Generic Models for Topological Evidence Logics

MSc Thesis *(Afstudeerscriptie)*

written by

Saúl Fernández González
(born on December 10th 1992 in Uviéu, Spain)

under the supervision of **Nick Bezhanishvili** and **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:
20th of September 2018

Members of the Thesis Committee:
Nick Bezhanishvili (supervisor)
Alexandru Baltag (supervisor)
Yde Venema
Benno van den Berg
Aybüke Özgün

Institute for Logic, Language and Computation

Abstract

This thesis studies several aspects of the topological semantics for evidence-based belief and knowledge introduced by Baltag, Bezhanishvili, Özgün, and Smets (2016).

Building on this work, we introduce a notion of *generic models*, topological spaces whose logic is precisely the sound and complete logic of topological evidence models. We provide generic models for the different fragments of the language.

Moreover, we give a multi-agent framework which generalises that of single-agent topological evidence models. We provide the complete logic of this framework together with some generic models for a fragment of the language. Finally, we define a notion of group knowledge which differs conceptually from previous approaches.

Contents

Acknowledgements

I wish to thank my supervisors Alexandru Baltag and Nick Bezhanishvili. Alexandru, working with you has been mesmerising. Thanks for throwing ideas at me at such a speed that I had to push my boundaries to be able to catch them. I also wish to thank you for the six 25cl pilseners, one 50cl pilsener, two IPA's and one white beer that you bought me in the process of writing this thesis. Thank you Nick for your expertise in topology but also for being a mentor to me. I thank you above all for trusting and believing in me as a mathematician. It has really made a difference.

I had very enlightening discussions with Eric Pacuit and Adam Bjorndahl about ideas which unfortunately didn't make it in this thesis. Special thanks go to Guram Bezhanishvili with whom I had two astoundingly fruitful meetings in which the ideas behind chapter 4 and some other parts of this thesis were conceived. He was answering my questions before I even had time to ask them. Section 2.1 also owes a lot to Guram.

I extend my gratefulness to the members of my thesis committee Yde Venema, Benno van den Berg and specially Aybüke Özgün, without whose contributions to the field my thesis would have probably been on some much less fun topic.

Other than that, I wish to thank Paßmann, Benjamins, O'Connell and the Adam-Days for letting me pick their brains about the most minute mathematical details. Such beautiful brains you people have.

Other acknowledgements in increasing order of cheesiness:

Pesetsky, I'm letting you proofread these acknowledgements merely to have the chance once again of yelling at each other about my excessive use of commas. Like most unhealthy things, your passion for research is fun and contagious, and I owe you a great deal for that. Thank you from the bottom of my heart for all the loudness. McHugh, you cheeky leprechaun, thank you for occasionally getting me close, despite my reticence, to seeing life the way you do. Möller Kalpak, thank you for the countless hours of reality TV, your invaluable friendship and your roommate-ness. You were always there. By 'there' I mean the living room.

Hoeks, I've reached peak cheesiness here, so I'm going with vagueness now. What I have to thank you for doesn't fit in these lines. You're incredible.

All errors are mine, includign the typo in this sentence.

When, in a given bedroom, you change the position of the bed,
can you say you are changing rooms, or else what?

(cf. topological analysis.)

– Georges Perec, *Species of Spaces*

Introduction

The present Master's thesis aims to take a further step towards bringing together topology and epistemic logic. In particular, it focuses on a certain topological semantics (the *dense interior semantics*) defined on a certain class of models based on topological spaces (*topological evidence models*), introduced in Baltag et al. (2016).

Epistemic logics (meaning the family of modal logics which study that which an epistemic agent *believes* or *knows*) found a modelisation in Hintikka (1962) in the form of Kripke frames, i.e. sets of possible worlds connected by (epistemic or doxastic) accessibility relations. Knowledge (K) and belief (B) are thus modal operators which are interpreted via standard possible worlds semantics.

Hintikka (1962) claims that the accessibility relation for knowledge must be (minimally) reflexive and transitive. On the syntax level, this demand translates to the fact that any logic for knowledge based on these frames must contain the axioms of $\mathsf{S4}$. And this, paired with the fact, famously proven by McKinsey and Tarski (1944), that $\mathsf{S4}$ is the logic of topological spaces under a certain semantics, lays the ground for a topological treatment of knowledge.

The semantics outlined in McKinsey and Tarski (1944) treats the "knowledge" modality as the interior operator, which, if one thinks of the open sets as "pieces of evidence", adds an evidential dimension to the notion of knowledge that one could not get within the framework of Kripke frames.

Under this interpretation, knowing a proposition amounts to having evidence for it. This can be an undesirable property. Depending on the properties one appoints to knowledge, belief and the relation thereof, one can get different epistemic logics, each with their axioms and rules. Following the precepts of Stalnaker (2006), a logic that allows us to talk about knowledge and belief, evidence (both "basic" and "combined") and a notion of *justification* is introduced in Baltag et al. (2016) and explored in depth in Özgün (2017), along with a class of models for this logic based on topological spaces: topo-e-models. The present pages study several aspects of this framework.

The novel contributions of this thesis can be sorted in two groups: on the one hand, we introduce a notion of *generic models* over a language $\mathcal{L}$, which are topological spaces whose logic is precisely the sound and complete $\mathcal{L}$-logic of topo-e-models, and provide several of these for the different fragments of the language. On the other hand, we bring to the table a notion of multi-agent topological evidence models which generalises the single-agent case and differs substantially from prior approaches. In this sense, we provide the several logics of multi-agent models and give some conceptual and theoretical contributions for a notion of group knowledge in this framework.

This thesis is structured as follows: in chapter 1, we introduce the notion of a topological space and provide the relevant technical results, leading up to the introduction of the framework of topological evidence models that is studied in the rest of the thesis.

