Topological Models for Group Knowledge and Belief

MSc Thesis *(Afstudeerscriptie)*

written by

Aldo Iván Ramírez Abarca
(born March 5, 1987 in México Distrito Federal, México)

under the supervision of **Dr. Alexandru Baltag**, and submitted to the Board of Examiners 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:**
August 14, 2015	Dr. Nick Bezhanishvili
	Prof. Dr. Dick de Jongh
	Prof. Dr. Ronald de Wolf

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

The advantage of using Topology in formal models of knowledge and belief comes from the spatial intuitions it endorses. In the possible-worlds framework, standard topological semantics for knowledge and belief are given in terms of the *interior* operator, and the interpretation is unavoidably connected with traditional Kripke semantics. A recent development in the field is the use of Topology to model *evidence*, and in turn evidence-based belief and knowledge. In contrast to the standard approach, the evidential one does not characterize knowledge in terms of interior. In both approaches, single-agent semantics are fairly consolidated. The present work inaugurates the use of the second approach -evidential in nature- to model epistemic group notions, and in particular *distributed knowledge*. We offer a sound and complete axiomatization of a topological logic of distributed knowledge that accounts for a notion of belief based on evidence.

Acknowledgements

All my thanks go to Alexandru Baltag. "When it's time to face the music, he is the one who knows where the volume knob is. He knows that you know he knows you know".

I would also like to express my gratitude to Dr. Nick Bezhanishvili, Prof. Dr. Dick de Jongh, and Prof. Dr. Ronald de Wolf, for accepting to read this work.

Whatever the following pages are–whatever they turn out to be–they are dedicated to Ana, Nigel, and Konstantinos.

Contents

Introduction

The following pages deal with formal representations of knowledge and belief. As such, we like to think that they mean some contribution to the research field of Formal Epistemology.

Formal representations of knowledge and belief

Our line of work is part of a huge tradition that searches for a formal theory of the properties of rational agents within a given environment. Such agents are typically assumed to be persons, computer programs, robots, or the like. They are best identified with what philosophers have come to present as *intentional systems* (see [52]), whose actions in the environment are determined by certain mental states. Even though there are many components to such mental states, research has favored a convenient categorization of them. Among said categories, we will be focusing on the so-called 'information attitudes,' i.e., the attitudes that an agent has towards the *information* perceivable in the environment ([52]). The most prominent members in the category of information attitudes are the concepts of *knowledge* and *belief*.

Hintikka ([31]) was the first to use ideas of Kripke's Modal Logic to formalize (and axiomatize) *knowledge* and *belief*. Ever since this breakthrough (in the early 60's), many epistemologists have followed his example and offered a plethora of results concerning the logics that ensue from interpreting the formulas of formal epistemic languages over Kripke frames. The underlying intention has always been to use the power and reach of Logic and Mathematics to describe with utmost precision whatever properties of knowledge and belief one would want to convey.

In this line, it was only natural that Topology, a powerful mathematical theory of space, was eventually applied in the modelling. Such application was highly relevant, especially ever since McKinsey and Tarski laid forward a topological interpretation of Modal Logic in the 40's (see [48]). Topological modelling in Epistemic Logic can be recognized as a field by itself nowadays, and the variety of 'topological' tools that one can use to tackle problems of formal epistemology has led to a lush literature.

The standard—and relatively prevalent—view is that *knowledge* corresponds to an **S4** modality given by the operator *int* (standing for *interior*) over arbitrary topological spaces. Its prevalence is due to several reasons, some of them conceptual, some of them purely technical. The fact of the matter is that *int* behaves as an **S4** necessity modal operator, and thus satisfies axioms that are agreeable—even desirable—in the philosophical concept of knowledge. In principle, we can say that the association of **S4** to topological spaces has nothing to do with epistemology. It is only through an epistemological reading that a notion of topological knowledge arises. From there, the options are, if not endless, surely quite vast. We could try to use Topology and topological models to interpret virtually any epistemic notion, and thus obtain results and new intuitions phrased in spatial terminology. It is all about using the tools at hand with sensible philosophical intuition.

For the epistemic atittude of *belief*, for example, there has been more than one candidate proposed for its topological interpretation, and the motivations are diverse. However, one could say that all of them stem from the desire to capture *some* inherent philosophical intuition, even if taken to its most formal phrasing (in terms of axioms).

Since any frame based on a preorder is in fact a topological space (an *Alexandrov* space), the topological interpretations of these individual information attitudes are linked to their standard relational depictions. Since standard relational depictions of knowledge (and belief) open the door to epistemic *group attitudes* in multi-agent settings—such as common knowledge, distributed knowledge, and common belief—these group notions have also been coupled with various topological interpretations. This is where we come in.

In this thesis

The starting point of this thesis is the aforementioned confluence between the streams of Kripke semantics for Epistemic Logic, on one hand, and the topological semantics for Modal Logic, on the other. Whereas in Epistemic Modal Logic we have that the accessibility relations on a state space reflect the agents' information attitudes, in the topological setting these are reflected by topologies on the state space.

With a considerable amount of work already done in the single-agent cases (where there are only individual operators for knowledge and/or belief), we took it as our job to explore logics for *group* attitudes by use of topological models, always keeping close at hand the relational counterpart and usual axiomatizations. Thus, our primary target is 'the multi-agent case.' In particular, we focus on the concept of *distributed knowledge*.

As one would expect, there already exist topological semantics for *distributed knowledge* in the literature, that are logically well-behaved and interesting in themselves (see [50], [48]). The originality of our approach, then, comes from a relatively new—based on the concept of *evidence*.

The topological semantics that pervade throughout our original contribution are *not* traditional, in the sense that they do not identify knowledge with *interior*. They are what we could call a variation on the evidential (single-agent) semantics introduced by Van Benthem and Pacuit in [49]. The latter authors devised so-called *evidence models* to evaluate the formulas of a language having modal evidential operators, designed to encode the notion of *having evidence for* something. However, it turns out that these evidence models naturally give rise to particular topological models for evidence, which we call *topological evidential models, topo-evidential models*, or *topo-e-models*, in short. Provided with these new tools, it is possible to construct a rich single-agent logic of knowledge and belief based on evidence, and this is what Baltag *et alia* do in their unpublished work [4]. It is precisely this work that motivates our whole enterprise, which we can state in (hopefully) clear terms as follows:

The main objective of the present thesis is to provide a sound and complete axiomatization for a logic of distributed knowledge with respect to topological evidential models that account for a clear notion of belief.

The background material for our work is provided mainly by two groups of sources. One—concerning the (standard) topological *interior semantics* for knowledge and belief—is given by [5], [50], [41], [48], and [42]. The other—concerning the evidence models and topo-evidential models—is given by [4].

The thesis is organised as follows:

- (Part I) The whole of Part I is dedicated to standard topological semantics (i.e., in terms of '*knowledge as interior*.')
  - Chapter 1 provides a sturdy background of (standard) relational and topological semantics for basic logics of knowledge and belief, in the single-agent case. First, we introduce the basic Kripke semantics for usual axiomatizations of knowledge, of belief, and for Stalnaker's combined system of knowledge and belief. Secondly, we review the topological interior semantics for knowledge, and the ensuing topological semantics for belief as rendered by Stalnaker's axioms. Here, we compile the results of soundness and completeness of the systems **S4.2** (for the knowledge-language), **KD45** (for the belief-language), and Stalnaker's **KB** (for the combined language) with respect to the class of extremally disconnected topological spaces. Such results provide substantial technical backup for the according 'soundness and completeness' proof regarding our new logic. The driving force of this chapter comes from [5] and [41].
  - Chapter 2 is dedicated to the review of the multi-agent (standard) relational and topological semantics, both for common and distributed knowledge. Once again, we first introduce the basic Kripke semantics, and then proceed to address the topological counterpart. The definitions and observations supply technical intuition for several mathematical subtleties as to our own—topological evidential—definition of distributed knowledge. We address soundness and completeness of the system $\mathbf{S4}_2^{\mathcal{DC}}$ with respect to the class of bi-topological spaces. The result is significant and, once again, technically linked to our 'soundness and completeness' proof. The main sources used are [50], [27], [43], and [48].
- (Part II) The whole of Part II is dedicated to topological evidential models for knowledge and belief.

– In Chapter 3, we introduce *evidence models* for a language of evidence-based belief. This is done as motivation for the subsequent presentation of Baltag *et alia*'s topo-evidential models for basic logics of knowledge and belief (single-agent case) ([4]). These models provide the most specific background for our original work. The 'soundness and completeness' results for the systems **S4.2** (for the knowledge-language), **KD45** (for the belief-language), and Stalnaker's **KB** (for the combined language) with respect to the class of topo-evidential models, are also recovered, in accordance with later needs.

– Chapter 4 launches the original part of the thesis. We introduce *two-agent topological (evidential) models* for a language with individual and distributed knowledge operators. The logic rendered is named TKD (standing for Topological Logic of Knowledge and Distributed Knowledge). We justify our choice of semantics with arguments inspired by the evidential paradigm. After that, we present an axiom system $\mathbf{\Lambda}_{TKD}$ for the logic TKD. We briefly discuss its axioms, and state the main Theorem of the thesis: that $\mathbf{\Lambda}_{TKD}$ is sound and weakly complete for the pertinent language with respect to the class of two-agent topological (evidential) models.

– Chapter 5 is dedicated in its entirety to proving the main Theorem mentioned above. The proof of soundness relies on many arguments expounded in the previous chapters. The proof of (weak) completeness is technically the hardest result of this thesis. It involves many phases, which are all properly addressed, and carefully exposed. This proof can be singled out as the most important contribution of the present work.

– Finally, Chapter 6 closes the thesis with some further directions for our proposal, as well as a succinct discussion regarding the conceptual aspect of our original constructions.

Since our work comprises a wide variety of intuitions, interpretations, constructions, and conceptual arguments, each chapter includes an introduction section that intends both to motivate its content and to help the reader keep track of the structure of intent.

We stress the fact that all original contributions of this thesis are part of an ongoing investigation, the bedrock of which is set by the unpublished work by Baltag *et alia* in [4].

Part I

Standard Topological Semantics for Knowledge and Belief

Chapter 1

Background. Single-agent models

In this chapter, we present a survey of standard relational and topological models for logics of knowledge and belief in the single-agent case. This is meant as background for the particular topological (multi-agent) interpretation that we will introduce in Chapter 4. Our discussion lies within a specific branch of formal epistemology, where epistemic/doxastic languages are used as the syntactic counterpart to semantic representations of knowledge and belief. The most standard (relational) interpretation employs Kripke-frames based on sets of possible worlds, or state spaces. In turn, the topological interpretation generalizes relational semantics by use of topological spaces over sets of possible worlds.

Though we will not go very deep into the *ontology* of the possible-worlds framework—or of the relational and topological semantics—it should be of service to motivate the content of this chapter with some of the conceptual notions supporting the formalization effort that we are focusing on.

According to [5] and [41], formal epistemologists have tackled the problem of finding the correct relation between knowledge and belief from two main angles:

1. Belief first. Provided with a sensible definition for *belief*, the concept of *knowledge* is interpreted as a kind of powerful belief, i.e., a belief that meets certain conditions with respect to the believer's environment and her/his development inside it. Ever since Plato addressed the problem of describing the nature of the relation between knowledge and belief, there has been overall acceptance of the principle that knowledge should be—at the very least—*true* and *justified* belief. The so-called 'JTB' *characterization* of knowledge (knowledge as *justified true belief*), though, was shattered by the Gettier counterexamples ([24]), which offer a scenario in which an agent believes that ϕ, (s)he is justified in believing that ϕ, but the belief cannot be considered as knowledge.[1] The counterexamples generated a lively debate among philosophers that goes on to this day. Searching for the missing ingredient that would make a justified true belief also knowledge, many have proposed interesting ways of addressing the problem. One of the most profitable accounts is due to Lehrer and Paxson in the *Defeasibility Analysis* of Knowledge ([33]). Their main tenet is that knowledge is justified true belief that cannot be defeated by the addition of any new true evidence; i.e., knowledge as *undefeated justified true belief.*

2. Knowledge first. Starting from a chosen notion of knowledge, one weakens it to obtain a well-behaved interpretation of belief. According to [5], this approach has not been given a lot of attention among epistemologists, but is highly favored by Williamson (see [55]), and by Stalnaker (see [46]). Stalnaker's formalization of belief in terms of knowledge serves as basis for the (standard) topological semantics of knowledge and belief that we review in this chapter.

[1] A typical Gettier-style counterexample goes as follows. Suppose that every morning I hear live guitar music coming from my neighbor Konstantinos's apartment. Moreover, I have seen a guitar inside his apartment, and many-a-time I have ran into him in the street while he is carrying a guitar case. Therefore, I have strong evidence to believe that Konstantinos plays guitar, and thus that *one of my neighbors plays guitar.* However, Konstantinos does not play guitar. Unbeknownst to me, it is his girlfriend Kriss who plays the guitar every morning in their apartment, and every time that I saw Konstantinos in the street with the guitar case, it was because he was bringing the guitar back home from Kriss's guitar lesson, who does not like to carry the guitar case while cycling back home. In this scenario, I am justified in having the true belief that *one of my neighbors plays guitar*, but it cannot amount to knowledge, since it was derived from the false premise that Konstantinos plays guitar.

1.1 Kripke semantics

1.1.1 Knowledge

Kripke frames allow us to evaluate the formulas of a modal language. They are tools commonly used to model *knowledge* and *belief*. In what follows, we will proceed with order, first addressing a formal language $\mathcal{L}_K$ having only a knowledge-operator, then a formal language $\mathcal{L}_B$ having only a belief-operator, and finally we will incorporate both operators into a full epistemic/doxastic language $\mathcal{L}_{KB}$.

Let $\mathcal{L}_K$ consist of a countable set *Prop* of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operator K. The grammar of this language is given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid K\phi \tag{1.1}$$

for $p \in \mathit{Prop}$. Naturally, the abbreviations for the connectives $\vee$, $\rightarrow$, $\leftrightarrow$ are the ones inherited from propositional logic. The (epistemic) possibility operator $\langle K\rangle$ is defined as $\neg K \neg$.

Definition 1.1.1. (Epistemic Kripke models) A *Kripke frame* consists of a tuple (X, R_K), where X is a non-empty set of possible worlds or states, and $R_K \subseteq X \times X$ is a binary relation on X (called the *accessibility* relation). An *epistemic Kripke model* is a tuple $(X, R_K, \mathcal{V})$, where $\mathcal{V} : \mathit{Prop} \rightarrow \mathcal{P}(X)$ is a valuation function. The semantics for the formulas of $\mathcal{L}_K$ is given recursively by the following rules of model-*satisfaction*:

For a given world $w \in X$,

$$\begin{array}{lll}
\mathcal{M}, w \Vdash \phi & \text{iff} & w \in \mathcal{V}(p) \\
\mathcal{M}, w \Vdash \neg\phi & \text{iff} & \mathcal{M}, w \nVdash \phi \\
\mathcal{M}, w \Vdash \phi \wedge \psi & \text{iff} & \mathcal{M}, w \Vdash \phi \text{ and } \mathcal{M}, w \Vdash \psi \\
\mathcal{M}, w \Vdash K\phi & \text{iff} & \forall v \in X, \quad [wR_K v \;\Rightarrow\; \mathcal{M}, v \Vdash \phi]
\end{array}$$

A formula ϕ is said to be *true* at w in $\mathcal{M}$ if $\mathcal{M}, w \Vdash \phi$. It is said to be *true* in $\mathcal{M}$ if $\mathcal{M}, v \Vdash \phi$ for every $v \in X$. A formula ϕ is said to be *valid* on a frame (X, R) if it is true in every model based on that frame. It is said to be *valid* on a class $\mathbb{C}$ of frames if it is valid on every frame in $\mathbb{C}$.

We define the *extension* of our formulas in the traditional way: for a given formula ϕ of $\mathcal{L}_K$, $[\![\phi]\!]^{\mathcal{M}} := \{w \in X \mid \mathcal{M}, w \Vdash \phi\}$. If the context of the models is clear, we shall avoid indexing these sets with $^{\mathcal{M}}$.

Notice that a given formula ϕ of $\mathcal{L}_K$ is true in a Kripke model $(X, R_K, \mathcal{V})$ if it has a *global* extension, i.e., if $[\![\phi]\!] = X$.

It is rather obvious that we use the epithet *epistemic* because K, and its corresponding accessibility relation R_K, are intended to represent knowledge. As such, the particular properties that we ask of R_K are going to have to reflect the desired qualities of the concept. For example, if the relation R_K is reflexive, then K reflects *factivity*, meaning that a proposition that is known at a given world w must be true at w ($\mathcal{M}, w \Vdash K\phi$ implies that $\mathcal{M}, w \Vdash \phi$). If R_K is transitive, then K reflects *positive introspection*, ($\mathcal{M}, w \Vdash K\phi$ implies that $\mathcal{M}, w \Vdash KK\phi$).

It is well established that some relational properties correspond to particular modal axioms, and are even characterized by them: any frame that validates the axiom will have the corresponding relational property, and viceversa. Also, it turns out that such axioms are sound and complete with respect to an according class of Kripke frames. Following [41], we present a table that includes some relevant relational properties and their respective defining modal axioms, as well as the epistemological reading of such qualities, where applicable.

Name	Axiom	Relational Property	Epistemic Property
(K)	$K(p \to q) \to (Kp \to Kq)$	Modal relation	*Logical Omniscience*
(T)	$Kp \to p$	Reflexivity	*Factivity*
(4)	$Kp \to KKp$	Transitivity	*Positive Introspection*
(.2)	$\langle K \rangle Kp \to K\langle K \rangle p$	Directedness	*Stalnaker's condition* (to be addressed later)
(5)	$\langle K \rangle p \to K\langle K \rangle p$	Euclideanness	*Negative introspection*
(B)	$p \to K\langle K \rangle p$	Symmetry	Implies "*Infallibility*"

Table 1.1: Some axioms for knowledge

This table deserves some commentary. Some of the most common axiomatizations for knowledge come from well-known modal proof systems. Typically, the proof systems associated to knowledge are **S4**, **S4.2**, and **S5**. There is no small amount of debate as to which one is better, and we shall only scratch the surface of the dilemma. Recall that the proof system **S4** is comprised by the axioms $(K)+(T)+(4)$, all propositional tautologies, and the rules of inference *Modus Ponens* and *Necessitation* for K. **S4.2** corresponds to a system with the same rules of inference and including all propositional tautologies, but generated by the axiomatic extension $(K)+(T)+(4)+(2)$. **S5** is given by the axiomatic extension $(K)+(T)+(4)+(5)$.

It is clear that the (K) axiom implies that agents are idealized logical thinkers. The axiom entails that they will know all the logical consequences of their knowledge. This issue is controversial, to say the least, and we refer the reader to [44] for an interesting exposition of the problem. *Negative introspection* is quite debatable as well. It seems unlikely that someone would know that they do not know something. This characteristic is contended by virtually all philosophers, and the Voorbrak's paradox ([53]) shows how it implies that knowledge becomes the same as 'belief in knowledge' in traditional epistemic/doxastic representations, which is undesirable. Once again, we refer the reader to [35] for germane examination of the matter. We should also keep in mind that, though much more appealing than its negative counterpart, *positive introspection* has also been taken to test by philosophers. Hintikka proposed 'logical'—rather than 'philosophical'—reasons for supporting its pertinence in the possible-worlds framework, since it is implied by the consistency of one's knowledge with those things unknown (see [46], p. 172). On the other front, Williamson, for example, advocates interesting reasons for rejecting the principle (see [55]).

From Table 1.1, we can see that if we axiomatize our logic of knowledge with the proof system **S4**, then we are talking about a *factive*, *positively introspective* type of knowledge. Since in the context of *reflexivity* and *transitivity*, *euclideanness* actually implies *symmetry*, a choice of **S5** means that we are dealing with an equivalence accessibility relation.[2] Knowledge will turn out to be *fully introspective*, then. Quoting Goldbach in [25], the agents will be "fully aware of the extent of their knowledge" (p. 13). From another philosophical angle, **S5** knowledge can be considered as absolute, *hard* knowledge—infallible in the sense of not being subject to defeat. The concept of *(in)fallibility* is connected to the *defeasibility analysis* of knowledge. Traditionally, grounding the concept of knowledge upon an equivalence relation yields that the possible worlds that make up the domain of our models are in fact epistemically indistinguishable. Knowledge of a proposition will be belief that cannot be revised upon receiving new information, no matter whether the information received is *soft* (possibly false) or *hard* (true) (see [11]). In this sense, it can be regarded as infallible.[3] As for **S4.2**, whose singular characteristic is that it defines *directed* frames, it will correspond to a kind of knowledge that we will call Stalnaker's knowledge. The reason is that Robert Stalnaker affirms that the true underlying logic of knowledge is given by **S4.2**. But this seemingly bold claim does not come from nowhere. Since it is closely tied to his conception of *belief*, we first introduce a formal way of referring to beliefs before discussing Stalnaker's take. For the time being, we just mention that **S4.2** is also an interesting system to axiomatize a logic of knowledge. Without going any further, the axiom that singles out this system, (.2), can be read as 'if it is possible to come to know something, then the epistemic possibility (of that something) is certain.'

[2] The **S5** brand of knowledge is favored by most theoretical computer scientists and economists.

[3] We will briefly explore some of these topics in successive chapters, when we set on a conceptual exploration of our own topological evidential interpretation of knowledge.

1.1.2 Belief

As done with knowledge, we can use Kripke frames to try and interpret *belief*. *Doxastic* Kripke models are structures of the form $(X, R_B, \mathcal{V})$. They are implemented to validate formulas formed in a comparable language $\mathcal{L}_B$, which is built up exactly as in (1.1), but having the unary modal operator B instead of K.

Here, the standard axiomatization for the logic of belief is the proof system **KD45**, whose axioms reflect the relational properties given in the following table:

Name	Axiom	Relational Property	Doxastic Property
(K)	$B(p \to q) \to (Bp \to Bq)$	Modal relation	*Logical Omniscience*
(D)	$Bp \to \langle B \rangle p$	Seriality	*Consistency of Belief*
(4)	$Bp \to BBp$	Transitivity	*Positive Introspection*
(5)	$\langle B \rangle p \to B\langle B \rangle p$	Euclideanness	*Negative introspection*

Table 1.2: Some axioms for belief

Once again, we briefly review the selection. The (K) axiom for belief implies that the agent will believe all the logical consequences of her/his beliefs. Axiom (D) means that nobody can believe something and its negation at the same time (recall that the operator $\langle B \rangle$ is defined as $\neg B \neg$). Thus, the property rendered is precisely 'consistency of belief.' Consistency of belief is fairly uncontroversial, as are positive and negative introspection, for that matter. Their colloquial—mundane—readings seem far too reasonable to repeal, and this is why **KD45** is called the standard axiomatic system for belief. Notice that beliefs, thus expounded, are not assumed to be factive, which reflects the ordinary tenet that people can have false beliefs.

1.1.3 Putting knowledge and belief together

Let $\mathcal{L}_{KB}$ consist of a countable set *Prop* of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators K and B. The grammar of this language is given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid K\phi \mid B\phi$$

for $p \in Prop$. Once again, the abbreviations for the connectives $\vee$, $\to$, $\leftrightarrow$ are the ones inherited from propositional logic. The possibility operators $\langle K \rangle$ and $\langle B \rangle$ are defined as $\neg K \neg$ and $\neg B \neg$, respectively.

Kripke semantics allows for a way in which to model knowledge and belief at the same time. Moreover, this can be done in terms of the *same* accessibility relation on a set of possible worlds. From our discussion above, it is evident that we cannot identify the knowledge-accessibility relation with the belief-accessibility relation; their standard interpretations satisfy different axioms. What can be done is to define a new belief-accessibility relation in terms of the knowledge-one. Such a technique comes from what is usually referred to as *plausibility models* for belief, which we briefly address in the following paragraphs.

A tuple $(X, \preccurlyeq, \mathcal{V})$ is called a plausibility model for $\mathcal{L}_B$ if $\preccurlyeq$ is a preorder on X, called the plausibility preorder, and $\mathcal{V}$ is a valuation function. For all practical purposes, we can take $\preccurlyeq$ as being the **S4** knowledge-accessibility relation for a given agent. Therefore, plausibility models can in fact be seen as models for the formulas of the extended language $\mathcal{L}_{KB}$.

If knowledge is already taken care of by the plausibility preorder, there is mainly one way in which epistemologists have interpreted belief in plausibility models, which is ascribed to Grove ([26]). However, for matters of convenience that will become apparent when we introduce our topo-evidential semantics, we will distinguish between two types of this plausibility-belief:

1. *(Grove-belief).*[4] For a given plausibility model $\mathcal{M} = (X, \preccurlyeq, \mathcal{V})$ and $w \in X$, we define $\mathcal{M}, w \Vdash B^{Gr}p$ iff $MAX_{\preccurlyeq}(X) \subseteq \mathcal{V}(p)$, where we have used the index ^{Gr} to distinguish this interpretation of belief from the one below (not because we mean to change the language). Therefore, Grove-belief amounts to truth in all the *most plausible worlds*. Notice that this renders Grove-belief as a *global* modality, in

[4] We use the name 'Grove-Belief' because this interpretation is essentially the same as belief in a Grove's sphere model. For details, see [26] and [10].

the sense that it is independent from the world w in which the formula is evaluated. In other words, if $MAX_{\preccurlyeq}(X) \subseteq \mathcal{V}(p)$, then $B^{Gr}p$ is true in $\mathcal{M}$ (in the whole model).

2. *(Lewis-belief)*.[5] For a given plausibility model $\mathcal{M} = (X, \preccurlyeq, \mathcal{V})$ and $w \in X$, we take $\mathcal{M}, w \Vdash B^{Le}p$ iff $\forall x \in X$, $\exists y \succcurlyeq x$ such that $\forall z \succcurlyeq y, z \in \mathcal{V}(p)$. In other words, Lewis-belief amounts to truth in all the worlds that are *plausible enough.*

   Notice that Grove-belief will not necessarily be *consistent* always, since there is no reason to assume that $MAX_{\preccurlyeq}(X) \neq \emptyset$. Thus, we can get situations in which $B^{Gr}p$ and $B^{Gr}\neg p$ hold at the same time. Lewis-belief, on the contrary, is always consistent for all plausibility models, and it also coincides with Grove-belief whenever Grove-belief is consistent.

   As for the axiomatization, then, it is well-known and easy to show that **KD45** is sound and complete for $\mathcal{L}_B$ with respect to plausibility models with the Lewis-belief semantics. Moreover, the system made up by the **S4** axioms for K, the **KD45** axioms for B, all propositional tautologies, and whose rules of inference are *Modus Ponens* and *Necessitation* for both operators, is sound and complete for $\mathcal{L}_{KB}$ with respect to plausibility models (using the preorder as K's accessibility relation and taking B as B^{Le}).

Stalnaker's combined system of knowledge and belief

As mentioned earlier, Stalnaker vouches for a type of knowledge which is not of the **S5** kind, since according to his standards, knowledge and belief collapse to one and the same in this case. However, he does not side with **S4** as the system for knowledge, either. Working with the underlying intuition that knowledge should be *correctly* justified true belief, Stalnaker presents a formal relation between the operators K and B, and by doing so he obtains a notion of knowledge which turns out to be axiomatizable by the system **S4.2**. Notice that we are not implying that he defined knowledge in terms of belief. As mentioned previously, he takes the other way around. By weakening an idealized—primitive—notion of knowledge into something resembling a standard depiction of belief—in a reasonable manner—he is able to conclude that the true logic for knowledge is rendered by **S4.2**.[6]

The axioms of his system, which we will refer to as **KB**, are given in Table 1.3.

	Stalnaker's Epistemic-Doxastic Axioms	
(K)	$K(p \to q) \to (Kp \to Kq)$	*Logical Omniscience*
(T)	$Kp \to p$	*Factivity of Knowledge*
(4)	$Kp \to KKp$	*Positive Introspection for K*
(D)	$Bp \to \langle B \rangle p$	*Consistency of Belief*
(SPI)	$Bp \to KBp$	*(Strong) Positive Introspection*
(SNI)	$\neg Bp \to K(\neg Bp)$	*(Strong) Negative Introspection*
(KB)	$Kp \to Bp$	*Knowledge implies Belief*
(FB)	$Bp \to BKp$	*Full Belief*
	Rules of Inference	
(MP)	From p and $p \to q$ infer q	Modus Ponens
(Nec_K)	From p infer Kp	*Necessitation for K*

Table 1.3: Stalnaker's proof system **KB**

As seen, Stalnaker starts off with the **S4** axioms for knowledge and the agreeable *consistency of belief.* No comment necessary. *Strong Positive Introspection* and *Strong Negative Introspection* speak of full introspection about one's own beliefs, which is not far-fetched. In principle, an idealized thinker could know what does (s)he believe and does not believe, in a customary—all-encompassing—sense of knowledge. Moving

[5]We use the name 'Lewis-Belief' because this interpretation is essentially taken from Lewis definition of beliefs in terms of counterfactuals. For details, see [37] and [10].

[6]As pointed out by Stalnaker himself, Lenzen ([34]) was the first that campaigned for the definability of belief in terms of knowledge, and presented the assumptions between the relation of both concepts that would imply that the logic of knowledge should be given by **S4.2** instead of **S4**.

on, we should say that the axiom (KB) is not only quite appealing but also philosophically necessary, for (conscious) knowledge, however it is rendered, should imply belief. As for (FB), it conveys the idea that belief is subjective certainty: belief is "subjectively indistinguishable from knowledge" ([5], p. 28).

The system **KB** yields a **KD45** pure belief logic and an **S4.2** pure knowledge logic. This is largely due to the fact that $B\phi \leftrightarrow \langle K\rangle K\phi$ is a theorem of the proof system. Such a theorem is the exact reason because of which we say that Stalnaker's framework promotes a definition of belief in terms of knowledge. This is very relevant in both the conceptual and technical dimensions. On the one hand, it endorses the aforementioned interpretation of belief as subjective certainty: believing p is equivalent to not knowing that you do not know p. On the other, it opens the opportunity for a translation from the formulas of $\mathcal{L}_B$ and $\mathcal{L}_{KB}$ into the formulas of $\mathcal{L}_K$. For example, the main results that Ozgün collects in [41] for both relational and topological completeness of these systems with respect to their respective models take great advantage from such translation. The following proposition formalizes everything that was said in this paragraph.

Proposition 1.1.2. *For all formulas ϕ, ψ of $\mathcal{L}_{KB}$, the following are theorems of KB:*

- $B(\phi \to \psi) \to (B\phi \to B\psi)$ *(The (K) axiom for B).*
- $B\phi \to \langle B\rangle\phi$ *(The (D) axiom for B).*
- $B\phi \to BB\phi$ *(The (4) axiom for B).*
- $\neg B\phi \to B\neg B\phi$ *(The (5) axiom for B).*
- $B\phi \leftrightarrow \langle K\rangle K\phi$.
- $\langle K\rangle K\phi \to K\langle K\rangle\phi$ *(The $(.2)$ axiom for K).*

For a proof of all these items, see [41] pp. 27-28. □

Stalnaker's system speaks of a particular relation between belief and knowledge. In the context of the *defeasibility analysis*, which equates knowledge with justified true belief that will not be defeated upon receiving any new information, his **S4.2** knowledge falls short of being the same as undefeated justified true belief. In order to show this, he incorporates an AGM Belief Revision paradigm into his system (for the basic review on the AGM theory of Belief Revision, we refer the reader to [1]). He concludes that the system supporting a kind of knowledge that is equal to undefeated justified true belief is given by **S4.3** (**S4.2**+ $(K(Kp \to Kq) \vee K(Kq \to Kp))$). Stalnaker then concludes that the defeasibility analysis "provides a *sufficient* condition for knowledge (in [his] idealized setting), [...] [b]ut it does not seem to be a plausible *necessary* and sufficient condition for knowledge" ([46], p. 191). However, as we will mention in due time, Baltag *et alia* show in [5] that the topological semantics for Stalnaker's knowledge, having **S4.2** as it complete system, coincides with undefeated justified true belief for a particular Belief Revision paradigm that *generalizes* AGM theory.[7]

So much for this syntactic and relational-semantic introduction for the logics of knowledge and belief. Let us get down to the business of addressing the topological models for them.

1.2 Topological semantics

1.2.1 Interior semantics

We start by reminding the reader of some basic definitions from General Topology. For any other basic definitions that we might be taking for granted, we refer the reader to [54] or [19] as proper background textbooks.

Definition 1.2.1. (Topological spaces) Let X be a set. $\tau \subseteq \mathcal{P}(X)$ is called a *topology* on X if it meets the following requirements:

- $X, \emptyset \in \tau$.

[7] The Belief Revision paradigm employed by [5] makes use of the notions of *conditional beliefs* (see [10] and [11]) and the therein introduced *topological update* operator.

- If $U, V \in \tau$, then $U \cap V \in \tau$.
- For a family $\mathcal{G} \subseteq \tau$, $\bigcup \mathcal{G} \in \tau$.

A *topological space* is a pair (X, τ), where X is a set and τ is a topology on X. The elements of τ are called *open sets*. Complements of open sets are called *closed sets*.

For $A \subseteq X$, the *interior* of A is defined as the $\subseteq$-largest open set included in A, and will be denoted by $int\ A$. The *closure* of A is defined as the $\subseteq$-least closed set including A, and will be denoted by $Cl(A)$. A standard result in General Topology, we have that for a given $x \in X$ and $A \subseteq X$, $x \in int\ A$ iff there exists an open set U such that $x \in U \subseteq A$; $x \in Cl(A)$ iff every open set U such that $x \in U$ intersects A (i.e., $U \cap A \neq \emptyset$).

For a given $x \in X$, we will often refer to the open sets of τ that include x as *τ-neighborhoods* of x.

Provided with the basic definitions, we introduce the topological semantics for the formulas of the languages we have reviewed so far. Once again, we will proceed with alleged order. First, we will deal with $\mathcal{L}_K$ and exhibit the basic logical results for the models given. Then we will address the topological models for $\mathcal{L}_B$ and $\mathcal{L}_{KB}$ that constitute sensible background to our framework, and proceed accordingly.

Definition 1.2.2. (Single-agent standard topological models. Interior semantics)

A tuple $(X, \tau, \mathcal{V})$ is called a topological model if and only if τ is a topology on X and $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

The semantics for the formulas of $\mathcal{L}_K$ is given recursively by

$$
\begin{array}{lcl}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X \backslash \|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|K\phi\| & = & \{x \in X \mid \exists U \in \tau \text{ s.t. } x \in U \subseteq \|\phi\|\} = int\ \|\phi\|.
\end{array}
$$

Truth and *Validity* are defined the same way as for standard Kripke semantics.

We have defined the semantics for our formulas through their so-called *extension*. In this sense, we will say that a topological model *satisfies* a given formula ϕ of $\mathcal{L}_K$ at world $x \in X$ iff $x \in \|\phi\|$. Such satisfaction will be marked as it is customary in Kripke semantics, but with a different terminology (to keep it apart from relational satisfaction): $\mathcal{M}, x \vDash \phi$. Thus, we have that

$$\mathcal{M}, x \vDash \phi \text{ iff } x \in \|\phi\|.$$

Notice that for every topological space (X, τ) and $A \subseteq X$, we have that $int\ X \backslash A = X \backslash Cl(A)$, so that $\|\langle K \rangle \phi\| = Cl(\|\phi\|)$.

The modern approach to the interior semantics for modal logic is highly profitable, on many levels. According to [48], this sort of "[...] modelling was particularly vivid and attractive for the language of *intuitionistic logic*, where open sets may be viewed as information stages concerning some underlying point [...]" (p. 222). In the field of epistemic logic, the idea points in the direction of treating the open sets as 'pieces of evidence' available to a given agent at some possible world, which is exactly what we will talk about formally as of Chapter 3.

Topological spaces can actually be seen as generalizations of relational structures, and hence speak of modal logics (including epistemic/doxastic logics) from a wider, or at least interesting, angle. For any space, the operator int behaves in a way that reminds of the modality of a reflexive and transitive relation. To shed light on the matter, let us recall the Kuratowski axioms for the operator int:

For every topological space (X, τ) and $A \subseteq X$, we have that

- $int\ X = X$.
- $int\ A \subseteq A$.
- $int\ (A \cap B) = int\ A\ \cap\ int\ B$.
- $int\ (int\ A) = int\ A$.

The Kuratowski axioms can be seen as the topological formulation of the axioms in the proof system **S4** (see [48], p. 237). This means to say that **S4** is actually sound with respect to the class of topological models with the *interior semantics*. It is well-known that the proof system is also complete with respect to the class of all topological models. A very elegant proof of completeness can be found in [48], through the so-called *canonical topo-models*. However, we introduce a humbler approach, that revolves around a celebrated connection between the topological semantics, on one side, and the standard Kripke relational semantics, on the other, for the formulas of $\mathcal{L}_K$. This association is of the utmost importance for our intentions ahead.

Definition 1.2.3. (*Alexandrov* spaces) A topological space (X, τ) is said to be an *Alexandrov* space if and only if the intersection of any collection of open sets of X is an open set as well.

Notice that a space is *Alexandrov* iff every point $x \in X$ has a $\subseteq$-smallest open set including it, namely the intersection of all the open sets around x.

Definition 1.2.4. For a topological space (X, τ), the *specialization preorder* on X is a relation $\leq_\tau$ on X defined by $x \leq_\tau y$ iff $O_x \subseteq O_y$, where we have taken $O_z := \{U \in \tau \mid z \in U\}$ for any $z \in X$. Another way of putting it would be to say that $x \leq_\tau y$ iff $x \in Cl_\tau(\{y\})$.

It is easy to check that the order defined above is in fact reflexive and transitive.

Definition 1.2.5. For a given **S4** frame $(X, \leq)$, we call a set $A \subseteq X$ *upward-closed* iff for every $x \in A$, if $x \leq y$ for some $y \in X$, then $y \in A$ as well. It is customary to mark $x\uparrow_\leq$ to refer to the set $\{y \in X \mid x \leq y\}$, which is clearly upward closed.

Observation 1. For a given **S4** frame $(X, \leq)$, the set of all $\leq$-upward-closed sets forms an *Alexandrov* topology on X, which we denote $\tau_\leq$. For $x \in X$, the $\subseteq$-smallest open set including x is precisely $x\uparrow_\leq$. This implies that $\{x\uparrow_\leq \mid x \in X\}$ is a basis for the topology $\tau_\leq$.

Observation 2.

- For a given **S4** frame $(X, \leq)$, $\leq = \leq_{\tau_\leq}$. To ascertain this, regard that the inclusion $\leq \subseteq \leq_{\tau_\leq}$ comes from the fact that $x\uparrow_\leq$ is the $\subseteq$-least open set including x in the topology $\tau_\leq$, so that $y \in x\uparrow_\leq$ implies that $x \in Cl_{\tau_\leq}(\{y\})$, and thus that $x \leq_{\tau_\leq} y$. The other inclusion is straightforward.
- For a given *Alexandrov* space (X, τ), $\tau = \tau_{\leq_\tau}$. That $\tau \subseteq \tau_{\leq_\tau}$ can be seen from the following argument: for any $U \in \tau$ and $x \in U$, we have that if $x \leq_\tau y$, then $y \in U$ as well; this means that $x\uparrow_{\leq_\tau} \subseteq U$; the other inclusion comes from noticing that, since τ makes the space *Alexandrov*, then for every $x \in X$, $\bigcap\{U \in \tau \mid x \in U\} \subseteq x\uparrow_{\leq_\tau}$.

