$ and $\Pi'_i = \Pi_i$ (Definition 3.17). Therefore, $\tau'_M = \tau_M$ and $\Pi'_M = \Pi_M$.

We apply Fact 4.4 to rewrite τ'_I in terms of N and M : we have

$$\begin{aligned}
 \tau'_I &= \tau'_N \vee \tau'_M & \Pi'_I &= \{\pi_1 \cap \pi_2 \mid \pi_1 \in \Pi'_N, \pi_2 \in \Pi'_M\} \\
 &= \tau_j \vee \tau_M & &= \{\pi_1 \cap \pi_2 \mid \pi_1 \in \Pi_j, \pi_2 \in \Pi_M\} \\
 &= \tau_{M \cup J} & &= \Pi_{M \cup J} \\
 &= \tau_{I \cup J} & &= \Pi_{I \cup J}
 \end{aligned}$$

where the last step follows from the observations that $J = J \cup N$ and $I = N \cup M$, which give us that $J \cup M = (J \cup N) \cup M = J \cup I = I \cup J$. \square

A.3 Proofs from Chapter 4

Proof of Proposition 4.3. Fix a multi-agent topo-e-model $\mathfrak{M} = (X, (\mathcal{E}_i^0)_{i \in A}, (\tau_i)_{i \in A}, (\Pi_i)_{i \in A}, \pi)$ and a subgroup $I \subseteq A$. Let $I \subseteq A$.

- (K-Distributivity) $\mathfrak{M} \models K_I(\varphi \rightarrow \psi) \rightarrow (K_I\varphi \rightarrow K_I\psi)$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models K_I(\varphi \rightarrow \psi)$, and that $(\mathfrak{M}, x) \models K_I\varphi$. We show that $(\mathfrak{M}, x) \models K_I\psi$, that is, we show that (1) $(\mathfrak{M}, x) \models \Box_I\psi$ and that (2) $(\mathfrak{M}, x) \models B_I\psi$.

For (1), observe that by our assumptions, there exist $U, V \in \tau_I|_{\Pi_I(x)}$ such that $x \in U \cap V$; $U \subseteq \llbracket \neg\varphi \vee \psi \rrbracket = \llbracket \neg\varphi \rrbracket \cup \llbracket \psi \rrbracket$; and $V \subseteq \llbracket \varphi \rrbracket$. As topologies are by definition closed under finite intersections, $U \cap V \in \tau_I$. Furthermore, the set supports φ : we have $U \cap V = (\llbracket \neg\varphi \rrbracket \cup \llbracket \psi \rrbracket) \cap \llbracket \varphi \rrbracket = \llbracket \psi \rrbracket \cap \llbracket \varphi \rrbracket \subseteq \llbracket \varphi \rrbracket$. Thus, $(\mathfrak{M}, x) \models \Box_I\psi$.

For (2), we use the fact that both U and V are locally dense in $\Pi_I(x)$, that is, both sets intersect with every open $W \in \tau_I|_{\Pi_I(x)}$. It remains to show that $U \cap V$ is locally dense. So let $W \in \tau_I|_{\Pi_I(x)}$. Recall that $(U \cap W) \neq \emptyset$, by density of U . Furthermore, by definition of a topology, $U \cap W \in \tau_I|_{\Pi_I(x)}$. So since $U \cap W$ is also open, it has nonempty intersection with V , by density of V . This gives us that $(U \cap V) \cap W \neq \emptyset$, and thus we conclude that $U \cap V$ is locally dense in $\Pi_I(x)$. Therefore, $(\mathfrak{M}, x) \models B_I\psi$.

- (T) $\mathfrak{M} \models K_I\varphi \rightarrow \varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models K_I\varphi$. This is equivalent to $(\mathfrak{M}, x) \models \Box_I\varphi \wedge B_I\varphi$. In particular, $(\mathfrak{M}, x) \models \Box_I\varphi$ gives us that there exists $U \in \tau_I|_{\Pi_I(x)}$ such that $x \in U \subseteq \llbracket \varphi \rrbracket$. Thus, $(\mathfrak{M}, x) \models \varphi$.

- (KK) $\mathfrak{M} \models K_I\varphi \rightarrow K_IK_I\varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models K_I\varphi$. Then $\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$ is locally dense and contains x . We need to show that $\text{Int}_{\tau_I}(\llbracket K_I\varphi \rrbracket) = \text{Int}_{\tau_I}(\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket))$ is locally dense and contains x ; but this follows from Proposition 2.3: $\text{Int}_{\tau_I}(\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)) = \text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$. Thus, $(\mathfrak{M}, x) \models K_IK_I\varphi$.

- (CB) $\mathfrak{M} \models B_I\varphi \rightarrow \neg B_I\neg\varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models B_I\varphi$. For contradiction, suppose that we also have $(\mathfrak{M}, x) \models B_I\neg\varphi$. Then $\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$ and $\text{Int}_{\tau_I}(\llbracket \neg\varphi \rrbracket)$ are both locally dense in $\Pi_I(x)$. But then, by definition of density and by Proposition 2.3, $\llbracket \varphi \rrbracket \cap \llbracket \neg\varphi \rrbracket \neq \emptyset$: we have reached a contradiction. Therefore, $(\mathfrak{M}, x) \not\models B_I\neg\varphi$.

- (Strong PI) $\mathfrak{M} \models B_I\varphi \rightarrow K_IB_I\varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models B_I\varphi$. Recall from the semantics of belief on multi-agent topo-e-models that $(\mathfrak{M}, x) \models B_I\varphi$ if and only if I believes φ in every state in $\Pi_I(x)$, i.e. if $\Pi_I(x) \subseteq \llbracket B_I\varphi \rrbracket$. We need to show that $\text{Int}_{\tau_I}(\llbracket B_I\varphi \rrbracket)$ is locally dense in $\Pi_I(x)$ and contains x . But both follow immediately from the fact that $x \in \Pi_I(x) \subseteq \llbracket B_I\varphi \rrbracket$. Thus, $(\mathfrak{M}, x) \models K_IB_I\varphi$.

- (Strong NI) $\mathfrak{M} \models \neg B_I \varphi \rightarrow K_I \neg B_I \varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models \neg B_I \varphi$. Recall from the semantics of belief on multi-agent topo-e-models that $(\mathfrak{M}, x) \models \neg B_I \varphi$ if and only if there exists no state in $\Pi_I(x)$ in which I believes φ , that is, $\Pi_I(x) \subseteq \llbracket \neg B_I \varphi \rrbracket$. So $\Pi_I(x)$ is a locally dense open subset of $\llbracket \neg B_I \varphi \rrbracket$ which contains x , and therefore we have $(\mathfrak{M}, x) \models K_I \neg B_I \varphi$.

- (KB) $\mathfrak{M} \models K_I \varphi \rightarrow B_I \varphi$.

The claim is immediate from the semantics of $K_I \varphi$ being defined as $\Box_I \varphi \wedge B_I \varphi$.

- (FB) $\mathfrak{M} \models B_I \varphi \rightarrow B_I K_I \varphi$.

Let $x \in X$ and suppose that $(\mathfrak{M}, x) \models B_I \varphi$. Then $\text{Int}_{\tau_I}(\llbracket \varphi \rrbracket)$ is locally dense in $\Pi_I(x)$, i.e. $(\mathfrak{M}, x) \models B_I K_I \varphi$.