Chapter 2 starts off by recalling the well-known theorem by McKinsey and Tarski (1944) relating the logic $\mathsf{S4}$ and the topological interior semantics on $\mathbb{R}$ and, following this spirit, introduces the notion of a *generic model.* We show that any dense-in-itself metrisable space such as $\mathbb{R}$ is a generic model for the knowledge fragment of the logic of topo-e-models and that certain such spaces, such as $\mathbb{Q}$ or the Cantor space, are generic models for other fragments.

In chapter 3 we introduce and explain our multi-agent models and provide their sound and complete logics. More saliently, we show that the logic of knowledge of these models is the fusion logic $\mathsf{S4.2} + \mathsf{S4.2}$.

Chapter 4 is concerned with providing generic models for this multi-agent logic. We show that the infinite quaternary tree $\mathcal{T}_{2,2}$ and the rational plane $\mathbb{Q} \times \mathbb{Q}$ are examples of such models.

Chapter 5 is an account of the notions of distributed and common knowledge. We introduce our proposal for the interpretation of distributed knowledge within this framework and outline how it differs conceptually from previous attempts. We provide a sound and complete logic of distributed knowledge and make some brief comments on common knowledge as well.

Finally, in chapter 6 we discuss the results obtained throughout the thesis, address some shortcomings and point out some directions in which to take the present work.

Chapter 1

Preliminaries

The present chapter introduces the technical preliminaries, concepts, definitions and results that will be used in the remainder of the thesis.

Section 1 introduces topological spaces, explains in which way these spaces are suited to model knowledge via the *interior semantics* and talks about McKinsey and Tarski's (1944) result relating this semantics with a widely accepted logic for knowledge.

In Section 2 we introduce several logics of knowledge and belief, and explain some of the shortcomings of the interior semantics, paving the way to the introduction, in section 3, of the framework that will be studied throughout the rest of the present thesis, namely that of *topological evidence models* with the *dense interior semantics*.

1.1 Topology and the interior semantics

Let us start off by recalling some of the basic concepts of topology.

Definition 1.1.1 (Topological space)**.** A *topological space* is a pair (X, τ), where X is a set and τ is a collection of subsets of X with the following properties:

i. $X \in \tau$ and $\varnothing \in \tau$;

ii. τ is closed under binary intersections: $U, V \in \tau$ implies $U \cap V \in \tau$;

iii. τ is closed under arbitrary unions: $\sigma \subseteq \tau$ implies $\bigcup \sigma \in \tau$.

A set $A \subseteq X$ is called *open* whenever $A \in \tau$, *closed* whenever $X \backslash A \in \tau$ and *clopen* whenever it is closed and open.

The *interior* of a set $A \subseteq X$ is the largest open set contained in A or equivalently

$$\operatorname{Int} A = \bigcup \{U \in \mathcal{P}(X) : U \in \tau \,\&\, U \subseteq A\}.$$

The *closure* of A is the least closed set which contains A, or

$$\operatorname{Cl} A = \bigcap \{U \in \mathcal{P}(X) : X \backslash U \in \tau \,\&\, A \subseteq U\}.$$

Note that the interior is always an open set, the closure is always a closed set and, moreover, A is open if and only if $A = \operatorname{Int} A$ and closed if and only if $A = \operatorname{Cl} A$. The following characterisation of interior and closure will be useful in what follows:

$$x \in \operatorname{Int} A \text{ iff } \exists U \in \tau(x \in U \subseteq A);$$
$$x \in \operatorname{Cl} A \text{ iff } \forall U \in \tau(x \in U \text{ implies } U \cap A \neq \varnothing).$$

Definition 1.1.2 (Bases and subbases)**.** A collection of sets $\mathcal{B} \subseteq \tau$ is a *basis of* τ if every open set can be expressed as a union of elements of $\mathcal{B}$ or, equivalently, if for every $U \in \tau$ and $x \in U$ there exists $B \in \mathcal{B}$ with $x \in B \subseteq U$.

A collection $\mathcal{S} \subseteq \tau$ is a *subbasis of* τ if the collection of finite intersections of elements of $\mathcal{S}$ forms a basis or, equivalently, if for every $U \in \tau$ and $x \in U$ there exist $S_1, ..., S_n \in \mathcal{S}$ with $x \in S_1 \cap ... \cap S_n \subseteq U$.

1.1.1 Some examples of topological spaces

The following are some particular topological spaces that will be used in subsequent chapters.

Example 1.1.3 (The real line)**.** Let $\mathbb{R}$ be the set of real numbers. We can define the *natural topology* on $\mathbb{R}$, $\tau_{\mathbb{R}}$ as the topology generated by the basis of open intervals

$$\mathcal{B} = \{(a, b) : a, b \in \mathbb{R}, a < b\}.$$

Equivalently, $U \subseteq \mathbb{R}$ is an open set if, for each $x \in U$, there exists some $\epsilon > 0$ such that $(x - \epsilon, x + \epsilon) \subseteq U$.

Definition 1.1.4 (Subspace topology)**.** Given a topological space (X, τ) and a set $Y \subseteq X$, we can define the *subspace topology* $\tau|_Y$ on Y as the set

$$\tau|_Y := \{U \cap Y : U \in \tau\}.$$

Note that $(Y, \tau|_Y)$ is trivially a topological space. We can use this in our next example:

Example 1.1.5 (The (ir)rational numbers)**.** The *natural topology* $\tau_{\mathbb{Q}}$ on the set of rational numbers $\mathbb{Q}$ is simply $\tau_{\mathbb{R}}|_{\mathbb{Q}}$ or, equivalently, the topology generated on $\mathbb{Q}$ by the basis of open intervals $\{(a, b) : a, b \in \mathbb{R}, a < b\}$, where $(a, b) = \{x \in \mathbb{Q} : a < x < b\}$.