These observations provide a one-to-one correspondence between *Alexandrov* spaces and **S4** Kripke frames. The correspondence is truth-preserving, in the sense that every relational model based on an **S4** frame will satisfy the same formulas of $\mathcal{L}_K$ as its associated topological space, and viceversa (for a nice proof of this, we refer the reader to [42] ,p. 306. This is a capital result—and we ourselves use its underlying intuition and basic technique in order to prove the topological completeness of the logic we introduce in Chapter 4. A formal statement of it is the following:

Proposition 1.2.6. *Let $(X, \leq)$ be an* **S4** *frame and $\mathcal{V} : Prop \to \mathcal{P}(X)$. For every formula ϕ of $\mathcal{L}_K$ and $x \in X$, we have that*

$$(X, \leq, \mathcal{V}), x \Vdash \phi \text{ iff } (X, \tau_\leq, \mathcal{V}), x \vDash \phi.$$

□

As an important byproduct, we get

Proposition 1.2.7. *The proof system* **S4** *is sound and complete for $\mathcal{L}_K$ with respect to the class of topological models with the interior semantics.*

□

Consequently, the system **S4** renders the logic of all topological spaces underlying models for $\mathcal{L}_K$ with the *interior semantics*. So the topological interpretation of knowledge will turn out to be at least factive and positively introspective. But what about **S4.2** and **S5**, then? They give rise to the logics of some particular classes of topological spaces. For example, we have that **S5** is sound and complete for $\mathcal{L}_K$ with respect to the class of *clopen* topological spaces (i.e., those spaces in which each open set is also closed). In turn, **S4.2** is sound and complete for $\mathcal{L}_K$ with respect to the class of *extremally disconnected* spaces. Since these provide sensible background to what we intend to assemble in Chapter 4, we introduce them and mention the handsome results delivered for their epistemic logics. After some pertinent discussion, we will also approach their doxastic logics and the comparably nice properties Ozgün found for them in [41].

Definition 1.2.8. A topological space (X, τ) is *extremally disconnected* if for every $A \subseteq X$, $Cl(int\ A) \in \tau$.

From [48] (p.253), we know that for every $\phi \in \mathcal{L}_K$, the formula $\langle K\rangle K\phi \to K\langle K\rangle\phi$ is valid in a topological space (X, τ) iff (X, τ) is extremally disconnected.

Moreover, we have that

Proposition 1.2.9. *The one-to-one correspondence between* **S4** *Kripke frames and Alexandrov topological spaces maps* **S4.2** *frames to extremally disconnected Alexandrov spaces.*

Proof. See [41] (p. 22). □

This proposition yields that the system **S4.2** is sound and (weakly) complete for $\mathcal{L}_K$ with respect to the class of extremally disconnected spaces with the interior semantics. However, as suggested earlier, extremally disconnected spaces will also constitute important interpretational ground for *belief*. In [41], Ozgün affords a sly definition for a topological semantics of logics of belief, based on Stalnaker's combined system for knowledge and belief.

Recall the language $\mathcal{L}_B$ and the combined language $\mathcal{L}_{KB}$. Let us expand Definition 1.2.2 and give a topological semantics for their formulas.

Definition 1.2.10. (Standard topological semantics for knowledge and belief) A tuple $(X, \tau, \mathcal{V})$ is called a *topological model* if and only if τ is a topology on X and $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

The semantics for the formulas of $\mathcal{L}_{KB}$ is given recursively by:

$$\begin{array}{lcl} \|p\| & = & \mathcal{V}(p) \\ \|\neg\phi\| & = & X \backslash \|\phi\| \\ \|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\ \|K\phi\| & = & int\ \|\phi\| \\ \|B\phi\| & = & Cl(int\ \|\phi\|). \end{array}$$

Notice that, whatever the conceptual explanation for the choice of topological interpretation of belief might be, it rests upon the philosophical advantages already exposed by Stalnaker, since the combined operator $Cl(int)$ is just the topological translation of $\langle K\rangle K$, itself a syntactic equivalent of B in **KB**. All these associations, together with the bridge between topological and relational semantics, point out to according generalizations of relational soundness-and-completeness results for **KD45** and **KB**. Unsurprisingly, it is Ozgün ([41]) that provides such outcomes. We gather them in a bundle of propositions next.

Proposition 1.2.11. *(already mentioned) The system* **S4.2** *is sound and weakly complete for* $\mathcal{L}_K$ *with respect to the class of extremally disconnected models with the (interior) semantics of Definition 1.2.10.*

□

Proposition 1.2.12. *The system* **KD45** *is sound and weakly complete for* $\mathcal{L}_B$ *with respect to the class of extremally disconnected models with the (closure-interior) semantics of Definition 1.2.10.*

□

It is worth talking about the strategy used in the completeness proof given in [41] for the last proposition. The reason is that it is essentially the same technique as the one we will employ for our own completeness endeavors in Chapter 5. The method can be described in simple terms as 'going through relational completeness.' For each consistent formula ϕ of the language for the system under discussion, one builds a relational structure that satisfies ϕ. Then Proposition 1.2.6 is used to render a topological model that also satisfies ϕ.

In this particular case, there are relevant addendums to the basic method described above: the proof goes through completeness of **KD45** with respect to the class of relational structures called *brushes*, and in particular with respect to *finite pins*. For each finite pin, we can find an **S4.2** model (based on a reflexive and transitive *cofinal frame*[8]) that satisfies the translations into $\mathcal{L}_K$ of all the formulas of $\mathcal{L}_B$ that the **KD45** model based on that finite pin satisfies. Thus, if ϕ is not a theorem of **KD45**, we can find a finite pin that falsifies ϕ, and then an **S4.2** model which falsifies the translation of ϕ into $\mathcal{L}_K$. Finally, Proposition 1.2.6 and Proposition 1.2.9 ensure that we can find an extremally disconnected model that falsifies ϕ, thus showing weak completeness of **KD45** with respect to extremally disconnected spaces.

We can say something very strong about this particular topological semantics for the formulas of $\mathcal{L}_{KB}$. Any *extensional* semantics[9] for the formulas of $\mathcal{L}_{KB}$ that validates the axioms of Stalnaker's system **KB** on a set X of possible worlds is actually a topological semantics, as presented in Definition 1.2.10, for which the pertinent topology on X is *extremally disconnected* (see [41], p.29). This comes from the fact that any extensional semantics that validates the **S4** axioms for K is actually the interior semantics, and, as we have seen, B can be defined as $\langle K\rangle K$ in **KB**, so that $\|B\phi\|$ coincides with $Cl(int\ \|\phi\|)$. The satisfaction of the axiom (CB), then, guarantees that the relevant topology is *extremally disconnected.* As a corollary, we get

Proposition 1.2.13. *Stalnaker's system* **KB** *is sound and complete for* $\mathcal{L}_{KB}$ *with respect to the class of extremally disconnected models with the full semantics of Definition 1.2.10.*

Proof. For completeness, we just allude to the fact that **KB** is complete with respect to its canonical model, which is an extensional model for the formulas of $\mathcal{L}_{KB}$ that clearly validates the axioms of **KB**. Thus, it is in fact an extensional topological model given by an extremally disconnected topology on the canonical domain. □

All these results show that there is a very interesting topological semantics for an axiomatic logic of knowledge and belief. As we mentioned earlier, the topological semantics comes with new intuitions about the philosophy of the relation between the two concepts, particularly within the frame of the *Defeasibility Analysis.* In [5], Baltag *et alia* formalize a Belief Revision method for this topological semantics, which abides by the most conventional AGM standards. They introduce a *topological update* operator that allows for the evaluation of beliefs after information is received. This is done by restricting the topology to the extension of the formula that is received, or with which the update is performed. In this setting, then, one can prove that knowledge coincides with belief that will not be given up upon receiving any new *true* hard information. Therefore, it coincides with undefeated justified true belief, while keeping its axiomatization in **S4.2**. So in this sense, the topological models vouch for the identification of Stalnaker's knowledge with the undefeated justified true belief of Lehrer and Paxson. For the subtleties of the topological Belief Revision paradigm, we refer the reader to [5] and [41].

We should mention that there are *other* topological interpretations for the formulas of $\mathcal{L}_B$, which give rise to their own doxastic logics. The most famous one was devised by Steinsvold in [47], taking $\|B\phi\|$ as the *co-derived* set of $\|\phi\|$. Safe to say, his inspiration must have come from McKinsey and Tarski's interpretation in [39] of the modal operator $\Diamond$ as the *derived set.*[10] Recall that, for a given topological space (X, τ) and $A \subseteq X$, we say that $x \in X$ is a *limit* point of A if for every τ-neighborhood U of x, $U\backslash\{x\} \cap A \neq \emptyset$. The set of all limit points of A is called the *derived set* of A. Similarly, we say that $x \in X$ is a *co-limit point* of A iff there exists $U \in \tau$ such that $x \in U$ and $U\backslash\{x\} \subseteq A$. The set of all co-limit points of A, then, is referred to as A's *co-derived set*, with the notation $t(A)$. With the topological interpretation $\|B\phi\| = t(|\phi\|)$, we get that the system **wK4**$= (K) + ((p \wedge Bp) \to BBp)$ is complete for $\mathcal{L}_B$ with respect to the class of

[8]The precise definitions of the relational structures *brushes, finite pins,* and *cofinal frames* do not matter so much for our purposes. What matters is the underlying scheme designed for the completeness proof. Nevertheless, we refer the reader to [41] (Chapter 4) for the detailed account.

[9]An extensional semantics gives the same meaning to sentences with the same extension (see [5] and [41]).

[10]This interpretation which was largely developed by Esakia (see [20]).

all topological spaces. Similarly, the standard **KD45** is complete with respect to the class of DSO-spaces, i.e., dense-in-themselves spaces that satisfy the TD-axiom and such that the derived set of every subset is open (see [47] and [48] for the proofs). An important shortcoming of this interpretation, however, is that if we use the interior semantics for K, then knowledge is equated with true belief ($Kp \leftrightarrow p \wedge Bp$), making it vulnerable to "Gettier counterexamples."

Chapter 2

Background. Multi-agent models

In this chapter, we present a brief summary of the multi-agent semantic representations for logics of group knowledge (and belief). We focus on relational and topological interpretations, so that the arguments expounded serve as preparation and groundwork for the logic of distributed knowledge that we will present in Chapter 4, as well as for the soundness and completeness results for such logic with respect to its topological (evidential) models. Though there has been recent development in the field of multi-agent doxastic systems (see [43]) and their logics of belief and *common belief*, we will restrict our discussion to languages having only operators for individual and group *knowledge*. This is no reason to think that we have left beliefs out of the picture, since our far reaching goal is to incorporate them as elements of a formal setup for multi-agent systems. Our proposal for a logic of distributed knowledge *implicitly* accounts for the opportunity of such an incorporation. We invoke the idea that any formal portrayal of knowledge also implies a formal portrayal of belief. As a disclaimer note, we mention that the expert reader can, in principle, skip this chapter.

As done in Chapter 1, we start by motivating our discourse with an overview of the conceptual perspective on multi-agent epistemic systems, taking care in mentioning the transcendence of the concepts addressed for the field that binds us, that of formal epistemology and theory of agency.

2.1 Group knowledge

When knowledge representation deals with environments in which there is more than one agent, "things get much more interesting—and challenging" ([52] p. 887). An explanation for calling the representational effort a challenge might be in that it involves depicting very complicated situations. The different agents reason not only about the properties of their own place in the environment, but also about each of the other agent's personal reasoning-strategies. Therefore, their rational states end up including views about each other, on one hand, and views about the group as a whole, on the other.

Here, the relevant idea is that of *interaction* between the agents. We strive to grasp how their conceptions of the environment are shaped by interacting with each other and as a group. In interactive multi-agent settings, we distinguish *individual* attitudes from those which are *group-related*. Individual attitudes are typically inherited from the single-agent paradigm, and they correspond to any given agent's knowledge and beliefs. As for the collective attitudes, Van Benthem and Sarenac say in [50] that "[p]erhaps the most interesting topic in an interactive epistemic setting has been the discovery of various notions of what may be called *group knowledge*" (p. 2). In the philosophical, economic, computer scientific, and linguistic literature, there are two main concepts of group knowledge for multi-agent interactive systems: *distributed knowledge* and *common knowledge*.[1]

Common knowledge

[1]Other very interesting notions of group knowledge have been introduced in the epistemological tradition. For a sensible overview of them, as well as interesting proposals to address the philosophical content of the concept of collective knowledge, we refer the reader to [18].

Common knowledge is arguably the predominant notion of group knowledge as far as research goes. It was first treated by David Lewis in the context of conventions ([36]). Lewis pointed out that in order for something to be a *convention*, it must be common knowledge among the members of a group. An example that almost any epistemologist knows is that of 'green lights while driving.' All drivers know that green lights express preference in a crossroads section. Moreover, all drivers know that drivers know this, they know that drivers know that all drivers know this, and so on. Some other basic examples—e.g. of common knowledge arising in dialogue understanding—can be found in [21] (Chapter 1), which speaks of the concept's most colloquial characterization, ascribed to John McCarthy: common knowledge is what "any fool" would know. As pointed out in [50], this instance of group knowledge is likely to be studied as a necessary precondition for coordinating actions between agents—a prerequisite for achieving agreement, so to speak. *Agreement* and *coordination* are central issues for multi-agent systems; game theorists, decision theorists, and economists have all feasted on the pertinent literature and analysis of common knowledge, adding significant contributions to its comprehensive formalization.

Informally, common knowledge of ϕ is defined as the infinite conjunction 'everybody knows ϕ and everybody knows that everybody knows ϕ, and so on.' The infinitary behavior has been tried to be captured formally in many ways, and a standard one views common knowledge of ϕ as the greatest fixed-point of λx(everybody knows ϕ and everybody knows x), where we follow [50] in the use of the terminology from modal μ-calculus. As we will briefly mention, common knowledge as a fixed-point does not always coincide with the infinite conjunction of modal formulas expressed above (see [14] and [50]), but both depictions should be taken as standard.

Distributed knowledge.

According to Fagin *et al*, if common knowledge is what "any fool" knows, then distributed knowledge can be viewed as what a "mysterious wise man," that has complete knowledge of what each member of the relevant group knows, would know ([21], p. 3).

The concept of distributed knowledge in a group of agents is linked to the process of *sharing* information. Its motivation comes from an interest in reasoning about what agents would end up knowing if they could combine (or aggregate) their information or their knowledge. A traditional example of such a situation (see, e.g., [23] and [27]) is constructed as follows. Suppose that Alice has the information that ϕ is the case (or that Alice *knows* ϕ). Bob, on the other hand, knows that ϕ implies ψ; but *neither* of them individually knows that ψ is the case. There is a sense in which the information that ψ is the case is already present in the information states of Alice and Bob, if taken together. One way of phrasing this would be to say that the information that ψ is the case prevails in a distributed form over the information states of Alice and Bob ([23], P. 111). Notice that we can think of distributed knowledge in two ways: either as the knowledge that is implicitly present in the group of agents, or as the knowledge that everybody would get if they shared their information. In this line, we can combine these two views in another typical illustration of the concept: suppose that Alice and Bob communicated everything they know to a third agent. Then, distributed knowledge of the group can be identified with this third agent's subsequent knowledge.

As for its significance in other fields, Van der Hoek points out in [51] that, since the notion of distributed knowledge stems from reasoning about what knowledge a group can attain through communication, it must be clear that it is also crucial when reasoning about the efficacy of speech acts and about communication protocols in distributed systems.

In the next section, we begin our formal approach to multi-agent systems. This thesis is mainly concerned with topological models for distributed knowledge, but it should be beneficial to review the semantic definitions (both in the relational and topological paradigms) for both common and distributed knowledge. The logic results that have been produced in this line of research, and that we shall review presently, will help us maintain a clear picture about our margin of discussion. At the same time, they provide support for many of the arguments and notions that we will use when introducing our logic in Chapter 4, and thereafter. Once again, we stress the fact that the multi-agent epistemic logic that we are shooting for has clear basis on a desirable relation between knowledge and belief. Though no explicit formal representation of belief will be addressed in this chapter, the notions of knowledge addressed are meant to be coupled with an underlying account of beliefs.

2.2 Syntax and Kripke semantics

In the single-agent case, an agent's epistemic capacities are modeled with accessibility relations in a possible worlds-framework. We can easily extend our definition of Kripke frames to include accessibility relations for different agents on the same state space. We have but to consider multi-relational frames, i.e., tuples of the form $(X, R_1, \dots, R_m)$ for some $m \in \mathbb{Z}^+$. Each relation, then, will correspond to the modality for the knowledge of its respective indexed agent. For the ease of exposition, **we will restrict our formal representations to two-agent models**. Therefore, we focus on bi-relational Kripke frames, of the form (X, R_1, R_2). In keep with the notation, we will refer to the anonymous agents as Agent 1 and Agent 2, respectively. When we talk about group knowledge, then, it is implicitly obvious that the group is comprised by Agent 1 and Agent 2.

As before, we first introduce the formal language that we will use to build up the formulas of the logics of group knowledge.

Definition 2.2.1. (Language) Let $\mathcal{L}_K^{\mathcal{DC}}$ consist of a countable set $Prop$ of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators K_1, K_2, D, and C. The grammar of our language is given by

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

for $p \in Prop$. Naturally, the abbreviations for the connectives $\vee$, $\rightarrow$, $\leftrightarrow$ are the ones inherited from propositional logic. The possibility operators $\langle K_1 \rangle$, $\langle K_2 \rangle$, $\langle D \rangle$, and $\langle C \rangle$ are defined as $\neg K_1 \neg$, $\neg K_2 \neg$, $\neg D \neg$, and $\neg C \neg$, respectively.

Once endowed with a two-agent formal language, we can introduce the most usual models to evaluate its formulas in, the so-called two-agent Kripke models (for the sake of coherence with the definitions of successive chapters, we will refer to these models as two-agent *bi-relational* models). The underlying intuition for their construction is a natural progression from the single-agent case. The interesting part, then, comes from the choice of semantics for the formulas involving the group-knowledge operators.

Just so that notation does not get in the way of a clear exposition, we first mention that for a given bi-relational frame (X, R_1, R_2), we will use the notation $(R_1 \cup R_2)^*$ to refer to the reflexive transitive closure of the union $R_1 \cup R_2$.[2]

Definition 2.2.2. (Two-agent bi-relational models) A tuple $(X, R_1, R_2, \mathcal{V})$ will be called a *two-agent bi-relational model* if and only if X is a non-empty set of possible worlds, R_1, R_2 are binary preorders on X, and $\mathcal{V} : Prop \rightarrow \mathcal{P}(X)$ is a valuation function. The semantics for the formulas of $\mathcal{L}_K^{\mathcal{DC}}$ is given recursively by the following rules of model-*satisfaction*:

For a given world $w \in X$,

$$\begin{array}{lll}
\mathcal{M}, w \Vdash \phi & \text{iff} & w \in \mathcal{V}(p) \\
\mathcal{M}, w \Vdash \neg\phi & \text{iff} & \mathcal{M}, w \nVdash \phi \\
\mathcal{M}, w \Vdash \phi \wedge \psi & \text{iff} & \mathcal{M}, w \Vdash \phi \text{ and } \mathcal{M}, w \Vdash \psi \\
\mathcal{M}, w \Vdash K_1\phi & \text{iff} & \forall v \in X,\ [wR_1v \Rightarrow \mathcal{M}, v \Vdash \phi] \\
\mathcal{M}, w \Vdash K_2\phi & \text{iff} & \forall v \in X,\ [wR_2v \Rightarrow \mathcal{M}, v \Vdash \phi] \\
\mathcal{M}, w \Vdash D\phi & \text{iff} & \forall v \in X,\ [wR_1 \cap R_2v \Rightarrow \mathcal{M}, v \Vdash \phi] \\
\mathcal{M}, w \Vdash C\phi & \text{iff} & \forall v \in X,\ [w(R_1 \cup R_2)^*v \Rightarrow \mathcal{M}, v \Vdash \phi]
\end{array}$$

The definition provided for $\|D\phi\|$ should be self-explanatory. The group is said to have distributed knowledge of ϕ at x iff ϕ holds at all worlds that are both R_1-accessible and R_2-accessible from x. In a sense,

[2] Recall that for any relation R on a set X, the transitive closure R^* of R is defined as the intersection of all the transitive relations including R. Alternatively, we can define R^* by the following rule: xR^*y iff there exists a finite chain $\{z_1, \dots, z_n\} \subseteq X$ such that

- $z_1 = x$, $z_n = y$,
- for every $k \in \{1, \dots, n-1\}$, z_kRz_{k+1}.

we can see that in order to compute the distributed knowledge of a group at a given world x, we consider the worlds that are related to x by *both* relations. It is intuitively clear that this corresponds to a semantic merge of information or knowledge. In our Alice-and-Bob example, we can see that if Alice and Bob communicated all their knowledge to a third agent, then (s)he would intuitively exclude all the possible worlds by virtue of which Alice and Bob admit the possibility of $\neg\psi$ at x, since in all the worlds that are accessible *by both of their relations*, ψ is always the case. As Van Benthem and Sarenac point out in [50], the resulting notion of distributed knowledge is technically different from common and individual knowledge, in the sense that it will not be invariant under modal bisimulation. Therefore, its mathematical treatment demands new ways to speak about invariance.

In the case of common knowledge, we can see from the definition provided for $\|C\phi\|$ that the group is said to have common knowledge of ϕ at x if for every world y that is accessible from x by a finite sequence of successive steps taken from either of the accessibility relations R_1 and R_2, ϕ holds at y. Notice that this intends to capture the intuition that something will not be common knowledge if one member of the group admits the possibility of any member admitting the possibility of any member admitting the possibility that... ϕ is false. If there is a world y *reachable* from x by some combination of the epistemic relations, such that ϕ is not true in y, then ϕ is not common knowledge at x. It turns out to be the case that this semantic definition of common knowledge over bi-relational frames works fairly good, in the sense that it can be identified with the greatest fixed-point interpretation. On a related note, we mention that in the case of relational semantics, the fixed-point interpretation coincides with the infinite conjunction of iterated individual knowledge modalities.

Notice that since the relations $(R_1 \cup R_2)^*$ and $R_1 \cap R_2$ are also preorders, then we can think about these notions of group knowledge as new agents. Turns out, then, that both distributed and common knowledge are *factive* and *positively introspective*.

2.3 Axiomatization

As pointed out in Chapter 1, there are standard axiomatizations for the logics that arise from relational (Kripke) semantics for the formulas of epistemic languages. The multi-agent case is no exception. The extension of the language to incorporate operators destined to be interpreted as group knowledge offers many different philosophical roads to take. It should be clear that the debate as to which system is better for characterizing knowledge carries over when we are speaking about the individual knowledge of two or more agents. For formulas of the fragment of the language having only the operators K_1 and K_2, fusion systems (see [22]) such as $\mathbf{S4}+\mathbf{S4}$ and $\mathbf{S5}+\mathbf{S5}$ are good candidates for purposes of axiomatization. Recall that these fusion systems are defined as the $\subseteq$-least sets of formulas (of the aforementioned fragment) that contain the **S4**, resp. **S5**, axioms for K_1 and for K_2 at the same time, and that are closed under *Modus Ponens*, *Necessitation* for K_1, and *Necessitation* for K_2. In this line, notice that, for example, if we were to use $\mathbf{S5}+\mathbf{S5}$ as an axiomatization, then soundness and completeness results for said fragment of the language would apply to the class of bi-relational models where R_1 and R_2 are not only preorders, but also equivalence relations.

As for group knowledge, the same story applies. As we reviewed in the previous section for relational semantics, the usual representations interpret distributed and common knowledge as a third agent's knowledge, in the sense that they behave as the necessity operators for modal relations on the state space. As such, the basic requirements both for D and for C are the **S4** axioms. But what has been much more interesting to study is the relation between the group-knowledge operators and the individual-knowledge operators. Axiomatizing such interaction has also been spurred by lively philosophical debate. Regardless of the philosophical advantages of some axiom system over any other, we run into certain principles that are widely accepted as standard.

Since our main focus of discussion concerns distributed knowledge, we start by reviewing the standard requisite for the relation between D and K_i. Here, there is only one characteristic axiom: $K_ip \to Dp$, which we will refer to as the (DsK) axiom from here on. As one can see, it is meant to capture the idea that if an agent of the group knows ϕ, then the whole group would know ϕ after the sharing of information is done. Informally, we can think of an agent revealing her/his information to the rest of the group.

In the case of common knowledge, as we have seen, the most standard semantic clause governing the relation between C and K_i is the following. If ϕ is common knowledge, then ϕ is true, every member of the group knows ϕ, every member of the group knows that every member of the group knows ϕ, and so on. The traditional way in which to formalize such demand is by imposing two conditions in the proof system. One is called the *equilibrium* axiom—which we will mark as (Eq)—and the other is a particular rule of inference called the *induction rule*. The equilibrium axiom is given by $Cp \leftrightarrow (p \wedge (K_1 Cp \wedge K_2 Cp))$, and it is meant to capture formally the condition that common knowledge of ϕ can only occur when every member of the group knows that ϕ is common knowledge. I.e., common knowledge is an epistemic fixed-point. As for the induction rule, it is given by $\dfrac{p \to (K_1(p \wedge q) \wedge K_2(p \wedge q))}{p \to Cq}$.

This rule gives us a way to derive common knowledge in whatever the resulting system might be.

Because of all of this, we can agree that the most standard proof system for a traditional logic for knowledge, distributed knowledge, and common knowledge would be given by the following table:

	Axioms
(K)	$K_i(p \to q) \to (K_i p \to K_i q)$
(T)	$K_i p \to p$
(4)	$K_i p \to K_i K_i p$
(K_D)	$D(p \to q) \to (Dp \to Dq)$
(T_D)	$Dp \to p$
(4_D)	$Dp \to DDp$
(K_C)	$C(p \to q) \to (Cp \to Cq)$
(T_C)	$Cp \to p$
(4)	$Cp \to CCp$
(DsK)	$K_i p \to Dp$
(Eq)	$Cp \leftrightarrow (p \wedge (K_1 Cp \wedge K_2 Cp))$
	Rules of Inference
Modus Ponens	From p and $p \to q$ infer q
Induction Rule	From $p \to (K_1(p \wedge q) \wedge K_2(p \wedge q))$ infer $p \to Cq$
(Nec_K)	From p infer Kp
(Nec_C)	From p infer Cp

Table 2.1: The proof system $\mathbf{S4}_2^{\mathcal{DC}}$

We will refer to this proof system as $\mathbf{S4}_2^{\mathcal{DC}}$. Notice that *Necessitation* for D follows from (DsK) and *Necessitation* for K_i.

The main result regarding this proof system and the two-agent relational semantics reviewed in Definition 2.2.2 was given by Halpern and Moses in [29] (see also [51], [21], and [43]). We introduce it by the following proposition, which we recognize as very important. For its detailed proof, we refer the reader to the aforementioned sources.

Proposition 2.3.1. *The system $\mathbf{S4}_2^{\mathcal{DC}}$ is sound and weakly complete for $\mathcal{L}_K^{\mathcal{DC}}$ with respect to the class of two-agent bi-relational models where the two relations are reflexive and transitive.*

□

2.4 Topological semantics

We proceed to introduce the topological models for the formulas of $\mathcal{L}_K^{\mathcal{DC}}$. As usual, we first review the basic concepts from General Topology that we need for the construction of our models.

Let X be a set, and let $\mathcal{F}$ be a family of topologies on X.

- For any $\tau \in \mathcal{F}$, we will write int_τ, resp. Cl_τ, to refer to the *interior* operator, resp. *closure*, operator *with respect the topology* τ.

- For $\tau, \tau' \in \mathcal{F}$, $\tau \vee \tau'$ will be used to denote the topology *generated* on X by $\tau \cup \tau'$ (i.e., the smallest topology on X including $\tau \cup \tau'$). Recall that a standard basis for $\tau \vee \tau'$ is the set $\{U \cap V \mid U \in \tau_1, V \in \tau_2\}$.

Definition 2.4.1. (Two-agent bi-topological models) A tuple $(X, \tau_1, \tau_2, \mathcal{V})$ is called a bi-topological model for a logic of distributed and common knowledge if and only if τ_1, τ_2 are topologies on X and $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

The topological semantics for the formulas of $\mathcal{L}_K^{\mathcal{DC}}$ is given recursively by:

$$
\begin{array}{lll}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X \backslash \|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|K_1\phi\| & = & \{x \in X \mid \exists U \in \tau_1 \text{ s.t. } x \in U \subseteq \|\phi\|\} = int_{\tau_1}\|\phi\| \\
\|K_2\phi\| & = & \{x \in X \mid \exists U \in \tau_2 \text{ s.t. } x \in U \subseteq \|\phi\|\} = int_{\tau_2}\|\phi\| \\
\|D\phi\| & = & \{x \in X \mid \exists U \in \tau_1, V \in \tau_2 \text{ s.t. } x \in U \cap V \subseteq \|\phi\|\} = int_{\tau_1 \vee \tau_2}\|\phi\| \\
\|C\phi\| & = & \{x \in X \mid \exists U \in \tau_1 \cap \tau_2 \text{ s.t. } x \in U \subseteq \|\phi\|\} = int_{\tau_1 \cap \tau_2}\|\phi\|.
\end{array}
$$

Satisfaction is denoted as usual: $\mathcal{M}, x \vDash \phi$ iff $x \in \|\phi\|$. *Truth* and *Validity* are defined the same way as for standard Kripke semantics.

Notice that we could already start talking about *evidence*. Since in the topological case, opens sets can be interpreted as pieces of evidence that serve as justifications for the agents' knowledge, we can phrase the semantics given in Definition 2.4.1 in terms of open sets of evidence. In this line, notice that the definitions for individual knowledge come from standard interior semantics, which we have already discussed. As for distributed and common knowledge, the motivation comes from the relational interpretation, in its likely translation to topological concepts. From the conceptual angle, distributed knowledge is equated to the kind of group knowledge that the group *would* acquire if its members were to pool together all their respective pieces of evidence, and thus justify *new* (or *implicit*, if one pleases) beliefs. Similarly, ϕ is taken to be common knowledge in the group if each of its members has common (evidential) justification supporting ϕ. That the justification is common is a strong requisite, and, as we will see, it is meant to address the depth of the concept of common knowledge in its usually acceptable formal interpretations. In other words, the requisite is strong enough to match a plausible axiomatization for common knowledge. Though this is a very casual use of the word 'evidence,' Part II of this thesis revolves around the formalization of topological evidence in a much more specific manner.

As of yet, it remains to be seen if the topological semantics fits with the axiomatization by $\mathbf{S4}_2^{\mathcal{DC}}$. However, one can anticipate that the fact that it somewhat matches the relational counterpart will play a role in the arguments of soundness, and much more importantly, in those of completeness. The following section is dedicated in its entirety to the careful exposition of the arguments that render $\mathbf{S4}_2^{\mathcal{DC}}$ as sound and weakly complete for $\mathcal{L}_K^{\mathcal{DC}}$ with respect to the class of bi-topological spaces. Similarly as we did in Chapter 1, we emphasize the importance of the association between *Alexandrov* spaces and **S4** frames. As we will see, there is a natural extension for such association in the multi-agent case. It is important to mention that we will focus on soundness and completeness *with respect to the class of bi-topological spaces*. Here, we deviate a little bit from Van Benthem and Sarenac ([50], [48]), whose results involve different topological models for the multi-agent epistemic logics. Their treatment concerns a particular class of product spaces, and for the sake of completion, we will shortly address such models later on.

2.5 Soundness and completeness (topological semantics)

Recall that for every **S4** frame, we can construct an associated *Alexandrov* topological space. This association carries over to bi-topological spaces without problem. For every bi-relational **S4** frame $(X, \leq_1, \leq_2)$, each preorder induces a topology on X, which we denote $\tau_{\leq_1}, \tau_{\leq_2}$, respectively. The tuple $(X, \tau_{\leq_1}, \tau_{\leq_2})$ is a bi-topological space.

Lemma 2.5.1. *For every bi-relational* **S4** *frame* $(X, \leq_1, \leq_2)$ *we have that* *a)* $\tau_{\leq_1\cap\leq_2} = \tau_{\leq_1} \vee \tau_{\leq_2}$

b) $\tau_{(\leq_1\cup\leq_2)^*} = \tau_{\leq_1} \cap \tau_{\leq_2}$

Proof.

a) For the $\supseteq$ inclusion, let $x \in X$ and $V \in \tau_{\leq_1} \vee \tau_{\leq_2}$ such that $x \in V$. It suffices to show that $x \uparrow_{\leq_1\cap\leq_2} \subseteq V$. By the definition of $\tau_{\leq_1} \vee \tau_{\leq_2}$, we know that there exist $V_1 \in \tau_{\leq_1}$, $V_2 \in \tau_{\leq_2}$ such that $x \in V_1 \cap V_2 \subseteq V$. Now, we also know that the relevant topologies were built in such a way that $x \uparrow_{\leq_1} \subseteq V_1$ and $x \uparrow_{\leq_2} \subseteq V_2$ (the $\leq_i$-upward-closed sets form a basis for $\tau_{\leq_i}$). Thus, $x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2} \subseteq V_1 \cap V_2$. But $x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2}$ is exactly the same as $x \uparrow_{\leq_1\cap\leq_2}$, so that $x \uparrow_{\leq_1\cap\leq_2} \subseteq V$. For the $\subseteq$ inclusion, we just point out that the basic open sets of $\tau_{\leq_1\cap\leq_2}$ are open sets of $\tau_{\leq_1} \vee \tau_{\leq_2}$. This comes from the fact—also used above—that for every $x \in X$, $x \uparrow_{\leq_1\cap\leq_2}$ is exactly $x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2}$, which is obviously an open set of $\tau_{\leq_1} \vee \tau_{\leq_2}$.

b) For the $\supseteq$ inclusion, let $x \in X$ and $V \in \tau_{\leq_1} \cap \tau_{\leq_2}$ such that $x \in V$. It suffices to show that $x \uparrow_{(\leq_1\cup\leq_2)^*} \subseteq V$. Let $y \in X$ such that $x(\leq_1 \cup \leq_2)^* y$. This means that there exists a finite chain $\{z_1, \ldots, z_n\} \subseteq X$ such that

- $z_1 = x$, $z_n = y$,
- for every $k \in \{1, \ldots, n-1\}$, either $z_k \leq_1 z_{k+1}$ or $z_k \leq_2 z_{k+1}$.

We show that $z_k \in V$ for every $k \in \{1, \ldots, n-1\}$ by induction on k. The base case ($k = 1$) holds because we have taken $z_1 = x \in V$. Assume that $z_{k-1} \in V$. We know that either $z_{k-1} \leq_1 z_k$ or $z_{k-1} \leq_2 z_k$ holds. If $z_{k-1} \leq_1 z_k$, then $z_k \in z_{k-1} \uparrow_{\leq_1}$, but this upward closed set is included in V due to the fact that V is open in $\tau_{\leq_1}$. Therefore, $z_k \in V$. Analogously, we can see that if $z_{k-1} \leq_2 z_k$, then z_k actually lies in V. By this inductive argument, we have that $y = z_n \in V$. Therefore, $x \uparrow_{(\leq_1\cup\leq_2)^*} \subseteq V$ and the $\supseteq$ inclusion is shown.

For the $\subseteq$ inclusion, we point out that the basic open sets of $\tau_{(\leq_1\cup\leq_2)^*}$ are open sets of $\tau_{\leq_1} \cap \tau_{\leq_2}$. Notice that for every $x \in X$, the inclusions $x \uparrow_{\leq_1} \subseteq x \uparrow_{(\leq_1\cup\leq_2)^*}$ and $x \uparrow_{\leq_2} \subseteq x \uparrow_{(\leq_1\cup\leq_2)^*}$ are straightforward. Therefore, for every $x \in X$, $x \uparrow_{(\leq_1\cup\leq_2)^*} \in \tau_{\leq_i}$ for both $i \in \{1, 2\}$, with which the desired inclusion is shown.

$\square$

The correspondence between bi-relational **S4** frames and bi-topological spaces is also truth-preserving, for formulas of $\mathcal{L}_K^{\mathcal{DC}}$. We show this by a natural correlate of Proposition 1.2.6.

Proposition 2.5.2. *Let $(X, \leq_1, \leq_2)$ be a bi-relational **S4** frame and $\mathcal{V} : Prop \to \mathcal{P}(X)$. For every formula ϕ of $\mathcal{L}_K^{\mathcal{DC}}$ and $x \in X$, we have that*

$$(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash \phi \text{ iff } (X, \tau_{\leq_1}, \tau_{\leq_2}, \mathcal{V}), x \vDash \phi.$$

Proof. By induction on ϕ. The base case and the Boolean cases are routine. Let us address the cases for our modal operators:

- ("K_i") Fix $i \in \{1, 2\}$. Assume that the claim holds for ϕ. We have that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash K_i\phi$ iff $x \uparrow_{\leq_i} \subseteq [\![\phi]\!] = \|\phi\|$ (where this last equality of extensions holds by induction hypothesis). Thus, $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash \phi$ iff $x \in int_{\tau_{\leq_i}} \|\phi\| = \|K_i\phi\|$, as witnessed by the $\tau_{\leq_i}$-open set $x \uparrow_{\leq_i}$. This means that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash K_i\phi$ iff $(X, \tau_{\leq_1}, \tau_{\leq_2}, \mathcal{V}), x \vDash K_i\phi$.
- ("D") Assume that the claim holds for ϕ. In this case, we have that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash D\phi$ iff $x \uparrow_{\leq_1\cap\leq_2} \subseteq [\![\phi]\!] = \|\phi\|$ (where this last equality of extensions holds by induction hypothesis). But recall that $x \uparrow_{\leq_1\cap\leq_2} = x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2}$, which is an open set of $\tau_{\leq_1} \vee \tau_{\leq_2}$ by Lemma 2.5.1 *a*). Therefore, $x \uparrow_{\leq_1\cap\leq_2} \subseteq \|\phi\|$ iff $x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2} \subseteq \|\phi\|$ iff $x \in int_{\tau_{\leq_1} \vee \tau_{\leq_2}} \|\phi\| = \|D\phi\|$. So we have shown that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash D\phi$ iff $(X, \tau_{\leq_1}, \tau_{\leq_2}, \mathcal{V}), x \vDash D\phi$.

- ("C") Assume that the claim holds for ϕ. We have that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash C\phi$ iff $x \uparrow_{(\leq_1 \cup \leq_2)^*} \subseteq [\![\phi]\!] = \|\phi\|$ (where this last equality of extensions holds by induction hypothesis). But Lemma 2.5.1 *b*) tells us that $x \uparrow_{(\leq_1 \cup \leq_2)^*}$ is an open set of $\tau_{\leq_1} \cap \tau_{\leq_2}$. Therefore, $x \uparrow_{(\leq_1 \cup \leq_2)^*} \subseteq \|\phi\|$ iff $x \in int_{\tau_{\leq_1} \cap \tau_{\leq_2}} \|\phi\| = \|C\phi\|$. With this, we get that $(X, \leq_1, \leq_2, \mathcal{V}), x \Vdash C\phi$ iff $(X, \tau_{\leq_1}, \tau_{\leq_2}, \mathcal{V}), x \vDash C\phi$.

□

Proposition 2.5.3. *The system* $\mathbf{S4}_2^{\mathcal{DC}}$ *is sound and weakly complete with respect to the class of bi-topological models.*

Proof.