□

A.4 Proofs from Chapter 5

Proof of Proposition 5.8. For soundness of $\Box[\forall]_I$, let $\mathbf{S} = (S, (\leq_I)_{\emptyset \neq I \subseteq A}, (\sim_I)_{\emptyset \neq I \subseteq A}, \pi)$ be a pseudo-model. We show that the axioms and rules from $\Box[\forall]_I$ (Definition 4.8) are valid on \mathbf{S} . Because the axioms and rules of propositional logic are straightforward, we focus on the others. Let $s \in S$ be a state and let $H \subseteq J \subseteq A$ be nonempty subgroups of agents.

S4 axioms and rules for \Box_I . The proofs for \Box_I -Necessitation and \Box_I -Distribution are trivial. The reflexivity and transitivity axioms follow immediately from the fact that \leq_I is a pre-order.

S5 axioms and rules for $[\forall]_I$. The proofs for $[\forall]_I$ -Necessitation and $[\forall]_I$ -Distribution are trivial. The reflexivity, transitivity, and symmetry axioms follow immediately from the fact that \sim_I is an equivalence relation.

We prove \Box -monotonicity by contraposition. Let $J \subseteq I$ be a nonempty subgroup and let $(\mathbf{S}, s) \not\models \Box_I \varphi$, that is, suppose there is $t \in S$ such that $s \leq_I t$ and $(\mathbf{S}, t) \not\models \varphi$. By the anti-monotonicity condition in Definition 5.1 of a pseudo-model, $s \leq_I t$ implies $s \leq_J t$. This gives us $(\mathbf{S}, s) \not\models \Box_J \varphi$, as required.

We prove $[\forall]$ -Monotonicity by contraposition. Let $J \subseteq I$ be a nonempty subgroup and let $(\mathbf{S}, s) \not\models [\forall]_I \varphi$. Then there is $t \in S$ such that $s \sim_I t$ and $(\mathbf{S}, t) \not\models \varphi$. By the anti-monotonicity condition in Definition 5.1 of a pseudo-model, $s \sim_I t$ implies $s \sim_J t$. We have $(\mathbf{S}, s) \not\models [\forall]_J \varphi$, as required.

Inclusion. Suppose $(\mathbf{S}, s) \models [\forall]_I \varphi$ and let $t \in S$ with $s \leq_I t$. By the inclusion condition in Definition 5.1 of a pseudo-model, we have $s \sim_I t$ and thus, $(\mathbf{S}, t) \models \varphi$. Since t was arbitrary, we have $(\mathbf{S}, s) \models \Box_I \varphi$, as required.

We prove soundness of $\Box[\forall]_{i,A}$ by contraposition. Let \mathbf{S} be a $\mathbf{-model}$ for $\Box[\forall]_{i,A}$ that does not satisfy all of the axioms and rules in $\Box[\forall]_{i,A}$. We can trivially map \mathbf{S} to a pseudo-model \mathbf{S} for $\Box[\forall]_I$ by adding empty relations for the subgroups: let $\leq_J := \sim_J := \emptyset$ for all nonempty $J \subsetneq A$ with $J \neq \{i\}$ for all $i \in A$. By Proposition 5.7, \mathbf{S} and \mathbf{S} agree on all formulas. Then at least one axiom or rule from $\Box[\forall]_{i,A}$ is not sound with respect to pseudo-models for $\Box[\forall]_I$. This contradicts Proposition 5.8, therefore, $\Box[\forall]_{i,A}$ must be sound with respect to pseudo-models for $\Box[\forall]_{i,A}$. □

Proof of Lemma 5.10. Before proving Lemma 5.10, we introduce the following notation ([BS20]). Let $\pm \varphi$ denote any of the formulas $\{\varphi, \sim \varphi\}$ and, given a set C of formulas over $\Box[\forall]_I(V)$, let $C^\pm :=$

$\{\pm\varphi \mid \varphi \in C\}$ be the closure of φ under single negations. Additionally, let $sub(\varphi)$ denote the set of subformulas of φ .

Let $\varphi_0 \in \Box[\forall]_I(V)$. We build the closure of φ_0 in a finite number of steps: let

1. $\Phi_0 := (sub(\varphi_0))^\pm$
2. $\Phi_1 := (\Phi_0 \cup \{[\forall]_I\varphi \mid \emptyset \neq I \subseteq A \text{ and } [\forall]_I\varphi \in \Phi_0\})^\pm$.
3. $\Phi_2 := (\Phi_1 \cup \{\Box_I[\forall]_I\varphi \mid [\forall]_I\varphi \in \Phi_0\} \cup \{\Box_I\neg[\forall]_I\varphi \mid \neg[\forall]_I\varphi \in \Phi_0\})^\pm$.
4. $\Phi(\varphi_0) := (\Phi_2 \cup \{\Box_I\varphi \mid \emptyset \neq I \subseteq A \text{ and } ([\forall]_I\varphi \in \Phi_2 \text{ or } (\Box_J\varphi \in \Phi_2 \text{ for any } \emptyset \neq J \subseteq A))\})^\pm$.

It is clear that the construction of $\Phi(\varphi_0)$ has finitely many steps, namely four, and that every step adds only finitely many formulas, given that the set of agents A is finite. It remains to show that $\Phi(\varphi_0)$ indeed satisfies the conditions from Definition 5.9.

Condition 1. We have $\varphi_0 \in \Phi(\varphi_0)$ by step 1.

Condition 2. $\Phi(\varphi_0)$ is closed under subformulas. First, note that Φ_0 is closed under subformulas by step 1: let $\psi \in (sub(\varphi_0))^\pm$. We show that $sub(\psi) \subseteq (sub(\varphi_0))^\pm$. We distinguish the following two cases: If (1) $\psi = \varphi_0$, then the claim is immediate. If (2) ψ is of the form $\sim \psi'$, then without loss of generality, assume that ψ is a subformula of $(\sim \psi)$, i.e. that $(\sim \psi) = \neg\psi$. Either $\sim \psi \in sub(\varphi_0)$, or $\psi \in sub(\varphi_0)$. If $\sim \psi \in sub(\varphi_0)$, then we automatically have $\psi \in sub(\varphi_0)$ and we are done. So instead, suppose $\psi \in sub(\varphi_0)$. Then we have that $sub(\psi) \subseteq sub(\varphi_0)$, and $sub(\sim \psi) = (\{\sim \psi\} \cup sub(\psi)) \subseteq (sub(\varphi_0))^\pm$, which proves our claim. It remains to show that $\Phi(\varphi_0)$ is closed under subformulas. This follows from the fact that for each $i \leq 3$, for every formula ψ that gets added to Φ_i , we have that $(sub(\psi) \setminus \{\psi\}) \subseteq \Phi_i$. Thus, $sub(\psi) \subseteq \Phi_{i+1}$.

Condition 3. $\Phi(\varphi_0)$ is closed under single negations. First, observe that the claim holds for Φ_0 . Next, note that at every step, we close the resulting set under single negations. As a result, for each added formula ψ , we add $(\sim \psi)$ in the same step. Observe that the previous claim, regarding closure under subformulas, ensures that $sub(\psi)$ is also closed under single negations. Thus, the claim also holds for the resulting set $\Phi(\varphi_0)$.