The *natural topology* on the set of irrational numbers $\mathbb{I}$ can be defined in an analogous manner.

Example 1.1.6 (The Baire space and the Cantor space)**.** Let ω^ω be the set of infinite sequences of natural numbers, and ω^* be the set of finite such sequences. For $s \in \omega^*$ and $\alpha \in \omega^\omega$ we say $s \lhd \alpha$ whenever s is an *initial segment* of α, i.e., whenever $s = \langle s_1, ..., s_n \rangle$ with $s_i = \alpha(i)$ for $1 \leq i \leq n$. For $s \in \omega^*$, let $O(s)$ denote the set of sequences of natural numbers that have s as an initial segment, i.e. $O(s) = \{\alpha \in \omega^\omega : s \lhd \alpha\}$. The *Baire space* $\mathfrak{B} = (\omega^\omega, \tau_{\mathfrak{B}})$ is the topological space that has ω^ω as its underlying set together with the topology $\tau_{\mathfrak{B}}$ generated by the basis

$$\mathcal{B}_{\mathfrak{B}} = \{O(s) : s \in \omega^*\}.$$

We can analogously define the *Cantor space* $\mathfrak{C}$ on the set $\mathbf{2}^\omega$ of countable sequences of zeros and ones. The Cantor space has a nice visual representation in the form of the infinite binary tree. This is a tree whose nodes are the finite sequences of zeros and ones. It has the empty sequence as the root and each node $\langle i_1, ..., i_n\rangle \in \mathbf{2}^*$ has exactly two successors, namely $\langle i_1, ..., i_n, 0\rangle$ as its left successor and $\langle i_1, ..., i_n, 1\rangle$ as its right successor. The elements of the Cantor space can be identified with *branches* of this tree, where a branch is a countable collection of nodes $\{s_0, s_1, s_2, ...\}$ such that s_0 is the empty sequence (i.e. the root of the tree) and each s_{k+1} is an immediate successor of s_k. The basic open sets $O(s)$ are identified with "fans", each fan being the subtree that spurs from one node. An open set is any union of some of these fans. $\alpha \in \mathbf{2}^\omega$ is in a basic open set $O(s)$ whenever the corresponding branch "enters" the fan.

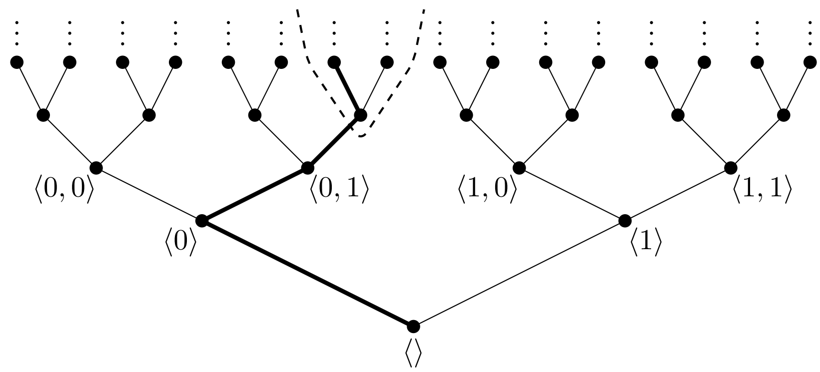

Figure 1.1: The Cantor space $\mathfrak{C}$. The dashed line represents the fan corresponding to $O(s)$, where $s = \langle 0, 1, 1\rangle$. The thick line represents a branch of the tree, corresponding to an infinite sequence $\alpha = \langle 0, 1, 1, 0, ...\rangle$. As we can see, the branch enters the fan and thus $\alpha \in O(s)$.

Example 1.1.7 (The binary tree $\mathcal{T}_2$)**.** If we consider the nodes of the binary tree in figure 1.1 instead of its branches to be the points of our space, we can give it a topology by setting the basic open sets to be $O(s)$, where $s = \langle a_0, ..., a_n\rangle$ and $t \in O(s)$ if t is a finite sequence of length greater than or equal to $n + 1$ with its $n + 1$ first elements being $a_0, ..., a_n$.

1.1.2 Preorders and Alexandroff topologies

Let us show the relation existing between topological spaces and preordered frames.

Definition 1.1.8 (Preordered frames). A *preordered frame* is a pair $(X, \leq)$ where X is a set and $\leq$ is a *preorder* on X, i.e. a reflexive and transitive binary relation defined on X.

A set $U \subseteq X$ is an *upwards closed set* (or an *upset* for short) if for all $x \in U$ and all $y \geq x$, we have that $y \in U$. It is *downwards closed* (or a *downset*) if $x \in U$ and $y \leq x$ imply $y \in U$.

The *upwards*- and *downwards-closure* of U are, respectively, the sets

$$\begin{aligned} \uparrow_{\leq} U :=& \{y \in X : x \leq y \text{ for some } x \in U\}; \\ \downarrow_{\leq} U :=& \{y \in X : y \leq x \text{ for some } x \in U\}. \end{aligned}$$

In the remainder of this text we will use the notation $\mathrm{Up}_{\leq}(X)$ to refer to the collection of all upsets of $(X, \leq)$. For a point $x \in X$, we will write $\uparrow_{\leq} x = \uparrow_{\leq}\{x\}$ and $\downarrow_{\leq} x = \downarrow_{\leq}\{x\}$. If there is no risk of ambiguity we will drop the $\leq$ in the subindices and write, for instance, $\mathrm{Up}(X), \uparrow x$.