For *soundness*, we have the following observations:

- Any bi-topological space (X, τ_1, τ_2) validates the **S4** axioms for K_1 and K_2 according to the same argument given in Chapter 1 (for the single-agent case), namely the satisfaction of the Kuratowski axioms for int_{τ_1} and int_{τ_2}. Similarly, any bi-topological space validates the **S4** axioms for D and C, since $int_{\tau_1 \vee \tau_2}$, resp. $int_{\tau_1 \cap \tau_2}$, also satisfy the Kuratowski axioms for the *interior* operator.
- For the (DsK) axiom, we observe that, since $\tau_1, \tau_2 \subseteq \tau_1 \vee \tau_2$, then for every $A \in X$ and a fixed $i \in \{1, 2\}$, $int_{\tau_i} A \subseteq int_{\tau_1 \vee \tau_2} A$ (for any $x \in int_{\tau_i} A$, there exists $U \in \tau_i$ such that $U \subseteq A$, but U is also a member of $\tau_1 \vee \tau_2$, so that $x \in int_{\tau_1 \vee \tau_2} A$ as well). Therefore, for any bi-topological model $(X, \tau_1, \tau_2, \mathcal{V})$ and a fixed $i \in \{1, 2\}$, we have that $x \in \|K_i\phi\|$ implies that $x \in \|D\phi\|$ for every $x \in X$. Therefore, $\|K_i\phi \to D\phi\| = X$, which implies that (DsK) is sound for bi-topological spaces (see Chapter 5, Lemma 5.1.1).
- As for the *equilibrium* axiom for C, we refer the reader to [15] (p. 30) for a proof of its validation in the interpretation of C as interior of the intersection of topologies.
- With respect to the rules of inference, it is well-known that *Modus Ponens* preserves validity; *Necessitation* for K_1, K_2, D, and C comes from the standard equality $int_{\tau_1} X = int_{\tau_2} X = int_{\tau_1 \vee \tau_2} X = int_{\tau_1 \cap \tau_2} X = X$. Though it strays from our focus, we sketch a proof of how the *C-induction rule* preserves validity: if for every space (X, τ_1, τ_2) and valuation function $\mathcal{V}$, the inclusion $\|p\| \subseteq (int_{\tau_1} \|p \wedge q\| \cap int_{\tau_2} \|p \wedge q\|)$ holds, then for every for space (X, τ_1, τ_2), $\|p\|$ is an open both in τ_1 and τ_2 included in $\|q\|$, so that $x \in \|p\|$ implies that $x \in int_{\tau_1 \cap \tau_2} \|q\| = \|Cq\|$ in the given space.

For completeness, we have observed in Proposition 2.3.1 that $\mathbf{S4}_2^{\mathcal{DC}}$ is complete for $\mathcal{L}_K^{\mathcal{DC}}$ with respect to the class of bi-relational **S4** frames, interpreting the modality for D as the intersection of the relations, and the modality for C as the transitive closure of the union of the relations. So assume ϕ_0 is a formula of $\mathcal{L}_K^{\mathcal{DC}}$ such that $\mathbf{S4}_2^{\mathcal{DC}} \nvdash \phi_0$. We get that there exists a bi-relational **S4** model $(X^0, \leq_1^0, \leq_2^0, \mathcal{V})$ and a world $x_0 \in X^0$ such that $(X^0, \leq_1^0, \leq_2^0, \mathcal{V}) \nVdash \phi_0$. According to Proposition 2.5.2, the associated bi-topological model $(X^0, \tau_{\leq_1^0}, \tau_{\leq_2^0}, \mathcal{V})$ is such that $(X^0, \tau_{\leq_1^0}, \tau_{\leq_2^0}, \mathcal{V}), x_0 \nvDash \phi_0$.

□

At the end of the previous section, we mentioned that Van Benthem and Sarenac offer in [50] a slightly different approach to topological models for multi-agent epistemic logics, based on particular product spaces. As it turns out, these also comprise a very interesting class of models, and we briefly sketch their construction.

Definition 2.5.4. (Two-agent product models)

Let (X, τ_1'), (Y, τ_2') be two topological spaces. We say that $A \subseteq X \times Y$ is H-open if for every $(x, y) \in A$, there exists $U \in \tau_1'$ such that $x \in U$ and $U \times \{y\} \subseteq A$. Similarly, we say that A is V-open if for every $(x, y) \in A$, there exists $V \in \tau_2'$ such that $y \in V$ and $\{x\} \times V \subseteq A$. The family of all H-open sets forms a topology on $X \times Y$, and the family of all the V-open sets forms a topology on $X \times Y$. Let these topologies be denoted by τ_1 and τ_2.

A tuple $(X \times Y, \tau_1, \tau_2, \mathcal{V})$ will be called a *two-agent product model* for $\mathcal{L}_K^{\mathcal{DC}}$ if and only if (X, τ_1'), (Y, τ_2') are two topological spaces such that τ_1 is the topology of all the H-open sets associated to τ_1' and τ_2 is the topology of all the V-open sets associated to τ_2'.

The product semantics for the formulas of $\mathcal{L}_K^{\mathcal{DC}}$ is given recursively by:

$$
\begin{array}{lll}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X\backslash\|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|K_1\phi\| & = & \{(x,y) \in X \times Y \mid \exists U \in \tau_1 \text{ s.t. } (x,y) \in U \subseteq \|\phi\|\} \\
\|K_2\phi\| & = & \{(x,y) \in X \times Y \mid \exists U \in \tau_2 \text{ s.t. } (x,y) \in U \subseteq \|\phi\|\} \\
\|D\phi\| & = & \{(x,y) \in X \times Y \mid \exists U \in \tau_1, V \in \tau_2 \text{ s.t. } (x,y) \in U \cap V \subseteq \|\phi\|\} \\
\|C\phi\| & = & \{(x,y) \in X \times Y \mid \exists U \in \tau_1 \cap \tau_2 \text{ s.t. } (x,y) \in U \subseteq \|\phi\|\}.
\end{array}
$$

Observe that the definition is quite similar to the one we have presented in Definition 2.4.1 (Two-agent topological models). Van Benthem and Sarenac show in [50] that $\mathbf{S4}_2^{\mathcal{DC}}$ is also sound and weakly complete with respect to product models. In fact, they provide a very nice completeness result: $\mathbf{S4}_2^{\mathcal{DC}}$ is complete with respect to product models based on $\mathbb{Q} \times \mathbb{Q}$. Moreover, they show that these product models distinguish between common knowledge as a fixed point and common knowledge as infinite conjunction of iterated modalities of individual knowledge, something that we have mentioned in passing throughout this chapter. For the details, we refer the reader to either [50] or [48].

2.6 Multi-agent doxastic logics

Up to now, the concept of belief has remained only implicit in our discussion for multi-agent epistemic logics. However, we can easily think of a language having modal operators B_1 and B_2, meant to be interpreted as the belief modalities of Agent 1 and Agent 2, respectively, over a class of relevant models. Accordingly, notions of group belief have been proposed in the literature, among which the most promising concept is that of *common belief*. For a pertinent discussion on the subject, we refer the reader to [43], [45], and [38].

Common belief is meant to capture an essentially weaker precondition for agreement than common knowledge. In this case, it is evident that the notion is trying to embody beliefs that are ubiquitous within the members of an epistemic group.

As for the formal representations of such multi-agent doxastic languages, we mention that in [43], Pearce and Uridia provide a topological semantics for the formulas of a multi-agent doxastic language having modal operators for individual belief (B_1 and B_2), and a modal operator C_b for common belief. Their topological interpretation of individual belief is in terms of the co-derived set, which we reviewed in its single-agent case in the last part of Chapter 1. The soundness and completeness results Pearce and Uridia achieve in [43] involve, on the syntactic side, a system with the usual rules of inference, the **K4** axioms for the individual belief operators, and the equilibrium axiom and *induction rule* for the operator C_b; on the topological semantic side, the models are made up by bi-topological spaces that are TD-intersection closed, in the sense that the two relevant topologies, and their intersection, all satisfy the TD-axiom.

It would seem that this line of research strays a bit from our intentions, since even though it is a topological interpretation, it does not have much to do with the topo-semantics for belief that we will introduce in Chapter 3. However, it is more than worth mentioning, since it is in fact just a different approach to using topology for modeling comprehensive multi-agent systems that would in principle account both for knowledge and belief. As it should be obvious by now, this is exactly what this thesis is all about.

Part II

Topological Evidential Semantics for Knowledge (and Belief)

Chapter 3

Background. Single-agent e-models

The present chapter provides the *specific* conceptual framework for *our* proposal of a topological semantics for distributed knowledge (Chapter 4). Cutting to the chase, this semantics is based on the idea of *evidence* and *evidence*-based information attitudes. Though there are important results addressed, this chapter is meant to serve as *an explanation* concerning the road we took in our topological-modelling efforts. We stress the fact that the original contribution of this thesis is part of an ongoing investigation, the bedrock of which is set by the unpublished work by Baltag *et alia* in [4]. Section 3.2 lays out all the fundamental preliminaries that we will require from this investigation.

3.1 Evidence models

In [49], Van Benthem and Pacuit present models for evidence and evidence-management, and they use them to evaluate formulas of a language with a modal evidential-operator. The objective is to capture the idea of *'having evidence for ϕ,'* and producing a logic of evidence. Let us refer to these models as *evidence models.* With them, Van Benthem and Pacuit offer a particular interpretation of what we can call 'evidence-based belief.' Let us start with the suitable definitions.

Definition 3.1.1. (Evidence language) Let $\mathcal{L}_{ev}$ consist of a countable set $Prop$ of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators E^0 and B^v. The grammar of this evidential language is given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid E^0\phi \mid B^v\phi$$

for $p \in Prop$. The abbreviations for the connectives $\vee$, $\rightarrow$, $\leftrightarrow$ are the ones inherited from propositional logic. The possibility operators $\langle E^0 \rangle$ and $\langle B^v \rangle$ are defined as $\neg E^0 \neg$ and $\neg B^v \neg$, respectively.

Definition 3.1.2. (Evidence models) A tuple $(X, E, \mathcal{V})$ will be called an *evidence model* if and only if X is a set of possible worlds and $\mathcal{E} \subseteq \mathcal{P}(X)$ is a family of subsets of X such that $X \in \mathcal{E}$. This family is meant to be interpreted as the evidence available to a given agent at a world in X. The elements of $\mathcal{E}$ are informally called 'pieces of evidence.'

Let $(\mathcal{E})^{Max} = \{\mathcal{F} \subseteq \mathcal{E} \mid \mathcal{F}$ has the FIP and no proper $\subseteq$-extension of $\mathcal{F}$ has the $FIP\}$.[1]

The semantics for the formulas of $\mathcal{L}_{ev}$ is given recursively by:

[1]Recall that the FIP is the 'Finite Intersection Property.' A family of sets $\mathcal{F}$ has the FIP iff $\bigcap \mathcal{G} \neq \emptyset$ for every finite $\mathcal{G} \subseteq \mathcal{F}$ (i.e., every finite subfamily of $\mathcal{F}$ has non-empty intersection). Notice then that $(\mathcal{E})^{Max}$ is the collection of all the maximally FIP evidence families.

$$
\begin{array}{lcl}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X \backslash \|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|E^0\phi\| & = & \{x \in X \mid \exists G \in \mathcal{E} \text{ such that } G \subseteq \|\phi\|\} \\
\|B^v\phi\| & = & \{x \in X \mid \forall \mathcal{F} \in (\mathcal{E})^{Max}\ ,\ \bigcap \mathcal{F} \subseteq \|\phi\|\}.
\end{array}
$$

E^0 is the evidential-operator we mentioned before, so that '$E^0\phi$' expresses *'having evidence for ϕ.'* Observe that it has global extension, in the sense that $\|E^0\phi\| = X$ if there is some piece of evidence G supporting ϕ, and $\|E^0\phi\| = \emptyset$ otherwise.

B^v encodes a notion of evidence-based belief, and its motivation can actually be linked to Grove-belief (see Chapter 1, section 1.1.3). We will not go deep into the details of Van Benthem and Pacuit's evidence-based belief, but notice that the role of the worlds in Grove's definition is here played by FIP families of evidence, while the plausibility preorder is replaced by inclusion; the role of the most plausible worlds is played by the maximally-FIP families of evidence, and the role of (relational) local satisfaction of a given ϕ is played by the inclusion of $\bigcap \mathcal{F}$ in $\|\phi\|$. In fact, when the family $\mathcal{E}$ is closed under arbitrary intersections, one can show that evidence-based belief is equivalent to Grove-belief, where the plausibility preorder for the relevant space is given in the same way as the *specialization preorder* (i.e., defining $x \preccurlyeq y$ iff for every $G \in \mathcal{E}$, $x \in G$ implies that $y \in G$). However, this connection poses a problem for Van Benthem and Pacuit's evidence-based belief: like Grove-belief, it can be inconsistent.

The discussion above should remind the reader of a topological flavored scenario. In this line, notice that the definition of evidence-based belief implies that whatever is believed has to be consistent with *every* finite intersection of pieces of evidence.[2] This is a key fact for us. *Such a condition resembles very much the concept of topological density, which will dominate our topo-evidential definitions as of the next section.*

3.2 Topo-evidential semantics

According to Baltag *et alia* ([4]), we can use *topological* models to convey the same representations as Van Benthem and Pacuit's evidence models. What we will refer to as 'topological evidential models' or 'topo-evidential models' can in turn be used to formalize both knowledge and belief in a profitable manner. Such formalization offers a precise evidence-based interpretation for K and B, on the one hand, while bearing a relation with the (standard) topological semantics that we reviewed in Part I, on the other. In the present section, we will explore all of this.

Let $(X, \mathcal{E}, \mathcal{V})$ be an evidence model. Let $\tau_{\mathcal{E}}$ stand for the topology generated on X by the finite intersections of elements of $\mathcal{E}$. Trivially, the evidential feature captured by the operator E^0 (*'having a piece of evidence for ϕ'*) here coincides with the existence of a sub-basic open set of $\tau_{\mathcal{E}}$ supporting ϕ. However, the shift into a topological setting allows for the introduction of new evidence-related notions such as, for example, *'having combined evidence for ϕ'* (an open set supporting ϕ), or *'having conclusive evidence for ϕ at x'* (an open set supporting ϕ *and* including x). The questions at hand, then, are the following:

a) What can the topological setting contribute to the evidence-based definition of belief?

b) How can knowledge be taken to an evidential interpretation?

c) In what way will these topo-evidential accounts for knowledge and belief relate to standard topological semantics?

In [4], Baltag *et alia* provide satisfactory answers to all these questions. As we will see, once endowed with a topology of evidence, it is possible to generalize the concept of *Lewis-belief* (which is **KD45**) from

[2] This comes from the fact that every finite family pieces of evidence with non-empty intersection can be extended to a maximally FIP family.

its standard definition in *plausibility* models (see Chapter 1, section 1.1.3). After that, one can produce a topo-evidential interpretation of knowledge, according to whatever relation between belief and knowledge one would want to obtain. In the particular case of [4], the knowledge rendered is of the Stalnaker's kind, satisfying the **S4.2** axioms. Thus, we can say that the resulting semantics successfully amalgamates three different views:

1. The standard **KD45** conception of belief from plausibility models (i.e., what we have called Lewis-Belief), which, unlike Van Benthem and Pacuit's B^v, is always consistent.

2. The evidential paradigm, where agents' information attitudes are dependent on the evidence available.

3. The relation between knowledge and belief axiomatized by Stalnaker's combined system **KB**.

It seems that there are many different concepts lying around, and not only concepts, but also many different epistemological trends being addressed. Let us carefully elucidate the proposal. The motivation and intuitions can be summarized as follows:

- We can use topological models, built from Van Benthem and Pacuit's evidence models, to address evidence.

- *Alexandrov* topological models have clear relational counterparts (see Definition 1.2.3 and subsequent Observations).

- Relational plausibility models offer an adequate interpretation of belief (**KD45**).

- We can use the ideas of relational plausibility models of belief to polish the evidence-based definition of B^v and turn it into a purely topological notion (for the evidential topology, of course). This polished version of evidence-based belief will be a **KD45** kind of belief.

- With the topo-evidential notion of belief and Stalnaker's intuitions about the true relation between belief and knowledge, we can produce a topo-evidential notion of knowledge, that is correlated with the interior semantics of the concept but is *not the same*.

- The topo-evidential semantics for K and B validate Stalnaker's axioms, on the one hand. On the other, they have fertile evidential readings that allow for formalization of new attitudes and epistemic scenarios.

The details of this scheme can of course be found in [4], but it should be beneficial to review the basic definitions and results.

Definition 3.2.1. (Topological evidence language for knowledge and belief) Let $\mathcal{L}_{KBev}$ consist of a countable set $Prop$ of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators E^0, E, $\Box$, K, and B. The grammar of this evidential language is given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid E^0\phi \mid E\phi \mid \Box\phi \mid K\phi \mid B\phi$$

for $p \in Prop$. The abbreviations for the connectives $\vee$, $\rightarrow$, $\leftrightarrow$ are the ones inherited from propositional logic. The possibility operators $\langle E^0\rangle$, $\langle E\rangle$, $\Diamond$, $\langle K\rangle$, and $\langle B\rangle$, are defined as $\neg E^0\neg$, $\neg E\neg$, $\neg\Box\neg$, $\neg K\neg$, and $\neg B\neg$, respectively.

Definition 3.2.2. (Single-agent topo-e-models)

Let X be a set and $\mathcal{E} \subseteq \mathcal{P}(X)$. A tuple $(X, \mathcal{E}, \tau_{\mathcal{E}}, \mathcal{V})$ is called a *topo-evidential model* (topo-e-model, in short) if and only if $\mathcal{E} \subseteq \mathcal{P}(X)$ is a sub-basis for $\tau_{\mathcal{E}}$ (i.e., $\tau_{\mathcal{E}}$ is generated by the family of finite intersections of elements of $\mathcal{E}$) and $\mathcal{V} : Prop \rightarrow \mathcal{P}(X)$ is a valuation function. The elements of the family $\tau_{\mathcal{E}}$ will be called *sub-basic evidence* sets.

Before the explicit definition of the semantics for the formulas, we present the evidential reading of the models:

- The sets $G \in \mathcal{E}$ are interpreted as *sub-basic pieces of evidence*, possessed by the agent. These can be seen as propositions that (s)he obtained either by direct observation, by the testimony of others, etc. However, they are not necessarily *factive*, they might be false in the actual world. This fits with the common usage of the term *evidence* in everyday-life (where we talk about 'fake evidence,' 'misleading evidence,' or 'evidence that points in the wrong direction'). However, it is in contrast with the most widespread usage of this term in epistemology, according to which evidence is always factive.

- A *derived evidence* (or just evidence, in short) is any non-empty intersection of finitely many pieces of evidence; i.e., any set of the form $G = \bigcap \mathcal{F}$, where $\mathcal{F} \subseteq \mathcal{E}$ is a finite, non-empty family of sub-basic evidence sets such that $\bigcap F \neq \emptyset$. Notice that all the evidence sets (both sub-basic and derived) are open in $\tau_{\mathcal{E}}$. Indeed, open sets will be the unions of derived evidence sets, so that the collection of derived evidence sets constitute a basis for the topology $\tau_{\mathcal{E}}$.

- An *argument for* p is a union of a non-empty family of derived evidences that entail p (i.e., a union of the form $\bigcup_{i \in I} G_i$, where $G_i \subseteq \|p\|$ for all $i \in I$). Thus, an argument may provide multiple evidential paths G_i to support a common conclusion p. In topological terms, an argument for p amounts to a non-empty open set U such that $U \subseteq \|p\|$.

- In order to refer to the always factive type of evidence, we will use the expression 'conclusive evidence.'

  In other words, we say that $G \in \mathcal{E}$ is a *conclusive piece of evidence* at world $x \in X$ whenever it is true at x (i.e. $x \in G$). Similarly, for any finite $\mathcal{F} \subseteq \mathcal{E}$, we say that $\bigcap \mathcal{F}$ is *derived conclusive evidence* at world $x \in X$ whenever it is derived evidence that is true at x (i.e. $x \in \bigcap \mathcal{F}$).

  Finally, an open set $U \in \tau$ is a *conclusive argument for* p at world x iff $x \in U \subseteq \|p\|$.

- In accordance with the use we have been giving to the word 'support,' we say that an evidence (be it sub-basic or derived) G *supports* a proposition p (or that 'G is evidence for p') whenever $G \subseteq \|p\|$. In this case, we also say that *the agent has evidence for* p. If the evidence for p is factive at world $x \in X$ (meaning that x lies within the evidence open set), then we say that *the agent has conclusive evidence for* p. Similarly, we say that an argument supports p if it is an *'argument for* p*'* in the sense of the treatment addressed above.

The association between evidence and topology can be given by:

sub-basic evidence	–	sub-basic opens of $\tau_{\mathcal{E}}$
derived evidence	–	basic opens of $\tau_{\mathcal{E}}$
arguments	–	(non-empty) opens of $\tau_{\mathcal{E}}$.

Remark 3.2.1. Notice that, according to the evidential treatment of terms, we are going to say that an agent *has conclusive evidence for* p at $x \in X$ if x lies within the interior of $\|p\|$.

Let us address the construction of the topo-evidential notion of belief. Recall that in section 3.1 we mentioned that Van Benthem and Pacuit's evidence-based belief need not be consistent always, and that it can be seen as a kind of Grove-belief. Baltag *et alia*'s proposal for belief, then, can be seen as an evidential analogue of the concept of Lewis-belief that we presented in Chapter 1, section 1.1.3. The intuition is that something is believed iff *it is entailed by all the sufficiently strong derived evidence.* Since the definition will only imply derived evidence, we can state it in purely topological terms. Once again, it will be related to topological density.[3]

We will say that p is believed at $x \in X$ iff every argument (or non-empty open) $V \in \tau_{\mathcal{E}}$ can be strengthened to an argument for p (i.e., iff for every $V \in \tau_{\mathcal{E}}$ such that $V \neq \emptyset$, there exists $V' \in \tau_{\mathcal{E}}$ such that $V' \subseteq V \cap \|p\|$). We can show, then, that in every topo-evidential model of the form $(X, \mathcal{E}, \tau_{\mathcal{E}}, \mathcal{V})$, believing p is equivalent to the existence of $U \in \tau_{\mathcal{E}}$ such that $U \subseteq \|p\|$ and U is consistent (or has non-empty intersection) with *every* other evidence open set.

[3]Recall that for any topological space (X, τ) and $A \subseteq X$, we say that A is τ-*dense* in X iff $Cl(A) = X$, or, equivalently, if for every non-empty open set $O \in \tau$, $O \cap A \neq \emptyset$.

Therefore, we will take $\|B\phi\| := \{x \in X \mid \exists U \in \tau_{\mathcal{E}} \text{ such that } U \subseteq \|\phi\| \wedge Cl(U) = X\}$.

Any dense argument $U \subseteq \|p\|$ is called a *justification for* p. In other words, a justification for p is an argument for p which is not contradicted by any other evidence (or argument). Therefore, a justification for p is nothing more than a union of derived evidences that supports p and is consistent with every derived evidence. We observe, then, that Baltag *et alia*'s notion of belief can be seen as 'justified belief': a proposition p is believed iff the agent has a justification for p.

As for knowledge, we can then define a topo-evidential notion of knowledge, that better approximates the common usage of the word than the infallible knowledge (as truth in all the possible worlds). For reasons that will become clear, we will call this Stalnaker's knowledge (see subsection 3.2.2). Essentially, this concept corresponds to having a *correct (i.e., true) justification*.

We will say that an agent *knows* p at x iff there exists some *justification* $U \in \tau_{\mathcal{E}}$ for p such that $x \in U$. In other words, p is known iff there exists some correct (i.e., true) *argument for* p that is consistent with all the available evidence.

After this presentation, we are entitled to introduce the full semantics for the formulas of $\mathcal{L}_{KBev}$, given recursively by:

$$
\begin{array}{lcl}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X \setminus \|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|E^0\phi\| & = & \{x \in X \mid \exists G \in \mathcal{E} \text{ such that } G \subseteq \|\phi\|\} \\
\|E\phi\| & = & \{x \in X \mid int\ \|\phi\| \neq \emptyset\} \\
\|\Box\phi\| & = & \{x \in X \mid \exists U \in \tau_{\mathcal{E}} \text{ such that } x \in U \subseteq \|\phi\|\} \\
\|B\phi\| & = & \{x \in X \mid \exists U \in \tau_{\mathcal{E}} \text{ such that } U \subseteq \|\phi\| \wedge Cl(U) = X\} \\
\|K\phi\| & = & \{x \in X \mid \exists U \in \tau_{\mathcal{E}} \text{ such that } x \in U \subseteq \|\phi\| \wedge Cl(U) = X\}.
\end{array}
$$

Truth and *validity* are defined the same way as for standard Kripke semantics. As usual, satisfaction at a given world is marked as $\mathcal{M}, x \vDash \phi$.

A bit of discussion about the semantics seems appropriate.

First of all, the semantics for operator $\Box$ is quite interesting. It is meant to capture the notion of '*having conclusive evidence (derived evidence, or arguments) for* ϕ.' Observe that the agent has conclusive evidence for ϕ at x iff $x \in int\ \|\phi\|$. In other words, the modality $\Box$ coincides with the interior operator:

$$\|\Box\phi\| = int\ \|\varphi\|.$$

This note points to a major difference between Baltag *et alia*'s framework and the standard topological semantics: while in the standard semantics the interior operator represents knowledge of something, in this interpretation the interior represents only 'having true evidence for' something. The difference arises from the fact that an agent may be in possession of some evidence that happens to be true, without the agent necessarily knowing, or even believing, that this evidence is true.

Observation 3. Notice that we can restate the definitions for the knowledge and belief operators in the following terms:

- $\|K\phi\| = \{x \in X \mid x \in int\ \|\phi\| \wedge Cl(int\ \|\phi\|) = X\}$
- $\|B\phi\| = \{x \in X \mid Cl(int\ \|\phi\|) = X\}$

It is easy to see, then, that there is a particular relation between the topo-evidential belief and the topo-evidential knowledge:

$$K\phi \Leftrightarrow (\Box\phi \wedge B\phi).$$

So this Stalnaker's knowledge amounts to having conclusive evidence for ϕ and believing that ϕ (with justification). In other words, this knowledge can be seen 'correctly justified belief.'

On a different note, regard that B turns out to be a global modality. It has global extension, meaning that if there is $x \in X$ such that $x \in \|B\phi\|$, then $\|B\phi\| = X$. Similarly, if $Cl(int\ \|\phi\|) \neq X$, then $\|B\phi\| = \emptyset$.

Operators E^0 and E are inspired by Van Benthem and Pacuit's evidential-operator. They are used to track down the possession of sub-basic evidence, resp. of derived evidence and arguments. We presented them here just for the sake of completion, since the language that we will be dealing with in the next chapters does not include evidential operators.

3.2.2 Soundness and completeness results for fragments of the language $\mathcal{L}_{KBev}$

The fragments of the language $\mathcal{L}_{KBev}$ that matter to us are $\mathcal{L}_K$, $\mathcal{L}_B$, and $\mathcal{L}_{KB}$. Here, we compile the basic 'soundness and completeness' results for said fragments that Baltag *et alia* produce in [4]. Such results concern the epistemic/doxastic systems addressed in Chapter 1. In what follows, we will use the abbreviated locution topo-e-model to refer to our topo-evidential models. Moreover, since we will focus only on the topology of evidence, without referring directly to the family of sub-basic evidence $\mathcal{E}$, *the signature of our topo-e-models will not include the family* $\mathcal{E}$ from here on.

Soundness

The proofs for soundness are almost redundant, but they already involve the exact type of methods that we will use in the proof of soundness for our logic of distributed knowledge, so it should be beneficial to include them in full display.[4]

We will use a convenient preliminary Lemma:

Lemma 3.2.3. *Let (X, τ) be a topological space. For every $U, V \in \tau$ such that U and V are both τ-dense, we have that $U \cap V$ is also τ-dense.*

Proof. Let $U, V \in \tau$ such that $Cl(U) = X$ and $Cl(V) = X$. For any open set $O \in \tau$, we have that $O \cap V \neq \emptyset$ and $O \cap V \in \tau$. Thus, $O \cap V$ is a non-empty open set of τ. As such, it must intersect U: $(U \cap V) \cap O = U \cap (O \cap V) \neq \emptyset$. Therefore, we have shown that $Cl(U \cap V) = X$. □

Likewise, we point out that in these soundness proofs, we will be using the fact that for any topo-e-model $\mathcal{M}$ (with a set X as its domain), we have that $\|\phi \to \psi\|^{\mathcal{M}} = X$ iff for every $x \in X$, $x \in \|\phi\|^{\mathcal{M}}$ implies that $x \in \|\psi\|^{\mathcal{M}}$.

*($\mathcal{L}_K$ – **S4.2**)*

Proposition 3.2.4. *The system* **S4.2** *is sound for $\mathcal{L}_K$ with respect to the class of all topo-e-models with the knowledge-only semantics of Definition 3.2.2.*

Proof. We proceed as it is usual for soundness proofs.

We show that the given axioms are semantic validities, and that the rules of inference preserve validity. In this sense, we just mention that it is well-known that the propositional tautologies are valid, and that *Modus Ponens* preserves validity.

Let $(X, \tau, \mathcal{V})$ be any topo-e-model.

- (T) Let $x \in X$. If $x \in \|Kp\|$, then $x \in int\ \|p\| \subseteq \|p\|$ (recall Observation 3).
- (K) Let $x \in X$. Assume that $x \in \|K(p \to q)\|$ and $x \in \|Kp\|$. Then, on one hand, there exists $U \in \tau$ such that $x \in U$, $U \subseteq \|(p \to q)\|$, and $Cl(U) = X$; on the other, there exists $V \in \tau$ such that $x \in V$, $V \subseteq \|p\|$, and $Cl(V) = X$. We have that $U \cap V$ is an open of τ such that $x \in U \cap V$, $U \cap V \subseteq \|q\|$ and $Cl(U \cap V) = X$ (Lemma 3.2.3), which means that $x \in \|Kq\|$.

[4] Notice that soundness of **S4.2** for $\mathcal{L}_K$ and soundness of **KD45** for $\mathcal{L}_B$ are actually implied by soundness of **KB** for $\mathcal{L}_{KB}$ (Recall from Chapter 1, section 1.1.3, that the system **KB** yields a **KD45** pure belief logic and an **S4.2** pure knowledge logic.)

- (4) Let $x \in X$. Assume that $x \in \|Kp\|$. Since for every $y \in \|p\|$ we have that $y \in int\ \|p\| \subseteq \|p\|$ and that $Cl(int\ \|p\|) = X$, then for every $y \in int\ \|p\|$, $y \in \|Kp\|$. Therefore, there exists a straightforward open set $U \in \tau$ such that $x \in U$, $U \subseteq \|Kp\|$, and $Cl(U) = X$, namely $U = int\ \|p\|$. Therefore, $x \in \|KKp\|$.

- (.2) Assume that $x \in \|\langle K \rangle Kp\|$.

  First of all, we observe that this assumption implies that $Cl(int\ \|p\|) = X$. We show this by contraposition. Suppose that $Cl(int\ \|p\|) \neq X$; then, for all $y \in X$, $y \in \|\neg Kp\|$; thus, there exists a straightforward open set $U \in \tau$ such that $x \in U$, $U \subseteq \|\neg Kp\|$, and $Cl(U) = X$ (namely $U = X$); therefore, $x \in \|K\neg Kp\| = X \backslash \|\langle K \rangle Kp\|$.

  Next, we show that $Cl(int\ \|p\|) = X$ implies that $x \in \|K\langle K \rangle p\|$. Notice that if $Cl(int\ \|p\|) = X$, then $Cl(\|p\|) = X$. It is easy to realize that this implies that $int\ \|\neg p\| = int\ (X \backslash \|p\|) = \emptyset$. Therefore, for all $y \in X$, $y \in \|\neg K\neg p\|$. As before, this renders that there exists a straightforward open set $U \in \tau$ such that $x \in U$, $U \subseteq \|\neg K\neg p\|$, and $Cl(U) = X$ (namely $U = X$). This last thing shows that $x \in \|K\neg K\neg p\| = \|K\langle K \rangle p\|$.

Finally, we show that the *Necessitation* rule for K preserves validity. Assume that for every topo-e-model $(X, \tau, \mathcal{V})$, $\|\phi\| = X$. This means that for every $x \in X$, $x \in \|\phi\|$. Observe then that there exists a straightforward open set $U \in \tau$ such that $x \in U \subseteq \|\phi\|$ and $Cl(U) = X$, namely $U = X$. Therefore, $x \in \|K\phi\|$, so that $\|K\phi\| = X$. Since the space is arbitrary, we get the desired property. ☐

*($\mathcal{L}_B$ – **KD45**)*

Proposition 3.2.5. *The system* **KD45** *is sound for $\mathcal{L}_B$ with respect to the class of all topo-e-models with the belief-only semantics of Definition 3.2.2.*

Proof. Let $(X, \tau, \mathcal{V})$ be any topo-e-model.

- (K) Let $x \in X$. Assume that $x \in \|B(p \to q)\|$ and $x \in \|Bp\|$. Then, on one hand, there exists $U \in \tau$ such that $U \subseteq \|(p \to q)\|$, and $Cl(U) = X$; on the other, there exists $V \in \tau$ such that $V \subseteq \|p\|$ and $Cl(V) = X$. We have that $U \cap V$ is an open of τ such that $U \cap V \subseteq \|q\|$ and (by Lemma 3.2.3), $Cl(U \cap V) = X$. Therefore, $x \in \|Bq\|$.

- (D) Let $x \in X$. Assume that $x \in \|Bp\|$. Then $Cl(int\ \|p\|) = X$ (recall Observation 3). This implies that $Cl(\|p\|) = X$, which in turn renders that $int\ \|\neg p\| = \emptyset$. But this means that for every $y \in X$, $y \in \|\neg B\neg p\|$. In particular, $x \in \|\neg B\neg p\| = \|\langle B \rangle p\|$.

- (4) Let $x \in X$. Assume that $x \in \|Bp\|$. We know that this implies that $\|Bp\| = X$. Therefore, there exists a straightforward open set $U \in \tau$ such that $U \subseteq \|Bp\|$ and $Cl(U) = X$ (namely $U = X$). With this, we have that $x \in \|BBp\|$.

- (5) Let $x \in X$. Assume that $x \in \|\neg Bp\|$. Then $Cl(int\ \|p\|) \neq X$ (Observation 3). Thus, $\|\neg Bp\| = X$. Therefore, there exists a straightforward open set $U \in \tau$ such that $U \subseteq \|\neg Bp\|$ and $Cl(U) = X$ (namely $U = X$). With this, we have that $x \in \|B\neg Bp\|$.

That *Necessitation* for B preserves validity can be shown in the exact same way as we did above, for K. ☐

*($\mathcal{L}_{KB}$ – **KB**)*

Proposition 3.2.6. *The system* **KB** *is sound for $\mathcal{L}_{KB}$ with respect to the class of all topo-e-models with the knowledge-belief semantics of Definition 3.2.2.*

Proof. Let $(X, \tau, \mathcal{V})$ be any topo-e-model.

- Above, we showed that the axioms (T), (K), (4), and (CB) = (5) are all semantic validities in the proposed semantics.

- (KB) Straight from definitions, $\|Kp \to Bp\| = X$.

- (SPI) Let $x \in X$. Assume that $x \in \|Bp\|$. We know that this implies that $\|Bp\| = X$. Therefore, there exists a straightforward open set $U \in \tau$ such that $x \in U$, $U \subseteq \|Bp\|$, and $Cl(U) = X$, namely $U = X$. Therefore, $x \in \|KBp\|$.
- (SNI) Let $x \in X$. Assume that $x \in \|\neg Bp\|$. Then $Cl(int\ \|p\|) \neq X$ (Observation 3). Thus, $\|\neg Bp\| = X$. Therefore, there exists a straightforward open set $U \in \tau$ such that $x \in U$, $U \subseteq \|\neg Bp\|$, and $Cl(U) = X$, namely $U = X$. Therefore, $x \in \|K\neg Bp\|$.
- (FB) Let $x \in X$. Assume that $x \in \|Bp\|$. We get that there exists an open set $U \in \tau$ such that $U \subseteq \|p\|$ and $Cl(U) = X$. Notice that for every $y \in U$, $y \in \|Kp\|$. Thus, $U \subseteq \|Kp\|$, and since $Cl(U) = X$, then $x \in \|BKp\|$.

We have already shown that the *Necessitation* rule for K, B preserves validity. Therefore, we are done. □

Completeness

Notice that *any* topological space serves as basis for two (essentially different) models for the formulas of $\mathcal{L}_{KB}$, under the *same* valuation function. One is given by the semantics of Definition 1.2.10 (Topological models), and the other is given by the semantics of Definition 3.2.2 (Topo-e-models). As pointed out by Baltag *et alia* in [4], there is a special bond between the alternate semantics. Because of this, the results of completeness for the epistemic/doxastic systems with respect to topo-e-models revolve around the corresponding propositions that we reviewed in the case for extremally disconnected models (see Chapter 1, section 1.2).

We will not address the subtleties of the completeness proofs. Rather, we provide brief arguments to outline their basic methodology, which, again, involves 'going through relational completeness.'

*($\mathcal{L}_K$ – **S4.2**)*

For **S4.2**, the completeness proof *(for the fragment $\mathcal{L}_K$)* presented by Baltag *et alia* in [4] makes use of the fact that every topological space is endowed with an extremally disconnected topology, the topology of all its open dense sets (plus $\emptyset$). Let us sketch the argument, for it serves as appropriate background.

By virtue of Lemma 3.2.3, we can see that the set $\{O \in \tau \mid Cl(O) = X\} \cup \{\emptyset\}$ forms a topology on X. We use the notation $\overline{\tau}$ to refer to this topology. Notice then that Observation 3 implies that, for a given topo-e-model $(X, \tau, \mathcal{V})$, $\|K\phi\| = int_{\overline{\tau}}\|\phi\|$ (where we have used the subindex $_{\overline{\tau}}$ to make explicit note that it is the interior in the topology $\overline{\tau}$). For a given topo-e-model $\mathcal{M} = (X, \tau, \mathcal{V})$, we will define $\overline{\mathcal{M}} := (X, \overline{\tau}, \mathcal{V})$.

For the moment, let us use different notation for the extensions of the formulas, according to whether we are using Definition 1.2.10 (Topological models) or Definition 3.2.2 (Topo-e-models). Let $[\phi]^{\mathcal{M}}$ stand for the extension in the interior semantics, while $\|\phi\|^{\mathcal{M}}$ is kept as the extension in topo-e-models, with a dense-interior semantics. As Baltag *et alia* do in [4], one can easily show that

Proposition 3.2.7. *For every space (X, τ), $x \in X$, and ϕ of $\mathcal{L}_K$, $\|\phi\|^{\mathcal{M}} = [\phi]^{\overline{\mathcal{M}}}$.*

□

Once this has been settled, Baltag *et alia* show that for a formula ϕ_* of $\mathcal{L}_K$ which is not a theorem of **S4.2**, there exists a very particular Kripke model $(X, \leq, \mathcal{V})$ (based on a finite rooted reflexive and transitive cofinal frame) that falsifies ϕ_*. After applying Proposition 1.2.6, we get a *topological model* (in the sense of Definition 1.2.10) $\mathcal{M} = (X, \tau_{\leq}, \mathcal{V})$ that falsifies ϕ_* (i.e., such that $[\phi_*]^{\mathcal{M}} \neq X$). Incidentally, for the particular class of Kripke models used (based on finite rooted reflexive and transitive cofinal frames), the topology $\tau_{\leq}$ turns out to be *exactly* equal to $\overline{\tau_{\leq}}$. So, using the terminology introduced, $\mathcal{M} = \overline{\mathcal{M}}$. By Proposition 3.2.7 above, then, we get that $\|\phi_*\|^{\mathcal{M}} = [\phi_*]^{\overline{\mathcal{M}}} = [\phi_*]^{\mathcal{M}} \neq X$. But this means that we have found a topo-e-model that falsifies ϕ_*. Therefore, we have weak completeness of **S4.2** for $\mathcal{L}_K$ *with respect to topo-e-models.*

*($\mathcal{L}_B$ – **KD45**)*

For **KD45**, the proof presented in [4] uses the relational structures called *brushes* that we mentioned in Chapter 1 (see Proposition 1.2.12). We know that **KD45** is complete for $\mathcal{L}_B$ with respect to the class of brushes. Now, it can be shown that for every relational Kripke model $\mathcal{M}^{rel} = (X, <, \mathcal{V})$ based on a brush, the (relational) extension $[\![\phi]\!]^{\mathcal{M}^{rel}}$ of a given formula ϕ of $\mathcal{L}_B$ coincides with the topo-e-extension $\|\phi\|^{\mathcal{M}^{top}}$

in the *Alexandrov* topo-e-model $\mathcal{M}^{top} := (X, \tau_{(<)^*}, \mathcal{V})$, where $(<)^*$ is the reflexive transitive closure of the (brush's) irreflexive $<$. Therefore, for a formula ϕ_* of $\mathcal{L}_B$ which is not a theorem of **KD45**, we can find a brush model that falsifies ϕ_* such that the 'associated' *Alexandrov* topo-e-model (built as above) also falsifies ϕ_*. With this, we have weak completeness of **KD45** for $\mathcal{L}_B$ with respect to topo-e-models.