Condition 4. Let $[\forall]_I\varphi \in \Phi(\varphi_0)$ and let $\emptyset \neq I \subseteq A$. We show that $[\forall]_I\varphi \in \Phi(\varphi_0)$. Observe that by construction of $\Phi(\varphi_0)$ it must be that (1) $[\forall]_I\varphi \in \Phi_0$, or (2) $[\forall]_I\varphi \in \Phi_1 \setminus \Phi_0$. In the case of (1), $[\forall]_I\varphi$ was added in step 2. In the case of (2), we know that either $[\forall]_I\varphi$ was already in Φ_0 , or both $[\forall]_I\varphi$ and $[\forall]_I\varphi$ were added in step 2, which concludes our proof.

Condition 5. Let $[\forall]_I\varphi \in \Phi(\varphi_0)$. We show that $\Box_I[\forall]_I\varphi \in \Phi(\varphi_0)$. Recall from the proof of item 4 that either (1) $[\forall]_I\varphi \in \Phi_0$, or (2) $[\forall]_I\varphi \in \Phi_1 \setminus \Phi_0$. In either case, $\Box_I[\forall]_I\varphi$ gets added in step 3 and thus, $\Box_I[\forall]_I\varphi \in \Phi(\varphi_0)$.

Condition 6. Let $\neg[\forall]_I\varphi \in \Phi(\varphi_0)$. We show that $\Box_I\neg[\forall]_I\varphi \in \Phi(\varphi_0)$. Recall that for any $\psi \in \Phi(\varphi_0)$, the formula $(\sim \psi)$ is added in the same step as ψ , if it is not already included. It follows that either (1) $[\forall]_I\varphi \in \Phi_0$ and $\neg[\forall]_I\varphi \in \Phi_0$, or (2) $[\forall]_I\varphi \in \Phi_1 \setminus \Phi_0$ and $\neg[\forall]_I\varphi \in \Phi_1 \setminus \Phi_0$ (see item 4). In either case, $[\forall]_I\varphi \in \Phi_1$ and $\neg[\forall]_I\varphi \in \Phi_1$. Then $\Box_I\neg[\forall]_I\varphi$ gets added in step 3, which concludes our proof.

Condition 7. Let $\Box_I\varphi \in \Phi(\varphi_0)$ and let $\emptyset \neq J \subseteq A$. Then $\Box_J\varphi$ gets added in step 4.

Condition 8. Let $[\forall]_I \in \Phi(\varphi_0)$. Then $\Box_I\varphi$ gets added in step 4.

In conclusion, $\Phi(\varphi_0)$ is the finite closure of φ_0 , as required. \square

Proof of Lemma 5.12. We check the conditions from Definition 5.1. Fix a nonempty subgroup $I \subseteq A$. To see that \leq_I is a preorder, note that reflexivity is immediate from our definition of \leq_I on S^C . For

transitivity, let $T, U, V \in S^C$ with $T \leq_I U \leq_I V$. Then for any nonempty $J \subseteq I$, and $\Box_J \varphi \in T$, we have by our definition of \leq_I that $\Box_J \varphi \in U$ (by $T \leq_I U$) and $\Box_J \varphi \in V$ (by $U \leq_I V$). Thus, $T \leq_I V$.

It is immediate from our definition of \sim_I on S^C that \sim_I is an equivalence relation. As for the conditions on the model, anti-monotonicity of the \leq and \sim relations is baked into our definitions of \leq_I and \sim_I on S^C . It remains to prove inclusion: we show that $\leq_I \subseteq \sim_I$.

Let $T, W \in S^C$, suppose $T \leq_I W$. We have $T \sim_I W$ if and only if for all nonempty $J \subseteq I$, and for all $([\forall]_J \varphi) \in \Phi$, we have

$$[\forall]_J \varphi \in T \text{ iff } [\forall]_J \varphi \in W.$$

Let $J \subseteq I$ and let $[\forall]_J \varphi \in \Phi$ be arbitrary.

For the left-to-right direction, suppose $([\forall]_J \varphi) \in T$. By Positive Introspection for \Box , we have $\vdash [\forall]_J \varphi \rightarrow [\forall]_I [\forall]_J \varphi$. From this validity, using inclusion we obtain $\vdash [\forall]_J \varphi \rightarrow \Box_J [\forall]_J \varphi$. Applying Definition 5.9 of closure, we obtain $\Box_J [\forall]_J \varphi \in T$. Recall from the definition of \leq_I that $\Box_J [\forall]_J \varphi \in T$ implies $\Box_J [\forall]_J \varphi \in W$ and that, applying Veracity for \Box_J , this implies $[\forall]_J \varphi \in W$, as required.

For the converse direction, suppose $[\forall]_J \varphi \notin T$. By our definition of closure and by maximal consistency of T , we have $\neg [\forall]_J \varphi \in T$. Again by the definition of closure, we have $\Box_J \neg [\forall]_J \varphi \in T$. Now, by the assumption that $T \leq_I W$, we know that $\Box_J \neg [\forall]_J \varphi \in W$. By Veracity for \Box_J , we have that $\vdash \Box_J \neg [\forall]_J \varphi \rightarrow [\forall]_J \varphi$. Now recall that W is maximally consistent with respect to Φ and that Φ is closed under subformulas. Therefore, it must be that $\neg [\forall]_J \varphi \in W$. We conclude using consistency of W that $[\forall]_J \varphi \notin W$, as required. \square

Proof of Lemma 5.13 (Truth Lemma, $\Box[\forall]_I$). To prove the Truth Lemma, we first need to prove the Diamond Lemmas for the evidence relations on the model:

Lemma A.1 (Diamond Lemma, \leq_I). *Fix a nonempty $I \subseteq A$. If $\neg \Box_I \varphi \in T \in S^C$, then there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \leq_I W$.*

Proof. Suppose $\neg \Box_I \varphi \in T \in S^C$. Let $C := \{\Box_J \psi \in T \mid J \subseteq I\}$.

Claim A.2. *The set $W_0 = \{\sim \varphi\} \cup C$ is consistent.*

Proof of claim. For contradiction, suppose not. Then we have $\vdash (\bigwedge C) \rightarrow \varphi$. Applying B_I -Necessitation and B_I -Distribution to $\vdash (\bigwedge C) \rightarrow \varphi$, we obtain that $\vdash \Box_I (\bigwedge C) \rightarrow \Box_I \varphi$.

By \Box_I -monotonicity, we have for any $J \subseteq I$ that $\vdash \Box_J \psi \rightarrow \Box_I \psi$. We apply this to the Positive Introspection axiom for \Box_J , to obtain that $\vdash \Box_J \psi \rightarrow \Box_I \Box_J \psi$ for all nonempty $J \subseteq I$. Applying this validity to all formulas in C gives us

$$\vdash \bigwedge \{\Box_J \psi \in T \mid J \subseteq I\} \rightarrow \bigwedge \{\Box_I \Box_J \psi \mid \Box_J \psi \in T, J \subseteq I\}$$

i.e. $\vdash (\bigwedge C) \rightarrow \Box_I (\bigwedge C)$.