A topology need not be closed under arbitrary intersection. If it is, (X, τ) is called an *Alexandroff* space. There is a 1-to-1 correspondence between preordered sets and Alexandroff spaces:

Lemma 1.1.9 (For details, see e.g. van Benthem and Bezhanishvili, 2007). *Let $(X, \leq)$ be a preordered set and $\tau_{\leq} := \mathrm{Up}(X)$. Then $\tau_{\leq}$ is an Alexandroff topology on X. Conversely, let (X, τ) be a topological space. The relation*

$$x \leq_{\tau} y \textit{ iff for all } U \in \tau (x \in U \textit{ implies } y \in U)$$

defines a preorder on X, called the specialisation preorder.

If $\leq$ is a preorder we have that $\leq_{\tau_{\leq}} = \leq$ and, if τ is an Alexandroff topology, $\tau_{\leq_{\tau}} = \tau$. Moreover, for any $U \subseteq X$

$$\mathrm{Int}_{\tau_{\leq}} U = \{x : \uparrow_{\leq} x \subseteq U\} \ \& \ \mathrm{Cl}_{\tau_{\leq}} U = \downarrow_{\leq} U.$$

For this reason, in the remainder of this thesis we will sometimes refer to preordered sets $(X, \leq)$ as topological spaces and, in this context, we will use the terms "open set" and "upset" interchangeably.

Before we proceed, let us use lemma 1.1.9 to look at p-morphisms from a topological lens. Given two relational structures (W, R) and (V, R') a *$p-$morphism* is a map $f : W \to V$ that satisfies:

i. *The forth condition*: Rxy implies $R'(fx)(fy)$;

ii. *The back condition*: $R'(fx)v$ implies that there exists $y \in W$ such that Rxy and $fy = v$.

But what are the topological counterparts of these conditions if the relations we are dealing with are preorders? It is straightforward to check that f satisfies the forth condition if and only if the preimage of an R'-open set is an R-open set. Similarly, the forth condition is satisfied if and only if the image of an R-open set is an R'-open set. These properties of a function f, in the more general context of topological spaces, are respectively called *continuity* and *openness.* A map which is open and continuous is called an *interior map.* In what follows, we will use the terms p-morphism and interior map interchangeably when we are dealing with Alexandroff topologies.

Definition 1.1.10 (Continuous map, open map, homeomorphism)**.** Let (X, τ) and (Y, σ) be two topological spaces. A map $f : X \to Y$ is called *continuous* if $A \in \sigma$ implies $f^{-1}A \in \tau$, and *open* if $A \in \tau$ implies $f[A] \in \sigma$.

We call f an *interior map* if it is open and continuous, and we call f a *p-morphism* if it is an interior map between Alexandroff spaces.

A bijective interior map is called a *homeomorphism.*

1.1.3 Interior semantics for modal logics

Let us work with a language $\mathcal{L}_\Box$ that includes a modal operator $\Box$ and can be defined recursively as follows:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box\phi,$$

where $p \in \mathsf{Prop}$, a countable set of propositional variables.

Given a topological space (X, τ) along with a *valuation* $V : \mathsf{Prop} \to \mathcal{P}(X)$ assigning sets of worlds to propositional variables, we have ourselves a *topological model* $\mathfrak{X} = (X, \tau, V)$ and we can define a notion of a formula ϕ being true at a world x recursively as follows:

$$\begin{aligned}
&\mathfrak{X}, x \vDash p \text{ iff } x \in V(p);\\
&\mathfrak{X}, x \vDash \neg\phi \text{ iff } \mathfrak{X}, x \nvDash \phi;\\
&\mathfrak{X}, x \vDash \phi \wedge \psi \text{ iff } \mathfrak{X}, x \vDash \phi \text{ and } \mathfrak{X}, x \vDash \psi;\\
&\mathfrak{X}, x \vDash \Box\phi \text{ iff } \exists U \in \tau(x \in U \ \&\ \mathfrak{X}, y \vDash \phi \text{ for all } y \in U).
\end{aligned}$$

By setting $\|\psi\|$ to be the set of worlds in which ψ holds, this definition gives us that $\|\Box\phi\| = \operatorname{Int}\|\phi\|$.

Note that, if we read $\Box\phi$ with the usual Kripke semantics on a relational model $\mathfrak{M} = (W, R, V)$, i.e.

$$x \in \|\Box\phi\| \text{ iff } y \in \|\phi\| \text{ for all } y \text{ such that } Rxy,$$

and if R happens to be a preorder on W, lemma 1.1.9 gives us that the set of worlds in which a formula is true is the same as if we read it with the interior semantics on $(W, \mathrm{Up}_R(W), V)$. And conversely, if $\mathfrak{X} = (X, \tau, V)$ is an Alexandroff topological model and $\leq_\tau$ is the specialisation preorder, then the interior semantics on $\mathfrak{X}$ coincides with the Kripke semantics on $(X, \leq_\tau, V)$.

This means that the interior semantics on topological spaces generalises the Kripke semantics on preordered frames. If we are reading $\Box$ as an epistemic operator, we can translate the semantics of Hintikka (1962) into this topological framework, with the addition that having a topological space allows us to have an *evidential* view of knowledge. Indeed, if we read $\Box$ as a knowledge modality, interpret the open sets in the topology to be pieces of evidence the agent has, and we say that P entails Q whenever $P \subseteq Q$, then the interior semantics defined above gives us that the agent *knows* ϕ whenever she has a piece of evidence which entails ϕ.

Let us get more specific and revisit some of the examples in subsection 1.1.1.