*($\mathcal{L}_{KB}$ – **KB**)*

For **KB**, the argument goes as follows. For a formula ϕ_* of $\mathcal{L}_{KB}$ which is not a theorem of **KB**, ϕ_* can be reduced to a formula ψ_* of $\mathcal{L}_K$ by the correspondence given in $\vdash_{\mathbf{KB}} B\phi \leftrightarrow \langle K \rangle K\phi$ (see Proposition 1.1.2 from Chapter 1). Since we also know that $\mathbf{S4.2} \subseteq \mathbf{KB}$, we get that the corresponding ψ_* is not a theorem of **S4.2**. By completeness of **S4.2** with respect to topo-e-models, we can find a topo-e-model that falsifies ψ_*, and hence falsifies ϕ_*. With this, we get weak completeness of **KB** for $\mathcal{L}_{KB}$ with respect to topo-e-models.

Chapter 4

A logic for distributed knowledge

As of this point, all that follows constitutes the original contribution of this thesis. In this chapter, we introduce a proposal for a topo-evidential logic of distributed knowledge, with a traditional syntax, on one hand, and a semantics given by very particular topological models within the evidence-model paradigm, on the other.

4.1 Motivation

If our goal is to model epistemic/doxastic logics of group knowledge (for two-agent groups) within tradition, a likely candidate for the formal language would be constructed as follows. A countable set *Prop* of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators K_1, K_2, B_1, B_2, D, and C, such that the grammar would be given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid K_1\phi \mid K_2\phi \mid B_1\phi \mid B_2\phi \mid D\phi \mid C\phi \tag{4.1}$$

for p in *Prop*. It is a considerable task to find *nice* models for such a rich language. But the literature on the profitability of topological semantics is also considerable. So, in colloquial terms, it is definitely worth a shot. The ambitious first goal of this thesis was precisely to use the single-agent topo-evidential semantics developed by Baltag *et alia* in [4] as basis for two-agent topological models that would provide an interesting—if possible 'well-behaved'—semantic counterpart to the syntax mentioned above.

Recall that in single-agent topo-e-models, the semantics for knowledge and belief are given in terms of dense interiors (informally, conclusive evidential arguments that are consistent with all the other available evidence). As such, the relevant extensions are

$$\begin{aligned}
\|K\phi\| &= \{x \in X \mid \exists U \in \tau \text{ s.t.} x \in U \subseteq \|\phi\| \wedge \ Cl(U) = X\} \\
\|B\phi\| &= \{x \in X \mid \exists U \in \tau \text{ s.t. } U \subseteq \|\phi\| \wedge \ Cl(U) = X\}.
\end{aligned}$$

When trying to take this topo-e-framework to the multi-agent case—*working with the underlying ideas addressed in Chapter 2*—we were faced by two problems. One is of a relatively conceptual kind, and it concerns the degree of mutual knowledge one would want to admit between the given agents, and the other constitutes a technical mismatch with the intuitive direction in which the relevant research pointed us.

Let us entertain the scenario. Call the language whose grammar is given in (4.1) $\mathcal{L}_2$. Suppose we build two-agent topological models for its formulas from bi-topological spaces (X, τ_1, τ_2) such that, for a valuation function $\mathcal{V} : Prop \to \mathcal{P}(X)$, the relevant semantics was given (in an extensional form) by

$$
\begin{aligned}
[K_i\phi] &\stackrel{?}{=} \{x \in X \mid \exists U \in \tau_i \text{ s.t.} x \in U \subseteq [\phi] \wedge \ Cl_{\tau_i}(U) = X\} \\
[B_i\phi] &\stackrel{?}{=} \{x \in X \mid \exists U \in \tau_i \text{ s.t. } U \subseteq [\phi] \wedge \ Cl_{\tau_i}(U) = X\} \\
[D\phi] &\stackrel{?}{=} \{x \in X \mid \exists U \in \tau_1 \vee \tau_2 \text{ s.t. } x \in U \subseteq [\phi] \wedge \ Cl_{\tau_1 \vee \tau_2}(U) = X\} \\
[C\phi] &\stackrel{?}{=} \{x \in X \mid \exists U \in \tau_1 \cap \tau_2 \text{ s.t. } x \in U \subseteq [\phi] \wedge \ Cl_{\tau_1 \cap \tau_2}(U) = X\}.
\end{aligned}
$$

The motivation for the hypothetical definitions for K_i and B_i should be clear. As for C and D, recall that, according to the usual interpretation of knowledge as topological interior, distributed and common knowledge are just third-agent kinds of knowledge, namely the knowledge in the join, resp. intersection, of the agents' topologies (see Chapter 2, section 2.4). Thus, it does not seem far-fetched to propose D and C as above.

However, let us see what happens with these models. A survey of the predicaments they led us into will help the reader understand the choice of topological semantics we introduce in section 4.2.

The conceptual problem: belief is common knowledge!

As seen in Chapter 3, the topo-e-semantics for B yield a global modality. In the hypothetical bi-topological models that we have presented, we would have that if $x' \in [B_i\phi]$ for some $i \in \{1, 2\}$ and $x' \in X$, then $[B_i\phi]$ would be the whole space X. Let us say, for example, that Agent 1 believes that ϕ at world x'. Therefore, at any world $y \in X$, Agent 1 still believes that ϕ. But there is something more. It turns out that, in this setting, Agent 2 will *know* that Agent 1 believes that ϕ, at any world in X. If $[B_1\phi] \neq \emptyset$, then $[B_1\phi] = X$, so that for every $y \in X$, we have that $X (\in \tau_2)$ is such that $y \in X = [B_1\phi]$ and $Cl_{\tau_2}(X) = X$. Hence, $y \in [K_2B_1\phi]$. This happens for the beliefs of Agent 2 as well. Agent 1 would end up knowing them all. Moreover, both agents would know this, they would know that they know this, and so on. Regardless of any formal hypothetical conception of common knowledge in this paradigm, the situation is highly undesirable from the epistemological perspective, to say the least.

The technical problem: density is not preserved under join (or intersection) of topologies.

Chapter 2 collects plausible ways in which to topologically model distributed and common knowledge. As discussed, this is done in terms of the join, resp. intersection, of the agents' topologies (which are are meant to embody the agents' epistemic justifications). In the framework we wish to incorporate, these justifications are interpreted as conclusive evidence that is consistent with all the other evidence available. Now, if we choose to interpret D and C as above, some of the customary—widely accepted—axioms for both concepts fail to hold. The reason for the collapse comes from the fact that density of the open sets of τ_i is preserved neither in $\tau_1 \vee \tau_2$ nor in $\tau_1 \cap \tau_2$. The following examples should clarify a bit this issue, and at the same time provide a warm-up for the sort of topological arguments we will be employing in our proofs later on.

Example 4.1.1. ($K_i\phi \not\Rightarrow D\phi$) Let $X = \{a, b\}$, $\tau_1 = \{X, \emptyset, \{a\}\}$, $\tau_2 = \{X, \emptyset, \{b\}\}$. Clearly, $\tau_1 \vee \tau_2 = \{X, \emptyset, \{a\}, \{b\}\}$. Let $\mathcal{V} : \{p\} \to \mathcal{P}(X)$ be such that $\mathcal{V}(p) = \{a\}$. Here, $a \vDash K_1p$, because $\{a\} \in \tau_1$, $\{a\} \subseteq [p]$, and $Cl_{\tau_1}(\{a\}) = X$. But $\{b\} \in \tau_1 \vee \tau_2$ and $\{b\} \cap \{a\} = \emptyset$. Therefore, every $(\tau_1 \vee \tau_2)$-neighborhood of a included in $[p]$, namely $\{a\}$, is not $(\tau_1 \vee \tau_2)$-dense, which means that $a \nvDash Dp$.

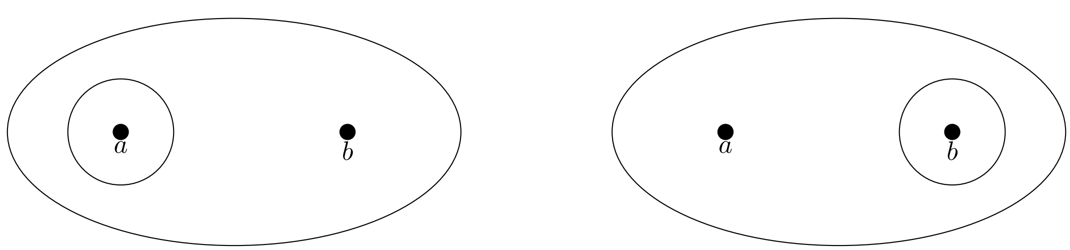

(X, τ_1) (X, τ_2)

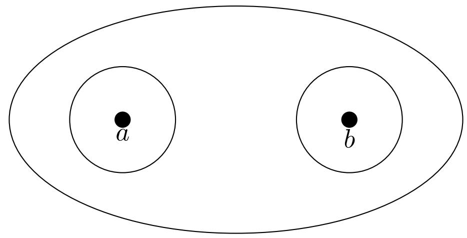

$(X, \tau_1 \vee \tau_2)$

Example 4.1.2. ($C\phi \not\Rightarrow K_1\phi \wedge K_2\phi$). Let $X = \{a, b\}$, $\tau_1 = \{X, \emptyset, \{a\}\}$, $\tau_2 = \{X, \emptyset, \{a\}, \{b\}\}$. Clearly, $\tau_1 \cap \tau_2 = \tau_1$. Let $\mathcal{V} : \{p\} \to \mathcal{P}(X)$ be such that $\mathcal{V}(p) = \{a\}$. We observe that $a \vDash Cp$: $\{a\} \in \tau_1 \cap \tau_2$, $\{a\} \subseteq [p]$, and $Cl_{\tau_1 \cap \tau_2}(\{a\}) = X$. But $\{b\} \in \tau_2$ and $\{b\} \cap \{a\} = \emptyset$. Therefore, every ($\tau_2$)-neighborhood of a included in $[p]$, namely $\{a\}$, is not (τ_2)-dense, which means that $a \nvDash K_2 p$.

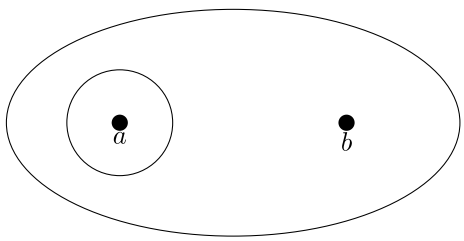

(X, τ_1)

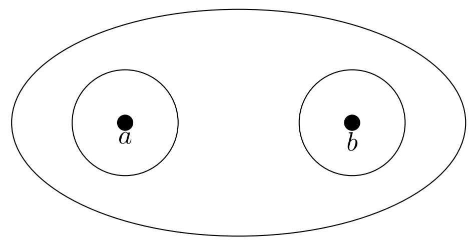

(X, τ_2)

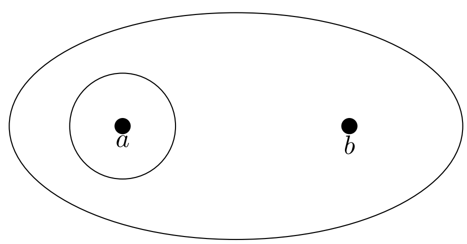

$(X, \tau_1 \cap \tau_2)$

Because of these problematic arguments and examples, it is clear that, in order to implement the topo-e-framework in a multi-agent scenario, several adjustments must be made. We took it as our job to go step-by-step in such an endeavor. As mentioned in the Introduction, the proposal that we will be dealing with *is restricted* to the fragment of $\mathcal{L}_2$ having only the operators K_1, K_2, and D. It should be worthwhile to briefly sketch, then, the virtual solutions to both of the problems addressed.

The solutions

For the first problem, we introduce *partitioned* topological spaces. Even when we do not tackle thoroughly the formal aspect of beliefs in this thesis, the topological models that we introduce open the possibility of interpreting beliefs in a way that does not collapse them into common knowledge. By assigning to each agent a partition of the state space, we circumscribe their epistemic scope to *information cells.* At a given world $x \in X$, the agent will consider as epistemically possible only the worlds lying in the same class as x. The idea is borrowed from epistemic game theory. Though we are not implementing it, we are certainly inspired by 'Aumann's structure model' (see [2]), in which each agent's knowledge is coupled with an information partition. As opposed to his take, we *do not* support an **S5** brand of knowledge. We adopt (and adapt) the information-partitions criterion under the philosophical tradition that differentiates between *information* and *knowledge* (see [23] and [14]).

For the second problem, though in principle technical, the solution stems from a clever conceptual reinterpretation of the notion of topological evidence. Such reinterpretation is ascribed to Baltag, and implies immersing topo-evidential knowledge into a slightly different scheme, that of topological *Learning Theory* (see [32]). Based on ideas from Baltag *et alia* (see [7]), we introduce a *topology of learnable evidence*.[1] This topology is meant to represent all the pieces of evidence that that agents can *potentially* come to learn (within their respective information cells at a given world, of course). As such, this topology of learnable evidence must be taken to include *both* agents' evidential topologies. In this light, for example, 'having conclusive evidence for ϕ' now becomes 'having *current* conclusive evidence,' The conceptual shift allows us to define—in a very precise way—an agreeable alternative for distributed knowledge which is a) based on the usual topological interpretation of D, b) consistent with the information-partitions criterion, c) evidential in nature, and d) satisfactory in terms of the usually accepted axioms for distributed knowledge.

In the next section, we will formalize these solutions with precision.

4.2 Syntax and semantics (the logic TKD)

We present a topological semantics for the formulas of a formal language with operators for knowledge and distributed knowledge (i.e. the fragment of $\mathcal{L}_2$ having only the operators K_1, K_2, and D). In this manner, we are endowed with a logic for distributed knowledge. We will refer to this logic as TKD (standing for Topological Logic for Knowledge and Distributed Knowledge).

Before the statement of our language, we recover some conventions about our topological notation. We also introduce key topological concepts that we will use constantly from here on.

Let X be a set, and let $\mathcal{F}$ be a family of topologies on X.

- For any $\tau \in \mathcal{F}$, we will write int_τ, resp. Cl_τ, to refer to the *interior* operator, resp. *closure*, operator *with respect to the topology* τ.
- For $A \subseteq X$ and $\tau \in \mathcal{F}$, τ^A will denote the sub-space topology on A (i.e., the family $\{U \cap A \mid U \in \tau\}$).
- For $\tau, \tau' \in \mathcal{F}$, $\tau \vee \tau'$ will be used to denote the topology generated by $\tau \cup \tau'$.
- For any $A \subseteq X$ and $\tau \in \mathcal{F}$, we will say that A is $\tau - dense$ in X iff A is dense in X *with respect to the topology* τ. I.e., A is $\tau - dense$ in X iff $Cl_\tau(A) = X$ iff every non-empty open set $O \in \tau$ is such that $O \cap A \neq \emptyset$.

Definition 4.2.1. (Language for the logic TKD)

Let $\mathcal{L}_K^\mathcal{D}$ consist of a countable set $Prop$ of propositional letters, the traditional Boolean operators $\neg$, $\wedge$, and the unary modal operators K_1, K_2, and D. The grammar of our language is given by

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

for $p \in Prop$. The abbreviations for the connectives $\vee$, $\rightarrow$, $\leftrightarrow$ are the ones inherited from propositional logic. The possibility operators $\langle K_1\rangle$, $\langle K_2\rangle$, and $\langle D\rangle$ are defined as $\neg K_1\neg$, $\neg K_2\neg$, and $\neg D\neg$, respectively.

Provided with the precise definition of our formal language, we introduce our models. It is very important to stress the fact that *these models are topo-evidential in nature.* For matters of clarity in notation and presentation, we will *not* use the epithet 'topo-e-model' in what follows. We will refer to our models simply as 'topological models,' However, we emphasize that their underlying motivation is *evidential*, and their place lies within the topo-evidentialist framework introduced in Chapter 3.

Definition 4.2.2. (Two-agent topological models)

A tuple $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ will be called a *two-agent topological model* for the formulas of $\mathcal{L}_K^\mathcal{D}$ if and only if the following conditions are met:

1. X is a set of possible worlds. Π_1, Π_2 are two agents' partitions on X. τ_1, τ_2, τ are topologies on X.
2. $\tau_1, \tau_2 \subseteq \tau$.
3. $\Pi_1 \subseteq \tau_1$, $\Pi_2 \subseteq \tau_2$.

[1] For a more detailed exposition of the concept, see Chapter 6.

4. $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

For a given $x \in X$ and $i \in \{1,2\}$, the notation $x[i]$ will be used to refer to the class of x in the corresponding partition Π_i. The topological semantics for the formulas of $\mathcal{L}_K^{\mathcal{D}}$ is defined recursively by:

$$\begin{aligned}
\|p\| &= \mathcal{V}(p)\\
\|\neg\phi\| &= X \backslash \|\phi\|\\
\|\phi \wedge \psi\| &= \|\phi\| \cap \|\psi\|\\
\|K_1\phi\| &= \{x \in X \mid \exists U \in \tau_1^{x[1]} \text{ s.t. } x \in U \subseteq \|\phi\| \wedge \ Cl_{\tau^{x[1]}}(U) = x[1]\}\\
\|K_2\phi\| &= \{x \in X \mid \exists U \in \tau_2^{x[2]} \text{ s.t. } x \in U \subseteq \|\phi\| \wedge \ Cl_{\tau^{x[2]}}(U) = x[2]\}\\
\|D\phi\| &= \{x \in X \mid \exists U \in \tau_1^{x[1]}, V \in \tau_2^{x[2]} \text{ s.t. } x \in (U \cap V) \subseteq \|\phi\| \wedge Cl_{\tau^{x[1]\cap x[2]}}(U \cap V) = x[1] \cap x[2]\}.
\end{aligned}$$

Truth and *Validity* are defined the same way as for standard Kripke semantics.

As seen, we have defined the semantics for our formulas through their *extension*. In this sense, we will say that a topological model *satisfies* a given formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ at world $x \in X$ iff $x \in \|\phi\|$. Such satisfaction will be marked as it is customary: $\mathcal{M}, x \vDash \phi$. Thus, we have that

$$\mathcal{M}, x \vDash \phi \text{ iff } x \in \|\phi\|.$$

Our definition might benefit from some discussion. As mentioned in the first section of this chapter, we mean to capture a notion of learnable evidence confined to information cells. For a given $i \in \{1,2\}$ and $x \in X$, $x[i]$ can be seen as Agent i's inherent epistemic *reach* at world x. The worlds lying within $x[i]$ are *conceivable* by i at x (but not necessarily *accessible*). Those worlds outside of $x[i]$ could be taken as virtually *impossible* for Agent i at world x. As for the evidence, $\tau_i^{x[i]}$ could be best interpreted as the *current* evidence that is available to Agent i at world x. Similarly, $\tau^{x[i]}$ should be taken to represent *all* the learnable evidence of Agent i at world x. Notice that *current* evidence is also *learnable* evidence. Hence the condition $\tau_i \subseteq \tau$. As it must, such inclusion holds when focusing on the agents' information cells at a given world x, i.e., when relativizing the topologies τ_i, τ to $x[i]$. In principle, the available evidence should be limited to the domain of *conceivability*, and this is why we have stated all of our definitions in terms of sub-space topologies. Another important constraint regarding the nature of cells is that they are themselves evidence. The whole margin of the agents' epistemic reach should constitute available evidence. Hence the condition $\Pi_i \subseteq \tau_i$. As it turns out, this condition is rather important on the technical side, for it will ensure that our operators are well-behaved in terms of the standard axiomatization for logics of distributed knowledge.

As for the definitions of individual knowledge, notice that they are very much in keep with the topo-e-semantics for knowledge in the single-agent case, but relativized to the information cells. On the conceptual side, we can say that it is a strong kind of knowledge. At a given world x, Agent i will know ϕ iff a) (s)he has conclusive evidence supporting ϕ (there is a τ_i-neighborhood U of x included in $\|\phi\|$), and b) *any* learnable evidence which is available to i at x will be consistent with i's evidential argument for ϕ (the aforementioned U is $\tau^{x[i]}$-dense). We stress the fact that, in order for something to constitute knowledge, the according argument should be consistent not only with all *current* evidence, but also with all *potentially* learnable evidence (e.g., learned in the future).

In the case of distributed knowledge, we intend to model two things at the same time. One is the hypothetical sharing of the group's full *information* at a given world, thus allowing its members to rule out each other's impossible worlds.[2] The other is the pooling of the group's current evidence in order to justify new—or implicit, if one pleases—beliefs (i.e., beliefs that the group would attain after sharing).

So first of all, if the two agents were able to share all of their *information* (based on learnable evidence) at a given $x \in X$, the resulting group knowledge should in principle be bounded to the intersection of the respective information cells (at x). Hence the definition of $\|D\phi\|$ in terms of topologies restricted to $x[1] \cap x[2]$. Informally, at a given $x \in X$, the unrealistic third agent to which the other two have communicated *all* of their information (both current and potential), should rule out both 1 and 2's classes of inaccessible worlds at x.

[2]In this particular context, *information* amounts to all learnable evidence within a relevant domain of conceivability.

Secondly, the distributed knowledge should be rendered from a refinement of justifications: if communication was possible, then the agents would be able to justify beliefs by pooled (current) evidence. Hence the use of $\tau_1 \vee \tau_2$ (restricted to $x[1] \cap x[2]$) as the relevant conclusive evidence for the group. Analogously to the cases for individual knowledge, the group will be said to have distributed knowledge of ϕ iff a) there is conclusive pooled evidence supporting ϕ ($U \in \tau_1^{x[1]}$ and $V \in \tau_2^{x[2]}$ are neighborhoods of x such that $U \cap V$ is included in $\|\phi\|$), and b) *any* learnable evidence within the pertinent $x[1] \cap x[2]$ will be consistent with the group's (implicit) justification for ϕ ($U \cap V$ is $\tau^{x[1] \cap x[2]}$-dense).

We can already draw some conclusions about our choice of semantics in terms of the interior of propositions on the state space. The following observation is very important, for many of our proofs will revolve around it.

Observation 4.

- $\|K_1\phi\| = \{x \in X \mid x \in int_{\tau_1^{x[1]}}\|\phi\| \wedge Cl_{\tau^{x[1]}}(int_{\tau_1^{x[1]}}\|\phi\|) = x[1]\}$
- $\|K_2\phi\| = \{x \in X \mid x \in int_{\tau_2^{x[2]}}\|\phi\| \wedge Cl_{\tau^{x[2]}}(int_{\tau_2^{x[2]}}\|\phi\|) = x[2]\}$
- $x \in \ \|D\phi\| \Rightarrow \ \begin{cases} x \in int_{(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}}\|\phi\| \ \textit{and} \\ Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1\vee\tau_2)^{x[1]\cap x[2]}}\|\phi\|) = x[1] \cap x[2]. \end{cases}$

The first two items are straightforward. The third item might need clarification. Assume that $x \in \|D\phi\|$. We know that there exist $U \in \tau_1^{x[1]}, V \in \tau_2^{x[2]}$ such that $x \in (U \cap V) \subseteq \|\phi\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(U \cap V) = x[1] \cap x[2]$. Notice that, since the i-cells are open in τ_i, we actually have that $U \in \tau_1$ and $V \in \tau_2$. But then $(U \cap V) \cap (x[1] \cap x[2])$ is an open set of $(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}$ such that $x \in (U \cap V) \cap (x[1] \cap x[2]) \subseteq \|\phi\|$, which means that $x \in int_{(\tau_1\vee\tau_2)^{x[1]\cap x[2]}}\|\phi\|$. Moreover, we have that $(U \cap V) \cap (x[1] \cap x[2]) \subseteq int_{(\tau_1\vee\tau_2)^{x[1]\cap x[2]}}\|\phi\|$, so that if $Cl_{\tau^{x[1]\cap x[2]}}(U \cap V) = x[1] \cap x[2]$, it is clear that $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1\vee\tau_2)^{x[1]\cap x[2]}}\|\phi\|) = x[1] \cap x[2]$ as well.

As for the technical layout, we observe that, in two-agent topological models, the assumptions $\Pi_1 \subseteq \tau_1$ and $\Pi_2 \subseteq \tau_2$ imply that the topological spaces we are looking at are *disconnected* for the three relevant topologies:

Remark 4.2.1. Let $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ be a two-agent topological model. If the partitions Π_1, Π_2 are non-trivial, then, for a fixed $i \in \{1, 2\}$ and every $x \in X$, $x[i]$ and $X \backslash x[i]$ are disjoint open sets of $\tau_i (\subseteq \tau)$ such that $x[i] \cup X \backslash x[i] = X$. Therefore, if the partitions are non-trivial, X is disconnected for τ_1, τ_2, and τ. Observe, then, that for every $x \in X$, $x[i]$ is clopen in τ_i. Similarly, for every $x \in X$, $x[1] \cap x[2]$ is clopen in τ.

To illustrate our setting, we introduce a comprehensive example.

Example 4.2.3. Let $X = \{x, y, z, w, v, u\}$ be a set. We take

- $\Pi_1 := \{\{x, y, x\}, \{w, v, u\}\}$,
  $\Pi_2 := \{\{x, w, v\}, \{y, z, u\}\}$.
- $\tau_1 := \{X, \emptyset, \{x, y\}, \{u\}\} \vee \Pi_1$,
  $\tau_2 := \{X, \emptyset, \{x, w\}, \{y, z\}\} \vee \Pi_2$,
  $\tau := \{\{z\}, \{w, v\}\} \vee (\tau_1 \vee \tau_2)$.
- $\mathcal{V} : \{p, q\} \to \mathcal{P}(X)$ given by $\mathcal{V}(p) = \{x, y\}$, $\mathcal{V}(q) = \{x, w, u\}$.

Let $\mathcal{M} := (X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$. With the aid of the diagram below, on can check, first and foremost, that $\mathcal{M}$ is a two-agent topological model:
-Π_1, Π_2 were constructed so that they are partitions on X;
-τ_1, τ_2, τ were constructed so that they are topologies on X;
-the pertinent inclusions $\tau_1 \subseteq \tau, \tau_2 \subseteq \tau$, $\Pi_1 \subseteq \tau_1$, $\Pi_2 \subseteq \tau_2$ all hold by definition.

Agent 1's partition of the space is represented by the horizontal blue line. Agent 2's partition is represented by the vertical red line. Accordingly, current-evidence open sets for Agent 1 are outlined in blue, and current-evidence open sets for Agent 2 are outlined in red. Non-current learnable evidence is outlined in purple. We point out that *all* blue, red, and purple open sets are part of the learnable evidence.

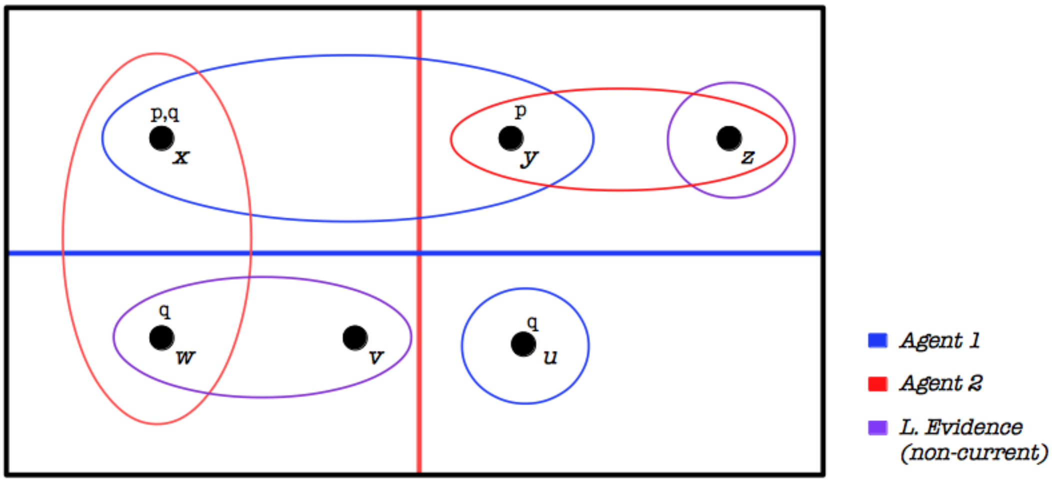

Example 4.2.3

So we have that, for example,

a) $\mathcal{M}, x \nvDash K_1 p$: though the open set $\{x, y\} \in \tau_1^{x[1]}$ is such that $\{x, y\} \subseteq \|p\|$, notice that $\{x, y\}$ does not intersect the (non-current-evidence) open set $\{z\} \in \tau^{x[1]}$.

However, we have that $\mathcal{M}, x \vDash Dp$: the set $\{x\}$ is open in $(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}$, it supports p, and it has non-empty intersection with all the other non-empty open sets of $\tau^{x[1] \cap x[2]}$, namely $x[1] \cap x[2]$ and $\{x\}$. This gives us an example of distributed knowledge that does not constitute individual knowledge (at a given world). Moreover, neither agent knows p at x. $\mathcal{M}, x \nvDash K_2 p$ because there is no $\tau_2^{x[2]}$-neighborhood of x included in $\|p\|$. We point out that $\mathcal{M}, x \vDash Dp$ comes from pooling together evidence.

Similarly, $\mathcal{M}, x \nvDash K_1 q$ because there is no $\tau_1^{x[1]}$-neighborhood of x included in $\|q\|$. On the other side, $\mathcal{M}, x \vDash Dq$, as witnessed by the open $\{x\}$ in $(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}$.

b) $\mathcal{M}, w \vDash K_2 q$: the set $\{x, w\}$ is open in $\tau_2^{w[2]}$, it supports q, and it has non-empty intersection with all the other non-empty open sets of $\tau^{w[2]}$, namely $w[2], w[1] \cap w[2], x[1] \cap x[2], \{x\}, \{w, v\}$, and $\{w\}$. In this line, notice that $\{w, v\}$ and $\{w\}$ are open sets of $\tau^{w[2]}$ that are not current evidence for Agent 2 at w. Informally, Agent 2 has not learned such pieces of evidence.

As one would hope for, $\mathcal{M}, w \vDash Dq$ as well: the set $\{w\}$ is open in $(\tau_1 \vee \tau_2)^{w[1] \cap w[2]}$, it supports q, and it has non-empty intersection with all the other non-empty open sets of $\tau^{w[1] \cap w[2]}$, namely $w[1] \cap w[2], \{w, v\}$, and $\{w\}$.

c) Though $\{y, z\} \subseteq \|\neg q\|$, we have that $\mathcal{M}, y \nvDash K_2 \neg q$ and $\mathcal{M}, z \nvDash K_2 \neg q$, for $\{y, z\} \cap \{u\} = \emptyset$ and $\{u\} \in \tau_1^{y[2]} \subseteq \tau^{y[2]} = \tau^{z[2]}$. On the other hand, it should be straightforward to realize that $\mathcal{M}, y \vDash D\neg q$ and that $\mathcal{M}, z \vDash D\neg q$. In both cases, the witnessing set is $\{y, z\}$. Thus, this implicit knowledge comes from ruling out impossible worlds. Not from pooling together evidence.

d) $\mathcal{M}, v \vDash K_1 \neg p$. The open set $v[1] \in \tau_1^{v[1]}$ supports $\neg p$, and evidently intersects all the learnable open sets of $\tau^{v[1]}$. Likewise $\mathcal{M}, v \vDash D\neg p$, and the witnessing set is $v[1] \cap v[2]$.

e) $\mathcal{M}, u \nvDash K_1 q$, but $\mathcal{M}, u \vDash Dq$. $\mathcal{M}, u \vDash K_1 \neg p$ (regard that the witnessing open set is $u[1]$).

Example 4.2.4. Let ρ stand for the euclidean topology on $\mathbb{R}$. Consider the cartesian product $X := \mathbb{R} \times \{1, 2, 3\}$. We take

- $\Pi_1 := \{\mathbb{R} \times \{1, 2\}, \mathbb{R} \times \{3\}\}$,

  $\Pi_2 := \{\mathbb{R} \times \{1\}, \mathbb{R} \times \{2\}, \mathbb{R} \times \{3\}\}$.

- $\tau_1 := \{\emptyset\} \cup [\{(\mathbb{R} \times \{1\})\backslash F \mid F \text{ is finite}\} \vee \{(\mathbb{R} \times \{2\})\backslash F \mid F \text{ is finite}\} \vee \{(\mathbb{R} \times \{3\})\backslash F \mid F \text{ is finite}\}]$,
  $\tau_2 := \{\emptyset\} \cup [\{(\mathbb{R} \times \{1\})\backslash C \mid C \text{ is count.}\} \vee \{(\mathbb{R} \times \{2\})\backslash C \mid C \text{ is count.}\} \vee \{(\mathbb{R} \times \{3\})\backslash C \mid C \text{ is count.}\}]$,
  $\tau := \{U \times \{1\} \mid U \in \rho\} \vee \{U \times \{2\} \mid U \in \rho\} \vee \{U \times \{3\} \mid U \in \rho\}$.
- $\mathcal{V} : Prop \to \mathcal{P}(X)$ any valuation function.

There are several comments to be made, but first, let us introduce some visual aid. A picture of the state space X, with the given partitions, would look like this:

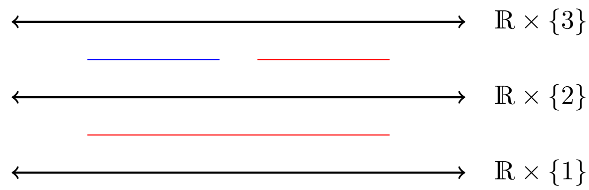

Example 4.2.4

Once again, we have represented Agent 1's partition of X with a blue horizontal line, and Agent 2's partition with red horizontal lines. This is an interesting, albeit simple, example. First of all, notice that Π_1, Π_2 are partitions straight from their construction (observe as well that Π_2 is a refinement of Π_1). τ_1 can be seen as the join of three disjoint *cofinite* topologies on the components $\mathbb{R} \times \{1\}$, $\mathbb{R} \times \{2\}$, and $\mathbb{R} \times \{3\}$, respectively. Similarly, τ_2 is constructed as the join of three *cocountable* topologies on the pertinent components, and τ is constructed as the join of three euclidean topologies on the pertinent components. It is well-known that the *cofinite* topology on $\mathbb{R}$ is included in the *cocountable* topology on $\mathbb{R}$, and that the latter is in turn included in the euclidean topology. Because of this, we get that $\tau_1 \subseteq \tau_2 \subseteq \tau$ (see [54] (Chapter 2)). As for the relation between the agents' topologies and their respective partitions, it is easy to see that $\tau_1 \subseteq \Pi_1$ and $\tau_2 \subseteq \Pi_2$ (their construction guarantees it). Therefore, we are fully entitled the call the tuple $\mathcal{M} := (X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ a two-agent topological model. Moreover, it is not a trivial example, since the partitions and topologies are all different from each other.

Now, we observe that in this example, knowledge will amount to interior: for every $x \in X$ and a fixed $i \in \{1, 2\}$, every non-empty open set in $\tau_i^{x[i]}$ is $\tau^{x[i]}$-dense. The reason is that the *cofinite* and *cocountable* topologies are dense in the euclidean topology, meaning that every non-empty open set of the euclidean topology intersects *any* open set of either the *cofinite* or *cocountable* topologies. This is not hard to realize, and again we refer the reader to [54] (Chapter 3) for details. We mention, then, that Observation 4 will imply that having individual (or distributed) knowledge of ϕ at a given $x \in X$ is the same as lying in the interior of ϕ in the corresponding relativized topology. We will show this explicitly in Chapter 5, because our example model is a member of a class of models that will be extremely important in our completeness proof, and which we call *dense* models (see Chapter 5, Definition 5.2.1).

Conceptually, it should be evident that we have modeled a very particular situation, in which Agent 2 can be seen as being both wiser and better-informed than Agent 1. First of all, Agent 2's information-partition is more refined than 1's. Secondly, no matter in which information cell we end up, the amount of evidence available for Agent 2 exceeds that of Agent 1 at the intersection of the cells, which is the only place where such a comparison is even significant. However, as implied by this discussion, our topological models allow for peculiar epistemic situations. For example, the evidential topology of a given agent can be more refined than that of her/his peer at a given world, but within a much broader domain of conceivability. The opposite situation is also possible, in which someone is endowed with a highly refined partition, but with not much learnable evidence within some information cell. Though our example presents a plausible scenario, we encourage the reader to entertain different layouts.

We have dealt with the syntax and the semantics of the logic TKD. In the next section, we proceed to axiomatize it.

4.3 The proof system $\mathbf{\Lambda}_{TKD}$

Let $\mathbf{\Lambda}_{TKD}$ denote a proof system for the logic TKD, constructed as follows:

- *(Axioms)*
  - All classical tautologies from propositional logic.
  - The **S4** axioms for K_i ($i \in \{1,2\}$):

$$(T) \quad K_i p \to p$$

$$(K) \quad K_i(p \to q) \to (K_i p \to K_i q)$$

$$(4) \quad K_i p \to K_i K_i p$$

  - The **S4.2** axioms for D:

$$(T_D) \quad Dp \to p$$

$$(K_D) \quad D(p \to q) \to (Dp \to Dq)$$

$$(4_D) \quad Dp \to DDp$$

$$(.2_D) \quad \langle D\rangle Dp \to D\langle D\rangle p$$

  - The *Distributed Knowledge* axioms:

$$(DsK) \begin{cases} K_1 p \to Dp \\ K_2 p \to Dp \end{cases}$$

$$(WSD) \begin{cases} \langle K_1\rangle K_1 p \to K_1\langle D\rangle p \\ \langle K_2\rangle K_2 p \to K_2\langle D\rangle p \end{cases}$$

- *(Rules of inference)*
  - *Modus Ponens*
  - Substitution
  - *Necessitation* for K_1, K_2, and D.

From Chapter 1, we know that the **S4** axioms are the basic requirements for standard individual knowledge. Therefore, it should be clear that we intend for our models to validate them in the semantics given. Similarly, from Chapter 2 we know that standard representations of distributed knowledge must conform to the **S4** axioms for operator D. Additionally, they must satisfy what we have called the (DsK) axioms: a pair of rules meant to capture the intuition that if some agent knows ϕ, then ϕ is distributed knowledge within the whole group. These considerations imply that, if $\mathbf{\Lambda}_{TKD}$ is sound for $\mathcal{L}_K^{\mathcal{D}}$ with respect to our two-agent topological models, then we will at least have done a good job in keeping to tradition. We will end up with topological knowledge that is *factive*, *positively introspective*, not *negatively introspective*, and whose behavior regarding distribution within a group of agents is adequate.