Now, since $C \subseteq T$, we have $\vdash (\bigwedge T) \rightarrow (\bigwedge C)$. But that gives us that $\vdash (\bigwedge T) \rightarrow \Box_I \varphi$. Since T is maximally consistent with respect to Φ , and the closure conditions imply that $\Box_I \varphi \in T$, we have reached a contradiction to the assumption that $\neg \Box_I \varphi \in T$. Therefore, W_0 must be consistent. \square

Given the consistency of W_0 , we can use the Lindenbaum Lemma to extend W_0 to a set $W \subseteq \Phi$ which is maximally consistent with respect to Φ (i.e. $W \in S^C$). Then we have $(\sim \varphi) \in W$, and $T \leq_I W$, since we have for all nonempty $J \subseteq I$, that $\Box_J \psi \in W$ for all ψ such that $\Box_J \psi \in T$. \square

Lemma A.3 (Diamond Lemma, \sim_I). *For all nonempty $I \subseteq A$, if $\neg[\forall]_I \varphi \in T \in S^C$, then there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \sim_I W$.*

Proof. We use the notation that we introduced for the proof of Lemma 5.10: letting $\pm\varphi$ denote any of the formulas $\{\varphi, \sim\varphi\}$, given a set C of formulas over $\Box[\forall]_I(V)$ and a nonempty subgroup $I \subseteq A$, let

$$T_I^\pm := \{\pm[\forall]_I \varphi \mid \varphi \in T, I \subseteq I\}$$

such that we can characterize \sim_I in terms of one-way inclusion: we get that $T \sim_I W$ if and only if $T_I^\pm \subseteq W$. Now suppose that $\neg[\forall]_I \varphi \in T \in S^C$ and consider the set $W_0 = \{\sim\varphi\} \cup T_I^\pm$.

Claim A.4. *The set W_0 is consistent.*

Proof of claim. Suppose not. Then we have $\vdash (\bigwedge T_I^\pm) \rightarrow \varphi$. Applying $[\forall]_I$ -Necessitation and $[\forall]_I$ -Distribution, we obtain that $\vdash [\forall]_I(\bigwedge T_I^\pm) \rightarrow [\forall]_I \varphi$.

By \Box_I -monotonicity, we have for any $J \subseteq I$ that $\vdash [\forall]_J \psi \rightarrow [\forall]_I \psi$. Together with the axiom $\pm[\forall]_J \psi \rightarrow [\forall]_J \pm[\forall]_J \psi$, which is derivable from the proof system $\Box[\forall]_I^1$, this gives us $\vdash \pm[\forall]_J \psi \rightarrow [\forall]_I \pm[\forall]_J \psi$. Applying this to each formula in T_I^\pm gives us

$$\vdash \bigwedge \{\pm[\forall]_J \psi \in T_I^\pm \mid \emptyset \neq J \subseteq I\} \rightarrow \bigwedge \{[\forall]_I(\pm[\forall]_J \psi) \mid \pm[\forall]_J \psi \in T_I^\pm, \emptyset \neq J \subseteq I\}$$

and thereby, we have $\vdash (\bigwedge T_I^\pm) \rightarrow [\forall]_I(\bigwedge T_I^\pm)$.

With $T_I^\pm \subseteq T$, we also have that $\vdash (\bigwedge T) \rightarrow (\bigwedge T_I^\pm)$. But recall that, as a consequence, $\vdash (\bigwedge T) \rightarrow [\forall]_I \varphi$; by our definition of closure, this implies that $[\forall]_I \varphi \in T$, contradicting our assumption that $\neg[\forall]_I \varphi \in T$. \square

Given the consistency of W_0 , we can use the Lindenbaum Lemma to extend W_0 to a set $W \subseteq \Phi$ that is maximally consistent with respect to Φ (i.e. $W \in S^C$). Then we have $(\sim \varphi) \in W$, and $T \sim_I W$, since $T_I^\pm \subseteq W$. \square

We can now prove the Truth Lemma (Lemma 5.13).

Proof of Lemma 5.13. By induction on the complexity of φ . The atomic case is taken care of by the canonical valuation and the Boolean cases are trivial. So we focus on the cases involving the modalities \Box_I and $[\forall]_I$.

- Inductive step for \Box_I . For the left-to-right direction, suppose $T \models \Box_I \varphi$, for some $\Box_I \varphi \in \Phi$. Now suppose for contradiction that $\Box_I \varphi \notin T$. Then, since T is maximally consistent with respect to Φ , we have $\neg\Box_I \varphi \in T$. By the Diamond Lemma for \leq_I (Lemma A.1), there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \leq_I W$. Now $T \leq_I W$ and $T \models \Box_I \varphi$ together imply that $W \models \varphi$, which by the induction hypothesis gives us that $\varphi \in W$. But given the consistency of W , this directly contradicts the fact that $(\sim \varphi) \in W$: we conclude that $\Box_I \varphi \in T$, after all.

¹It is equivalent to the conjunction of the two validities $[\forall]_J \psi \rightarrow [\forall]_J [\forall]_J \psi$ (which is included in the S5 axioms for $[\forall]_J$) and $\neg[\forall]_J \psi \rightarrow [\forall]_J \neg[\forall]_J \psi$ (which can be proved from the S5 axioms for $[\forall]_J$).

For the converse direction, suppose $\Box_I \varphi \in T$. Recall that by our definition of \leq_I in the filtration of the canonical pseudo-model (Definition 5.11), we have $\Box_I \varphi \in W$ for all $W \in S^C$ such that $T \leq_I W$. Let such W be arbitrary. To prove that $T \models \Box_I \varphi$, it suffices to show that $W \models \varphi$. But this follows directly from Veracity for \Box_I : we get that $\varphi \in W$, and the induction hypothesis then gives us that $W \models \varphi$, as required.

- Inductive step for $[\forall]_I$. For the left-to-right direction, suppose $T \models [\forall]_I \varphi$, for some $[\forall]_I \varphi \in \Phi$. Now suppose for contradiction that $[\forall]_I \varphi \notin T$. Then, since T is maximally consistent with respect to Φ , we have $\neg[\forall]_I \varphi \in T$. By the Diamond Lemma for \sim_I (Lemma A.3), there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \sim_I W$. Now $T \sim_I W$ and $T \models [\forall]_I \varphi$ together imply that $W \models \varphi$, which by the induction hypothesis gives us that $\varphi \in W$. But given the consistency of W , this directly contradicts the fact that $(\sim \varphi) \in W$: we conclude that $[\forall]_I \varphi \in T$, after all.

For the converse direction, suppose $[\forall]_I \varphi \in T$. Recall that by our definition of \sim_I in the canonical pseudo-model, we have $[\forall]_I \varphi \in W$ for all $W \in S^C$ such that $T \sim_I W$. Let such W be arbitrary. To show that $T \models [\forall]_I \varphi$, it suffices to show that $W \models \varphi$. Since $[\forall]_I \varphi \in W$, we know from Veracity for $[\forall]$ that $\varphi \in W$. Then it follows from the induction hypothesis that $W \models \varphi$.

□

□

Proof of Proposition 5.38. Let $\mathbf{M} = (S, (\leq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \leq_A, \rightarrow_A, \pi)$ be a pseudo-model for $KB_{i,A}$. We show that the axioms and rules given in Definition 4.11 are valid on \mathbf{M} . The proofs for the axioms and rules of propositional logic, and for the inference rules, are a routine check, therefore we omit them.

For Stalnaker's axioms, let $s \in S$ be a state and let $\varphi \in KB_{i,A}$. We only show the individual variants, as the proofs for the full group A are identical, so let $i \in A$ be any agent.