Example 1.1.11. Let us suppose our epistemic agent is an underfunded ornithologist attempting to determine the weight of a certain bird. Her devices of measurement are not particularly precise and produce results with a margin of error of ± 10g. Let us code the set of possible worlds with the positive real numbers $(0, \infty)$, where at world x the weight of the bird is precisely x grams. Now, suppose the actual world is $x_0 = 509$ and the ornithologist obtains a measurement of 500g $\pm$ 10g. Then the open interval $(490, 510)$ is her piece of evidence. With this, there are things she knows and things she does not know. She does not know, for instance, the proposition "the bird is heavier than 500g" to be true. She knows, however, that the bird is heavier than 400g. This proposition can be interpreted as the set of worlds $P = (400, \infty)$ and she has a piece of evidence which includes the actual world and entails this proposition: $x_0 \in (490, 510) \subseteq P$.

Example 1.1.12. Let us equate a world with an infinite stream of data, represented by a sequence of natural numbers. We are thus in our Baire space. Our epistemic agent this time is a scientist, and her evidence comes in the form of *observations*, which are finite streams of data that the scientist is able to grasp. A world is compatible with her observation whenever the stream of data is an initial segment of said world. If she observes $s = \langle a_1, ..., a_n \rangle$, then the set of worlds compatible with it (the corresponding *piece of evidence* in our sense) is precisely the basic open set $O(s)$.

In this setting, open sets correspond with *verifiable* propositions: if P is an open set and the actual world x_0 is in P, then there exist a basic open set $O(s)$ such that $x \in O(s) \subseteq P$, thus this scientist can potentially make an observation, s, which will allow her to know P. Similarly, closed sets correspond to *refutable* propositions and clopen sets to *decidable* propositions. For more details on this interpretation, see Kelly (1996).

1.1.4 McKinsey and Tarski: S4 as a topological logic of knowledge

Modelling knowledge as topological interior gives us an intuitive, evidence-based idea of what knowledge amounts to. But what is the logic of topological spaces?

Definition 1.1.13 (S4)**.** The logic S4 is the least set of formulas in the language $\mathcal{L}_\Box$ which contains all the propositional tautologies, is closed under uniform substitution and the rules of modus ponens (from $\phi \to \psi$ and ϕ infer ψ) and necessitation (from ϕ infer $\Box\phi$) and contains the axioms:

(K) $\Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$;

(T) $\Box\phi \to \phi$;

(4) $\Box\phi \to \Box\Box\phi$.

And the relevant result to this respect:

Theorem 1.1.14 (McKinsey and Tarski, 1944)**.** S4 *is sound and complete with respect to topological spaces under the interior semantics.*

This result sheds light on another reason why one would want to use the topological interior to model knowledge. The resulting logic S4 gives knowledge some desirable (and in general philosophically accepted) properties. The (T) axiom gives us factivity of knowledge: an agent only knows things which are true. Axiom (4) gives us positive introspection: if an agent knows something, she knows that she knows it.

Note that the soundness of these axioms corresponds to certain properties of the topological interior: for instance, axiom (T) corresponds to the fact that the interior of a set is contained in the set, axiom (K) to the fact that interior distributes over finite intersections and axiom (4) to the idempotence of the Int operator. Completeness can be derived from relational completeness with respect to preordered frames plus lemma 1.1.9.

But McKinsey and Tarski proved something else. We do not need to consider the class of all topological spaces to get the logic S4. There need not be a straightforward conceptual interpretation of knowledge relative to every possible topology. They showed that, instead, we can take some particular, "natural" topological space used to model knowledge whose logic is S4.

Definition 1.1.15 (Dense-in-itself space)**.** A topological space (X, τ) is *dense-in-itself* if no singleton is an open set, i.e., if $\{x\} \notin \tau$ for all $x \in X$.

Definition 1.1.16 (Metrisable space)**.** Given a set X a *metric* on X is a map $d : X \times X \to [0, \infty)$ satisfying for all $x, y, z \in X$:

i. $d(x, y) = 0$ iff $x = y$;

ii. $d(x, y) = d(y, x)$;

iii. $d(x, z) \leq d(x, y) + d(y, z)$.

A metric d on X induces a topology τ_d: we say that a set $U \subseteq X$ is open if, for every $x \in U$, there exists some $\epsilon > 0$ such that $d(x, y) < \epsilon$ implies $y \in U$.

A topological space (X, τ) is *metrisable* if there exists a metric d on X such that $\tau = \tau_d$.

Remark 1.1.17. All the spaces presented as examples in subsection 1.1.1 are both dense-in-itself and metrisable. The corresponding metric for $\mathbb{R}$, $\mathbb{Q}$ and $\mathbb{I}$ is $d(x, y) = |x - y|$, and clearly no singleton contains an open interval in these spaces. The binary tree $\mathcal{T}_2$ clearly has no open singletons and it is a regular space with a countable basis and thus metrisable. The fact that $\mathfrak{B}$ has these properties is a consequence of the fact that $\mathfrak{B}$ is homeomorphic to $\mathbb{I}$. Similarly, $\mathfrak{C}$ is homeomorphic to a dense-in-itself metrisable subspace of $\mathbb{R}$ (for details on these claims, see Munkres, 2000; Engelking, 1989).

Theorem 1.1.18 (McKinsey and Tarski, 1944). S4 *is the logic of any dense-in-itself metrisable space.*[1]

This is to say: whatever is provable in S4 is true at any world in any model based on (for example) $\mathbb{R}$ with the interior semantics and, conversely, whatever is true of the topology of $\mathbb{R}$ is provable in S4.

We thus have a semantics based on evidence that allows us to talk about knowledge and whose logic is a philosophically felicitous epistemic logic. Moreover, we have some specific spaces which provide "nice" ways to conceptualise knowledge and whose logic is still S4.

This semantics, however, is not the topic of this thesis. Instead, we will be working with the *dense interior* semantics. Understanding the conceptual reasons to move away from the interior and introducing this semantics is the aim of the next section.