However, as frequently stated throughout this thesis, our proposal of topological semantics does not come without a reason. We want our models to account for a notion of *belief*, whose relation to knowledge conveys the intuitions that Stalnaker posed in [46] (see Chapter 1, section 1.1). Let us not forget, then, that we want an **S4.2** axiomatization of knowledge. As it turns out to be the case, the (.2) axiom for our individual K_i operators can be derived in $\mathbf{\Lambda}_{TKD}$:[3]

Remark 4.3.1. The (.2) axiom for K_1 and K_2 can be derived from the other principles. We present a brief $\mathbf{\Lambda}_{TKD}$-proof of this result:

Fix $i \in \{1, 2\}$. We have the following:

(1) $\vdash_\Lambda \langle D \rangle p \to \langle K_i \rangle p$	Substitution and Contraposition on (DsK)
(2) $\vdash_\Lambda K_i \langle D \rangle p \to K_i \langle K_i \rangle p$	K_i-*Necessitation* on (1), substitution on (K)
(3) $\vdash_\Lambda \langle K_i \rangle K_i p \to K_i \langle D \rangle p$	(WSD)
(4) $\vdash_\Lambda \langle K_i \rangle K_i p \to K_i \langle K_i \rangle p$	Substitution on *transitivity of implication* (3), (2), and *Modus Ponens*.

The (WSD) axioms might seem mysterious at first glance, and we admit that in a way they are. But they are not gratuitous, and they are meant to encode an important density condition inherent to the particular choice of topological semantics we favored.[4] Though their clear purpose *can only be fully ascertained when coupled to our proof of completeness*, we take notice of the resemblance with the (.2) axioms for the individual knowledge operators. Recall that in Stalnaker's combined system of knowledge and belief (Chapter 1, Section 1.1), belief is syntactically equated with $\langle K \rangle K$. In turn, $\langle K \rangle K$ turns out to be a global modality in the topo-e-framework. In our setting, $\langle K_i \rangle K_i$ plays the role of an operator encoding a kind of robust belief, in the sense that it is necessary for knowledge, and that it will imply the density of the justifications for knowledge, with respect to the share-able evidence. Its colloquial reading—in a relational interpretation—is neither satisfactory nor utterly abnormal, for a way of phrasing it can be by saying that 'if knowledge is possible for someone, then the possibility of distributed knowledge in the group is certain for him/her.'

So much for the discussion of our axioms. Once in hold of a proof system, we present the main theorem of this thesis:

Theorem 4.3.1. *(Main Theorem) The proof system $\mathbf{\Lambda}_{TKD}$ is sound and weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of two-agent topological models.*

Chapter 5 will be dedicated in its entirety to proving this result.

[3]Notice that, nevertheless, axiom $(.2_D)$ is still independent from the rest. In fact, the derivation of the (.2) axioms in $\mathbf{\Lambda}_{TKD}$ requires $(.2_D)$. In this line, we have to say that the topological models—and semantics for the formulas—that we presented in Definition 4.2.2 were constructed so that (DsK) and the **S4.2** axioms for K_i and D were satisfied. Though this is not evident from the discussion so far, we will properly show that the axioms are valid in the next chapter, when we prove soundness.

[4]During the research stage of this work, we colloquially referred to the property as 'weak super-density,' Hence the use of the acronym (WSD).

Chapter 5

Soundness and completeness of $\mathbf{\Lambda}_{TKD}$

In this chapter, we prove Theorem 4.3.1: we show that $\mathbf{\Lambda}_{TKD}$ is sound and weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of two-agent topological models that we introduced in Chapter 4.

5.1 Soundness

Recall that a formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ is *true* in a two-agent topological model $\mathcal{M} = (X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ iff $\|\phi\|^{\mathcal{M}} = X$.

Just to keep things clear, we explicitly note that we will be using the following easy result in virtually every item of our proof of soundness.

Lemma 5.1.1. *Let $\mathcal{M} = (X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ be a two-agent topological model. For all ϕ, ψ of $\mathcal{L}_K^{\mathcal{D}}$, we have that $\|\phi \to \psi\|^{\mathcal{M}} = X$ iff for every $x \in X$, $x \in \|\phi\|^{\mathcal{M}}$ implies that $x \in \|\psi\|^{\mathcal{M}}$.*

Proof. For the left-to-right direction, take $x \in X$ such that $x \in \|\phi\|^{\mathcal{M}}$. This means that $\mathcal{M}, x \vDash \phi$. By assumption, we have that $\|\phi \to \psi\|^{\mathcal{M}} = X$, so that $\mathcal{M}, x \vDash \phi \to \psi$. Therefore, $\mathcal{M}, x \vDash \psi$ and hence $x \in \|\psi\|^{\mathcal{M}}$.

For the right-to-left direction, let $x \in X$ such $\mathcal{M}, x \vDash \phi$. This means that $x \in \|\phi\|^{\mathcal{M}}$, which by assumption implies that $x \in \|\psi\|^{\mathcal{M}}$ and thus that $\mathcal{M}, x \vDash \psi$. Therefore, we have that $\mathcal{M}, x \vDash \phi$ implies that $\mathcal{M}, x \vDash \psi$, which means that $\mathcal{M}, x \vDash \phi \to \psi$, and hence that $x \in \|\phi \to \psi\|^{\mathcal{M}}$. Since x was arbitrarily chosen, we have that $\|\phi \to \psi\|^{\mathcal{M}} = X$.

$\square$

Another convenient result, analogous to Lemma 3.2.3, is the following:

Lemma 5.1.2. *Let Y be a set, and let $\upsilon_1, \upsilon_2, \upsilon$ be three topologies on Y such that $\upsilon_1, \upsilon_2 \subseteq \upsilon$. For a fixed $i \in \{1, 2\}$, take $\hat{\upsilon}_i := \{W \in \upsilon_i \mid Cl_\upsilon(W) = Y\} \cup \{\emptyset\}$. I.e., $\hat{\upsilon}_i = \{W \in \upsilon_i \mid W \text{ is } \upsilon - dense\} \cup \{\emptyset\}$.*

Then for every $U \in \hat{\upsilon}_k$, $V \in \hat{\upsilon}_j$ (with $k, j \in \{1, 2\}$ not necessarily different), we have that $U \cap V$ is υ-dense.

Proof. Let $k, j \in \{1, 2\}$, and let $U \in \hat{\upsilon}_k$, $V \in \hat{\upsilon}_j$. By assumption, for any non-empty set $O \in \upsilon$, we have that $O \cap V \neq \emptyset$. But $O \cap V \in \upsilon$ (because $\upsilon_j \subseteq \upsilon$). Thus, $O \cap V$ is a non-empty open set of υ as well. Hence, $(U \cap V) \cap O = U \cap (O \cap V) \neq \emptyset$. $\square$

Finally, we remind the reader of an important observation—reviewed in the previous chapter—regarding the semantics of our formulas. We will be using it constantly in the proof of soundness:

Observation 5.

- $\|K_1\phi\| = \{x \in X \mid x \in int_{\tau_1^{x[1]}}\|\phi\| \wedge Cl_{\tau^{x[1]}}(int_{\tau_1^{x[1]}}\|\phi\|) = x[1]\}$
- $\|K_2\phi\| = \{x \in X \mid x \in int_{\tau_2^{x[2]}}\|\phi\| \wedge Cl_{\tau^{x[2]}}(int_{\tau_2^{x[2]}}\|\phi\|) = x[2]\}$

- $x \in \ \|D\phi\| \Rightarrow \begin{cases} x \in int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|\phi\| \text{ and} \\ Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|\phi\|) = x[1] \cap x[2]. \end{cases}$

Proposition 5.1.3. *The system $\mathbf{\Lambda}_{TKD}$ is sound for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of two-agent topological models. That is, for every ϕ of $\mathcal{L}_K^{\mathcal{D}}$, $\vdash_{\Lambda_{TKD}} \phi$ implies that ϕ is valid in the pertinent class.*

Proof. We proceed as it is usual for soundness proofs. First, we show that the given axioms are semantic validities, and then we show that *validity* is preserved under the rules of inference. It is clear that all classical tautologies are valid in our pertinent class. Similarly, it is well-known that *Modus Ponens* preserves validity.

For the other items, we will make heave use of the items gathered in Observation 5.

Let $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau)$ be a space underlying a topological model. Let $\mathcal{V} : Prop \to \mathcal{P}(X)$ be any valuation function.

- The **S4** axioms for K_i $(i \in \{1, 2\})$:

  Fix $(i \in \{1, 2\})$.

  – (T) Let $x \in X$ such that $x \in \|K_i p\|$. Since $int_{\tau_1^{x[1]}} \|p\| \subseteq \|p\|$, we have that $x \in \|p\|$ (first items of Observation 5).

  – (K) Let $x \in X$. Assume that $x \in \|K_i(p \to q)\|$ and that $x \in \|K_i p\|$. Then, on one hand, there exists $U \in \tau_i^{x[i]}$ such that $x \in U$, $U \subseteq \|(p \to q)\|$, and $Cl_{\tau^{x[i]}}(U) = x[i]$; on the other, there exists $V \in \tau_i^{x[i]}$ such that $x \in V$, $V \subseteq \|p\|$ and $Cl_{\tau^{x[i]}}(V) = x[i]$. We have that $U \cap V$ is an open set of $\tau_i^{x[i]}$ such that $x \in U \cap V$, $U \cap V \subseteq \|q\|$ and $Cl_{\tau^{x[i]}}(U \cap V) = x[i]$ (Lemma 5.1.2), which means that $x \in \|K_i q\|$.

  – (4) Let $x \in X$. Assume $x \in \|K_i p\|$. Then there exists $U_* \in \tau_i^{x[i]}$ such that $x \in U_* \subseteq \|p\|$ and $Cl_{\tau^{x[i]}}(U_*) = x[i]$. Observe that for every $y \in U$, we have—rather obviously—that $y \in \|K_i p\|$. Therefore, there exists an open set $U \in \tau_i^{x[i]}$ such that $x \in U$, $U \subseteq \|K_i p\|$, and $Cl_{\tau^{x[i]}}(U) = x[i]$, namely $U = U_*$. Therefore, $x \in \|K_i K_i p\|$.

- The **S4.2** axioms for D:

  – (T_D) Let $x \in X$ such that $x \in \|Dp\|$. Since $x \in int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|p\| \subseteq \|p\|$, we have that $x \in \|p\|$.

  – (K_D) Let $x \in X$. Assume that $x \in \|D(p \to q)\|$ and that $x \in \|Dp\|$. Then, on one hand, there exist $O \in \tau_1^{x[1]}$, $O' \in \tau_2^{x[2]}$ such that $x \in O \cap O' \subseteq \|(p \to q)\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(O \cap O') = x[1] \cap x[2]$; on the other, there exist $U \in \tau_1^{x[1]}, U' \in \tau_2^{x[2]}$ such that $x \in U \cap U' \subseteq \|p\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(U \cap U') = x[1] \cap x[2]$. But then it is easy to see that $O \cap U \in \tau_1^{x[1]}$ and $O' \cap U' \in \tau_2^{x[2]}$, with $x \in (O \cap U) \cap (O' \cap U') \subseteq \|q\|$, and $Cl_{\tau^{x[1]\cap x[2]}}((O \cap U) \cap (O' \cap U')) = x[1] \cap x[2]$ (Lemma 5.1.2). This means that $x \in \|Dq\|$.

  – (4_D) Let $x \in X$. Assume that $x \in \|Dp\|$. Then there exist $U \in \tau_1^{x[1]}$, $V \in \tau_2^{x[2]}$ such that $x \in (U \cap V) \subseteq \|p\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(U \cap V) = x[1] \cap x[2]$. Observe that for every $y \in (U \cap V)$, we have that $y \in \|Dp\|$. Therefore, since $(U \cap V) \subseteq \|Dp\|$, we get that $x \in \|DDp\|$.

  – $(.2_D)$ Let $x \in X$. Assume that $x \in \|\langle D \rangle Dp\|$.

    First, we observe that this assumption implies that $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|p\|) = x[1] \cap x[2]$. We show this by contraposition. Suppose that $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|p\|) \neq x[1] \cap x[2]$. Then for all $y \in x[1] \cap x[2]$, $y \in \|\neg Dp\|$ (recall the third item of Observation 5); thus, $x[1] \cap x[2] \subseteq \|\neg Dp\|$, so that the open sets $x[1] \in \tau_1^{x[1]}$, $x[2] \in \tau_2^{x[2]}$ are such that $x \in x[1] \cap x[2] \subseteq \|\neg Dp\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(x[1] \cap x[2]) = x[1] \cap x[2]$; therefore, $x \in \|D\neg Dp\| = X \backslash \|\langle D \rangle Dp\|$.

    Next, we show that the fact that $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|p\|) = x[1] \cap x[2]$ implies that $x \in \|D\langle D \rangle p\|$. Notice that if $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|p\|) = x[1] \cap x[2]$, then $Cl_{\tau^{x[1]\cap x[2]}} \|p\| = x[1] \cap x[2]$. This implies that $int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|\neg p\| \subseteq int_{\tau^{x[1]\cap x[2]}} \|\neg p\| = int_{\tau^{x[1]\cap x[2]}}(X \backslash \|p\|) = \emptyset$. Therefore, $Cl_{\tau^{x[1]\cap x[2]}}(int_{(\tau_1 \vee \tau_2)^{x[1] \cap x[2]}} \|\neg p\|) = \emptyset$, which by the third item of Observation 5 renders that for all $y \in x[1] \cap x[2]$, $y \in \|\neg D \neg p\|$. As before, we choose $x[1] \in \tau_1^{x[1]}$, $x[2] \in \tau_2^{x[2]}$ such that $x \in x[1] \cap x[2] \subseteq \|\neg D \neg p\|$ and $Cl_{\tau^{x[1]\cap x[2]}}(x[1] \cap x[2]) = x[1] \cap x[2]$; therefore, $x \in \|D \neg D \neg p\| = \|D\langle D \rangle p\|$.

- (DsK). Let $x \in X$. Assume, without loss of generality, that $x \in \|K_1 p\|$. Then there exists $U \in \tau_1^{x[1]}$ such that $x \in U \subseteq \|p\|$ and $Cl_{\tau^{x[1]}}(U) = x[1]$. But this means that, in particular, $x \in U \cap x[2] \subseteq \|p\|$. By taking $x[2] \in \tau_2^{x[2]}$ as the other set we are looking for, we get that the existence of such U guarantees that $x \in \|Dp\|$ as well, and the thing that remains to be shown to convince ourselves of this is that $Cl_{\tau^{x[1]\cap x[2]}}(U \cap x[2]) = x[1] \cap x[2]$. So take a non-empty $O \in \tau^{x[1]\cap x[2]}$. O is of the form $O' \cap (x[1] \cap x[2])$ for some $O' \in \tau$. Notice that $O' \cap (x[1] \cap x[2]) = (O' \cap x[2]) \cap x[1]$. But we have that $\Pi_2 \subseteq \tau$, so that $x[2] \in \tau$ and therefore $O' \cap x[2]$ is a non-empty open set of τ also. This means that $(O' \cap x[2]) \cap x[1]$ is a non-empty open set of $\tau^{x[1]}$. Thus, by assumption, $((O' \cap x[2]) \cap x[1]) \cap U \neq \emptyset$, which implies that $O \cap (U \cap x[2]) \neq \emptyset$ (recall that $U \subseteq x[1]$).

- (WSD). Let $x \in X$. Assume, without loss of generality, that $x \in \|\langle K_1\rangle K_1 p\|$.

  Somewhat similarly as in the proof for the $(.2_D)$ axiom, we observe that this first assumption implies that $Cl_{\tau^{x[1]}}(int_{\tau_1^{x[1]}}\|p\|) = x[1]$. We show this by contraposition. Suppose that $Cl_{\tau^{x[1]}}(int_{\tau_1^{x[1]}}\|p\|) \neq x[1]$; then for all $y \in x[1]$, $y \in \|\neg K_1 p\|$ (recall the first item of Observation 5); thus, $x[1] \subseteq \|\neg K_1 p\|$, so that the open set $x[1] \in \tau_1^{x[1]}$ is such that $x \in x[1] \subseteq \|\neg K_1 p\|$ and $Cl_{\tau^{x[1]}}(x[1]) = x[1]$; therefore, $x \in \|K_1 \neg K_1 p\| = X \backslash \|\langle K_1\rangle K_1 p\|$.

  For the rest of the proof, let us use a neater terminology. By the discussion above, we know that there exists $U \in \tau_1^{x[1]}$ such that $U \subseteq \|p\|$ and $Cl_{\tau^{x[1]}}(U) = x[1]$ (where $int_{\tau_1^{x[1]}}\|p\|$ plays the role of the aforementioned "U").

  We will first show that for every $z \in x[1]$, $Cl_{\tau^{x[1]\cap z[2]}}(U \cap z[2]) = x[1] \cap z[2]$.

  To prove this, first notice that $U \cap z[2] \neq \emptyset$, since $z[2] \in \tau$ and hence $z[2] \cap x[1]$ is a non-empty open set of $\tau^{x[1]}$ (it is non-empty because $z \in z[2] \cap x[1]$); now, since U is taken to be $\tau^{x[1]}$-dense, then $U \cap (z[2] \cap x[1]) \neq \emptyset$, which implies that $U \cap z[2] \neq \emptyset$. Once that has been settled, take a non-empty $O \in \tau^{x[1]\cap z[2]}$. O is of the form $O' \cap (x[1] \cap z[2])$ for some $O' \in \tau$. Notice that $O' \cap (x[1] \cap z[2]) = (O' \cap z[2]) \cap x[1]$. But we have that $z[2] \in \tau$, so $O' \cap z[2]$ is a non-empty open set of τ also. But this means that $(O' \cap z[2]) \cap x[1]$ is a non-empty open set of $\tau^{x[1]}$. Since U is taken to be $\tau^{x[1]}$-dense, we have that $((O' \cap z[2]) \cap x[1]) \cap U \neq \emptyset$, which implies $O \cap (U \cap z[2]) \neq \emptyset$. Therefore, we have that seen that for every $z \in x[1]$, $Cl_{\tau^{x[1]\cap z[2]}}(U \cap z[2]) = x[1] \cap z[2]$.

  Secondly, regard that U is such that $U \subseteq (\|p\| \cap x[1])$. Therefore, for every $z \in x[1]$, the preliminary result from the above paragraph renders that $Cl_{\tau^{x[1]\cap z[2]}}((\|p\| \cap x[1]) \cap z[2]) = x[1] \cap z[2]$. This means that for every $z \in x[1]$, $int_{\tau^{x[1]\cap z[2]}}((\|\neg p\| \cap x[1]) \cap z[2]) = \emptyset$, which implies that for every $z \in x[1]$, $int_{(\tau_1 \vee \tau_2)^{x[1]\cap z[2]}}\|\neg\phi\| = \emptyset$. By the third item of Observation 5, then, we have that for every $z \in x[1]$, $z \in \|\neg D \neg p\|$. Thus, $x[1] \subseteq \|\neg D \neg p\|$.

  Observe, then, that $x[1] \in \tau_1^{x[1]}$ is such that $x \in x[1] \subseteq \|\neg D \neg p\|$ and $Cl_{\tau^{x[i]}}(x[1]) = x[1]$. With this, we get that $x \in \|K_1 \neg D \neg p\| = \|K_1 \langle D\rangle p\|$.

Finally, we show that the *Necessitation* rule for K_i preserves validity. Assume that for every topological space $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau)$ and every valuation function $\mathcal{V} : Prop \to \mathcal{P}(X)$, $\|\phi\| = X$. This means that for every $x \in X$, we have that $x \in \|\phi\|$. Then, and somewhat trivially, there exists an open set $U \in \tau_i^{x[i]}$ such that $x \in U$, $U \subseteq \|\phi\|$, and $Cl_{\tau^{x[i]}}(U) = x[i]$, namely $U = x[i]$. Therefore, $x \in \|K_i\phi\|$, and $\|K_i\phi\| = X$. Since the space is arbitrary, we get the desired property. Notice that, since $K_i\phi \to D\phi$ is valid (coming from the satisfaction of (DsK) axiom), we have that the *Necessitation* rule for D also preserves validity.

□

5.2 Completeness

5.2.1 Equivalence between relational and topological models

Our intention is to show that $\mathbf{\Lambda}_{TKD}$ is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of all topological models. We will do so by proving completeness with respect to a very particular class of them, which we introduce as *special* topological models.

Definition 5.2.1. We will call a topological model ***dense*** if and only if the following conditions are met:

($denk$) For every $x \in X$ and a fixed $i \in \{1, 2\}$, every non-empty open set in $\tau_i^{x[i]}$ is $\tau^{x[i]}$-dense.

($dend$) For every $x \in X$, every non-empty open set in $(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}$ is $\tau^{x[1]\cap x[2]}$-dense.

Observation 6. Notice that, in the particular case of *dense* topological models, we have that

- $\|K_1\phi\| = \{x \in X \mid x \in int_{\tau_1^{x[1]}}\}$
- $\|K_2\phi\| = \{x \in X \mid x \in int_{\tau_2^{x[2]}}\}$
- $\|D\phi\| = \{x \in X \mid x \in int_{(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}} \|\phi\|\}$.

The third item might be difficult to elucidate.

First, regard that the basic open sets of $(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}$ are of the form $(U' \cap V') \cap x[1] \cap x[2]$, with $U' \in \tau_1$ and $V' \in \tau_2$. Therefore, assuming that $x \in int_{(\tau_1 \vee \tau_2)^{x[1]\cap x[2]}} \|\phi\|$ yields that there exist $U \in \tau_1, V \in \tau_2$ such that $x \in (U \cap V) \cap (x[1] \cap x[2]) \subseteq \|\phi\|$. By ($dend$), $(U \cap V) \cap (x[1] \cap x[2])$ is assumed to be $\tau^{x[1]\cap x[2]}$-dense. But notice that $(U \cap x[1]) \in \tau_1^{x[1]}$ and $(V \cap x[2]) \in \tau_2^{x[2]}$. Moreover, $(U \cap x[1]) \cap (V \cap x[2]) = (U \cap V) \cap (x[1] \cap x[2])$. Therefore, we have found our witnessing open sets such that $x \in (U \cap x[1]) \cap (V \cap x[2]) \subseteq \|\phi\|$ and $Cl_{\tau^{x[1]\cap x[2]}}((U \cap x[1]) \cap (V \cap x[2])) = x[1] \cap x[2]$, thus rendering that $x \in \|D\phi\|$.

Definition 5.2.2. We will call a topological model ***special*** if and only if it is dense and, additionally, we have that:

S. $\tau = \tau_1 \vee \tau_2$.

In order to prove completeness with respect to this class of *special* topological models, we will go even further in our focus and show completeness with respect to the *special* topological models in which the relevant topologies are all *Alexandrov* topologies. The reason is that such models have clear relational counterparts according to our own discussion in Chapter 1. We briefly recover the appropriate definitions and main observations here, to keep closely at hand:

Definition 5.2.3. A topological space (X, τ) is said to be an *Alexandrov* space if and only if the intersection of any collection of open sets of X is an open set as well.

Notice that a space is *Alexandrov* iff every point $x \in X$ has a $\subseteq$-smallest open set including it, namely the intersection of all the open sets around x.

Definition 5.2.4. For a topological space (X, τ), the *specialization preorder* on X is a relation $\leq_\tau$ on X defined by $x \leq_\tau y$ iff $O_x \subseteq O_y$, where we have taken $O_z := \{U \in \tau \mid z \in U\}$ for any $z \in X$. Another way of putting it would be to say that $x \leq_\tau y$ iff $y \in Cl_\tau(\{x\})$.

It is easy to check that the order defined above is in fact reflexive and transitive.

Definition 5.2.5. For a given **S4** frame $(X, \leq)$, we call a set $A \subseteq X$ *upward-closed* iff for every $x \in A$, if $x \leq y$ for some $y \in X$, then $y \in A$ as well. On a related note, it is customary to mark $x \uparrow_\leq$ to refer to the set $\{y \in X \mid x \leq y\}$.

Observation 7. For a given **S4** frame $(X, \leq)$, the set of all $\leq$-upward-closed sets forms an *Alexandrov* topology on X, which we denote $\tau_\leq$. For $x \in X$, the $\subseteq$-smallest open set including x is precisely $x \uparrow_\leq$. This implies that $\{x \uparrow_\leq \mid x \in X\}$ is a basis for the topology $\tau_\leq$.

Observation 8.

- For a given **S4** frame $(X, \leq)$, $\leq = \leq_{\tau_\leq}$. To ascertain this, regard that the inclusion $\leq \subseteq \leq_{\tau_\leq}$ comes from the fact that $x \uparrow_\leq$ is the least open set including x in the topology $\tau_\leq$, so that $y \in x \uparrow_\leq$ implies that $y \in Cl_{\tau_\leq}(\{x\})$, and thus that $x \leq_{\tau_\leq} y$. The other inclusion is straightforward.

- For a given *Alexandrov* space (X,τ), $\tau = \tau_{\leq_\tau}$. That $\tau \subseteq \tau_{\leq_\tau}$ can be seen from the following argument: for any $U \in \tau$ and $x \in U$, we have that if $x \leq_\tau y$, then $y \in Cl_\tau(\{x\})$, and thus $y \in U$ as well. But this means that $x \uparrow_{\leq_\tau} \subseteq U$. The other inclusion comes from noticing that, since τ makes the space *Alexandrov*, then for every $x \in X$, $\bigcap\{U \in \tau \mid x \in U\} \subseteq x \uparrow_{\leq_\tau}$.

As shown in Chapter 1, these definitions allow us to draw a bridge between topological spaces and Kripke frames. In this sense, any topological model will have an associated relational structure that satisfies certain properties. While this bridge or association can be built for every topological model, we will concentrate our attention on *special* topological models, as said before.

Any *special* topological model in which the relevant topologies are *Alexandrov* will be henceforward called *special Alexandrov* model.

Let us elaborate a little bit on the matter at hand. For every *special Alexandrov* model, of the form $\mathcal{M}^{top} = (X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$, we can conceive the structure $\mathcal{M}^{rel} = (X, \Pi_1, \Pi_2, \leq_{\tau_1}, , \leq_{\tau_2}, \leq_\tau, \mathcal{V})$. All the topological constraints that we have demanded from $\mathcal{M}^{top}$ can be written down as relational conditions demanded from $\mathcal{M}^{rel}$. Moreover, we can restate the topological semantics for our formulas in purely relational terms. Such a reformulation will prove extremely helpful.

Definition 5.2.6. (Relational version of *special Alexandrov* models)

A tuple $\mathcal{M} = (X, \sim_1, \sim_2, \leq_1, \leq_2, \leq, \mathcal{V})$ is a *relational model* for TKD if and only if

i) X is a set. $\sim_1, \sim_2$ are two equivalence relations on X. $\leq_1, \leq_2, \leq$ are three preorders on X.

ii) $\leq_1 \subseteq \sim_1$, $\leq_2 \subseteq \sim_2$.

iii) For a fixed $i \in \{1,2\}$, for every $w, v \in X$ such that $w \sim_i v$, there exists $u \in X$ such that $w \leq_i u$ and $v \leq u$.

iv) For every $w, v \in X$ such that $w \sim_1 \cap \sim_2 v$, there exists $u \in X$ such that $w \leq u$ and $v \leq u$.

v) $\leq = \leq_1 \cap \leq_2$ (*special* condition S.)

vi) $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

The (relational) semantics for the formulas of $\mathcal{L}_K^{\mathcal{D}}$ is defined recursively by the following rules of model *satisfaction*:

$$
\begin{array}{lll}
\mathcal{M}, w \Vdash p & \text{iff} & w \in \mathcal{V}(p) \\
\mathcal{M}, w \Vdash \neg\phi & \text{iff} & \mathcal{M}, w \nVdash \phi \\
\mathcal{M}, w \Vdash \phi \wedge \psi & \text{iff} & \mathcal{M}, w \Vdash \phi \text{ and } \mathcal{M}, w \Vdash \psi \\
\mathcal{M}, w \Vdash K_1\phi & \text{iff} & \forall v \in X, \ \ [w \leq_1 v \ \Rightarrow \ \mathcal{M}, v \Vdash \phi] \\
\mathcal{M}, w \Vdash K_2\phi & \text{iff} & \forall v \in X, \ \ [w \leq_2 v \ \Rightarrow \ \mathcal{M}, v \Vdash \phi] \\
\mathcal{M}, w \Vdash D\phi & \text{iff} & \forall v \in X, \ \ [w \leq v \ \Rightarrow \ \mathcal{M}, v \Vdash \phi].
\end{array}
$$

Truth and *Validity* are defined the same way as for standard Kripke semantics.

We define the *extension* of our formulas the natural way: for a given formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$, $[\![\phi]\!] = \{w \in X \mid \mathcal{M}, w \Vdash \phi\}$. Just to avoid future complications, notice that this definition clearly entails the following equalities:

$$
\begin{aligned}
[\![\neg\phi]\!] &= X \setminus [\![\phi]\!] \\
[\![\phi \wedge \psi]\!] &= [\![\phi]\!] \cap [\![\psi]\!]
\end{aligned}
$$

As a terminological convention, we will use the following abbreviations:

- $[\leq_1]\phi = \{w \in X \mid \forall v \in X, \ \ [w \leq_1 v \ \Rightarrow \ \mathcal{M}, v \Vdash \phi]\}$.
- $[\leq_2]\phi = \{w \in X \mid \forall v \in X, \ \ [w \leq_2 v \ \Rightarrow \ \mathcal{M}, v \Vdash \phi]\}$.

- $[\leq]\phi = \{w \in X \mid \forall v \in X, \quad [w \leq v \ \Rightarrow\ \mathcal{M}, v \Vdash \phi]\}$

With this, we can rewrite our relational semantics in a much more friendly manner:

Observation 9.

- $w \in [\![K_1\phi]\!]$ iff $x \in [\leq_1]\phi$.
- $w \in [\![K_2\phi]\!]$ iff $x \in [\leq_2]\phi$.
- $w \in [\![D\phi]\!]$ iff $x \in [\leq]\phi$.

The following Proposition is very important, for it guarantees that the established link between these classes of models is actually meaningful. It is a multi-agent version of Proposition 1.2.6 from Chapter 1, the one that binds frames to topological structures in a truth-preserving manner.

Proposition 5.2.7. *There exists a bijective correspondence between the class of special Alexandrov models and the class of relational models. Such correspondence keeps the domain of the structures the same and preserves the truth of the formulas of $\mathcal{L}_K^{\mathcal{D}}$ at same states.*

Proof. Let $\mathcal{A}$ denote the class of *special Alexandrov* models, and $\mathcal{R}$ denote the class of relational models.

We define a map

$$(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V}) \overset{H}{\longmapsto} (X, \sim_1, \sim_2, \leq_{\tau_1}, \leq_{\tau_2}, \leq_\tau, \mathcal{V})$$

where

- $\sim_1, \sim_2$ are the equivalence relations induced by the partitions Π_1, Π_2, respectively,
- $\leq_{\tau_1}, \leq_{\tau_2}, \leq_\tau$ are the *specialization preorders* induced on X by τ_1, τ_2, τ, respectively.

Claim 1. $H[\mathcal{A}] \subseteq \mathcal{R}$.

Let $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ be a *special Alexandrov* model (i.e., *dense* and such that $\tau = \tau_1 \vee \tau_2$). We want to show that $(X, \sim_1, \sim_2, \leq_{\tau_1}, \leq_{\tau_2}, \leq_\tau, \mathcal{V})$ is a relational model.

i) By definition, we have that $\sim_1, \sim_2$ are equivalence relations. It is well-known that $\leq_{\tau_1}, \leq_{\tau_2}, \leq_\tau$ are reflexive and transitive relations on X iff τ_1, τ_2, τ are (*Alexandrov*) topologies on X (see Definition 5.2.4 and subsequent observations).

ii) For a fixed $i \in \{1, 2\}$, $\leq_{\tau_i} \subseteq \sim_i$ iff for every $x \in X$, $x \uparrow_{\leq_{\tau_i}} \subseteq x[i]$. But this happens iff for every $x \in X$, $x[i]$ is an open set of τ_i, which in turn happens iff $\tau_i \subseteq \Pi_i$.

iii) We want to show that if $w \sim_i v$, then there exists u such that $w \leq_{\tau_i} u$ and $v \leq_\tau u$.

By assumption, (*denk*) holds for $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$. Fix $i \in \{1, 2\}$. Our assumption means that for every $x \in X$, every open set in $\tau_i^{x[i]}$ is $\tau^{x[i]}$-dense. Let $w, v \in X$ such that $w \sim_i v$. Clearly, $w \uparrow_{\leq_{\tau_i}}$ is a non-empty open set of τ_i, and moreover, it is completely included in $w[i]$ (the cells are open). Therefore, $w \uparrow_{\leq_{\tau_i}} \in \tau_i^{w[i]}$. By a similar argument, $v \uparrow_{\leq_\tau}$ is a non-empty open set of $\tau^{v[i]} = \tau^{w[i]}$. Therefore, they have non-empty intersection. This renders the existence of u such that $w \leq_{\tau_i} u$ and $v \leq_\tau u$.

iv) We want to show that if $w \sim v$, then there exists u such that $w \leq_\tau u$ and $v \leq_\tau u$.

We have assumed that (*dend*) holds for $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ and that $\tau = \tau_1 \vee \tau_2$. Let $w, v \in X$ such that $w(\sim_1 \cap \sim_2)v$. $w \uparrow_{\leq_\tau}$ is actually $w \uparrow_{\leq_{\tau_1 \vee \tau_2}} = w \uparrow_{\leq_{\tau_1}} \cap w \uparrow_{\leq_{\tau_2}}$ (see the proof of item v) below). The latter is a non-empty open set of $(\tau_1 \vee \tau_2)^{w[1] \cap w[2]}$. Notice that $w(\sim_1 \cap \sim_2)v$ implies that $v[1] \cap v[2] = w[1] \cap w[2]$. Because of this equality, and by virtue of item ii) above, we get that $v \uparrow_{\leq_\tau} = v \uparrow_{\leq_{\tau_1}} \cap v \uparrow_{\leq_{\tau_2}}$ is a non-empty open set of $\tau^{w[1] \cap w[2]}$. Therefore, (*dend*) implies that $w \uparrow_{\leq_\tau}$ and $v \uparrow_{\leq_\tau}$ have non-empty intersection. This renders the existence of u such that $w \leq_\tau u$ and $v \leq_\tau u$.

v) (*S.*) Since we are dealing with *special* models, we have the underlying condition of $\tau = \tau_1 \vee \tau_2$. Therefore, $\leq_\tau = \leq_{\tau_1 \vee \tau_2}$. We show that $\leq_{\tau_1 \vee \tau_2}$ is exactly equal to $\leq_{\tau_1} \cap \leq_{\tau_2}$: the $\subseteq$ inclusion is pretty straightforward, since $y \in Cl_{\tau_1 \vee \tau_2}(\{x\})$ evidently implies that $y \in Cl_{\tau_1}(\{x\})$ and $y \in Cl_{\tau_2}(\{x\})$. For the $\supseteq$ inclusion, let $x(\leq_{\tau_1} \cap \leq_{\tau_2})y$. Take $U \in \tau_1 \vee \tau_2$ such that $x \in U$. There exist $U_1 \in \tau_1, U_2 \in \tau_2$ such that $x \in U_1 \cap U_2 \subseteq U$. But $x(\leq_{\tau_1} \cap \leq_{\tau_2})y$ implies that $y \in U_1$ and $y \in U_2$. Therefore, $y \in U$. Thus, we get that $y \in Cl_{\tau_1 \vee \tau_2}(\{x\})$ and hence that $x \leq_{\tau_1 \vee \tau_2} y$.

vi) The valuation function is the same for 'both' structures.

Claim 2. H is a bijection.

We indicate that the map

$$(X, \sim_1, \sim_2, \leq_1, \leq_2, \leq, \mathcal{V}) \overset{H^{-1}}{\longmapsto} (X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V})$$

is an inverse of H, where

- Π_1, Π_2 are the partitions induced by $\sim_1, \sim_2$, respectively,
- $\tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}$ are the *Alexandrov* topologies induced on X by the preorders $\leq_1, \leq_2, \leq$, respectively (recall Observation 7).

First of all, notice that for $(X, \sim_1, \sim_2, \leq_1, \leq_2, \leq, \mathcal{V})$, the structure $(X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V})$ is a *special Alexandrov* model. This can be seen by analogous arguments as the ones presented for *Claim 1.*: Π_1, Π_2, as defined, are clearly partitions. $\tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}$ are *Alexandrov* topologies iff $\leq_1, \leq_2, \leq$ are preorders. Condition *iii*) from Definition 5.2.6 (Relational models) implies that $(X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V})$ meets (*denk*). Condition *iv*) from Definition 5.2.6 implies that $(X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V})$ meets (*dend*). Finally, condition v) from Definition 5.2.6, that says $\leq = \leq_1 \cap \leq_2$, implies that $\tau_{\leq} = \tau_{\leq_1 \cap \leq_2}$, which is equal to $\tau_{\leq_1} \vee \tau_{\leq_2}$ by Lemma 2.5.1 *a*) from Chapter 2.

That H^{-1} is an inverse is implied by Observation 8:

$$(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V}) \overset{H}{\longmapsto} (X, \sim_1, \sim_2, \leq_{\tau_1}, \leq_{\tau_2}, \leq_\tau, \mathcal{V}) \overset{H^{-1}}{\longmapsto} (X, \Pi_1, \Pi_2, \tau_{\leq_{\tau_1}}, \tau_{\leq_{\tau_2}}, \tau_{\leq_\tau}, \mathcal{V})$$

where the equalities $\tau_{\leq_{\tau_1}} = \tau_1$, $\tau_{\leq_{\tau_2}} = \tau_2$, $\tau_{\leq_\tau} = \tau$ all hold by the second item of Observation 8.

Similarly,

$$(X, \sim_1, \sim_2, \leq_1, \leq_2, \leq, \mathcal{V}) \overset{H^{-1}}{\longmapsto} (X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V}) \overset{H}{\longmapsto} (X, \sim_1, \sim_2, \leq_{\tau_{\leq_1}}, \leq_{\tau_{\leq_2}}, \leq_{\tau_{\leq}}, \mathcal{V})$$

where the equalities $\leq_{\tau_{\leq_1}} = \leq_1$, $\leq_{\tau_{\leq_2}} = \leq_2$, $\leq_{\tau_{\leq}} = \leq$ all hold by the first item of Observation 8.

Claim 3. The correspondence is truth-preserving at same states.

Let $\mathcal{M}^{rel} = (X, \sim_1, \sim_2, \leq_1, \leq_2, \leq, \mathcal{V})$ be a relational model. Let us use $\mathcal{M}^{top}$ to refer to the *special Alexandrov* model given by $H^{-1}(\mathcal{M}^{rel}) = (X, \Pi_1, \Pi_2, \tau_{\leq_1}, \tau_{\leq_2}, \tau_{\leq}, \mathcal{V})$. We want to show that, for a given formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ and a point $x \in X$, we have that

$$\mathcal{M}^{rel}, x \Vdash \phi \text{ iff } \mathcal{M}^{top}, x \vDash \phi.$$

We proceed by induction on ϕ. The base case is a given, since the valuation function is the same for both models. The cases for the Boolean connectives are also pretty straightforward, since the extensions coincide in the two semantics.

It only remains to deal with the modal operators:

- ("K_i") Fix $i \in \{1, 2\}$. Assume that the claim holds for ϕ. We will show that

  "$x \in [\leq_i]\phi$ iff $x \in int_{\tau_{\leq_i}^{x[i]}} \|\phi\|$."

  We have that $x \in [\leq_i]\phi$ iff $x \uparrow_{\leq_i} \subseteq [\![\phi]\!]$. By induction hypothesis, this last thing holds iff $x \uparrow_{\leq_i} \subseteq \|\phi\|$. But recall that $x \uparrow_{\leq_i}$ is a basic open set in the topology $\tau_{\leq_i}$. Moreover, $x \uparrow_{\leq_i} \subseteq x[i]$ (the cells are open). Therefore, $x \uparrow_{\leq_i} \subseteq \|\phi\|$ iff $x \uparrow_{\leq_i} \subseteq (\|\phi\| \cap x[i])$ iff $x \in int_{\tau_{\leq_i}^{x[i]}} \|\phi\|$. Our double implication is correct.