- (K) Suppose $(\mathbf{M}, s) \models K_i(\varphi \rightarrow \psi)$, and suppose $(\mathbf{M}, s) \models K_i \varphi$. Let $t \in S$ be arbitrary such that $s \leq_i t$. We show that $(\mathbf{M}, t) \models \psi$. By assumption, $(\mathbf{M}, t) \models \varphi \rightarrow \psi$ and $(\mathbf{M}, t) \models \varphi$. But then $(\mathbf{M}, t) \models \psi$, as required.
- (T) Suppose $(\mathbf{M}, s) \models K_i \varphi$. By reflexivity of \leq_i , we have $s \leq_i s$ and therefore, $(\mathbf{M}, s) \models \varphi$.
- (KK) Suppose $(\mathbf{M}, s) \models K_i \varphi$. Let $t \in S$ be arbitrary such that $s \leq_i t$. We show that $(\mathbf{M}, t) \models K_i \varphi$ (thus giving us that $(\mathbf{M}, s) \models K_i K_i \varphi$, as required). So let $u \in S$ such that $t \leq_i u$. Then by transitivity of \leq_i , $s \leq_i u$, which gives us that $(\mathbf{M}, u) \models \varphi$, and so we are done.
- (CB) Suppose $(\mathbf{M}, s) \models B_i \varphi$. So let $t \in S$ with $s \rightarrow_i t$. Then $(\mathbf{M}, t) \models \varphi$, so it cannot be the case that $(\mathbf{M}, s) \models B_i \neg \varphi$, i.e. we have $(\mathbf{M}, s) \models \neg B_i \neg \varphi$.
- (Strong PI) By contraposition. Suppose $(\mathbf{M}, s) \not\models K_i B_i \varphi$. We show that $(\mathbf{M}, s) \not\models B_i \varphi$. So let $t, u \in S$ with $s \leq_i t \rightarrow_i u$ with $(\mathbf{M}, u) \not\models \varphi$. By the strong transitivity condition (Definition 5.35), $s \rightarrow_i u$. But then we have that $(\mathbf{M}, s) \not\models B_i \varphi$, as required.
- (Strong NI) Suppose $(\mathbf{M}, s) \not\models B_i \varphi$, i.e. suppose there is $t \in S$ with $s \rightarrow_i t$ and $(\mathbf{M}, t) \not\models \varphi$. Let $u \in S$ with $s \leq_i u$. By the strong Euclidean condition (Definition 5.35), we have $u \rightarrow_i t$. With $(\mathbf{M}, t) \not\models \varphi$, we obtain $(\mathbf{M}, u) \models \neg B_i \varphi$, as required.

- (KB) Suppose $(\mathbf{M}, s) \models K_i \varphi$. Then the fact that $(\mathbf{M}, s) \models B_i \varphi$ follows directly from the inclusion condition (Definition 5.35), which states that $s \rightarrow t$ implies $s \sqsubseteq_i t$ for all $t \in S$.
- (FB) By contraposition. Suppose $(\mathbf{M}, s) \not\models B_i K_i \varphi$. We show that $(\mathbf{M}, s) \not\models B_i \varphi$. So let $t, u \in S$ such that $s \rightarrow_i t \sqsubseteq_i u$ and $(\mathbf{M}, u) \not\models \varphi$. By the full belief condition (Definition 5.35), we have $s \rightarrow_i u$, giving us $(\mathbf{M}, s) \not\models B_i \varphi$.

For the group knowledge axioms, let $s \in S$ be a state and let $\varphi \in KB_{i,A}$. Again, let $i \in A$ be arbitrary.

(KBK) *Individual Knowledge and Group Belief imply Group Knowledge*. By contraposition. Suppose that $(\mathbf{M}, s) \not\models K_A \varphi$. We show that either $(\mathbf{M}, s) \not\models K_i \varphi$ or $(\mathbf{M}, s) \not\models B_A \varphi$. So by $(\mathbf{M}, s) \not\models K_A \varphi$, there is $t \in S$ with $s \sqsubseteq_A t$ and $(\mathbf{M}, t) \not\models \varphi$. By the KBK-condition of the pseudo-model for $KB_{i,A}$, having $s \sqsubseteq_A t$ gives us the following case distinction: either (1) $s \sqsubseteq_j t$ for all $j \in A$, or (2), $s \rightarrow_A t$. In the case of (1), we have in particular that $s \sqsubseteq_i t$, giving us that $(\mathbf{M}, s) \not\models K_i \varphi$. In the case of (2), we have $(\mathbf{M}, s) \not\models B_A \varphi$.

(KPB) *Group Knowledge of Positive Individual Beliefs*. By contraposition. Suppose $(\mathbf{M}, s) \not\models K_A B_i \varphi$, with $i \in A$. We show $(\mathbf{M}, s) \not\models B_i \varphi$. So let $t \in S$ with $s \sqsubseteq_A t$ and $(\mathbf{M}, t) \not\models B_i \varphi$. So there exists $u \in S$ with $t \rightarrow_i u$ and $(\mathbf{M}, u) \not\models \varphi$. For $(\mathbf{M}, s) \not\models B_i \varphi$, it suffices to show that $s \rightarrow_i u$. But this follows from the condition of group knowledge of individual beliefs on pseudo-models for $KB_{i,A}$: from $s \sqsubseteq_A t$ and $t \rightarrow_i u$, we conclude that $s \rightarrow_i u$.

(BDK) *Consistency of Group Belief with Distributed Knowledge*. Suppose $(\mathbf{M}, s) \models \bigwedge_{i \in A} K_i \varphi_i$, where $\{\varphi_i \mid i \in A\}$ are arbitrary formulas. We show that $(\mathbf{M}, s) \models \langle B_A \rangle (\bigwedge_{i \in A} \varphi_i)$. By the BDK-condition on pseudo-models for $KB_{i,A}$, there exists $t \in S$ such that both $s \rightarrow_A t$ and $s (\bigcap_{i \in A} \sqsubseteq_i) t$. Now let $i \in A$ be arbitrary. By assumption that $(\mathbf{M}, s) \models \bigwedge_{i \in A} K_i \varphi_i$, and by $s \sqsubseteq_i t$, we know that $(\mathbf{M}, t) \models \varphi_i$. But then, since i was arbitrary, $(\mathbf{M}, t) \models \bigwedge_{i \in A} \varphi_i$. So since $s \rightarrow_A t$, we have $(\mathbf{M}, s) \models \langle B_A \rangle (\bigwedge_{i \in A} \varphi_i)$, as required. \square

Proof of Corollary 5.39, $KB_{i,A}$. We prove completeness with respect to pseudo-models for $KB_{i,A}$ using the standard canonical-model construction (see e.g. [BRV01]). The canonical pseudo-model for $KB_{i,A}$ is defined as follows:

Definition A.5. Fix a finite set of agents A and a finite vocabulary V . The canonical pseudo-model for $KB_{i,A}$ is the structure $\mathbf{M}^C = (S^C, (\sqsubseteq_i)_{i \in A}, (\rightarrow_i)_{i \in A}, \sqsubseteq_A, \rightarrow_A, \pi)$ where S^C is given by the set of maximal consistent theories over $KB_{i,A}(V)$, i.e.