1.2 Topological accounts of belief and knowledge: a pre-history

Before explaining and introducing the framework, let us bring forward some of the logics that will be mentioned in this section and throughout the rest of the thesis.

1.2.1 Logics for knowledge and belief

We will mention the axioms and rules of several logics and provide a class of relational frames with respect to which they are sound and complete. For details about these soundness and completeness results, see e.g. Blackburn, De Rijke, and Venema (2001).

All of the logics in this subsection contain the propositional tautologies, are closed under uniform substitution and the rules of modus ponens and necessitation and moreover contain the following axioms:

- As introduced in the previous subsection, S4 has the additional axioms

[1] The original formulation of this theorem talked about dense-in-itself, metrisable, *separable* spaces. Later work (e.g. Rasiowa and Sikorski, 1970) dropped the separability condition.

(K) $\Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$;
(T) $\Box\phi \to \phi$;
(4) $\Box\phi \to \Box\Box\phi$.

S4 is sound and complete with respect to the class of finite rooted preordered frames.

- S4.2 has the axioms and rules of S4 plus the axiom

  (.2) $\Diamond\Box\phi \to \Box\Diamond\phi$.

  S4.2 is sound and complete with respect to the class of finite rooted frames $(W, \leq)$ where $\leq$ is reflexive, transitive and *weakly directed* (i.e. $y \geq x \leq z$ implies that there exists $t \in W$ such that $y \leq t \geq z$).

- S5 has the axioms and rules of S4 plus the axiom

  (5) $\neg\Box\phi \to \Box\neg\Box\phi$.

  It is sound and complete with respect to the class of finite frames $(W, \sim)$ where $\sim$ is an equivalence relation.

- KD45 has the (K), (4) and (5) axioms plus

  (D) $\Box\phi \to \neg\Box\neg\phi$.

  It is sound and complete with respect to the class of finite frames $(W, \leq)$ where $\leq$ is reflexive, transitive and *Euclidean* (i.e. $y \geq x \leq z$ implies $y \leq z$).

The following subsection is a recap of section 3.3 and chapter 4 in Özgün (2017).

1.2.2 The road to the dense interior semantics

The relation between belief and knowledge has historically been a main focus of epistemology. One would want to have a formal system that accounts for knowledge and belief together, which requires careful consideration regarding the way in which they interact.

Canonically, knowledge has been thought of as "true, justified belief". However, Gettier's (1963) counterexamples of cases of true, justified belief which do not amount to knowledge shattered this paradigm.

Epistemologists like Williamson (2002) and Stalnaker (2006) propose, proverbially having dessert before dinner, to think of this issue from a "knowledge first" perspective: instead of starting off with a notion of belief and strengthening it to what amounts to some felicitous idea of knowledge, one could start with some idea of knowledge and define belief from it.

Stalnaker (2006) argues that a relational semantics is insufficient to capture Gettier's (1963) considerations and, trying to stay close to most of the intuitions of Hintikka (1962), provides an axiomatisation for a system of knowledge and belief. This system, Stal, has two modal operators, B and

K, and on top of the $\mathsf{S4}$ axioms and rules for K it adds the axioms in table 1.1:

(PI)	$B\phi \to KB\phi$;
(NI)	$\neg B\phi \to K\neg B\phi$;
(KB)	$K\phi \to B\phi$;
(CB)	$B\phi \to \neg B\neg\phi$;
(FB)	$B\phi \to BK\phi$.

Table 1.1: Extra axioms for Stal

In this logic, knowledge is an $\mathsf{S4.2}$ modality, belief is a $\mathsf{KD45}$ modality and the following formula can be proven:

$$B\phi \leftrightarrow \neg K\neg K\phi.$$

Within this system, in the words of Baltag, Bezhanishvili, Özgün, and Smets (2013), "believing p" is the same as "not knowing you don't know p". The following can also be proven in Stal:

$$B\phi \leftrightarrow BK\phi.$$

Belief then becomes "subjective certainty", in the sense that the agent cannot distinguish whether she believes or knows p, and believing amounts to believing that one knows.

An attempt to introduce belief in a topological framework was conducted by Steinsvold (2006). Belief is seen as the dual of the *derived set* operator,

$$\neg B\neg P = d(P) := \{x : \forall U \in \tau(x \in U \text{ implies } \exists y \in (P \cap U)\backslash\{x\})\}.$$

In other words, $x \in BP$ if and only if there exists $U \in \tau$ such that $x \in U \subseteq P \cup \{x\}$. Steinsvold proves that $\mathsf{KD45}$ is sound and complete for this semantics defined on *DSO spaces* (i.e. dense-in-itself spaces in which every $d(P)$ is open, *DSO* meaning "derived sets are open"). A problem with this approach is that the epistemic agent at world x will always believe a false proposition, namely $X\backslash\{x\}$. An even bigger problem is that, if one tries to account for this notion of belief plus knowledge as interior in the same setting, one gets the equivalence $K\phi \leftrightarrow \phi \wedge B\phi$: knowledge amounts to true belief, which clearly falls short.

Baltag et al. (2013) take a Stalnakerian stand and observe that, if knowledge is interior and the principle $B\phi \leftrightarrow \neg K\neg K\phi$ is to be accounted for, then belief has to be modelled as the closure of the interior operator, i.e.

$$BP = \mathrm{Cl}\,\mathrm{Int}\,P.$$

Then Baltag et al. (2013) prove that the logic Stal is actually sound and complete with this semantics for a particular class of topological spaces, namely *extremally disconnected* spaces, i.e. those spaces with the property that the closure of an open set is an open set.