The first two items of Observation 6, by which $int_{\tau^{x[i]}_{\leq_i}}\|\phi\|$ coincides with $\|K_i\phi\|$ in *special* models, guarantee that $\mathcal{M}^{rel}, x \Vdash K_i\phi$ iff $\mathcal{M}^{top}, x \vDash K_i\phi$.

- ("D") Assume that the claim holds for ϕ. We will show that

  "$x \in [\leq]\phi$ iff $x \in int_{(\tau_{\leq_1} \vee \tau_{\leq_2})^{x[1]\cap x[2]}}\|\phi\|$."

  In this particular case, we have that $x \in [\leq]\phi$ iff $x \in [\leq_1 \cap \leq_2]\phi$ iff $x \uparrow_{\leq_1 \cap \leq_2} \subseteq [\![\phi]\!] = \|\phi\|$ (where this last equality of extensions holds by induction hypothesis). But recall that $x \uparrow_{\leq_1 \cap \leq_2} = x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2}$, which is an open set of $\tau_{\leq_1} \vee \tau_{\leq_2}$ included in $x[1] \cap x[2]$. Therefore, $x \uparrow_{\leq_1 \cap \leq_2} \subseteq \|\phi\|$ iff $x \uparrow_{\leq_1} \cap x \uparrow_{\leq_2} \subseteq (\|\phi\| \cap (x[1] \cap x[2]))$ iff $x \in int_{(\tau_{\leq_1} \vee \tau_{\leq_2})^{x[1]\cap x[2]}}\|\phi\|$.

  Once again, Observation 6 (third item) ensures that $\mathcal{M}^{rel}, x \Vdash D\phi$ iff $\mathcal{M}^{top}, x \vDash D\phi$.

$\square$

So we are endowed with a powerful association between *special Alexandrov* models and relational models. It must be crystal clear by now that if we succeed in proving weak completeness of the proof system $\mathbf{\Lambda}_{TKD}$ with respect to relational models, then we are done, since Proposition 5.2.7 above ensures that we get weak completeness with respect to *special Alexandrov* models.

5.2.2 Completeness of $\mathbf{\Lambda}_{TKD}$ with respect to relational models

Our proof is a step-by-step construction that involves many phases. It should be beneficial to keep in mind the structure of intent.

The plan

(Preliminaries.) We present a generalization of the concept of relational model for TKD, the so-called *relational pseudo-models* and *intermediary pseudo-models*. We intend to build a relational model directly from our syntax, with the well-known technique of *canonical structures*, where we have one relation for each one of our operators. Thus, we will get a relation for K_1, a relation for K_2, and a relation for D. However, it will not necessarily be the case that the relation for D is exactly the intersection of the relations for K_1 and K_2. Hence the notion of pseudo-model. Moreover, the syntactic side of our logic does not talk at all about our equivalence relations, so we will have to address such shortcoming by being very careful in our pertinent definitions and proofs. Though conceptually uninteresting, a notion of *intermediary pseudo-model* will help us deal with the properties we desire to attain from the equivalence relations of our structures. Defining everything rigorously will give us the opportunity to speak about satisfaction of the formulas of our language in the auxiliary relational tuples. In this preliminary phase, we also show explicitly that $\mathbf{\Lambda}_{TKD}$ is sound with respect to relational pseudo-models, since soundness will be convenient to have.

Step 1. We produce the canonical structure for $\mathbf{\Lambda}_{TKD}$, show that it is a pseudo-model *in all the sense of the word*, and prove weak completeness of $\mathbf{\Lambda}_{TKD}$ with respect to the class of relational pseudo-models.

Step 2. We prove weak completeness of $\mathbf{\Lambda}_{TKD}$ with respect to the class of finite relational pseudo-models. We use the key concept of *filtration* to do so.

Step 3. We show that any finite relational pseudo-model is a *bounded-morphic* image of a finite intermediary pseudo-model. Thus, we prove weak completeness with respect to the class of intermediary pseudo-models.

Step 4. We show that any finite intermediary pseudo-model is a *bounded-morphic* image of a finite relational model. Thus, we prove weak completeness with respect to the class of relational models.

Without further ado, we proceed with the execution of the plan.

(Preliminaries)

Definition 5.2.8. (Relational pseudo-models)

A tuple $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ will be called a *relational pseudo-model* if and only if

i) X is a set. $\sim_1, \sim_2, \sim$ are three equivalence relations on X. $\leq_1, \leq_2, \leq$ are three preorders on X.

ii)
- $\leq_1 \subseteq \sim_1$, $\leq_2 \subseteq \sim_2$, and $\leq \subseteq \sim$.
- $\leq \subseteq \leq_1 \cap \leq_2$.
- $\sim \subseteq \sim_1 \cap \sim_2$.

iii) (*denk*) For a fixed $i \in \{1, 2\}$, for every $w, v \in X$ such that $w \sim_i v$, there exists $u \in X$ such that $w \leq_i u$ and $v \leq u$.

iv) (*dend*) For every $w, v \in X$ such that $w \sim v$, there exists $u \in X$ such that $w \leq u$ and $v \leq u$.

v) $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

Moreover, such a tuple will evidently turn out to be a relational *model* if, additionally, we have that

I. $\sim = \sim_1 \cap \sim_2$,

S. $\leq = \leq_1 \cap \leq_2$.

Satisfaction and *extension* of the formulas of $\mathcal{L}_K^{\mathcal{D}}$ are taken as in Definition 5.2.6. Regard that this entails that Observation 9 applies here, as well.

As pointed out in the statement of our plan, we need another technical definition to make things work smoothly in our proof. The following structures are not quite models, but they are not as broad as pseudo-models.

Definition 5.2.9. We will say that a relational pseudo-model $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ is an *intermediary pseudo-model* if the equality $\sim = \sim_1 \cap \sim_2$ holds, but not necessarily $\leq = \leq_1 \cap \leq_2$.

Other very important tools that we will be using are truth-preserving transformations between pseudo-models.

Definition 5.2.10. (Bounded morphisms for pseudo-models) Let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ and $\mathcal{M}' = (X', \sim_1', \sim_2', \sim', \leq_1', \leq_2', \leq', \mathcal{V}')$ be two relational pseudo-models. A map $h : \mathcal{M} \to \mathcal{M}'$ is a *bounded morphism* if it meets:

- For every $w \in X$, w and $h(w)$ satisfy the same propositional letters.
- *(Forth Conditions)* For a fixed $i \in \{1, 2\}$, if $w \leq_i v$, then $h(w) \leq_i' h(v)$. If $w \leq v$, then $h(w) \leq' h(v)$.
- *(Back Conditions)* For a fixed $i \in \{1, 2\}$, if $h(w) \leq_i' v'$, then there exists $v \in X$ such that $w \leq_i v$ and $h(v) = v'$. If $h(w) \leq' v'$, then there exists $v \in X$ such that $w \leq v$ and $h(v) = v'$.

Proposition 5.2.11. *(Invariance of modal satisfaction under bounded morphisms)*

If $h : \mathcal{M} \to \mathcal{M}'$ is a bounded morphism between two relational pseudo-models $\mathcal{M}$ and $\mathcal{M}'$, then for every $w \in X$ and every formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$, we have that $\mathcal{M}, w \Vdash \phi$ iff $\mathcal{M}', h(w) \Vdash \phi$.

Proof. We proceed by induction on ϕ. The first item in the definition of bounded morphisms immediately gives us the base case. The Boolean cases are routine. Let us deal with the modalities:

- ("K_i") Fix $i \in \{1, 2\}$. We want to show that $\mathcal{M}, w \Vdash K_i\phi$ iff $\mathcal{M}', h(w) \Vdash K_i\phi$. For the left-to-right direction, assume that $\mathcal{M}, w \Vdash K_i\phi$. This means that $w \in [\leq_i]\phi$. Let $v' \in X'$ such that $h(w) \leq_i' v'$. We know that there exists $v \in X$ such that $w \leq_i v$ and $h(v) = v'$ (*Back Condition*). But our assumption implies that $\mathcal{M}, v \Vdash \phi$. By induction hypothesis, then, we get that $\mathcal{M}', v' \Vdash \phi$ as well, so that in fact we have shown that $\mathcal{M}', h(w) \Vdash K_i\phi$. For the right-to-left direction, assume that $\mathcal{M}', h(w) \Vdash K_i\phi$. Let $v \in X$ such that $w \leq_i v$. The *Forth Condition* implies that $h(w) \leq_i h(v)$, so that we get that $\mathcal{M}', h(v) \Vdash \phi$. By induction hypothesis, this last thing implies that $\mathcal{M}, v \Vdash \phi$. Thus, we have shown that $\mathcal{M}, w \Vdash K_i\phi$.

- ("D") The same argument can be used to show that $\mathcal{M}, w \Vdash D\phi$ iff $\mathcal{M}', h(w) \Vdash D\phi$. By marking "$\leq$" wherever it reads "$\leq_i$," and "$D$" wherever it reads "$K_i$" above, we get the desired property.

$\square$

Notice that our definition of bounded morphisms does not even touch the equivalence relations of our structures. Bounded morphisms are important because they are truth-preserving maps. Since the semantics for our formulas does not explicitly involve our equivalence relations, we only need the *Back* and *Forth Conditions* for the pertinent preorders to guarantee that two pseudo-models between which there exists a bounded morphism will satisfy the same formulas of our language.

Definition 5.2.12. (Filtrations of pseudo-models)

Let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be any relational pseudo-model for TKD, and let Σ be a sub-formula closed set of formulas of $\mathcal{L}_K^{\mathcal{D}}$. We define an equivalence relation on $\leftrightsquigarrow_\Sigma$ on X, given by the following rule:

$$w \leftrightsquigarrow_\Sigma v \text{ iff for all } \phi \in \Sigma,\ [\mathcal{M}, w \Vdash \phi \text{ iff } \mathcal{M}, v \Vdash \phi].$$

For a given $w \in X$, the class of w under $\leftrightsquigarrow_\Sigma$ will be denoted by $|w|$.

A *filtration* of $\mathcal{M}$ through Σ consists of a tuple $\mathcal{M}^F = (X^F, \sim_1^F, \sim_2^F, \sim^F, \leq_1^F, \leq_2^F, \leq^F, \mathcal{V}^F)$ for which the following conditions hold:

- $X^f = \{|w| \mid w \in X\}$.
- *("Forth" conditions)* For a fixed $i \in \{1, 2\}$, if $w \leq_i v$, then $|w| \leq_i^F |v|$; if $w \leq v$, then $|w| \leq^F |v|$.
- *("Back" conditions)* For a fixed $i \in \{1, 2\}$, if $|w| \leq_i^F |v|$, then for every $K_i\phi \in \Sigma$, $\mathcal{M}, w \Vdash K_i\phi$ implies that $\mathcal{M}, v \Vdash \phi$; if $|w| \leq^F |v|$, then for every $D\phi \in \Sigma$, $\mathcal{M}, w \Vdash D\phi$ implies that $\mathcal{M}, v \Vdash \phi$.
- $\mathcal{V}^F(p) = \{|w| \mid \mathcal{M}, w \Vdash p\}$.

Proposition 5.2.13. *(Invariance of modal satisfaction under filtrations)*

Let $\mathcal{M}$ be a relational pseudo-model for TKD, and let Σ be a sub-formula closed set of formulas of $\mathcal{L}_K^{\mathcal{D}}$. If $\mathcal{M}^F$ is a filtration of $\mathcal{M}$ through Σ, then for every $w \in X$ and every formula $\phi \in \Sigma$, we have that $\mathcal{M}, w \Vdash \phi$ iff $\mathcal{M}^F, |w| \Vdash \phi$.

Proof. By induction on ϕ. The base case comes straight from the definition of $\mathcal{V}^F$. The Boolean cases are easy, but it should be noted that we can apply the induction hypothesis only because Σ is closed under sub-formulas. Let us address the modalities, then.

- ("K_i") Fix $i \in \{1, 2\}$. Let $K_i\phi \in \Sigma$. We want to show that $\mathcal{M}, w \Vdash K_i\phi$ iff $\mathcal{M}^F, |w| \Vdash K_i\phi$. For the left-to-right direction, assume that $\mathcal{M}, w \Vdash K_i\phi$. This means that $w \in [\leq_i]\phi$. Let $v \in X$ such that $|w| \leq_i^F |v|$. By the appropriate *"Back" Condition*, our assumption then implies that $\mathcal{M}, v \Vdash \phi$. By induction hypothesis, then, we get that $\mathcal{M}^F, |v| \Vdash \phi$ as well, so that in fact we have shown that $\mathcal{M}^F, |w| \Vdash K_i\phi$. For the right-to-left direction, assume that $\mathcal{M}^F, |w| \Vdash K_i\phi$. Let $v \in X$ such that $w \leq_i v$. The *"Forth" Condition* implies that $|w| \leq_i^F |v|$, so we get that $\mathcal{M}^F, |v| \Vdash \phi$. By induction hypothesis, this last thing implies that $\mathcal{M}, v \Vdash \phi$. Thus, we have shown that $\mathcal{M}, w \Vdash K_i\phi$.
- ("D") Once again, a similar procedure can be shown to realize that, for every $D\phi \in \Sigma$, $\mathcal{M}, w \Vdash D\phi$ iff $\mathcal{M}^F, |w| \Vdash D\phi$.

$\square$

Proposition 5.2.14. *The proof system $\mathbf{\Lambda}_{TKD}$ is sound for $\mathcal{L}_K^{\mathcal{D}}$ with respect to relational pseudo-models.*

Proof. It is a matter of routine to show that the defined relational semantics validate the **S4** axioms both for K_i ($i \in \{1, 2\}$) and for D. Similarly, it is easy to see that the rules of inference preserve validity.

For the ($.2_D$) axiom, let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be any relational pseudo-model. For a given formula ϕ and an arbitrary $w \in X$, assume that $\mathcal{M}, w \Vdash \langle D \rangle D\phi$. This means that there exists a world $v \in X$ such that $w \leq v$ and $\mathcal{M}, v \Vdash D\phi$. Let $u \in X$ such that $w \leq u$. Trivially, we have that $v \sim u$, so condition

iv) throws the existence of $o \in X$ such that $v \leq o$ and $u \leq o$. From $\mathcal{M}, v \Vdash D\phi$ we get that $\mathcal{M}, o \Vdash \phi$. Thus, with $w \leq u \leq o$ we get that actually $\mathcal{M}, w \Vdash D\langle D\rangle\phi$.

As for the axiom (DsK), let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be any relational pseudo-model. For a given formula ϕ and an arbitrary $w \in X$, assume without loss of generality that $\mathcal{M}, w \Vdash K_1\phi$. If $w \leq v$, then we have that $w \leq_1 v$ as well. Therefore, we get that $\mathcal{M}, v \Vdash \phi$, which only means that $\mathcal{M}, w \Vdash D\phi$. Thus, we get that $\mathcal{M}, w \Vdash K_1\phi \to D\phi$.

For the (WSD) axioms, fix $i \in \{1, 2\}$, and let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be any relational pseudo-model. For a given formula ϕ and an arbitrary $w \in X$, assume that $\mathcal{M}, w \Vdash \langle K_i\rangle K_i\phi$. This means that there exists a world $v \in X$ such that $w \leq_i v$ and $\mathcal{M}, v \Vdash K_i\phi$. Let $u \in X$ such that $w \leq_i u$. Trivially, we have that $v \sim_i u$, so condition iii) guarantees that there exists $o \in X$ such that $v \leq_i o$ and $u \leq o$. From $\mathcal{M}, v \Vdash K_i\phi$ we get that $\mathcal{M}, o \Vdash \phi$. Thus, with $w \leq_i u \leq o$ we get that actually $\mathcal{M}, w \Vdash K_i\langle D\rangle\phi$.

□

Step 1. Completeness with respect to relational pseudo-models

In what follows, we will use the term $\mathbf{\Lambda}$ to refer to $\mathbf{\Lambda}_{TKD}$, solely for notational convenience.

Recall that a set Γ of formulas of $\mathcal{L}_K^{\mathcal{D}}$ is said to be $\mathbf{\Lambda}$-*consistent* if and only if $\Gamma \nvdash_{\mathbf{\Lambda}} \bot$, and $\mathbf{\Lambda}$-*inconsistent* otherwise. A formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ is $\mathbf{\Lambda}$-*consistent* if $\{\phi\}$ is, and otherwise it is $\mathbf{\Lambda}$-*inconsistent*. It is easy to see that a set Γ of formulas of $\mathcal{L}_K^{\mathcal{D}}$ is $\mathbf{\Lambda}$-*consistent* iff every finite subset of Γ is. Another simple but important fact to remember is that $\Gamma \nvdash_{\mathbf{\Lambda}} \phi$ iff $\Gamma \cup \{\neg\phi\}$ is $\mathbf{\Lambda}$-*consistent*.

For the trained eye, it should seem obvious by now that we will make heavy use of the following result.

Proposition 5.2.15. *$\mathbf{\Lambda}$ is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to a class of structures $\mathbb{C}$ iff every $\mathbf{\Lambda}$-consistent formula is satisfiable on some $\mathcal{C} \in \mathbb{C}$.*

□

In a similar fashion, we call a set Γ of formulas of $\mathcal{L}_K^{\mathcal{D}}$ $\mathbf{\Lambda}$-*maximally consistent* if it is consistent and, for every $\psi \notin \Gamma$, $\Gamma \cup \{\psi\}$ is $\mathbf{\Lambda}$-inconsistent.

An admittedly famous Lemma allows us to extend every $\mathbf{\Lambda}$-consistent set to a $\mathbf{\Lambda}$-maximally consistent set:

Lemma 5.2.16. *(Lindenbaum's Lemma)*
If Γ is a $\mathbf{\Lambda}$-consistent set of formulas of $\mathcal{L}_K^{\mathcal{D}}$, then there exists a $\mathbf{\Lambda}$-maximally consistent set Γ^+ such that $\Gamma \subseteq \Gamma^+$.

Proof. The language $\mathcal{L}_K^{\mathcal{D}}$ is countable. Let $\phi_0, \phi_1, \ldots$ be an enumeration of the formulas of $\mathcal{L}_K^{\mathcal{D}}$. We shall recursively define Γ^+ as a union of an $\subseteq$-ascending chain of $\mathbf{\Lambda}$-consistent sets:

$$\Gamma_0 = \Gamma$$
$$\Gamma_{n+1} = \begin{cases} \Gamma_n \cup \{\phi_n\} & \text{if this set is } \mathbf{\Lambda}\text{-consistent,} \\ \Gamma_n \cup \{\neg\phi_n\} & \text{otherwise} \end{cases}$$
$$\Gamma^+ = \bigcup_{n \in \mathbb{N}} \Gamma_n.$$

Clearly $\Gamma \subseteq \Gamma^+$, and it is relatively easy to show that Γ^+, thus defined, is actually a $\mathbf{\Lambda}$-maximally consistent set. □

From here on, we will abbreviate the locution '$\mathbf{\Lambda}$-maximally consistent set' with '$\mathbf{\Lambda}$-MCS.'

$\mathbf{\Lambda}$-MCS's have nice and usable properties:

Proposition 5.2.17. *If Γ is a $\mathbf{\Lambda}$-MCS, then*

a) Γ is closed under Modus Ponens

b) $\mathbf{\Lambda} \subseteq \Gamma$.

c) For every formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$, either $\phi \in \Gamma$ or $\neg\phi \in \Gamma$.

d) For all formulas ϕ, ψ, we have that $\phi \wedge \psi \in \Gamma$ iff both $\phi \in \Gamma$ and $\psi \in \Gamma$.

e)
- *For a fixed $i \in \{1,2\}$, if $\{K_i\psi \mid K_i\psi \in \Gamma\} \vdash_\Lambda \phi$, then $K_i\phi \in \Gamma$.*
- *If $\{D\psi \mid D\psi \in \Gamma\} \vdash_\Lambda \phi$, then $D\phi \in \Gamma$.*

Proof. The first four items are easy to show. They are particular instances of a claim introduced in [16] (p. 199). For the proof of the last item, we adapt an elegant method found in [25] (p. 26):

- Fix $i \in \{1,2\}$. Assume that $\{K_i\psi \mid K_i\psi \in \Gamma\} \vdash_\Lambda \phi$. This means that there exists a finite set of formulas $\{\psi_0, \ldots, \psi_m\}$ such that $K_i\psi_l \in \Gamma$ for every $l \in \{0, \ldots, m\}$ and such that $\vdash_{\mathbf{\Lambda}} K_i\psi_0 \wedge \ldots \wedge K_i\psi_m \to \phi$. By *Necessitation* for K_i, axiom (K), and *Modus Ponens*, we have that $\vdash_{\mathbf{\Lambda}} K_i(K_i\psi_0 \wedge \ldots \wedge K_i\psi_m) \to K_i\phi$.

  But notice that item $d)$ above implies that the formula $K_i\psi_0 \wedge \ldots \wedge K_i\psi_m$ lives in Γ. By *Necessitation* for K_i, we then get that $K_i(K_i\psi_0 \wedge \ldots \wedge K_i\psi_m) \in \Gamma$. Therefore, since by item $a)$ Γ is closed under *Modus Ponens*, we also get that $K_i\phi \in \Gamma$.
- With an analogous method as the one provided above, and using the according *Necessitation* for D and axiom (K_D), we see that the claim holds.

□

<u>The canonical relational pseudo-model</u>

Let $\mathcal{M}^{\mathbf{\Lambda}} = (X^{\mathbf{\Lambda}}, \leq_1^{\mathbf{\Lambda}}, \leq_2^{\mathbf{\Lambda}}, \leq^{\mathbf{\Lambda}}, \mathcal{V}^{\mathbf{\Lambda}})$ be the canonical structure of $\mathbf{\Lambda}$, where

- $X^{\mathbf{\Lambda}}$ is the set of $\mathbf{\Lambda}$-MCS's.
- $\leq_i^{\mathbf{\Lambda}}$ $(i \in \{1,2\})$ are binary relations on $X^{\mathbf{\Lambda}}$, defined by the following rule:
  $w \leq_i^{\mathbf{\Lambda}} v$ iff $\forall\phi(K_i\phi \in w \Rightarrow \phi \in v)$.
- $\leq^{\mathbf{\Lambda}}$ is a binary relation on $X^{\mathbf{\Lambda}}$, defined by the following rule:
  $w \leq^{\mathbf{\Lambda}} v$ iff $\forall\phi(D\phi \in w \Rightarrow \phi \in v)$.
- $\mathcal{V}^{\mathbf{\Lambda}}$ is the canonical valuation; i.e., $\mathcal{V}^{\mathbf{\Lambda}} : Prop \to \mathcal{P}(X^{\mathbf{\Lambda}})$ is defined by
  $\mathcal{V}^{\mathbf{\Lambda}}(p) = \{w \in X^{\mathbf{\Lambda}} \mid p \in w\}$.

Definition 5.2.18. For the canonical structure $\mathcal{M}^{\mathbf{\Lambda}}$ as constructed above, we consider a structural extension of the form $(X^{\mathbf{\Lambda}}, \sim_1^{\mathbf{\Lambda}}, \sim_2^{\mathbf{\Lambda}}, \sim^{\mathbf{\Lambda}}, \leq_1^{\mathbf{\Lambda}}, \leq_2^{\mathbf{\Lambda}}, \leq^{\mathbf{\Lambda}}, \mathcal{V}^{\mathbf{\Lambda}})$, where the new relations are defined the following way:

- For $i \in \{1,2\}$, $\sim_i^{\mathbf{\Lambda}}$ is the $\subseteq$-least equivalence relation including $\leq_i^{\mathbf{\Lambda}}$.[1]
- $\sim^{\mathbf{\Lambda}}$ will denote the $\subseteq$-least equivalence relation including $\leq^{\mathbf{\Lambda}}$.

From here on, when we write $\mathcal{M}^{\mathbf{\Lambda}}$, we will refer to this extension.

Lemma 5.2.19. *Let X be any set such that $\sim_1, \sim_2, \sim$ are equivalence relations on X and $\leq_1, \leq_2, \leq$ are preorders on X. If $\sim_1$ is the $\subseteq$-least equivalence relation containing $\leq_1$, $\sim_2$ is the $\subseteq$-least equivalence relation containing $\leq_2$, and $\sim$ is the $\subseteq$-least equivalence relation containing $\leq$, then condition $a)$ below is equivalent to condition* iii) *$(denk)$ of relational pseudo-models, while condition $b)$ below is equivalent to condition* iv) *$(dend)$ of relational pseudo-models:*

a) For a fixed $i \in \{1,2\}$, for every $w, x, y \in X$ such that $w \leq_i x$ and $w \leq_i y$, there exists z such that $x \leq_i z$ and $y \leq z$.

[1] Recall that for any relation R on a set X, the $\subseteq$-least equivalence relation $\sim_R$ including R is defined as the intersection of all the equivalence relations including R. Alternatively, we can define $\sim_R$ by the the following rule: $x \sim_R y$ iff there exists a finite chain $\{z_1, \ldots, z_n\} \subseteq X$ such that

– $z_1 = x$, $z_n = y$,

– for every $k \in \{1, \ldots, n-1\}$, either $z_k R z_{k+1}$ or $z_{k+1} R z_k$.

Observe that such a definition entails *reflexivity, symmetry, and transitivity* of $\sim_R$.

b) For every $w, x, y \in X$ such that $w \leq x$ and $w \leq y$, there exists z such that $x \leq z$ and $y \leq z$.

Proof.

[a) $\Leftrightarrow$ iii) ($denk$)]

The right-to-left direction is pretty straightforward, since $w \leq_i x$ and $w \leq_i y$ imply that $x \sim_i y$. For the left-to-right one, fix $i \in \{1, 2\}$. Let $w, v \in X$ such that $w \sim_i v$. We know that there exists a finite chain $\{w_1, \ldots, w_n\} \subseteq X$ such that

- $w_1 = w$, $w_n = v$,
- for every $k \in \{1, \ldots, n-1\}$, either $w_k \leq_i w_{k+1}$ or $w_k \geq_i w_{k+1}$ holds.

We will show that condition a) above implies that there exists $u \in X$ such that $w \leq_i u$ and $v \leq u$. In order to do so, we will use an inductive argument (on the length of the aforementioned chains).

Claim. For every finite chain $\{w_1, \ldots, w_n\}$ (of length n) such that $w_1 = w$ and either $w_k \leq_i w_{k+1}$ or $w_k \geq_i w_{k+1}$ (for every $k \in \{1, \ldots, n-1\}$), there exists $u \in X$ such that $w \leq_i u$ and $w_n \leq u$.

We proceed by induction on n:

The base case ($n = 1$) is straightforward, choosing $u = w$. Suppose the claim holds for chains of length $n - 1$. For the relevant chain of length n, there exists $u^* \in X$ such that $w \leq_i u^*$ and $w_{n-1} \leq u^*$. Notice that we have the following two cases:

(Case $w_{n-1} \leq_i w_n$) The assumptions $w_{n-1} \leq_i w_n$ and $w_{n-1} \leq_i u^*$ imply, with condition a), that there exists $u \in X$ such that $u^* \leq_i u$ and $w_n \leq u$. Regard that $w \leq_i u^* \leq_i u$ (IH), so that transitivity of $\leq_i$ implies that $w \leq_i u$ as well.

(Case $w_{n-1} \geq_i w_n$) In this case, we have that $w_n \leq_i w_{n-1} \leq_i u^*$, so that transitivity of $\leq_i$ implies that $w_n \leq_i u^*$. Condition a), then, implies that there exists $u \in X$ such that and $u^* \leq_i u$ and $w_n \leq u$, but then—exactly as above—$w \leq_i u^* \leq_i u$ (IH) and transitivity of $\leq_i$ imply that $w \leq_i u$ as well, since we have that $w \leq_i u^* \leq_i u$.

[b) $\Leftrightarrow$ iv) ($dend$)]

Once again, the right-to-left direction is straightforward, since $w \leq x$ and $w \leq y$ imply that $x \sim y$. For the left-to-right direction, it is easy to realize that we can use the same exact argument as in the previous item to show that for every $w, v \in X$ such that $w \sim v$, there exists $u \in X$ such that $w \leq u$ and $v \leq u$. The only change that has to be made in this analogous proof comes from marking "$\leq$" wherever it reads "$\leq_i$" in the according claim and proof of claim.

□

Lemma 5.2.20. *$\mathcal{M}^{\mathbf{\Lambda}} = (X^{\mathbf{\Lambda}}, \sim_1^{\mathbf{\Lambda}}, \sim_2^{\mathbf{\Lambda}}, \sim^{\mathbf{\Lambda}}, \leq_1^{\mathbf{\Lambda}}, \leq_2^{\mathbf{\Lambda}}, \leq^{\mathbf{\Lambda}}, \mathcal{V}^{\mathbf{\Lambda}})$ is a relational pseudo-model.*

Proof.

i) Preorders and equivalence relations

- It is clear that $\sim_1^{\mathbf{\Lambda}}$, $\sim_2^{\mathbf{\Lambda}}$ and $\sim^{\mathbf{\Lambda}}$ are equivalence relations, straight from their definition.
- Fix $i \in \{1, 2\}$.
  - *(Reflexivity of $\leq_i^{\mathbf{\Lambda}}$)* Let ϕ be a formula of $\mathcal{L}_K^{\mathcal{D}}$ and $w \in X^{\mathbf{\Lambda}}$. We want to show that $K_i\phi \in w$ implies that $\phi \in w$. Assume that $K_i\phi \in w$. Since w is a $\mathbf{\Lambda}$-MCS, $K_i\phi \to \phi \in w$ (a substitution on the (T) axiom). Therefore, since w is closed under *Modus Ponens*, we have that $\phi \in w$ and $w \leq_i^{\mathbf{\Lambda}} w$.

– *(Transitivity of $\leq_i^{\mathbf{\Lambda}}$)* Let $w, v, u \in X^{\mathbf{\Lambda}}$ such that $w \leq_i^{\mathbf{\Lambda}} v$ and $v \leq_i^{\mathbf{\Lambda}} u$. Let ϕ be a formula of $\mathcal{L}_K^{\mathcal{D}}$. We want to show that $K_i\phi \in w$ implies that $\phi \in u$. Assume that $K_i\phi \in w$. Since w is a $\mathbf{\Lambda}$-MCS, $K_i\phi \to K_iK_i\phi \in w$ (a substitution on the (4) axiom). Thus, we have that $K_iK_i\phi \in w$ and hence that $K_i\phi \in v$, which with $v \leq_i^{\mathbf{\Lambda}} u$ implies that $\phi \in u$. So $w \leq_i^{\mathbf{\Lambda}} u$ holds.

With this, we have that $\leq_1^{\mathbf{\Lambda}}$ and $\leq_2^{\mathbf{\Lambda}}$ are preorders on $X^{\mathbf{\Lambda}}$.

- *Reflexivity* and *Transitivity of* $\leq_i^{\mathbf{\Lambda}}$ can be shown in the exact same way, by marking "D" whenever it reads "K_i" above. The proofs would follow accordingly, but referring to axioms (T_D) and (4_D), respectively, which our logic includes as well. Thus, $\leq^{\mathbf{\Lambda}}$ is also a preorder on $X^{\mathbf{\Lambda}}$.

ii) Pertinent inclusions

- The inclusions $\leq_1^{\mathbf{\Lambda}} \subseteq \sim_1^{\mathbf{\Lambda}}$, $\leq_2^{\mathbf{\Lambda}} \subseteq \sim_2^{\mathbf{\Lambda}}$, and $\leq^{\mathbf{\Lambda}} \subseteq \sim^{\mathbf{\Lambda}}$ all hold by definition.
- We want to show that $\leq^{\mathbf{\Lambda}} \subseteq \leq_1^{\mathbf{\Lambda}} \cap \leq_2^{\mathbf{\Lambda}}$. Let $w, v \in X^{\mathbf{\Lambda}}$ such that $w \leq^{\mathbf{\Lambda}} v$. For a fixed $i \in \{1, 2\}$, assume that $K_i\phi \in w$. Since w is a $\mathbf{\Lambda}$-MCS and thus contains all the tautologies of $\mathbf{\Lambda}$, then $K_i\phi \to D\phi \in w$ (a substitution on the (DsK) axiom). Therefore, since w is closed under *Modus Ponens*, we have that $D\phi \in w$. With the assumption $w \leq^{\mathbf{\Lambda}} v$, this implies that $\phi \in v$. Thus, we have shown that $w \leq_i^{\mathbf{\Lambda}} v$ as well, so that $\leq^{\mathbf{\Lambda}} \subseteq \leq_1^{\mathbf{\Lambda}} \cap \leq_2^{\mathbf{\Lambda}}$.
- Notice that the above items imply that $\leq^{\mathbf{\Lambda}} \subseteq \sim_1^{\mathbf{\Lambda}} \cap \sim_2^{\mathbf{\Lambda}}$. Since $\sim^{\mathbf{\Lambda}}$ is defined as the $\subseteq$-least equivalence relation including $\leq^{\mathbf{\Lambda}}$, it is clear that $\sim^{\mathbf{\Lambda}} \subseteq \sim_1^{\mathbf{\Lambda}} \cap \sim_2^{\mathbf{\Lambda}}$ as well.

iii) Density condition iii) (*denk*)

We will show that $\mathcal{M}^{\mathbf{\Lambda}}$ meets condition *a*) in Lemma 5.2.19, so that in virtue of said Lemma, it will also satisfy condition iii) for relational pseudo-models.

Fix $i \in \{1, 2\}$. Let $w, v, u \in X^{\mathbf{\Lambda}}$ such that $w \leq_i^{\mathbf{\Lambda}} v$ and $w \leq_i^{\mathbf{\Lambda}} u$.

Let $a := \{\phi \mid K_i\phi \in v\}$, $b := \{\psi \mid D\psi \in u\}$, and $o' := a \cup b$. We affirm that o' is $\mathbf{\Lambda}$-consistent. Suppose it is not. Then $o' \vdash_{\mathbf{\Lambda}} \bot$. This means that there exists a finite set of formulas of the form $\{\phi_0, \ldots, \phi_n\} \cup \{\psi_0, \ldots, \psi_m\}$ such that $\phi_k \in a$ for every $k \in \{0, \ldots, n\}$, $\psi_l \in b$ for every $l \in \{0, \ldots, m\}$, and

$$\phi_0 \wedge \ldots \wedge \phi_n \wedge \psi_0 \wedge \ldots \wedge \psi_m \vdash_{\mathbf{\Lambda}} \bot. \tag{5.1}$$

Notice that, since v and u are $\mathbf{\Lambda}$-MCS's, we have that $K_i\phi_0 \wedge \ldots \wedge K_i\phi_n \in v$ and $D\psi_0 \wedge \ldots \wedge D\psi_m \in u$. Let us mark $\phi' := \phi_0 \wedge \ldots \wedge \phi_n$ and $\psi' := \psi_0 \wedge \ldots \wedge \psi_m$. Since both K_i and D are distributive over conjunction, and since v, u are $\mathbf{\Lambda}$-MCS's, we also get that $K_i\phi' \in v$ and $D\psi' \in u$, which in turn means that $\phi' \in a$ and $\psi' \in b$. From equation (5.1) we get that $\phi' \wedge \psi' \vdash_{\mathbf{\Lambda}} \bot$, which implies that $\vdash_{\mathbf{\Lambda}} \phi' \to \neg\psi'$, and hence that

$$\vdash_{\mathbf{\Lambda}} \langle D\rangle\phi' \to \langle D\rangle\neg\psi'. \tag{5.2}$$

We have seen that $K_i\phi' \in v$. Together with our assumption that $w \leq_i^{\mathbf{\Lambda}} v$, we get that $\langle K_i\rangle K_i\phi' \in w$. But we also know that $\langle K_i\rangle K_i\phi' \to K_i\langle D\rangle\phi' \in w$ (a substitution on the (WSD) axiom). Therefore, $K_i\langle D\rangle\phi' \in w$ (w is closed under *Modus Ponens*). With our other assumption that $w \leq_i^{\mathbf{\Lambda}} u$, this last thing entails that $\langle D\rangle\phi' \in u$. From equation (5.2) and the fact that u is a $\mathbf{\Lambda}$-MCS closed under *Modus Ponens*, then, we get that $\langle D\rangle\neg\psi' \in u$, but this contradicts our preliminary result of $D\psi' \in u$. Therefore, o' is $\mathbf{\Lambda}$-consistent.

By Lindenbaum's Lemma, we can extend o' to a $\mathbf{\Lambda}$-MCS, which we will name o. It is easy to see, now, that $v \leq_i^{\mathbf{\Lambda}} o$ and $u \leq^{\mathbf{\Lambda}} o$: for every ϕ, ψ, we have that, on one hand, $K_i\phi \in v$ iff $\phi \in a \subseteq o' \subseteq o$, while on the other, $D\psi \in u$ iff $\psi \in b \subseteq o' \subseteq o$.

iv) Density condition iv) (*dend*)

Same method as the previous item. We will show that $\mathcal{M}^{\mathbf{\Lambda}}$ meets condition *b*) in Lemma 5.2.19, so that in virtue of said Lemma, it will also satisfy condition iv) for relational pseudo-models. Let $w, v, u \in X^{\mathbf{\Lambda}}$ such that $w \leq^{\mathbf{\Lambda}} v$ and $w \leq^{\mathbf{\Lambda}} u$.

Let $a := \{\phi \mid D\phi \in v\}$, $b := \{\psi \mid D\psi \in u\}$, and $o' := a \cup b$. We affirm that o' is $\mathbf{\Lambda}$-consistent. Suppose it is not. Then $o' \vdash_{\mathbf{\Lambda}} \bot$. This means that there exists a finite set of formulas of the form

$\{\phi_0, \ldots, \phi_n\} \cup \{\psi_0, \ldots, \psi_m\}$ such that $\phi_k \in a$ for every $k \in \{0, \ldots, n\}$, $\psi_l \in b$ for every $l \in \{0, \ldots, m\}$, and

$$\phi_0 \wedge \ldots \wedge \phi_n \wedge \psi_0 \wedge \ldots \wedge \psi_m \vdash_{\mathbf{\Lambda}} \bot. \tag{5.3}$$

Notice that, since v and u are $\mathbf{\Lambda}$-MCS's, we have that $D\phi_0 \wedge \ldots \wedge D\phi_n \in v$ and $D\psi_0 \wedge \ldots \wedge D\psi_m \in u$. Let us mark $\phi' := \phi_0 \wedge \ldots \wedge \phi_n$ and $\psi' := \psi_0 \wedge \ldots \wedge \psi_m$. D is distributive over conjunction and v, u are $\mathbf{\Lambda}$-MCS's, so we also get that $D\phi' \in v$ and $D\psi' \in u$, which in turn means that $\phi' \in a$ and $\psi' \in b$. From equation (5.3) we get that $\phi' \wedge \psi' \vdash_{\mathbf{\Lambda}} \bot$, which implies that $\vdash_{\mathbf{\Lambda}} \phi' \to \neg\psi'$, and hence that

$$\vdash_{\mathbf{\Lambda}} \langle D\rangle\phi' \to \ \langle D\rangle\neg\psi'. \tag{5.4}$$

We have seen that $D\phi' \in v$. Then $w \leq^{\mathbf{\Lambda}} v$ implies that $\langle D\rangle D\phi' \in w$. But we also know that $\langle D\rangle D\phi' \to D\langle D\rangle\phi' \in w$ (a substitution on the $(.2_D)$ axiom). Therefore, $D\langle D\rangle\phi' \in w$ (w is closed under *Modus Ponens*). With our other assumption that $w \leq^{\mathbf{\Lambda}} u$, we get that $\langle D\rangle\phi' \in u$. From equation (5.4) and the fact that u is a $\mathbf{\Lambda}$-MCS closed under *Modus Ponens*, we get that $\langle D\rangle\neg\psi' \in u$, but this contradicts our preliminary result of $D\psi' \in u$. Therefore, o' is $\mathbf{\Lambda}$-consistent.