$$S^C := \{T \subseteq KB_{i,A}(V) \mid T \text{ is a maximally consistent set}\}$$

and the relations \sqsubseteq_i and \rightarrow_i for all $i \in A$, and \sqsubseteq_A and \rightarrow_A for the full group A , are given by

$$\begin{aligned} T \sqsubseteq_i W & \quad \text{iff} \quad \forall (K_i \varphi) \in KB_{i,A}(V) (K_i \varphi \in T \Rightarrow \varphi \in W) \\ T \sqsubseteq_A W & \quad \text{iff} \quad \forall (K_A \varphi) \in KB_{i,A}(V) (K_A \varphi \in T \Rightarrow \varphi \in W) \\ T \rightarrow_i W & \quad \text{iff} \quad \forall (B_i \varphi) \in KB_{i,A}(V) (B_i \varphi \in T \Rightarrow \varphi \in W) \\ T \rightarrow_A W & \quad \text{iff} \quad \forall (B_A \varphi) \in KB_{i,A}(V) (B_A \varphi \in T \Rightarrow \varphi \in W) \end{aligned}$$

for $T, W \in S^C$. Finally, we define for all $T \in S^C$

$$\pi(T) := \{p \in V \mid p \in T\}.$$

It can be checked that \mathbf{M}^C is indeed a pseudo-model for $KB_{i,A}$:

Proposition A.6. *The canonical model \mathbf{M}^C is a pseudo-model for $KB_{i,A}$.*

Proof. We check the conditions from Definition 5.35. For all conditions, except for the KBK-condition, the group knowledge of individual beliefs condition, and the BDK-condition, we only show the claim for \sqsubseteq_i : for \sqsubseteq_A , the proofs are analogous (since the Stalnaker conditions for i and A coincide, as well as the definitions of the relations \sqsubseteq_i and \sqsubseteq_A on \mathbf{M}^C).

- The relations \sqsubseteq_i for $i \in A$ are S4.2: let $i \in A$. To see that \sqsubseteq_i is a pre-order, observe that reflexivity follows directly from the (T) axiom and the definition of \sqsubseteq_i on S^C ; similarly, transitivity follows from the (KK) axiom.

It remains to show that \sqsubseteq_i is weakly directed. Let $T_1, T_2 \in S^C$ and suppose $S \sqsubseteq_i T_1$ and $S \sqsubseteq_i T_2$; we show that there exists theory U such that $T_1 \sqsubseteq_i U$ and $T_2 \sqsubseteq_i U$. We use the fact that strong Euclideanity and inclusion (see the proof below) are satisfied by \mathbf{M}^C , and that \rightarrow_i is serial (see the proof below). By seriality \rightarrow_i , there exists theory U such that $S \rightarrow_i U$. With the assumptions that $S \sqsubseteq_i T_1$ and $S \sqsubseteq_i T_2$, strong Euclideanity gives us that $T_1 \rightarrow_i U$ and $T_2 \rightarrow_i U$. Finally, inclusion gives us $T_1 \sqsubseteq_i U$ and $T_2 \sqsubseteq_i U$, as required.

- The relations \rightarrow_i for $i \in A$ are **KD45**: let $i \in A$. To see that \rightarrow_i is serial, let $S \in S^C$ be a theory. It suffices to show that $\{\varphi \mid B_i\varphi \in S\}$ is consistent (since this implies that there exists some maximally consistent set $T \subseteq KB_{i,A}(V)$ such that $T \supseteq \{\varphi \mid B_i\varphi \in S\}$; then $T \in S^C$. So by definition of \rightarrow_i on the canonical model, $S \rightarrow_i T$). So suppose for contradiction that it's not, i.e. that $\vdash \bigwedge \{\varphi \mid B_i\varphi \in S\} \rightarrow \perp$. Then, using B_i -Necessitation and B_i -Distribution, we obtain that $\vdash B_i(\bigwedge \{\varphi \mid B_i\varphi \in S\}) \rightarrow B_i\perp$. Clearly, $S \vdash \bigwedge \{B_i\varphi \mid B_i\varphi \in S\}$. Since $\vdash \bigwedge \{B_i\varphi \mid B_i\varphi \in S\} \rightarrow B_i\bigwedge \{\varphi \mid B_i\varphi \in S\}$, we get that $S \vdash B_i\bigwedge \{\varphi \mid B_i\varphi \in S\}$. Together with the previously obtained that $\vdash B_i\bigwedge \{\varphi \mid B_i\varphi \in S\} \rightarrow B_i\perp$, we have that $S \vdash B_i\perp$. Equivalently, for some $\psi \in KB_{i,A}(V)$ we have $S \vdash B_i(\psi \wedge (\sim \psi))^2$. But then $S \vdash B_i\psi \wedge B_i(\sim \psi)$, directly contradicting axiom (CB). Thus, $\{\varphi \mid B_i\varphi \in S\}$ must be consistent.

Now, to see that \rightarrow_i is transitive, let $S, T, U \in S^C$ and suppose that $S \rightarrow_i T \rightarrow_i U$. We show that $S \rightarrow_i U$. So let $B_i\varphi \in S$. Combining axioms (Strong PI) and (KB), we obtain that $\vdash B_i\varphi \rightarrow B_iB_i\varphi$. So $S \vdash B_iB_i\varphi$. Using our assumption that $S \rightarrow_i T \rightarrow_i U$, we get that $T \vdash B_i\varphi$, and therefore, that $U \vdash \varphi$, as required.

Finally, to see that \rightarrow_i is Euclidean, let $S, T, U \in S^C$ and suppose that $S \rightarrow_i T$ and $S \rightarrow_i U$. We show that $S \rightarrow_i U$, so let $B_i\varphi \in T$. For contradiction, suppose $\varphi \notin U$. By maximal consistency of U , $(\sim \varphi) \in U$. Since $S \rightarrow_i U$, we know that $B_i\varphi$ cannot be in S . So by maximal consistency of S , $(\sim B_i\varphi) \in S$. Now by the (Strong PI) and (Strong NI) axioms, $K_A(\sim B_i\varphi) \in S$. We apply the (KB) axiom to obtain $B_i(\sim B_i\varphi) \in S$. With $S \rightarrow_i T$, this gives us that $(\sim B_i\varphi) \in T$, which directly contradicts the (CB) axiom. Therefore, we must have $\varphi \in U$.

- Inclusion: let $i \in A$. Let $S, T \in S^C$ and suppose $S \rightarrow_i T$; we show that this implies $S \sqsubseteq_i T$, so let $K_i\varphi \in S$. Given maximal consistency of S and the (KB) axiom, $B_i\varphi \in S$. But then, by the (FB) axiom, we get $B_iK_i\varphi \in S$, which gives us by the assumption that $S \rightarrow_i T$ that $K_i\varphi \in T$. By the (T) axiom, $\varphi \in T$, as required.

²Like in Section 5.1, \sim denotes the single negation.

- Strong transitivity: let $i \in A$. Let $S, T, U \in S^C$ and suppose that $S \sqsubseteq_i T \rightarrow_i U$. We show that $S \rightarrow_i U$. So let $B_i\varphi \in S$. By the (Strong PI) axiom, $K_i B_i\varphi \in S$. By $S \sqsubseteq_i T$, we get $B_i\varphi \in T$. From the assumption that $T \rightarrow_i U$, we obtain $\varphi \in U$, as required.
- Strong Euclidean: let $i \in A$. Let $S \sqsubseteq_i T$ and $S \rightarrow_i U$. We show that $T \rightarrow_i U$. So let $B_i\varphi \in T$; we show that $\varphi \in U$. For contradiction, suppose not: then, by maximal consistency, $(\sim \varphi) \in U$. Then we know by $S \rightarrow_i U$ that $B_i\varphi$ cannot be in S , thus, by maximal consistency, $(\sim B_i\varphi) \in S$. Recall that $S \sqsubseteq_i T$, and therefore, $K_i(B_i\varphi)$ cannot be in T . But we know that $K_i(B_i\varphi) \in T$ by applying the (Strong PI) axiom to the assumption that $B_i\varphi \in T$. So we conclude that $\varphi \in U$.
- Full belief: let $i \in A$. Let $S, T, U \in S^C$ and suppose $S \rightarrow_i T \sqsubseteq_i U$; we show that $S \rightarrow_i U$. So let $B_i\varphi \in S$. By the (FB) axiom, $B_i K_i\varphi \in S$. Now, by our assumption that $S \rightarrow_i T \sqsubseteq_i U$, we obtain that $K_A\varphi \in T$ and $\varphi \in U$, as required.
- KBK-condition: by contraposition. Let $i \in A$, let $S, T \in S^C$, and suppose that $S \not\sqsubseteq_i T$, and $S \not\rightarrow_A T$. We show that this implies $S \not\sqsubseteq_A T$. By $S \not\sqsubseteq_i T$, there exists $K_i\varphi \in S$ such that $\varphi \notin T$. Since $S \not\rightarrow_A T$, there is also $B_I\psi \in S$ such that $\psi \notin T$. Now consider the formula $K_A\varphi \vee \psi$.