If one wants to take this setting to the realm of Dynamic Epistemic Logic and introduce a notion of public announcement in it, one runs into a problem. Executing a public announcement amounts (semantically) to erasing some possible worlds, therefore we would be dealing with subspaces of our original space. And there is no guarantee that these subspaces of an extremally disconnected space will themselves be extremally disconnected. To make this setting dynamic, one has to work with *hereditarily extremally disconnected spaces*, i.e., spaces with the property that any subspace of them is extremally disconnected. Despite there is a theory of h.e.d. spaces, these seem rather difficult to find "in the wild", not particularly comfortable to work with and none of the "natural" spaces provided above as examples are h.e.d.

Enter Baltag et al. (2016). In this paper a new semantics is introduced, improving on the idea of *evidence models* of van Benthem and Pacuit (2011) which exploits the notion of evidence-based knowledge allowing to account for notions as diverse as *basic evidence* versus *combined evidence*, *factual, misleading* and *nonmisleading evidence*, etcetera. It is a semantics whose logic maintains a Stalnakerian spirit with regards to the relation between knowledge and belief, which behaves well dynamically and which does not confine us to work with "weird" species of spaces.

This is the *dense interior semantics*, defined on *topological evidence models.*

1.3 The logic of topological evidence models

All the definitions and results in the present section appear in Baltag et al. (2016).

Definition 1.3.1. A *topological evidence model* or *topo-e-model* is a tuple (X, τ, E_0, V), where (X, τ) is a topological space, E_0 is a subbasis for τ and V is a valuation.

The elements $e \in E_0$ represent the *basic pieces of evidence* the agent has, i.e., the evidence the agent has acquired directly through observation, measurements, etc. A *combined evidence* or an *argument* is an evidence the agent can piece together from her basic evidence, i.e. a nonempty finite intersection $e_0 \cap ... \cap e_n$ of pieces of basic evidence. We say an agent has a basic piece of evidence for a proposition P at world x whenever there exists some $e \in E_0$ such that $x \in e \subseteq P$, and we denote the set of worlds in which this is true as $\Box_0 P$. Similarly, we say that an agent has evidence for P at x if she has a *factive argument for* P, i.e. if there is a combined evidence $e_0 \cap ... \cap e_n$ such that $x \in e_0 \cap ... \cap e_n \subseteq P$. We denote the set of worlds in which this holds by $\Box P$. Note that, since the combined evidence constitutes a topological basis, $\Box P$ is exactly the topological interior of P.

We can make a distinction between an argument and a justification. A *justification* is an argument which is sufficiently strong, in the sense that it cannot be contradicted by any other argument. This takes us to the topological notion of *density*.

Definition 1.3.2 (Density)**.** Let (X, τ) be a topological space. A set $U \subseteq X$ is *dense* if $U = \varnothing$[2] or $\mathrm{Cl}\, U = X$. In other words, a nonempty set U is dense if and only if it has a nonempty intersection with every nonempty open set in the topology.

Formally, a *justification* is simply a dense argument. This is where the notions of knowledge and belief come in: we say that our agent *believes* P *at* x if she has a justification for P, i.e., if there exists a dense piece of evidence $U \subseteq P$; we say that the agent *knows* P *at* x if she has a *factive* justification for P, i.e. if $x \in U \subseteq P$. Knowledge in this setting amounts to *correctly justified belief*. We represent the set of worlds in which the agent believes/knows P respectively by BP and KP. Note that

$$BP = X \text{ if } \mathrm{Int}\, P \text{ is dense and nonempty, } \varnothing \text{ otherwise;}$$
$$KP = \mathrm{Int}\, P \text{ if } \mathrm{Int}\, P \text{ is dense, } \varnothing \text{ otherwise.}$$

Finally, we can think of the whole space X as the set of worlds which are consistent with the agent's information. In this sense, we can represent a notion of *infallible knowledge* by a global modality $[\forall]$, in the sense that an agent knows P infallibly whenever P holds in every world she considers possible. So $[\forall]P = X$ if $P = X$, and $[\forall]P = \varnothing$ otherwise.

To summarise,

Definition 1.3.3 (Dense interior semantics for topo-e-models)**.** We have a language that includes a countable set of propositional variables, the Boolean connectives $\wedge$ and $\neg$ and the modalities $\Box_0, \Box, B, K$ and $[\forall]$. Given a topo-e-model $\mathfrak{X} = (X, \tau, E_0, V)$, let τ^+ be the collection of its dense open sets. We define truth in $x \in X$ recursively as follows:

$$
\begin{array}{lll}
\mathfrak{X}, x \vDash p & \text{iff} & x \in V(p); \\
\mathfrak{X}, x \vDash \neg\phi & \text{iff} & \mathfrak{X}, x \nvDash \phi; \\
\mathfrak{X}, x \vDash \phi \wedge \psi & \text{iff} & \mathfrak{X}, x \vDash \phi \text{ and } \mathfrak{X}, x \vDash \psi; \\
\mathfrak{X}, x \vDash \Box_0\phi & \text{iff} & \text{there exists } e \in E_0 \text{ such that } x \in e \subseteq \|\phi\|; \\
\mathfrak{X}, x \vDash \Box\phi & \text{iff} & x \in \mathrm{Int}\, \|\phi\|; \\
\mathfrak{X}, x \vDash B\phi & \text{iff} & \text{there exists } U \in \tau^+ \backslash \{\varnothing\} \text{ such that } U \subseteq \|\phi\| \\
 & \text{iff} & \varnothing \neq \mathrm{Int}\, \|\phi\| \in \tau^+; \\
\mathfrak{X}, x \vDash K\phi & \text{iff} & \text{there exists } U \in \tau^+ \text{ such that } x \in U \subseteq \mathrm{Int}\, \|\phi\| \\
 & \text{iff} & x \in \mathrm{Int}\, \|\phi\| \in \tau^+; \\
\mathfrak{X}, x \vDash [\forall]\phi & \text{iff} & \mathfrak{X}, y \vDash \phi \text{ for all } y \in X.
\end{array}
$$

Moreover we set $\Diamond_0, \Diamond, \hat{B}, \hat{K}, [\exists]$ to be, respectively, the duals of these modalities (i.e. $\Diamond_0\phi \equiv \neg\Box_0\neg\phi$ and so on).