By Lindenbaum's Lemma, we can extend o' to a $\mathbf{\Lambda}$-MCS, which we will name o. It is clear that $v \leq^{\mathbf{\Lambda}} o$ and $u \leq^{\mathbf{\Lambda}} o$: for every ϕ, ψ, we have that, on one hand, $D\phi \in v$ iff $\phi \in a \subseteq o' \subseteq o$, while on the other, $D\psi \in u$ iff $\psi \in b \subseteq o' \subseteq o$.

v) $\mathcal{V}^{\mathbf{\Lambda}}$ is defined as a valuation function.

□

Lemma 5.2.21. *(Truth Lemma) For every formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ and every $w \in X^{\mathbf{\Lambda}}$, we have that*

$$\mathcal{M}^{\mathbf{\Lambda}}, w \Vdash \phi \text{ iff } \phi \in w.$$

Proof. We proceed by induction on ϕ. The base case and the cases for the Boolean connectives are routine. Therefore, we proceed with the modalities.

- ("K_i") Fix $i \in \{1, 2\}$. Assume that the claim holds for ϕ. We want to show that $\mathcal{M}^{\mathbf{\Lambda}}, w \Vdash K_i\phi$ iff $K_i\phi \in w$.

  For the right-to-left direction, we work by contraposition. Assume that $\mathcal{M}^{\mathbf{\Lambda}}, w \nVdash K_i\phi$. This means that there exists $v \geq_i^{\mathbf{\Lambda}} w$ such that $\mathcal{M}^{\mathbf{\Lambda}}, v \nVdash \phi$, which by induction hypothesis entails that $\phi \notin v$. But then, since $v \geq_i^{\mathbf{\Lambda}} w$, this means that $K_i\phi \notin w$.

  For the left-to-right direction, assume that $\mathcal{M}^{\mathbf{\Lambda}}, w \Vdash K_i\phi$. We will first show that $w'' := \{K_i\psi \mid K_i\psi \in w\}$ is such that $w'' \vdash_{\mathbf{\Lambda}} \phi$. Suppose this does not hold. Then the set $w'' \cup \{\neg\phi\}$ is $\mathbf{\Lambda}$-consistent. By Lindenbaum's Lemma, we can extend w'' to a $\mathbf{\Lambda}$-MCS, which we will name w'. Notice that for every $K_i\psi \in w$, $K_i\psi \in w'$, but since w' is a $\mathbf{\Lambda}$-MCS, we have that for every $K_i\psi \in w$, actually $\psi \in w'$ (w' has every substitution of the (T) axiom and is closed under *Modus Ponens*). Therefore, we have that $w \leq_i^{\mathbf{\Lambda}} w'$. By assumption and induction hypothesis, then, we get that $\phi \in w'$, which is a contradiction. Therefore $w'' \vdash_{\mathbf{\Lambda}} \phi$, but then Proposition 5.2.17 *e*) implies that $K_i\phi \in w$.

- ("D") Assume that the claim holds for ϕ. We want to show that $\mathcal{M}^{\mathbf{\Lambda}}, w \Vdash D\phi$ iff $D\phi \in w$.

  Both right-to-left and left-to-right directions can be shown using the technique applied in the case for "K_i," accordingly appealing to Proposition 5.2.17 *e*).

□

Proposition 5.2.22. $\mathbf{\Lambda}$ *is (weakly) complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of relational pseudo-models.*

Proof. Let ϕ_* be a formula of $\mathcal{L}_K^{\mathcal{D}}$ such that ϕ_* is $\mathbf{\Lambda}$-consistent. By Lindenbaum's Lemma, we know that there exists a $\mathbf{\Lambda}$-MCS Γ_* such that $\{\phi_*\} \subseteq \Gamma_*$. By Lemma 5.2.21 (the Truth Lemma), we have that

$$\mathcal{M}^{\mathbf{\Lambda}}, \Gamma_* \Vdash \phi_*.$$

□

Step 2. Completeness with respect to finite relational pseudo-models

In this step, we will show that the system $\mathbf{\Lambda}$ is weakly complete with respect to the class of finite relational pseudo-models. In order to do so, we will use *filtrations* of relational pseudo-models.

For a given formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$, we define Σ_ϕ as the $\subseteq$-smallest set such that

i) $\phi \in \Sigma_\phi$.

ii) Σ_ϕ is closed under sub-formulas, conjunction, and negation.

iii) For every $K_i\psi \in \Sigma_\phi$ ($i \in \{1, 2\}$), $DK_i\psi \in \Sigma_\phi$ as well.

Observation 10. For any ϕ of $\mathcal{L}_K^{\mathcal{D}}$, Σ_ϕ is finite *up to logical equivalence*. We briefly sketch the proof of this. Let Σ_0 denote the set of all sub-formulas of ϕ, which is finite. Take $\Sigma_1 := \Sigma_0 \cup \{DK_i\psi \mid K_i\psi \in \Sigma_0, i \in \{1, 2\}\}$. Since Σ_0 is finite, then Σ_1 is finite as well. We notice that every formula of Σ_ϕ is logically equivalent to a Boolean combination of formulas in Σ_1. But it is well-known that, if Σ_1 is a finite set, then there are only finitely many non-logically-equivalent Boolean combinations of formulas in Σ_1 (a *Boolean algebra* generated by finitely many atoms is finite).

Now, let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be any relational pseudo-model for TKD and ϕ a formula of $\mathcal{L}_K^{\mathcal{D}}$. We define the structure $\mathcal{M}^f := (X^f, \sim_1^f, \sim_2^f, \sim^f, \leq_1^f, \leq_2^f, \leq^f, \mathcal{V}^f)$, where

- X^f is the quotient space of X under the equivalence relation $\leftrightsquigarrow_{\Sigma_\phi}$, which is given by the following rule: $w \leftrightsquigarrow_{\Sigma_\phi} v$ iff for all $\psi \in \Sigma_\phi$, $\mathcal{M}, w \Vdash \psi$ iff $\mathcal{M}, v \Vdash \psi$.
- For $i \in \{1, 2\}$, the relation $\leq_i^f$ is defined by the following rule: $|w| \leq_i^f |v|$ iff for every $K_i\psi \in \Sigma_\phi$, $\mathcal{M}, w \Vdash K_i\psi$ implies that $\mathcal{M}, v \Vdash K_i\psi$.
- The relation $\leq^f$ is defined by the following rule: $|w| \leq^f |v|$ iff for every $D\psi \in \Sigma_\phi$, $\mathcal{M}, w \Vdash D\psi$ implies that $\mathcal{M}, v \Vdash D\psi$.
- For $i \in \{1, 2\}$, $\sim_i^f$ is defined as the $\subseteq$-least equivalence relation including $\leq_i^f$.
- $\sim^f$ is defined as the $\subseteq$-least equivalence relation including $\leq^f$.
- $\mathcal{V}^f$ is a valuation function given by $\mathcal{V}^f(p) = \{|w| \mid w \in \mathcal{V}(p)\}$.

Proposition 5.2.23. *If $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ is a relational pseudo-model for TKD, and ϕ is a formula of $\mathcal{L}_K^{\mathcal{D}}$, then $\mathcal{M}^f$—as constructed above—is a filtration (of pseudo-models) of $\mathcal{M}$ through Σ_ϕ.*

Proof. We only have to check that the defined relations meet the filtrations' conditions that we reviewed earlier (Definition 5.2.12):

- *("Forth" conditions)* Fix $i \in \{1, 2\}$. Let $w, v \in X$ such that $w \leq_i v$. We want to show that $|w| \leq_i^f |v|$. For any $K_i\psi \in \Sigma_\phi$, assume that $\mathcal{M}, w \Vdash K_i\psi$. Since $\leq_i$ is transitive, we have that $\mathcal{M}, w \Vdash K_iK_i\psi$ (relational soundness of the (4) axiom). Therefore, we get that the assumption $w \leq_i v$ implies that $\mathcal{M}, v \Vdash K_i\psi$. With this, we have shown that for any $K_i\psi \in \Sigma_\phi$, $\mathcal{M}, w \Vdash K_i\psi$ implies that $\mathcal{M}, v \Vdash K_i\psi$, which means that $|w| \leq_i^f |v|$.

  The analogous property for $\leq^f$ can be shown by marking "$\leq$" wherever it reads "$\leq_i$," and "D" wherever it reads "K_i," in the previous paragraph. Since $\leq$ is transitive by assumption (the axiom (4_D) is also sound with respect to relational pseudo-models), we get the desired situation.

- *("Back" conditions)* Fix $i \in \{1, 2\}$. Let $w, v \in X$ such that $|w| \leq_i^f |v|$. For any $K_i\psi \in \Sigma_\phi$, assume that $\mathcal{M}, w \Vdash K_i\psi$. By our definition, we get that $\mathcal{M}, v \Vdash K_i\psi$, but since $\mathcal{M}$ is a relational pseudo-model and thus $\leq_i$ is reflexive, we have that $\mathcal{M}, v \Vdash \psi$ (relational soundness of the (T) axiom).

  Once again, if in the previous paragraph we mark "$\leq$" wherever it reads "$\leq_i$," and "D" wherever it reads "K_i," we can convince ourselves that the according condition holds for $\leq^f$, as well.

$\square$

The specific filtration that we have chosen has desirable features, and the definition of Σ_ϕ will play a role in ensuring that we have built a finite pseudo-model for TKD, as intended.

Lemma 5.2.24. *Let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be a relational pseudo-model for TKD, and ϕ a formula of $\mathcal{L}_K^{\mathcal{D}}$. The filtration $\mathcal{M}^f$ of $\mathcal{M}$ through Σ_ϕ, as constructed above, meets the following conditions:*

- *For a fixed $i \in \{1, 2\}$, for every $w, v \in X$ such that $|w| \leq_i^f |v|$, there exists $v' \in X$ such that $w \leq_i v'$ and $|v'| = |v|$.*
- *For every $w, v \in X$ such that $|w| \leq^f |v|$, there exists $v' \in X$ such that $w \leq v'$ and $|v'| = |v|$.*

Proof.

- Fix $i \in \{1, 2\}$. Assume that $|w| \leq_i^f |v|$. Let $v(\Sigma_\phi)' := \bigwedge\{\theta \in \Sigma_\phi \mid \mathcal{M}, v \Vdash \theta\}$.

  Since Σ_ϕ is finite up to logical equivalence, we have that $v(\Sigma_\phi)'$ is logically equivalent to a finite conjunction of formulas in Σ_ϕ. Let this finite conjunction be denoted by $v(\Sigma_\phi)$. Since Σ_ϕ is closed under conjunction, we have that $v(\Sigma_\phi)$ lies in Σ_ϕ, as well. Notice that $\mathcal{M}, v \Vdash v(\Sigma_\phi)$, by construction.

  Claim. $\mathcal{M}, w \Vdash \langle K_i \rangle v(\Sigma_\phi)$.

  In order to prove this claim, we assume towards a contradiction that $\mathcal{M}, w \nVdash \langle K_i \rangle v(\Sigma_\phi)$. Therefore, we have that $\mathcal{M}, w \Vdash \neg\langle K_i \rangle v(\Sigma_\phi)$, which implies that $\mathcal{M}, w \Vdash K_i \neg v(\Sigma_\phi)$. On the other side, we mentioned above that the fact that Σ_ϕ is taken to be closed under conjunction ensures that $v(\Sigma_\phi) \in \Sigma_\phi$; since Σ_ϕ is also closed under negation, we have that $\neg v(\Sigma_\phi) \in \Sigma_\phi$.

  Since $\neg v(\Sigma_\phi) \in \Sigma_\phi$, $\mathcal{M}, w \Vdash K_i \neg v(\Sigma_\phi)$ implies with $|w| \leq_i^f |v|$ that $\mathcal{M}, v \Vdash K_i \neg v(\Sigma_\phi)$. Since $\mathcal{M}$ is a relational pseudo-model and thus $\leq_i$ is reflexive, $\mathcal{M}, v \Vdash K_i \neg v(\Sigma_\phi)$ implies that $\mathcal{M}, v \Vdash \neg v(\Sigma_\phi)$ as well (relational soundness of the (T) axiom). But this contradicts the previously reviewed assumption that $\mathcal{M}, v \Vdash v(\Sigma_\phi)$.

  Therefore, the *Claim* is true and $\mathcal{M}, w \Vdash \langle K_i \rangle v(\Sigma_\phi)$.

  The *Claim* implies that there exists $v' \in X$ such that $w \leq_i v'$ and $\mathcal{M}, v' \Vdash v(\Sigma_\phi)$. It is easy to see that $\mathcal{M}, v' \Vdash v(\Sigma_\phi)$ implies that $|v'| = |v|$: for every $\psi \in \Sigma_\phi$, if $\mathcal{M}, v \Vdash \psi$, we know that there is a subformula ψ' of $v(\Sigma_\phi)$ which is logically equivalent to ψ and such that $\mathcal{M}, v' \Vdash \psi'$; on the other direction, if $\mathcal{M}, v \nVdash \psi$, then $\mathcal{M}, v \Vdash \neg\psi$, but $\neg\psi \in \Sigma_\phi$ (the latter is closed under negation), so that there exists a subformula θ of $v(\Sigma_\phi)$ which is logically equivalent to $\neg\psi$ and such that $\mathcal{M}, v' \Vdash \theta$, which happens iff $\mathcal{M}, v' \Vdash \neg\psi$ iff $\mathcal{M}, v' \nVdash \psi$. Therefore, for every $\psi \in \Sigma_\phi$, $\mathcal{M}, v \Vdash \psi$ iff $\mathcal{M}, v' \Vdash \psi$, which means that $|v'| = |v|$.

  With all this, we have shown that the desired condition holds.

- The proof of this item is analogous to the above one:

  Assume that $|w| \leq^f |v|$. Let $v(\Sigma_\phi)' := \bigwedge\{\theta \in \Sigma_\phi \mid \mathcal{M}, v \Vdash \theta\}$.

  Since Σ_ϕ is finite up to logical equivalence, we have that $v(\Sigma_\phi)'$ is logically equivalent to a *finite* conjunction of formulas in Σ_ϕ. Let this finite conjunction be denoted by $v(\Sigma_\phi)$. Since Σ_ϕ is closed under conjunction, we have that $v(\Sigma_\phi)$ lies in Σ_ϕ, as well. Notice that $\mathcal{M}, v \Vdash v(\Sigma_\phi)$, by construction.

  Claim. $\mathcal{M}, w \Vdash \langle D \rangle v(\Sigma_\phi)$.

  Once again, we assume towards a contradiction that $\mathcal{M}, w \nVdash \langle D \rangle v(\Sigma_\phi)$. Therefore, we have that $\mathcal{M}, w \Vdash \neg\langle D \rangle v(\Sigma_\phi)$, which implies that $\mathcal{M}, w \Vdash D \neg v(\Sigma_\phi)$. Since Σ_ϕ is closed under conjunction, we have that $v(\Sigma_\phi) \in \Sigma_\phi$; since Σ_ϕ is also closed under negation, we have that $\neg v(\Sigma_\phi) \in \Sigma_\phi$.

  Since $\neg v(\Sigma_\phi) \in \Sigma_\phi$, $\mathcal{M}, w \Vdash D \neg v(\Sigma_\phi)$ implies with $|w| \leq^f |v|$ that $\mathcal{M}, v \Vdash D \neg v(\Sigma_\phi)$. Since $\mathcal{M}$ is a relational pseudo-model, and thus $\leq$ is reflexive, $\mathcal{M}, v \Vdash D \neg v(\Sigma_\phi)$ implies that $\mathcal{M}, v \Vdash \neg v(\Sigma_\phi)$ (relational soundness of the (T_D) axiom). But this contradicts the previously reviewed assumption that $\mathcal{M}, v \Vdash v(\Sigma_\phi)$.

  Therefore, the *Claim* is true and $\mathcal{M}, w \Vdash \langle D \rangle v(\Sigma_\phi)$.

  $\mathcal{M}, w \Vdash \langle D \rangle v(\Sigma_\phi)$ implies there exists $v' \in X$ such that $w \leq v'$ and $\mathcal{M}, v' \Vdash v(\Sigma_\phi)$. By the same argument as in the previous item, $\mathcal{M}, v' \Vdash v(\Sigma_\phi)$ implies that $|v'| = |v|$.

$\square$

Lemma 5.2.25. *$\mathcal{M}^f$ is a finite relational pseudo-model for TKD.*

Proof. That X^f is finite comes precisely from the fact that Σ_ϕ is finite up to logical equivalence.

Once that has been settled, it all boils down to showing yet again that our structure meets the relevant conditions that characterize relational pseudo-models.

i) Preorders and equivalence relations

- $\sim_1^f, \sim_2^f$, and $\sim^f$ are equivalence relations, straight from their definition.
- It is easy to see that $\leq_1^f, \leq_2^f$ and $\leq^f$ are all reflexive and transitive.

ii) Pertinent inclusions

- The inclusions $\leq_1^f \subseteq \sim_1^f$, $\leq_2^f \subseteq \sim_2^f$, and $\leq^f \subseteq \sim^f$ all hold by definition.
- ($\leq^f \subseteq \leq_1^f \cap \leq_2^f$). Let $w, v \in X$ such that $|w| \leq^f |v|$. We want to show that $|w| \leq_1^f |v|$ and that $|w| \leq_2^f |v|$. For a fixed $i \in \{1, 2\}$, assume that $K_i\psi \in \Sigma_\phi$ and that $\mathcal{M}, w \Vdash K_i\psi$. By relational soundness of the (4) axiom, we have that $\mathcal{M}, w \Vdash K_iK_i\psi$. By relational soundness of the (DsK) axiom, we have that $\mathcal{M}, w \Vdash DK_i\psi$, but let us not forget that Σ_ϕ was constructed in such a way that $DK_i\psi \in \Sigma_\phi$ as well. Therefore, since we have assumed that $|w| \leq^f |v|$, $\mathcal{M}, w \Vdash DK_i\psi$ implies that $\mathcal{M}, v \Vdash DK_i\psi$. By relational soundness of the (T_D) axiom, $\mathcal{M}, v \Vdash DK_i\psi$ implies that $\mathcal{M}, v \Vdash K_i\psi$ as well. Thus, we have seen that for every $K_i\psi \in \Sigma_\phi$, $\mathcal{M}, w \Vdash K_i\psi$ implies that $\mathcal{M}, v \Vdash K_i\psi$, which shows that $|w| \leq_i^f |v|$. Since i was arbitrary, we have both that $|w| \leq_1^f |v|$ and that $|w| \leq_2^f |v|$.
- We want to show that $\sim^f \subseteq \sim_1^f \cap \sim_2^f$. From the item above, we know that $\leq^f \subseteq \leq_1^f \cap \leq_2^f$. But we have also mentioned that the inclusions $\leq_1^f \subseteq \sim_1^f$ and $\leq_2^f \subseteq \sim_2^f$ both hold. Therefore, we have that $\leq^f \subseteq \leq_1^f \cap \leq_2^f \subseteq \sim_1^f \cap \sim_2^f$. Since $\sim^f$ is defined as the $\subseteq$-least equivalence relation including $\leq^f$, $\leq^f \subseteq \sim_1^f \cap \sim_2^f$ implies that $\sim^f \subseteq \sim_1^f \cap \sim_2^f$.

iii) Density condition iii) $(denk)$

We will show that $\mathcal{M}^f$ meets condition a) in Lemma 5.2.19, so that in virtue of said Lemma, it will also satisfy condition iii) for relational pseudo-models.

Fix $i \in \{1, 2\}$. We want to show that for every $w, v, u \in X$ such that $|w| \leq_i^f |v|$ and $|w| \leq_i^f |u|$, there exists $o \in X$ such that $|v| \leq_i^f |o|$ and $|u| \leq^f |o|$.

Let $w, v, u \in X$ such that $|w| \leq_i^f |v|$ and $|w| \leq_i^f |u|$. By Lemma 5.2.24, we know that there exist $v', u' \in X$ such that $w \leq_i v'$, $w \leq_i u'$, and $|v'| = |v|$, $|u'| = |u|$. Since $\mathcal{M}$ is a relational pseudo-model and thus meets condition a) in Lemma 5.2.19 (see Proposition 5.2.20 iii)), we have that the assumptions $w \leq_i v'$ and $w \leq_i u'$ imply that there exists $o \in X$ such that $v' \leq_i o$ and $u' \leq o$. Since $\mathcal{M}^f$ is a filtration (of pseudo models) of $\mathcal{M}$ through Σ_ϕ, the *"Forth" condition* for $\leq_i^f$ yields that $|v'| \leq_i^f |o|$ and $|u'| \leq^f |o|$. But since $|v'| = |v|$, $|u'| = |u|$, we get the desired conclusion.

iv) Density condition iv) $(dend)$

Exactly the same as the previous item. We will show that $\mathcal{M}^f$ meets condition b) in Lemma 5.2.19, so that in virtue of said Lemma, it will also satisfy condition iv) for relational pseudo-models.

We want to show that for every $w, v, u \in X$ such that $|w| \leq^f |v|$ and $|w| \leq^f |u|$, there exists $o \in X$ such that $|v| \leq^f |o|$ and $|u| \leq^f |o|$.

Let $w, v, u \in X$ such that $|w| \leq^f |v|$ and $|w| \leq^f |u|$. By Lemma 5.2.24, we know that there exist $v', u' \in X$ such that $w \leq v'$, $w \leq u'$, and $|v'| = |v|$, $|u'| = |u|$. Since $\mathcal{M}$ is a relational pseudo-model and thus meets condition b) in Lemma 5.2.19 (see Proposition 5.2.20 iv)), we have that the assumptions $w \leq v'$ and $w \leq u'$ imply that there exists $o \in X$ such that $v' \leq o$ and $u' \leq o$. Since $\mathcal{M}^f$ is a filtration (of pseudo models) of $\mathcal{M}$ through Σ_ϕ, the *"Forth" condition* for $\leq^f$ yields that $|v'| \leq^f |o|$ and $|u'| \leq^f |o|$. But since $|v'| = |v|$, $|u'| = |u|$, we get the desired situation.

v) $\mathcal{V}^f$ is defined as a valuation function.

□

Directly from our Proposition 5.2.13, we have

Corollary 5.2.26. *For every $w \in X$ and every formula ψ of $\mathcal{L}_K^{\mathcal{D}}$, $\mathcal{M}, w \Vdash \psi$ iff $\mathcal{M}^f, |w| \Vdash \psi$.*

□

Corollary 5.2.27. *$\mathbf{\Lambda}$ is weakly complete with respect to the class of finite relational pseudo-models.*

Proof. Let ϕ_* be a formula of $\mathcal{L}_K^{\mathcal{D}}$ such that ϕ_* is $\mathbf{\Lambda}$-consistent. From *Step 1*, we know that there exists a relational pseudo-model $\mathcal{M}_*$ and a world w_* in its domain such that $\mathcal{M}_*, w_* \Vdash \phi_*$. We construct the set Σ_{ϕ_*} as before, and consider our particular filtration of $\mathcal{M}_*$ through Σ_{ϕ_*}, which we accordingly denote $\mathcal{M}_*^f$. By previous argumentation, we know that the domain of this filtration is finite (Observation 10). Then, from Corollary 5.2.26, we immediately get that

$$\mathcal{M}_*^f, |w_*| \Vdash \phi_*.$$

□

Step 3. Completeness with respect to (finite) intermediary pseudo-models

Lemma 5.2.28. *Any finite relational pseudo-model is a bounded morphic image of a (finite) intermediary pseudo-model.*

Proof. Let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be a finite relational pseudo-model.

Let $\{c_0, \ldots, c_{m-1}\}$ denote an enumeration without repetition of the set $X/\sim$ (the quotient of X under $\sim$).

We build up a structure

$$\mathcal{M}^p = (X \times \{c_0, \ldots, c_{m-1}\}, \sim_1^p, \sim_2^p, \sim^p, \leq_1^p, \leq_2^p, \leq^p, \mathcal{V}^p),$$

where the relations $\sim_1^p, \sim_2^p, \sim^p, \leq_1^p, \leq_2^p, \leq^p$ are defined by the following rules:

For two arbitrary elements $(w, c_s), (v, c_t)$ of $X \times \{c_0, \ldots, c_{m-1}\}$, we have

$(w, c_s) \sim_1^p (v, c_t)$ iff $s = t$ and $w \sim_1 v$.

$(w, c_s) \sim_2^p (v, c_t)$ iff $w \sim_2 v$ and $s - t \equiv lw - lv \pmod{m}$, where lw and lv are the unique indexes such that $w \in c_{lw}$ and $v \in c_{lv}$.

$(w, c_s) \sim^p (v, c_t)$ iff $s = t$ and $w \sim v$.

$(w, c_s) \leq_1^p (v, c_t)$ iff $(w, c_s) \sim_1^p (v, c_t)$ and $w \leq_1 v$.

$(w, c_s) \leq_2^p (v, c_t)$ iff $(w, c_s) \sim_2^p (v, c_t)$ and $w \leq_2 v$.

$(w, c_s) \leq^p (v, c_t)$ iff $(w, c_s) \sim^p (v, c_t)$ and $w \leq v$.

We define the valuation function $\mathcal{V}^p : Prop \to \mathcal{P}(X \times \{c_0, \ldots, c_{m-1}\})$ by $\mathcal{V}^p(p) = \{(w, c_l) \mid w \in \mathcal{V}(p)\}$.

Claim 1. $\mathcal{M}^p$ is a finite intermediary pseudo-model.

We have to show that $\mathcal{M}^p$ meets the five conditions of relational pseudo-models, that its domain is finite, and that, additionally, the equality $\sim^p = \sim_1^p \cap \sim_2^p$ holds.

Before anything, we observe that $X \times \{c_0, \ldots, c_{m-1}\}$ is finite, since X is finite, by assumption.

(Pseudo-Model Conditions)

i) Preorders and equivalence relations

- It is straightforward to realize that $\sim_1^p$ and $\sim^p$ are equivalence relations, since *reflexivity, symmetry,* and *transitivity* all follow easily from the according conditions for $\sim_1$ and $\sim$, respectively. What can cause a bit of doubt is whether $\sim_2^p$ is also an equivalence relation or not. But we show that it is in fact one:
  - *(Reflexivity of $\sim_2^p$)* For any (w, c_s), we have that $w \sim_2 w$. Pick the unique $\sim$-class of w. As above, let us denote it with c_{lw}. We have that $s - s \equiv lw - lw \pmod{m}$, rendering that $(w, c_s) \sim_2^p (w, c_s)$.
  - *(Symmetry of $\sim_2^p$)* Let $(w, c_s) \sim_2^p (v, c_t)$. Then $w \sim_2 v$ and $s - t \equiv lw - lv \pmod{m}$, where lw and lv are the unique indexes such that $w \in c_{lw}$ and $v \in c_{lv}$. Since $\sim_2$ is symmetric, we get that $v \sim_2 w$ as well. Observe that $-(s - t) \equiv -(lw - lv) \pmod{m}$, so that $t - s \equiv lv - lw \pmod{m}$. With these two, we get that $(v, c_t) \sim_2^p (w, c_s)$.
  - *(Transitivity of $\sim_2^p$)* Let $(w, c_s) \sim_2^p (v, c_t)$ and $(v, c_t) \sim_2^p (u, c_r)$. First and foremost, we can see that this implies that $w \sim_2 u$. Now, notice that our assumptions entail, on one hand, that $(s - t) - (lw - lv)$ is a multiple of m, let us say ma (for some $a \in \mathbb{Z}$), and that, on the other, $(t - r) - (lv - lu)$ is a multiple of m as well, let us say mb (for some $b \in \mathbb{Z}$). Behold that $(s-r)-(lw-lu) = s-t+t-r-lw+lv-lv+lu = (s-t)-(lw-lv)+(t-r)-(lv-lu) = ma + mb = m(a + b)$. Therefore, $(s - r) - (lw - lu)$ is a multiple of m as well, so that $(s - r) \equiv (lw - lu) \pmod{m}$ and hence $(w, c_s) \sim_2^p (u, c_r)$ holds.
- It is relatively easy to realize that $\leq_1^p, \leq_2^p, \leq^p$, as defined, are reflexive and transitive. Both conditions follow from reflexivity and transitivity of $\sim_1^p, \sim_2^p, \sim^p$, and reflexivity and transitivity of $\leq_1, \leq_2, \leq$.

ii) Pertinent inclusions

- The construction of $\leq_1^p, \leq_2^p, \leq_p$ guarantees the facts that $\leq_i^p \subseteq \sim_i^p$ for both $i \in \{1, 2\}$, and that $\leq^p \subseteq \sim^p$.
- For convenience, we will first show that $\sim^p \subseteq \sim_1^p \cap \sim_2^p$. Assume that $(w, c_s) \sim^p (v, c_t)$. This means that $s = t$ and $w \sim v$ (which obviously implies that $lw = lv$). Since $\sim \subseteq \sim_1 \cap \sim_2$ by assumption, we have that $w \sim_1 v$ and $w \sim_2 v$. Together with the fact that $s = t$ (and thus that $s - t \equiv lw - lv \pmod{m}$), we immediately get both that $w \sim_1^p v$ and that $w \sim_2^p v$.
- Now we will show that $\leq^p \subseteq \leq_1^p \cap \leq_2^p$. Assume that $(w, c_s) \leq^p (v, c_t)$. This means that $(w, c_s) \sim^p (v, c_t)$ and $w \leq v$. By the above item, we get that $(w, c_s) \sim_1^p (v, c_t)$ and that $(w, c_s) \sim_2^p (v, c_t)$. Since we have that $\leq \subseteq \leq_1 \cap \leq_2$ ($\mathcal{M}$ is a relational pseudo-model), we also have that $w \leq_1 v$ and that $w \leq_2 v$, so that retrieving the according pertinent information, we can see both that $w \leq_1^p v$ and that $w \leq_2^p v$.

iii) Density condition iii) (*denk*)

- *(For $\sim_1^P$)* Let $(w, c_s), (v, c_t) \in X \times \{c_0, \dots, c_{m-1}\}$ such that $(w, c_s) \sim_1^p (v, c_t)$. This means that $s = t$ and $w \sim_1 v$. By assumption, we have that there exists $u \in X$ such that $w \leq_1 u$ and $v \leq u$ ($\mathcal{M}$ is a relational pseudo-model). Notice that $w \leq_1 u$ implies that $w \sim_1 u$ ($\leq_1$ is included in $\sim_1$). Then the element (u, c_s) is such that $(w, c_s) \sim_1^p (u, c_s)$ and $w \leq_1 u$. Hence, $(w, c_s) \leq_1^p (u, c_s)$. Likewise, $v \leq u$ implies that $v \sim u$ ($\leq$ is included in $\sim$), so that (u, c_s) is such that $(v, c_t) \sim^p (u, c_s)$ and $v \leq u$, which means that $(v, c_t) \leq^p (u, c_s)$.
- *(For $\sim_2^P$)* Let $(w, c_s), (v, c_t) \in X \times \{c_0, \dots, c_{m-1}\}$ such that $(w, c_s) \sim_2^p (v, c_t)$. Then $w \sim_2 v$ and $s - t \equiv lw - lv \pmod{m}$. By assumption, we have that there exists $u \in X$ such that $w \leq_2 u$ and $v \leq u$ ($\mathcal{M}$ is a relational pseudo-model). First, notice that $v \leq u$ implies that $v \sim u$, which means that $lv = lu$. This trivially implies that $s - t \equiv lw - lu \pmod{m}$. Secondly, regard that $w \leq_2 u$ implies that $w \sim_2 u$ ($\leq_2$ is included in $\sim_2$). Therefore, (u, c_t) is such that $(w, c_s) \sim_2^p (u, c_t)$ and $w \leq_2 u$, but this means that $(w, c_s) \leq_2^p (u, c_t)$. On the other side, since $v \sim u$, we have that $(v, c_t) \sim^p (u, c_t)$, which with $v \leq u$ implies that $(v, c_t) \leq^p (u, c_t)$.

iv) Density condition iv) (*dend*)

Let $(w, c_s), (v, c_t) \in X \times \{c_0, \ldots, c_{m-1}\}$ such that $(w, c_s) \sim^p (v, c_t)$. This means that $s = t$ and $w \sim v$. By assumption we have that there exists $u \in X$ such that $w \leq u$ and $v \leq u$ ($\mathcal{M}$ is a relational pseudo-model). Notice that we have that $w \sim u$ and that $v \sim u$, ($\leq$ is included in $\sim$). Then, the element (u, c_s) is such that $(w, c_s) \sim^p (u, c_s)$ and $w \leq u$, which means that $(w, c_s) \leq^p (u, c_s)$. In the same manner, we can easily grasp that $(v, c_t) \sim^p (u, c_s)$, so that $v \leq u$ implies that $(v, c_t) \leq^p (u, c_s)$.

v) A redundant assertion, we point out that $\mathcal{V}^p$ is clearly a valuation function.

Up to now, we have shown that the structure $\mathcal{M}^p$ is a relational pseudo-model. The only thing that remains to acknowledge in order to convince ourselves that it is in fact an intermediary pseudo-model is that $\sim^p \supseteq \sim_1^p \cap \sim_2^p$, so that using ii) above we would get that $\sim^p = \sim_1^p \cap \sim_2^p$. Let us show such inclusion. Let $(w, c_s), (v, c_t) \in X \times \{c_0, \ldots, c_{m-1}\}$ such that $(w, c_s) \sim_1^p (v, c_t)$ and $(w, c_s) \sim_2^p (v, c_t)$. From the first assumption we get that $s = t$, but coupling this with $(w, c_s) \sim_2^p (v, c_t)$, we get that $|lw - lv|$ is a multiple of m bounded by $m - 1$. Therefore, $|lw - lv|$ must be equal to 0, so that $lw = lv$. But this means that w and v lie within the same $\sim$-class, so that $w \sim v$. Therefore, we have that $s = t$ and $w \sim v$, rendering that $(w, c_s) \sim^p (v, c_t)$.

Claim 2. The map $h : \mathcal{M}^p \to \mathcal{M}$ given by $h((w, c_s)) = w$ is a bounded morphism.

- Let $(w, c_s) \in X \times \{c_0, \ldots, c_{m-1}\}$. Then $\mathcal{M}^p, (w, c_s) \Vdash p$ iff $(w, c_s) \in \mathcal{V}^p(p)$ iff $w \in \mathcal{V}(p)$ iff $\mathcal{M}, w \Vdash p$. Thus, (w, s) and $h((w, s))$ satisfy the same propositional letters.
- *(Forth conditions)*

  Straight from our definitions, we have that for a fixed $i \in \{1, 2\}$, $(w, c_s) \leq_i^p (v, c_t)$ implies that $w \leq_i v$. Likewise, $(w, c_s) \leq^p (v, c_t)$ implies that $w \leq v$.
- *(Back conditions)*
  - Assume that $h((w, c_s)) = w \leq_1 v$ (and recall that this assumption implies that $w \sim_1 v$). Notice then that (v, c_s) is such that $(w, c_s) \sim_1^p (v, c_s)$ and $w \leq_1 v$. Therefore, we have that $(w, c_s) \leq_1^p (v, c_s)$, with $h((v, c_s)) = v$.
  - Assume that $h((w, c_s)) = w \leq_2 v$ (and recall that this assumption implies that $w \sim_2 v$). Let $r = s - lw + lv$. Notice then that $(v, c_{r \bmod m})$ is such that $s - (r \bmod m) \equiv s - r \pmod m$, but $s - r = lw - lv$. Therefore, $s - (r \bmod m) \equiv lw - lv \pmod m$, so that $w \sim_2 v$ implies that $(w, c_s) \sim_2^p (v, c_{r \bmod m})$. With the additional assumption of $w \leq_2 v$, we get that $(w, c_s) \leq_2^p (v, c_{r \bmod m})$, where $h((v, c_{r \bmod m})) = v$, clearly.
  - Assume that $h((w, c_s)) = w \leq v$ (and recall that this assumption implies that $w \sim v$). Notice that (v, c_s) is such that $(w, c_s) \sim^p (v, c_s)$ and $w \leq v$. Therefore, $(w, c_s) \leq^p (v, c_s)$, with $h((v, c_s)) = v$.

Therefore, $h : \mathcal{M}^p \to \mathcal{M}$ is a bounded morphism. With *Claim 1.* and Claim 2. thus shown, we have concluded the proof of our Lemma. □

Directly from our Proposition 5.2.11, we have

Corollary 5.2.29. *For a given finite relational pseudo-model for TKD, of the form $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$, we have that its according structure $\mathcal{M}^P$, built as in the proof of Lemma 5.2.28, is such that for every formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ and $w \in X$,*

$$\mathcal{M}^p, (w, c_s) \Vdash \phi \text{ iff } \mathcal{M}, w \Vdash \phi.$$

□

Corollary 5.2.30. $\mathbf{\Lambda}$ *is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of intermediary pseudo-models.*

Proof. Let ϕ_* be a formula of $\mathcal{L}_K^{\mathcal{D}}$ such that ϕ_* is $\mathbf{\Lambda}$-consistent. We know from *Step 2* that there exists a finite relational pseudo-model $\mathcal{M}_*$ and a world w_* in its domain such that $\mathcal{M}_*, w_* \Vdash \phi_*$. We construct the according structure $\mathcal{M}_*^p$. From Corollary 5.2.29, we immediately get that

$$\mathcal{M}_*^p, (w_*, c_0) \Vdash \phi_*.$$

□

Step 4. Completeness with respect to (finite) relational models

Before moving on, we introduce a couple of conventions regarding notation.
Let $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ be a finite relational pseudo-model.

- For any given relation R on X and any subset $A \subseteq X$, we write $MAX_R(A)$ to refer to the set of R-*maximal* elements in A; i.e., $MAX_R(A) = \{u \in A \mid \forall v \in A,\ uRv \ \Rightarrow\ vRu\}$. Likewise, we write $TOP_R(A)$ to refer to the set of R-*top* elements in X; i.e., $TOP_R(A) = \{u \in X \mid \forall v \in A,\ vRu\}$.
- For a fixed $i \in \{1, 2\}$ and $w \in X$, we will again use the notation $w[i]$ to refer to the class of w under the equivalence relation $\sim_i$. In this line, $w[d]$ will denote the class of w under the equivalence relation $\sim$.

Observation 11. If $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ is an intermediary pseudo-model, then $w[d] = w[1] \cap w[2]$ for every $w \in X$. It is very easy to realize this, because the assumption that $\mathcal{M}$ is an intermediary pseudo-model means that $\sim = \sim_1 \cap \sim_2$.

Lemma 5.2.31. *If $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$ is a finite intermediary pseudo-model, then the following conditions hold:*

a) For a fixed $i \in \{1, 2\}$ and a fixed $w \in X$,

$$MAX_{\leq}(w[i]) = TOP_{\leq_i}(w[i]) = MAX_{\leq_i}(w[i]). \tag{5.5}$$

b) For a fixed $w \in X$,

$$MAX_{\leq}(w[1] \cap w[2]) = TOP_{\leq}(w[1] \cap w[2]). \tag{5.6}$$

Proof. First and foremost, let us agree that the assumption that $\mathcal{M}$ is finite gives us that the sets of *maximal* elements addressed are non-empty.

a) Fix $i \in \{1, 2\}$ and $w \in X$. We observe that the condition "$\leq \subseteq \leq_i$" implies that $MAX_{\leq_i}(w[i]) \subseteq MAX_{\leq}(w[i])$. Now take $m \in MAX_{\leq}(w[i])$ and an arbitrary $w' \in w[i]$ ($w \sim_i w'$). By the assumed satisfaction of condition iii) (*denk*), we know that there exists $m' (\in w[i])$ such that $w' \leq_i m'$ and $m \leq m'$. But this means that $m = m'$, and the reason is that m is taken to be a $\leq$-maximal in $w[i]$. Therefore, we have shown that for all $m \in MAX_{\leq}(w[i])$, $m \geq_i w'$ for all $w' \in w[i]$, which means that $MAX_{\leq}(w[i]) \subseteq TOP_{\leq_i}(w[i])$. But it should be very clear that $TOP_{\leq_i}(w[i]) \subseteq MAX_{\leq_i}(w[i])$. Because of all of this, we get the following chain of inclusions:

$$MAX_{\leq}(w[i]) \subseteq TOP_{\leq_i}(w[i]) \subseteq MAX_{\leq_i}(w[i]) \subseteq MAX_{\leq}(w[i]).$$

So we have that

$$MAX_{\leq}(w[i]) = TOP_{\leq_i}(w[i]) = MAX_{\leq_i}(w[i]).$$

b) Fix $w \in X$. Take $m \in MAX_{\leq}(w[1] \cap w[2])$ and an arbitrary $w' \in w[1] \cap w[2]$ ($w \sim_1 w'$ and $w \sim_2 w'$). By Observation 11, we get that w' actually lies within $w[d]$, so that $w \sim w'$. By the assumed satisfaction of condition iv) (*dend*), we know that there exists $m' (\in w[d])$ such that $w' \leq m'$ and $m \leq m'$. But m is taken to be a $\leq$-maximal in $w[d] = w[1] \cap w[2]$, so that $m = m'$. Therefore, we have shown that for all $m \in MAX_{\leq}(w[1] \cap w[2])$, $m \geq_i w'$ for all $w' \in w[1] \cap w[2]$, which means that $MAX_{\leq}(w[1] \cap w[2]) \subseteq TOP_{\leq}(w[1] \cap w[2])$. The other inclusion is pretty straightforward.