Claim A.7. $K_A(\varphi \vee \psi) \in S$ and $\varphi \vee \psi \notin T$.

Proof of claim. First of all, we have $K_A(\varphi \vee \psi) \in S$. Observe that $K_i(\varphi \vee \psi) \in S$, since the validity $K_i\varphi \rightarrow_A K_i(\varphi \vee \psi)$ is derivable from Stalnaker's axioms. The claim then follows from having $K_i\varphi \in S$ and maximal consistency of S . Analogously, $B_I(\psi) \in S$ implies that $B_I(\varphi \vee \psi) \in S$. Now the (KBK) axiom gives us that $K_A(\varphi \vee \psi) \in S$.

However, we have that $\varphi \vee \psi \notin T$: suppose for contradiction that $\varphi \vee \psi \in T$. Then, by maximal consistency of T , either φ or ψ must be an element. But recall that $\varphi \notin T$ and $\psi \notin T$, therefore, we have reached a contradiction. \square

Having proved Claim A.7, we can conclude that $S \not\sqsubseteq T$, as required.

- Group knowledge of individual beliefs: Let $S, T \in S^C$ and suppose $S \sqsubseteq_A T$. We show that this implies that for all $i \in A$, for all $U \in S^C$, that $S \rightarrow_i U$ if and only if $T \rightarrow_i U$.
For the left-to-right direction, let $i \in A$ and let $U \in S^C$. Suppose that $S \rightarrow_i U$. For contradiction, let $B_i\varphi \in T$, but $\varphi \notin U$. By maximal consistency, $(\sim \varphi) \in U$. So we know from $S \rightarrow_i U$ that $B_i\varphi \notin S$. By maximal consistency, $(\sim B_i\varphi) \in S$. Now, by the (Strong PI) and (Strong NI) axiom, we have $K_A(\sim B_i\varphi) \in S$. Having assumed that $S \sqsubseteq T$, $(\sim B_i\varphi) \in T$, contradicting the assumption that $B_i\varphi \in T$. Thus, it must be that $\varphi \in U$.
For the converse direction, let $i \in A$ and let $U \in S^C$; suppose that $T \rightarrow_i U$. For contradiction, let $B_i\varphi \in S$, but $\varphi \notin U$. By maximal consistency, $(\sim \varphi) \in U$. So we know from $T \rightarrow_i U$ that $B_i\varphi \notin T$. By maximal consistency, $(\sim B_i\varphi) \in T$. We claim that $B_i\varphi$ is also in T : with $B_i\varphi \in S$, the (Strong PI) axiom gives us $K_A B_i\varphi \in S$. So by $S \sqsubseteq T$, we have $B_i\varphi \in T$, contradicting the fact that $(\sim B_i\varphi) \in T$. Thus, it must be that $\varphi \in U$.
- BDK-condition: let S be a theory. We show that there exists W such that $S (\rightarrow_A \cap \bigcap_{i \in A} \sqsubseteq_i) W$. We need the Diamond Lemma for the belief relation \rightarrow_A in order to prove this. It is presented in Lemma A.8.

Now, in order to prove that there exists W such that $S (\rightarrow_A \cap \bigcap_{i \in A} \leq_i) W$, we need to show that the set W_2 is consistent, which is defined as follows:

$$\begin{aligned} W_1 &:= \{\psi \mid B_A \psi \in S\} \\ W_2 &:= W_0 \cup W_1. \end{aligned}$$

We show that every finite subset of W_2 is consistent. First, let $U \subseteq W_0$ be any finite subset of W_0 . We show that U is consistent with W_1 . After proving that all finite subsets of W_0 are consistent with W_1 , we can conclude that W_2 is a consistent set.

Given $U \subseteq W_0$, we add the tautology \top to U to obtain $U' := U \cup \{\top\}$. Now, for each $\varphi_i \in U'$, there is some $i \in A$ such that $K_i \varphi \in S$. To see this, let $i \in A$. If there is no $\varphi_i \in U$ for which the claim already holds, then it holds by the fact that we have $K_i \top \in S$. As a result, U' is a subset of S by the (T) axiom. Furthermore, $U' = U \cup \{\top\}$ is finite, since U was assumed to be finite.

Let $i \in A$ be arbitrary and consider the set

$$U_i := \{\varphi_i \in U' \mid K_i \varphi \in S\}.$$

Each U_i is finite (being a subset of the finite set U') and therefore, $\bigwedge U_i$ is a formula. By the (T) axiom, $U_i \subseteq S$, so by maximal consistency of S , we have $(\bigwedge U_i) \in S$. Now consider the formula

$$\bigwedge_{i \in A} (\bigwedge U_i)$$

which must also be in S , and note that it is of the form $\bigwedge_{i \in A} \varphi_i$. It follows from the (KBK) axiom that the formula $\neg B_A \neg (\bigwedge_{i \in A} (\bigwedge U_i)) \in S$. We apply the Diamond Lemma (Lemma A.8) to obtain $V \in S^C$ such that $(\sim (\neg (\bigwedge_{i \in A} (\bigwedge U_i)))) = \bigwedge_{i \in A} (\bigwedge U_i) \in V$ and $S \rightarrow_A V$. By the semantics for B_A , we know that $W_1 \subseteq V$. Since V is a maximally consistent set, we must also have $U_i \subseteq V$ for each $i \in A$. Thus, $U' \cup W_1$ is consistent, and so is $U \cup W_1$.

Since U was an arbitrary finite subset of W_0 and it was consistent with W_1 , we conclude that $W_2 = W_0 \cup W_1$ must be consistent. Using the Lindenbaum Lemma, we extend W_2 into a maximally consistent set W . Since we have $W \in S^C$, we get by $W_0 \subseteq W$ that $S (\bigcap_{i \in A} \leq_i)$; and by $W_1 \subseteq W$ that $S \rightarrow_A W$. This concludes our proof. □

To prove completeness with respect to the canonical model, we use the Truth Lemma. As we previously mentioned, it is a standard lemma in canonical-model constructions (see e.g. [BRV01]). To prove the Truth Lemma, we first need to prove the Diamond Lemmas for each of the knowledge and belief relations, for individual $i \in A$ and for the group A . As the claim for the full group is equivalent to the claim for individual agents, we use the $\alpha \in \{A\} \cup A$ notation.

Lemma A.8 (Diamond Lemma, $\rightarrow_i, \rightarrow_A$). *Let $\alpha \in \{A\} \cup A$. If $\neg B_\alpha \varphi \in S \in S^C$, then there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $S \rightarrow_\alpha W$.*

Proof. We only prove the claim for individual agents, because the proof for the full group completely mirrors the proof for individual agents. Let $i \in A$, let $T \in S^C$ be a theory with $\neg B_i \varphi \in T$, and let

$S := \{\psi \mid B_i\psi \in T\}$. We show that $S \cup \{\sim \varphi\}$ is consistent. For contradiction, suppose not: then $\vdash (\wedge S) \rightarrow \varphi$. Applying B_i -Necessitation and B_i -Distribution, we obtain that $\vdash B_i(\wedge S) \rightarrow B_i\varphi$, i.e. $\vdash \wedge\{B_i\psi \mid B_i\psi \in T\} \rightarrow B_i\varphi$.

Now observe that, since $\{B_i\psi \mid B_i\psi \in T\} \subseteq T$, we have $\vdash (\wedge T) \rightarrow \wedge\{B_i\psi \mid B_i\psi \in T\}$. But this gives us $\vdash (\wedge T) \rightarrow B_i\varphi$. By maximal consistency, this implies that $B_i\varphi \in T$. But this directly contradicts the assumption that $\neg B_i\varphi \in T$, and therefore we conclude that $S \cup \{\sim \varphi\}$ must be consistent.

This means that we can use the Lindenbaum Lemma to construct a maximal consistent set $W \subseteq S^C$, such that $(S \cup \{\sim \varphi\}) \subseteq W$. Then $(\sim \varphi) \in W$, and having $\psi \in W$ for all $B_i\psi \in T$, we can conclude that $T \rightarrow_i W$. \square

Lemma A.9 (Diamond Lemma, $\trianglelefteq_i, \trianglelefteq_A$). *Let $\alpha \in \{A\} \cup A$. If $(\neg K_\alpha \varphi) \in T \in S^C$, then there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \trianglelefteq_\alpha W$.*

Proof. We only prove the claim for individual agents; let $i \in A$. Let $(\neg K_i \varphi) \in T \in S^C$ and let $S := \{\psi \mid K_i\psi \in T\}$. We show that $S \cup \{\sim \varphi\}$ is consistent. For contradiction, suppose not: then $\vdash (\wedge S) \rightarrow \varphi$. Applying K_i -Necessitation and K_i -Distribution, we obtain that $\vdash K_i(\wedge S) \rightarrow K_i\varphi$. Since $K_i(\wedge S) = \wedge\{K_i\psi \mid K_i\psi \in T\}$, this implies that $\vdash \wedge\{K_i\psi \mid K_i\psi \in T\} \rightarrow K_i\varphi$. But $\{K_i\psi \mid K_i\psi \in T\} \subseteq T$. So we have $\vdash (\wedge T) \rightarrow \wedge\{K_i\psi \mid K_i\psi \in T\}$. Combined with the fact that $\vdash \wedge\{K_i\psi \mid K_i\psi \in T\} \rightarrow K_i\varphi$, we obtain that $\vdash (\wedge T) \rightarrow K_i\varphi$. By maximal consistency of T , $K_i\varphi \in T$, contradicting our assumption that $(\neg K_i\varphi) \in T$. Thus, $S \cup \{\sim \varphi\}$ must be consistent.

This means that we can use the Lindenbaum Lemma to construct a maximal consistent set $W \subseteq S^C$, such that $(S \cup \{\sim \varphi\}) \subseteq W$. Then $(\sim \varphi) \in W$, and having $\psi \in W$ for all ψ such that $K_A\psi \in T$, we can conclude that $T \trianglelefteq_i W$. \square

The proof of the Truth Lemma is straightforward and similar to Truth Lemma 5.13 in the completeness proof for $\Box[\forall]_I$ with respect to the filtration of the canonical pseudo-model:

Lemma A.10 (Truth Lemma). *Given the canonical pseudo-model for $KB_{i,A} \mathbf{M}^C$, we have for all $\varphi \in KB_{i,A}$:*

$$T \models_{sc} \varphi \text{ iff } \varphi \in T, \text{ for every } T \in S^C.$$

Proof. By induction on the complexity of φ . The atomic case is taken care of by the canonical valuation and the Boolean cases are trivial. So we focus on the cases involving modalities B_i, B_A, K_i , and K_A . We only prove the claim for $B_i\varphi$ and $K_i\varphi$, as the proofs for $B_A\varphi$ and $K_A\varphi$ are completely analogous and are obtained by replacing each occurrence of B_i with B_A , and each occurrence of K_i with K_A , respectively.

Inductive step for B_i . For the left-to-right direction, suppose $T \models B_i\varphi$, for some $B_i\varphi \in KB_{i,A}$. Now suppose for contradiction that $B_i\varphi \notin T$. Then, by maximal consistency of T , we have $\neg B_i\varphi \in T$. By the Diamond Lemma for \rightarrow_i , there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \rightarrow_i W$. With $T \models B_i\varphi$, this implies that $W \models \varphi$; by the induction hypothesis, we obtain that $\varphi \in W$, which directly contradicts the fact that $(\sim \varphi) \in W$. We conclude that $B_i\varphi \in T$.

For the converse direction, suppose $B_i\varphi \in T$. Recall that by our definition of \rightarrow_i in the canonical pseudo-model for $KB_{i,A}$, we have $\varphi \in W$ for all $W \in S^C$ such that $T \rightarrow_i W$. Let such W be arbitrary. To prove that $T \models B_i\varphi$, it suffices to show that $W \models \varphi$. But this follows directly from the fact that $\varphi \in W$: applying the induction hypothesis gives us that $W \models \varphi$.

Inductive step for K_i . For the left-to-right direction, suppose $T \models K_i\varphi$, for some $K_i\varphi \in KB_{i,A}$. Now suppose for contradiction that $K_i\varphi \notin T$. Then, by maximal consistency of T , we have $\neg K_i\varphi \in T$. By

the Diamond Lemma for \leq_i , there exists $W \in S^C$ such that $(\sim \varphi) \in W$ and $T \leq_i W$. With $T \models K_i \varphi$, this implies that $W \models \varphi$; by the induction hypothesis, we obtain that $\varphi \in W$, which directly contradicts the fact that $(\sim \varphi) \in W$. We conclude that $K_i \varphi \in T$.

For the converse direction, suppose $K_i \varphi \in T$. Recall that by our definition of \leq_i in the canonical pseudo-model for $KB_{i,A}$, we have $\varphi \in W$ for all $W \in S^C$ such that $T \leq_i W$. Let such W be arbitrary. To prove that $T \models K_i \varphi$, it suffices to show that $W \models \varphi$. But this follows directly from the fact that $\varphi \in W$: the induction hypothesis gives us that $W \models \varphi$, as required. \square

This allows us to prove our claim:

Proof of Corollary 5.39. Soundness was established in Proposition 5.38. For completeness, let $\varphi_0 \in KB_{i,A}$ be any consistent formula and consider the canonical pseudo-model \mathbf{M}^C for $KB_{i,A}$. By the Lindenbaum Lemma, there exists some maximally consistent theory T_0 in \mathbf{S}^C with $\varphi_0 \in T_0$. By the Truth Lemma A.10, T_0 satisfies φ_0 in \mathbf{S}^C . This gives us weak completeness with respect to pseudo-models for $KB_{i,A}$. \square

\square