□

Lemma 5.2.32. *Any finite intermediary pseudo-model is a bounded morphic image of a (finite) relational model.*

Proof. For any finite intermediary pseudo-model, of the form $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$, we will consider the structure

$$\mathcal{M}^a = (X \times \{1,2\}, \sim_1^a, \sim_2^a, \sim^a, \leq_1^a, \leq_2^a, \leq^a, \mathcal{V}^a),$$

where the relations are defined by the following rules:

For a fixed $i \in \{1,2\}$, $(w,s) \sim_i^a (v,t)$ iff $w \sim_i v$.

$(w,s) \sim^a (v,t)$ iff $w \sim v$.

For a fixed $i \in \{1,2\}$, $(w,s) \leq_i^a (v,t)$ iff $\begin{cases} s = t \ \wedge \ w \leq v & \text{or} \\ t = i \ \wedge \ w \leq_i v & \text{or} \\ v \in MAX_{\leq}(w[i]). & \end{cases}$

$(w,s) \leq^a (v,t)$ iff $\begin{cases} s = t \ \wedge \ w \leq v & \text{or} \\ v \in MAX_{\leq}(w[1] \cap w[2]). & \end{cases}$

We define the valuation funtion $\mathcal{V}^a : \mathit{Prop} \to \mathcal{P}(X \times \{1,2\})$ by $\mathcal{V}^a(p) = \{(w,r) \mid w \in \mathcal{V}(p)\}$.

Claim 1. *$\mathcal{M}^a$ is a finite relational* ***model*** *for TKD.*

First of all, we have to show that $\mathcal{M}^a$ is a relational pseudo-model and that, moreover, it is an intermediary pseudo-model, so that $\sim^a$ is exactly equal to $\sim_1^a \cap \sim_2^a$. Finally, we will need to guarantee that the missing key ingredient that would make our structure an actual *model* also holds. Namely, we have to ascertain that $\leq^a$ is exactly equal to $\leq_1^a \cap \leq_2^a$.

Let us start by stressing the fact that $\mathcal{M}^a$ is finite, since it is built from a finite intermediary pseudo-model, so that X is finite and hence $X \times \{1,2\}$ is finite.

(Pseudo-model conditions)

i) <u>Preorders and equivalence relations</u>

- $\sim_1^a, \sim_2^a$ and $\sim^a$ are obviously equivalence relations, straight from their definition.
- Let us show that $\leq_1^a$ and $\leq_2^a$ are preorders. Fix $i \in \{1,2\}$.
  - *(Reflexivity of $\leq_i^a$)* Let $(w,s) \in X \times \{1,2\}$. Since $w \leq_i w$ by the reflexivity of $\leq_i$, we have that $(w,s) \leq_i^a (w,s)$.
  - *(Transitivity of $\leq_i^a$)* Let $(w,s), (v,t), (u,r) \in X \times \{1,2\}$ such that $(w,s) \leq_i^a (v,t)$ and $(v,t) \leq_i^a (u,r)$. We have the following cases:
    - (Case "$t = s$")
      - (Sub-case "$r = t$") Here, we know for sure both that $w \leq v$ and that $v \leq u$. Therefore, since $\leq$ is transitive, we have that $w \leq u$, and thus that $(w,s) \leq_i^a (u,r)$.
      - (Sub-case "$r = i$") Here, we know for sure that $w \leq v$ and that $v \leq_i u$. Since $\leq \subseteq \leq_i$, we also have that $w \leq_i v$. Since $\leq_i$ is transitive, we get that $w \leq_i u$, and thus that $(w,s) \leq_i^a (u,r)$.
      - (Sub-case "$u \in MAX_{\leq}(v[i])$") Notice that the assumptions $\leq_i \subseteq \sim_i$ and $(w,s) \leq_i^a (v,t)$ imply that $v[i] = w[i]$. Therefore, $u \in MAX_{\leq}(w[i])$, so we have that $(w,s) \leq_i^a (u,r)$ straightforwardly.
    - (Case "$t = i$")
      - (Sub-case "$r = t$") Here, we know for sure that $w \leq_i v$ and that $v \leq u$. Since $\leq \subseteq \leq_i$, we also have that $v \leq_i u$. Since $\leq_i$ is transitive, we get that $w \leq_i u$, and thus that $(w,s) \leq_i^a (u,r)$.

- (Sub-case "$r = i$") Here, we know for sure that $w \leq_i v$ and that $v \leq_i u$. Since $\leq_i$ is transitive, we get that $w \leq_i u$, and thus that $(w, s) \leq^a_i (u, r)$.
- (Sub-case "$u \in MAX_{\leq}(v[i])$") Once again, notice that $v[i] = w[i]$, so we have that $(w, s) \leq^a_i (u, r)$ straightforwardly.

∗ (Case "$v \in MAX_{\leq}(w[i])$")
Lemma 5.2.31 *a*) ensures that

$$MAX_{\leq_i}(w[i]) = MAX_{\leq}(w[i]).$$

It is not difficult to see, then, that the fact that $v \in MAX_{\leq}(w[i])$, together with the assumption that $(v, t) \leq^a_i (u, r)$, implies that $u \in MAX_{\leq}(w[i])$:

- If $v \leq u$ is the sub-case, then $\leq$-maximality of v implies the $\leq$-maximality of u;
- If $v \leq_i u$ is the sub-case, then Lemma 5.2.31 *a*) gives that the $\leq_i$-maximality of v implies the $\leq_i$-maximality of u, and thus the $\leq$-maximality of u;
- The last sub-case gives immediately that $u \in MAX_{\leq}(w[i])$.

Therefore, we get that in this case $u \in MAX_{\leq}(w[i])$, so that $(w, s) \leq^a_i (u, r)$.

- Now, let us show that $\leq^a$ is a preorder.
  - *(Reflexivity of $\leq^a$)* Let $(w, s) \in X \times \{1, 2\}$. Since $w \leq w$ by the reflexivity of $\leq$, we have that $(w, s) \leq^a (w, s)$.
  - *(Transitivity of $\leq^a$)* Let $(w, s), (v, t), (u, r) \in X \times \{1, 2\}$ such that $(w, s) \leq^a (v, t)$ and $(v, t) \leq^a (u, r)$. We have the following cases:
    ∗ (Case "$t = s$")
      - (Sub-case "$r = t$") Here, we know for sure both that $w \leq v$ and that $v \leq u$. Therefore, since $\leq$ is transitive, we have that $w \leq u$, and thus that $(w, s) \leq^a (u, r)$.
      - (Sub-case "$u \in MAX_{\leq}(v[1] \cap v[2])$") Once again, notice that the assumption $(w, s) \leq^a (v, t)$ implies that $v[d] = w[d]$, so that $v[1] \cap v[2] = w[1] \cap w[2]$. Then, we have that $u \in MAX_{\leq}(w[1] \cap w[2])$, and thus that $(w, s) \leq^a (u, r)$.

    ∗ (Case "$v \in MAX_{\leq}(w[1] \cap w[2])$")
    It is clear that the assumptions $v \in MAX_{\leq}(w[1] \cap w[2])$ and $(v, t) \leq^a (u, r)$ imply that $u \in MAX_{\leq}(w[1] \cap w[2])$ as well. Therefore, we get that $(w, s) \leq^a (u, r)$ straightforwardly.

ii) <u>Pertinent inclusions</u>

- If we carefully review the cases provided by our definitions, it is relatively easy to realize that the inclusions $\leq^a_1 \subseteq \sim^a_1$, $\leq^a_2 \subseteq \sim^a_2$, and $\leq^a \subseteq \sim^a$ all hold.
- ($\leq^a \subseteq \leq^a_1 \cap \leq^a_2$) Assume that $(w, s) \leq^a (v, t)$.
  - (Case "$t = s$") We know for sure that $w \leq v$ holds. Since $\leq \subseteq \leq_i$ for both $i \in \{1, 2\}$, we have that $w \leq_1 v$ and $w \leq_2 v$ hold as well. Therefore, $t = s$ implies both that $(w, s) \leq^a_1 (v, t)$ and that $(w, s) \leq^a_2 (v, t)$.
  - (Case "$v \in MAX_{\leq}(w[1] \cap w[2])$") For this case, it should be clear that $v \in MAX_{\leq}(w[1])$ and that $v \in MAX_{\leq}(w[2])$. Thus, we immediately get both that $(w, s) \leq^a_1 (v, t)$ and that $(w, s) \leq^a_2 (v, t)$.
- The assumed inclusion $\sim \subseteq \sim_1 \cap \sim_2$ implies that $\sim^a \subseteq \sim^a_1 \cap \sim^a_2$.

iii) <u>Density condition iii) (*denk*)</u>

Fix $i \in \{1, 2\}$. Let $(w, s), (v, t) \in X \times \{1, 2\}$ such that $(w, s) \sim^a_i (v, t)$. This implies that $w \sim_i v$, and therefore that $w[i] = v[i]$. Since X is finite, we know that $MAX_{\leq}(v[1] \cap v[2]) \neq \emptyset$. Let $u \in MAX_{\leq}(v[1] \cap v[2])$. Notice that $MAX_{\leq}(v[1] \cap v[2]) = MAX_{\leq}(v[1]) \cap MAX_{\leq}(v[2])$. Recall that i is fixed, so *in particular*, we have that $u \in MAX_{\leq}(v[i]) = MAX_{\leq}(w[i])$. Straight from our definitions, we get that $(w, s) \leq^a_i (u, r)$ and that $(v, t) \leq^a (u, r)$, for some arbitrarily chosen $r \in \{1, 2\}$.

iv) Density condition iv) (*dend*)
Let $(w,s),(v,t) \in X \times \{1,2\}$ such that $(w,s) \sim^a (v,t)$. This implies that $w \sim v$, and therefore that $w[1] \cap w[2] = w[d] = v[d] = v[1] \cap v[2]$. Since X is finite, we know that $MAX_{\leq}(w[1] \cap w[2]) \neq \emptyset$. Let $u \in MAX_{\leq}(w[1] \cap w[2])$. By the aforementioned equality of cells, we have that $u \in MAX_{\leq}(v[1] \cap v[2])$, so that obviously $(w,s) \leq^a (u,r)$ and $(v,t) \leq^a (u,r)$, for some arbitrarily chosen $r \in \{1,2\}$.

v) Granting that this is redundant, we point out that $\mathcal{V}^a$ is a valuation function.

(Model conditions)

I. That $\sim^a = \sim_1^a \cap \sim_2^a$ comes directly from the equality $\sim = \sim_1 \cap \sim_2$, which holds due to the assumption that these relations are part of an intermediary pseudo-model.

S. We have already shown that $\leq^a \subseteq \leq_1^a \cap \leq_2^a$. For the $\supseteq$ inclusion, let $(w,s),(v,t) \in X \times \{1,2\}$ such that $(w,s) \leq_1^a (v,t)$ and $(w,s) \leq_2^a (v,t)$. We have the following cases:

- (Case "$t = s$") Assume without loss of generality that $t = s = 1$. In this case, we get that the assumption $(w,1) \leq_1^a (v,1)$ implies either that $w \leq v$ or that $w \leq_1 v$. In any case, we have that $v \in w[1]$. On the other hand, the assumption $(w,1) \leq_2^a (v,1)$ implies either that $w \leq v$ or that $v \in MAX_{\leq}(w[2])$. If $w \leq v$ holds, then it is trivial to see that $(w,1) \leq^a (v,1)$ holds as well. If $v \in MAX_{\leq}(w[2])$, then the safe guarantee that v lies within $w[1]$ implies that actually $v \in MAX_{\leq}(w[1] \cap w[2])$ (for any u such that $v \leq u$, $u \in v[d] = v[1] \cap v[2] = w[1] \cap w[2] \subseteq w[2]$, so that $u \leq v$ as well). Therefore, we have that $(w,1) \leq^a (v,1)$.
- (Case "$t \neq s$") Assume without loss of generality that $t = 1$ and $s = 2$. The assumption $(w,1) \leq_1^a (v,2)$ implies that $v \in MAX_{\leq}(w[1])$. On the other hand, the assumption $(w,1) \leq_2^a (v,2)$ implies either that $w \leq_2 v$ or that $v \in MAX_{\leq}(w[2])$. In any case, we have that $v \in w[2]$, so that, similarly as above, $v \in MAX_{\leq}(w[1])$ implies that actually $v \in MAX_{\leq}(w[1] \cap w[2])$, thus rendering that $(w,1) \leq^a (v,2)$.

Claim 2. The map $h : \mathcal{M}^a \to \mathcal{M}$ given by $h((w,s)) = w$ is a bounded morphism.

- Let $(w,s) \in X \times \{1,2\}$. Then $\mathcal{M}^a,(w,s) \Vdash p$ iff $(w,s) \in \mathcal{V}^a(p)$ iff $w \in \mathcal{V}(p)$ iff $\mathcal{M}, w \Vdash p$. Thus, (w,s) and $h((w,s))$ satisfy the same propositional letters.
- *(Forth conditions)*
  - Fix $i \in \{1,2\}$. Assume that $(w,s) \leq_i^a (v,t)$.
    - (Case "$t = s$") In this case, we know for sure that $w \leq v$, but since $\leq \subseteq \leq_i$, we also have that $w \leq_i v$.
    - (Case "$t = i$") In this case, we immediately get that $w \leq_i v$.
    - (Case "$v \in MAX_{\leq}(w[i])$") We know from Lemma 5.2.31 *a)* that
$$MAX_{\leq}(w[i]) = TOP_{\leq_i}(w[i]),$$
so that $w \leq_i v$, straightforwardly.
  - Assume that $(w,s) \leq^a (v,t)$.
    - (Case "$t = s$") In this case, $w \leq v$ holds straightforwardly.
    - (Case "$v \in MAX_{\leq}(w[1] \cap w[2])$") Lemma 5.2.31 *b)* ensures that
$$MAX_{\leq}(w[1] \cap w[2]) = TOP_{\leq}(w[1] \cap w[2]),$$
so that $w \leq v$.
- *(Back conditions)*
  - Fix $i \in \{1,2\}$. Assume that $h((w,s)) = w \leq_i v$. By definition, we have that $(w,s) \leq_i^a (v,i)$. On the other hand, regard that $h((v,i)) = v$.

– Assume that $h((w,s)) = w \leq v$. By definition, we have that $(w,s) \leq^a (v,s)$. Once again, it is clear that $h((v,i)) = v$.

With this, we have shown that h is a bounded morphism. With *Claim 1.* and *Claim 2.* thus shown, we can conclude that the arbitrary finite intermediary pseudo-model with which we started is a bounded morphic image of a relational model. The Lemma has been proved.

□

By our Proposition 5.2.11, we have

Corollary 5.2.33. *For a given finite intermediary pseudo-model of the form $\mathcal{M} = (X, \sim_1, \sim_2, \sim, \leq_1, \leq_2, \leq, \mathcal{V})$, we have that its according structure $\mathcal{M}^a$, built as in the proof of Lemma 5.2.32, is such that for every formula ϕ of $\mathcal{L}_K^{\mathcal{D}}$ and $w \in X$,*

$$\mathcal{M}^a, (w,s) \Vdash \phi \text{ iff } \mathcal{M}, w \Vdash \phi.$$

□

Corollary 5.2.34. $\mathbf{\Lambda}$ *is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of (finite) relational models for TKD.*

Proof. Let ϕ_* be a formula of $\mathcal{L}_K^{\mathcal{D}}$ such that ϕ_* is $\mathbf{\Lambda}$-consistent. We know from *Step 3* that there exists a finite intermediary pseudo-model $\mathcal{M}_*$ and a world w_* in its domain such that $\mathcal{M}_*, w_* \Vdash \phi_*$. If we construct the according structure $\mathcal{M}_*^a$, then Corollary 5.2.33 immediately renders that

$$\mathcal{M}_*^a, (w_*, 1) \Vdash \phi_*.$$

□

5.2.3 Back to Topology

All the work has been done by now. We recover all of our results and their natural corollaries in this brief subsection.

Proposition 5.2.35. *The proof system $\mathbf{\Lambda}_{TKD}$ is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of finite special Alexandrov models.*

Proof. As tradition dictates, we proceed by contraposition. Let ϕ_0 be a formula of $\mathcal{L}_K^{\mathcal{D}}$ such that $\nvdash_{\mathbf{\Lambda}} \phi_0$. By our result of completeness with respect to finite relational models (Corollary 5.2.34), we have that there exists a finite relational model for TKD, of the form $\mathcal{M}_0^{rel} = (X^0, \sim_1^0, \sim_2^0, \leq_1^0, \leq_2^0, \leq^0, \mathcal{V}^0)$, and a world $w_0 \in X^0$ such that $\mathcal{M}_0^{rel}, w_0 \nVdash \phi_0$. We construct the topological structure $\mathcal{M}_0^{top} = (X^0, \Pi_1^0, \Pi_2^0, \tau_{\leq_1^0}, \tau_{\leq_2^0}, \tau_{\leq^0}, \mathcal{V}^0)$, which is in fact the image of $\mathcal{M}_0^{rel}$ under the correspondence H^{-1} defined in the proof of Proposition 5.2.7. By that same Proposition, we have that $\mathcal{M}_0^{top}$ is a finite *special Alexandrov* model such that

$$\mathcal{M}_0^{top}, w_0 \nvDash \phi_0.$$

□

Corollary 5.2.36. *The proof system $\mathbf{\Lambda_{TKD}}$ is weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of (finite) two-agent topological models.*

□

Corollary 5.2.37. *The logic TKD has the finite model property.*

□

With all this, we are now fully entitled to affirm that our Main (and only) Theorem (4.3.1) has been proved:

The proof system $\mathbf{\Lambda}_{TKD}$ is sound and weakly complete for $\mathcal{L}_K^{\mathcal{D}}$ with respect to the class of two-agent topological models. □

Chapter 6

Discussion & future work

This small chapter addresses some of the conceptual subtleties that were left unsaid in previous discourse. At the same time, it presents an intuitive road for our research to take.

6.1 Analysis of knowledge (connection to the *Defeasability Theory*).

In Chapter 5, we introduced a particular semantic interpretation for knowledge. As it turned out, such knowledge is not exactly the same as the topo-evidential knowledge we talked about in Chapter 3. They are related, but there is an essential difference. Regardless of the information-partitions scheme, we observe that this essential difference lies in the idea of *learnable evidence.* Elaborating on the matter, we present a much more formal version of the underlying idea that helped us solve the technical problem for our first approach: the so-called 'learning models.'

Definition 6.1.1 (Single-agent Learning Models)**.** We will say that a tuple $(X, \mathcal{E}, \mathcal{E}', \mathcal{V})$ is a *learning model* for $Prop$ if and only if

- X is a set of possible worlds.
- $\mathcal{E} \subseteq \mathcal{P}(X)$ is a family of non-empty sets such that $X \in \mathcal{E}$. We will refer to the sets in $\mathcal{E}$ as *sub-basic current-evidence sets.*
- $\mathcal{E}' \subseteq \mathcal{P}(X)$ is a family of non-empty sets such that $\mathcal{E} \subseteq \mathcal{E}'$. We will refer to the members of $\mathcal{E}'$ as *sub-basic potential-evidence sets*
- $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

As seen in Chapter 3, every learning model can in principle give rise to an associated *bi-topo-e-model* $(X, \tau, \tau', \mathcal{V})$, where $\tau \subseteq \tau'$. Indeed, here we have taken $\tau = \tau_{\mathcal{E}}$ as the topology generated by the family $\mathcal{E}$ of sub-basic current evidence, and $\tau' = \tau_{\mathcal{E}'}$ as the topology generated by the family $\mathcal{E}'$ of sub-basic potential evidence.

All the concepts introduced in Chapter 3 for evidence models can be presented on learning models as well (using only the current evidence $\mathcal{E}$ and its associated topology τ in the definitions). In particular, the concept that we referred to as 'Stalnaker's knowledge' in Chapter 3 can be defined as before (as $K^s p = \Box p \wedge Bp$, where we have used the superindex ^{s} to distinguish it from the kind of knowledge we put forward in Chapter 4).

In such a learning framework, *our* knowledge turns out to be much more robust than Stalnaker's knowledge, for it involves *all* the *learnable* potential evidence that is available for the agent in the set of worlds. We can define this brand of knowledge as *irrevocable evidential knowledge*:

Definition 6.1.2. (Irrevocable evidential knowledge) Essentially, we require our knowledge to be monotonic under evidence addition. More precisely, knowledge could be thought of as belief based on a correct *justification*, but with the added requisite that such justification remains to be consistent with *any* potential evidence, not just with the current.[1] Let us call these justifications *robust justifications.* Therefore, we will say that an agent (irrevocably) *knows* ϕ at x iff there exists some *argument* $U \in \tau$ such that $x \in U$ (it is correct), $U \subseteq \|p\|$ (it supports p), and $U \cap U' \neq \emptyset$ for all $U' \in \tau'$ (it is a *robust justification*). In topological terms, we have that $x \in \|Kp\|$ iff $x \in int_\tau \|p\|$ and $int_\tau \|p\|$ is τ'-dense.

Notice that this is exactly the single-agent version of the individual knowledge that we used in Chapter 4 (Definition 4.2.2). Therefore, our knowledge could be equated to a kind of robust belief, that will not be given up upon learning *any* new evidence. It is intuitively clear why we say that it is *irrevocable*: the agent will keep believing the object of knowledge no matter what evidence (s)he might come to learn. In the context of an analysis in the *Defeasability* tradition, it is a strong kind of knowledge. The interesting thing, though, is that the sound and complete system of this irrevocable knowledge is still **S4.2**. On the other hand, Stalnaker's (FB) principle fails for the notion: $Bp \not\Rightarrow BKp$, so we cannot identify it with 'subjective certainty' (see Chapter 1).

As for the *multi-agent* setting, notice that the aforementioned quality of robustness played an important role in producing a well-behaved notion of distributed knowledge. Recall that Example 4.1.1 offered a situation in which the distributed (non-robust) justifications of a group failed to justify each agent's Stalnaker's knowledge, in the sense of $K_i^s p \not\Rightarrow Dp$. In contrast, the *irrevocable group knowledge* (based on the distributed *robust* justifications of the group) *does* behave like the traditional distributed knowledge, as shown by the soundness result of Chapter 5.[2]

Therefore, though conceptually maybe too strong for one's taste, the notion of irrevocable knowledge is technically convenient, to say the least. In order to achieve a slightly weaker notion, we can modify Definition 6.1.2 above and think of an evidential version of the so-called (in)defeasible knowledge:

Definition 6.1.3. ((In)defeasible knowledge) We will say that an agent (indefeasibly) *knows* ϕ at x iff there exists some *argument* $U \in \tau$ such that $x \in U$ (it is correct), $U \subseteq \|p\|$ (it supports p), and $U \cap U' \neq \emptyset$ for all $U' \in \tau'$ *such that* $x \in U$ (it is a *justification* that is *robust* with respect to *conclusive* potential evidence).

This evidential version of indefeasible knowledge identifies knowledge with beliefs that are supported by true evidence, and that will not be dropped upon receiving any new (true) evidence. However, the agent might revise her/his belief after receiving *false* evidence. This is a very interesting notion by itself, and it is in keep with the notion of indefeasible knowledge of [33] (see also [8] and [12]). It opens the door to a different evidential interpretation of knowledge that is at the same time not as weak as what we have called Stalnaker's knowledge ever since Chapter 3, and not as strong as the irrevocable knowledge we introduced and worked with in the multi-agent evidential case. However, when exploring the multi-agent correlate of this indefeasible kind of knowledge under the information-partitions criterion that we favored, we ran into some trouble again for the distributed version. It turns out that the fact that this (in)defeasible knowledge is non-monotonic with respect to the addition of arbitrary learnable evidence (i.e., possibly non-conclusive or false) does not fit well with the traditional concept of distributed knowledge. Evidently, one could hope for amendment in terms of its definition, and this is a possible line of work to follow.

6.2 Extensions.

Recall the rich language $\mathcal{L}_2$ that we presented at the beginning of Chapter 4, whose grammar is given by

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid K_1\phi \mid K_2\phi \mid B_1\phi \mid B_2\phi \mid D\phi \mid C\phi$$

The original work in this thesis is restricted to the fragment of $\mathcal{L}_2$ having only the operators K_1, K_2, and D. One of the most obvious and immediate subsequent steps to take in the direction of our proposal is

[1] We use the word 'justification' in the sense of the evidential treatment we have given to topo-evidential models (see Chapter 3, section 3.2).

[2] On a related note, we should point out that we abode by an interesting underlying assumption, that could in principle be dropped in future work. For simplicity, we assumed that all potential evidence does not depend on which agent is learning. While there are agent-dependent current-evidence topologies τ_i, there is only one, common potential-evidence topology τ', including both τ_1 and τ_2. However, one could think of a model in which there were two *different* topologies of potential evidence τ_1' and τ_2', according to each agent. We anticipate that, in order for the *same* arguments to work for the likely analogue of distributed knowledge in this case, a sufficient condition for these two topologies would be that *both* of them included both τ_1 and τ_2.

to tackle the problem of providing an all-encompassing (multi-agent) topo-evidential semantics that would include belief, common knowledge, and even *common belief* (see Chapter 2, section 2.6).

Bringing belief back into the picture

Throughout the previous chapters, we frequently mentioned that our multi-agent proposal takes into account the notion of belief, and allows for a standard formalization of the concept without much complication. In this sense, we present a natural extension of the semantics given in Definition 4.2.2. Observe that the conceptual arguments addressed at the beginning of this chapter pervade through these definitions.

Let $\mathcal{L}_{KB}^{\mathcal{D}}$ stand for the fragment (of the full language $\mathcal{L}_2$) having only the operators K_1, K_2, B_1, B_2, and D.

Definition 6.2.1. (Two-agent topological models)
A tuple $(X, \Pi_1, \Pi_2, \tau_1, \tau_2, \tau, \mathcal{V})$ will be called a *two-agent topological model* for the formulas of $\mathcal{L}_{KB}^{\mathcal{D}}$ if and only if

1. X is a set of possible worlds. Π_1, Π_2 are two agents' partitions on X. τ_1, τ_2, τ are topologies on X.
2. $\tau_1, \tau_2 \subseteq \tau$.
3. $\Pi_1 \subseteq \tau_1$, $\Pi_2 \subseteq \tau_2$.
4. $\mathcal{V} : Prop \to \mathcal{P}(X)$ is a valuation function.

We recursively define the topological semantics for the formulas of $\mathcal{L}_{KB}^{\mathcal{D}}$ by:

$$
\begin{array}{lcl}
\|p\| & = & \mathcal{V}(p) \\
\|\neg\phi\| & = & X \backslash \|\phi\| \\
\|\phi \wedge \psi\| & = & \|\phi\| \cap \|\psi\| \\
\|B_1\phi\| & = & \{x \in X \mid \exists U \in \tau_1^{x[1]} \text{ s.t. } U \subseteq \|\phi\| \wedge \ Cl_{\tau_1^{x[1]}}(U) = x[1]\} \\
\|B_2\phi\| & = & \{x \in X \mid \exists U \in \tau_2^{x[2]} \text{ s.t. } U \subseteq \|\phi\| \wedge \ Cl_{\tau_2^{x[2]}}(U) = x[2]\} \\
\|K_1\phi\| & = & \{x \in X \mid \exists U \in \tau_1^{x[1]} \text{ s.t. } x \in U \subseteq \|\phi\| \wedge \ Cl_{\tau^{x[1]}}(U) = x[1]\} \\
\|K_2\phi\| & = & \{x \in X \mid \exists U \in \tau_2^{x[2]} \text{ s.t. } x \in U \subseteq \|\phi\| \wedge \ Cl_{\tau^{x[2]}}(U) = x[2]\} \\
\|D\phi\| & = & \{x \in X \mid \exists U \in \tau_1^{x[1]}, V \in \tau_2^{x[2]} \text{ s.t. } x \in (U \cap V) \subseteq \|\phi\| \wedge \ Cl_{\tau^{x[1] \cap x[2]}}(U \cap V) = x[1] \cap x[2]\}.
\end{array}
$$

As usual, *truth*, *validity*, and *satisfaction* are defined in the standard way.

First and foremost, notice that belief, unlike knowledge, is defined in terms of the current-evidence topologies. Because of this, the proof of soundness of **KD45** for $\mathcal{L}_B$ with respect to the class of topo-e-models that we presented in Chapter 3, section 3.2.2, can be easily adapted into this multi-agent case (abiding by the information-partitions criterion). In this sense, we would get soundness of the fusion system **KD45+KD45** (for the fragment of the language having only the operators B_1, B_2) with respect to the class of two-agent topological models. As for completeness, we anticipate that **KD45+KD45** is also complete for the aforementioned fragment with respect to the class of two-agent topological models, but a precise proof remains to be done.

Secondly, one can show that, for a fixed $i \in \{1, 2\}$, the relation between the individual K_i-operators and the individual B_i-operators is governed by every axiom of Stalnaker's system **KB** *except* for (FB), which is not sound with respect to these models, as mentioned in the first section of the present chapter. Once again, it should be worthwhile to explore all these considerations thoroughly in the future, and we encourage the reader to give a chance to such an endeavor.

Common knowledge and common belief?

At the present stage of our inquiry, the attempts to provide a decent semantics for *common knowledge* within the lines of our paradigm are unsatisfactory, so we will not include them. However, it is a problem worth tackling. As discussed, common knowledge is an important notion of group knowledge, and a (traditionally) full semantics for multi-agent systems would obviously account for it.

The same can be said about *common belief*. As of now, how to couple this concept with the topological interpretations of belief we laid forward throughout this thesis is an open question.

Conclusion

As a closing remark to the present work, we can say that its main objective was accomplished, to a certain degree. We succeeded in providing a sound and complete axiomatization for a topological logic of distributed knowledge, where the interpretation of such knowledge accounts for a notion of belief arising from *evidence* in a possible-worlds framework.

Moreover, the concept of knowledge and of group knowledge that we obtained in this effort vouches for a (mostly) customary and well-behaved formal relation between the interpretations for knowledge and belief. As such, we can safely say that our work fares sufficiently well with ordinary tenets of epistemology.

We mentioned in the Introduction that our main target was the multi-agent case for a recently conceived theory of evidence-based belief and knowledge over topological models (i.e., the work of Baltag *et alia* in [4]). As it was, the process of tackling such a challenge demanded necessary amendments to the already existing paradigm, as well as the incorporation of intuitions from other fields of formal epistemology.

There are a couple of points worth mentioning.

First of all, we should admit to the fact that our venture was mainly technical. The core of this thesis—and of the research done to support it—lies in the long and complicated proof of completeness from Chapter 5. As it occurs sometimes, it is hard to say if what drove the conceptual shift that allowed us to solve our two predicaments was philosophical epistemological intuition (regarding the behavior of evidence), or mathematical perspective about the structures used and its underlying possible logics. In all likelihood, it was a combination of both things.[3] As such, the ensuing formalism is groundwork on both the conceptual dimension and the technical one.

On the one hand, the topological spaces underlying two-agent topological models are interesting in themselves as purely mathematical structures. They are bi-partioned, tri-topological spaces where the cells are selectively open in the topologies. Such structures are hard to deal with at first glance, but they are also very stimulating as candidates for spatial representation.

On the other hand, the idea of treating both belief and knowledge as dependent on a subjective partition on the possibility space, but at the same time arising from evidential support within the information cells, lays out an enormous amount of interpretational ground, propelling interesting and exotic discussions about the ontology of evidence, knowledge, and belief.

In conclusion, the author would like to express that the virtue of the present work, if any, could lie in its being taken as an incentive for trusting Mathematics as a promoter of philosophical advancement.

[3] Since it was Alexandru Baltag that offered such solutions, one could probably ask him to provide elucidation on the matter.

Bibliography

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic*, 50(02):510–530, 1985.

[2] R. J. Aumann. Agreeing to disagree. *The annals of statistics*, pages 1236–1239, 1976.

[3] R. J. Aumann. Interactive epistemology i: knowledge. *International Journal of Game Theory*, 28(3):263–300, 1999.

[4] A. Baltag, N. Bezhanishvili, A. Özgün, and S. Smets. The topology of uncertain evidence, justified belief, and fallible knowledge. work in progress.

[5] A. Baltag, N. Bezhanishvili, A. Özgün, and S. Smets. The topology of belief, belief revision and defeasible knowledge. In *Logic, Rationality, and Interaction*, pages 27–40. Springer, 2013.

[6] A. Baltag, N. Bezhanishvili, A. Ozgün, and S. Smets. The topology of full and weak belief. 2014.

[7] A. Baltag, N. Gierasimczuk, and S. Smets. On the solvability of inductive problems: a study in epistemic topology. *ILLC Prepublication Series PP-2015-13 (to appear in Proceedings of TARK 2015)*, 2015.

[8] A. Baltag, B. Renne, and S. Smets. The logic of justified belief change, soft evidence and defeasible knowledge. In *Logic, Language, Information and Computation*, pages 168–190. Springer, 2012.

[9] A. Baltag, B. Renne, and S. Smets. The logic of justified belief, explicit knowledge, and conclusive evidence. *Annals of Pure and Applied Logic*, 165(1):49–81, 2014.

[10] A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 165:5–21, 2006.

[11] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. *Texts in logic and games*, 3:9–58, 2008.

[12] A. Baltag, H. van Ditmarsch, and L. Moss. Epistemic logic and information update. *Handbook of the Philosophy of Information, Elsevier Science Publishers, Amsterdam*, pages 361–456, 2008.

[13] A. Baltag, H. van Ditmarsch, and L. Moss. Epistemic logic and information update. *Handbook of the Philosophy of Information, Elsevier Science Publishers, Amsterdam*, pages 361–456, 2008.

[14] J. Barwise. On the model theory of common knowledge. *The situation in logic*, 17:201–220, 1989.

[15] N. Bezhanishvili and W. Van Der Hoek. Structures for epistemic logic. In *Johan van Benthem on Logic and Information Dynamics*, pages 339–380. Springer, 2014.

[16] P. Blackburn, M. De Rijke, and Y. Venema. *Modal Logic: Graph. Darst*, volume 53. Cambridge University Press, 2002.

[17] O. Board. Dynamic interactive epistemology. *Games and Economic Behavior*, 49(1):49–80, 2004.

[18] R. Boddy. *Epistemic Issues and Group Knowledge*. PhD thesis, Universiteit van Amsterdam, 2014.

[19] R. Engelking. General topology. 1989.

[20] L. Esakia. Intuitionistic logic and modality via topology. *Annals of Pure and Applied Logic*, 127(1):155–170, 2004.

[21] R. Fagin, Y. Moses, J. Y. Halpern, and M. Y. Vardi. *Reasoning about knowledge.* MIT press, 2003.

[22] D. M. Gabbay. Many-dimensional modal logics: theory and applications. 2003.

[23] J. Gerbrandy. Distributed knowledge. In *Twendial*, volume 98, pages 111–124, 1998.

[24] E. L. Gettier. Is justified true belief knowledge? *analysis*, pages 121–123, 1963.

[25] R. Goldbach. *Modelling Democratic Deliberation.* PhD thesis, Universiteit van Amsterdam, 2015.

[26] A. Grove. Two modellings for theory change. *Journal of philosophical logic*, 17(2):157–170, 1988.

[27] J. Y. Halpern and R. Fagin. Modelling knowledge and action in distributed systems. *Distributed computing*, 3(4):159–177, 1989.

[28] J. Y. Halpern and Y. Moses. A guide to the modal logics of knowledge and belief: Preliminary draft. In *Proceedings of the 9th international joint conference on Artificial intelligence-Volume 1*, pages 480–490. Morgan Kaufmann Publishers Inc., 1985.

[29] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence*, 54(3):319–379, 1992.

[30] J. Y. Halpern, D. Samet, and E. Segev. Defining knowledge in terms of belief: The modal logic perspective. *The Review of Symbolic Logic*, 2(03):469–487, 2009.

[31] J. Hintikka. Knowledge and belief. 1962.

[32] K. Kelly. The logic of reliable inquiry. 1996.

[33] K. Lehrer and T. Paxson. Knowledge: Undefeated justified true belief. *The Journal of Philosophy*, pages 225–237, 1969.

[34] W. Lenzen. Recent work in epistemic logic. 1978.

[35] W. Lenzen. Epistemologische betrachtungen zu [s4, s5]. *Erkenntnis*, 14(1):33–56, 1979.

[36] D. Lewis. Convention: A philosophical study. 1969.

[37] D. K. Lewis. Counterfactuals. 1973.

[38] L. Lismont and P. Mongin. *On the logic of common belief and common knowledge.* Springer, 1997.

[39] J. C. C. McKinsey and A. Tarski. The algebra of topology. *Annals of mathematics*, pages 141–191, 1944.

[40] J.-J. C. Meyer and W. van der Hoek. Epistemic logic for ai and computer science, volume 41 of cambridge tracts in theoretical computer science, 1995.

[41] A. Özgün. *Topological models for belief and belief revision.* PhD thesis, Universiteit van Amsterdam, 2013.

[42] R. Parikh, L. S. Moss, and C. Steinsvold. Topology and epistemic logic. In *Handbook of spatial logics*, pages 299–341. Springer, 2007.

[43] D. Pearce and L. Uridia. The topology of common belief. In *AT*, pages 246–259, 2012.

[44] R. Stalnaker. The problem of logical omniscience, i. *Synthese*, 89(3):425–440, 1991.

[45] R. Stalnaker. Common ground. *Linguistics and philosophy*, 25(5):701–721, 2002.

[46] R. Stalnaker. On logics of knowledge and belief. *Philosophical studies*, 128(1):169–199, 2006.

[47] C. Steinsvold. *Topological models of belief logics*. ProQuest, 2007.

[48] J. van Benthem and G. Bezhanishvili. Modal logics of space. In *Handbook of spatial logics*, pages 217–298. Springer, 2007.

[49] J. van Benthem and E. Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1-3):61–92, 2011.

[50] J. Van Benthem and D. Sarenac. The geometry of knowledge. 2004.

[51] W. Van der Hoek and J. Meyer. *A complete epistemic logic for multiple agents*. Springer, 1997.

[52] W. Van der Hoek and M. Wooldridge. Multi-agent systems. *Foundations of Artificial Intelligence*, 3:887–928, 2008.

[53] F. Voorbraak and R. te Utrecht. *As far as I know: Epistemic logic and uncertainty*. Department of Philosophy, University of Utrecht, 1993.

[54] S. Willard. *General topology*. Courier Corporation, 2004.

[55] T. Williamson. *Knowledge and its Limits*. Oxford University Press, 2002.