[2]Although considering the empty set dense is not standard, we will do so in this thesis for expositional clarity. Later we will define the topology of dense open sets and, for that, it is convenient to include the empty set in the definition.

One does not, in most cases, need to use all these modalities. Let us provide several fragments that will be used throughout this thesis and point out what the logic of each of them is.

The knowledge-only fragment $\mathcal{L}_K$. This is the fragment including only the knowledge modality K. It has $\mathsf{S4.2}$ as its logic. Similarly,

The belief-only fragment $\mathcal{L}_B$ has $\mathsf{KD45}$ as its logic.

The knowledge and belief fragment $\mathcal{L}_{KB}$. This is the fragment with the K and B modalities. Its logic is precisely Stal, introduced in subsection 1.2.2.

The knowledge fragment $\mathcal{L}_{\forall K}$. The logic of the fragment with the global "infallible knowledge" modality $[\forall]$ plus the "defeasible knowledge" modality K is $\mathsf{Logic}_{\forall K}$, which is the least logic including all the propositional tautologies, the $\mathsf{S4}$ axioms and rules for K, the $\mathsf{S5}$ axioms and rules for $[\forall]$ plus the axioms $[\forall]\phi \to K\phi$ and $[\exists]K\phi \to [\forall]\hat{K}\phi$.

The combined evidence fragment $\mathcal{L}_{\forall\Box}$. The logic of this fragment, $\mathsf{Logic}_{\forall\Box}$ contains the $\mathsf{S4}$ axioms and rules for $\Box$, the $\mathsf{S5}$ axioms and rules for $[\forall]$ plus the axiom $[\forall]\phi \to \Box\phi$.

Note that the following equality holds in topo-e-models:

$$\|K\phi\| = \|\Box\phi \wedge [\forall]\Diamond\Box\phi\|.$$

Indeed, $\Box\phi$ takes care of the "interior" part and $[\forall]\Diamond\Box\phi$ takes care of the "density" part. The following equality also holds (recall subsection 1.2.2):

$$\|B\phi\| = \|\hat{K}K\phi\|.$$

We see that K and B can actually be defined in this fragment.

The evidence fragment $\mathcal{L}_{\Box\Box_0\forall}$. The logic of the fragment including the $[\forall]$, $\Box$ and $\Box_0$ modalities is $\mathsf{Logic}_{\forall\Box\Box_0}$, including the propositional tautologies, the axioms and rules of $\mathsf{S5}$ for $[\forall]$ and of $\mathsf{S4}$ for $\Box$ plus the axioms:

$4_{\Box_0}$	$\Box_0\phi \to \Box_0\Box_0\phi$;
Universality	$[\forall]\phi \to \Box_0\phi$;
Factive evidence	$\Box_0\phi \to \Box\phi$;
Pullout	$(\Box_0\phi \wedge [\forall]\psi) \to \Box_0(\phi \wedge [\forall]\psi)$;

and the monotonicity rule: from $\phi \to \psi$, infer $\Box_0\phi \to \Box_0\psi$.

Remark 1.3.4. Unless we are dealing with a fragment which includes the $\Box_0$ modality, having a designated subbasis does not play a role in the semantics.

For this reason, whenever a notion of basic evidence is not involved, we will refer to models of the form (X, τ, V) as topo-e-models, with the understanding that they can be turned into a topo-e-model in the sense of definition 1.3.1 by setting $E_0 = \tau$.

Chapter 2

Generic spaces for the logic of topo-e-models

McKinsey and Tarski's (1944) theorem stating that S4 is the logic of any dense-in-itself metrisable space (such as the real line $\mathbb{R}$) under the interior semantics has a rather interesting implication. Not only do we have completeness with respect to the class of topological spaces, but moreover we have a space which gives a somewhat "natural" way of capturing knowledge yet it is "generic" enough so that its logic is precisely the logic of topological spaces. Whatever is not provable in the logic of knowledge S4 will find a refutation in $\mathbb{R}$ and whatever is true in S4 will hold in every model based on the topology of the real line.

Translating this idea to the framework of topo-e-models is the aim of this chapter. We wish to find topological evidence models which capture the logics presented in the preceding chapter, that is, special spaces whose logic under the dense interior semantics is exactly the logic of topo-e-models. Let us start this by formalising the idea of "generic".

Definition 2.0.1 (Generic models)**.** Let $\mathcal{L}$ be a language and (X, τ) a topological space. We will say that (X, τ) is a *generic model for* $\mathcal{L}$ if the sound and complete $\mathcal{L}$-logic over the class of all topological evidence models is sound and complete with respect to the family

$$\{(X, \tau, E_0) : E_0 \text{ is a subbasis of } \tau\}.$$

If $\Box_0$ is not in the language, then a generic model is simply a topological space which is sound and complete with respect to the corresponding $\mathcal{L}$-logic.

Since McKinsey and Tarski's theorem appeared in 1944, a number of simplified proofs of this result have been published. (For an overview, see van Benthem and Bezhanishvili, 2007.) The present chapter builds on one such proof, contained in Bezhanishvili, Bezhanishvili, Lucero-Bryan, and van Mill (2018), which uses the well-known fact that S4 is sound and complete with respect to finite rooted preorders (see e.g. Blackburn et al., 2001) and then constructs an interior map from a dense-in-itself metrizable space (X, τ) onto any such frame. That is, a surjective map $f : (X, \tau) \twoheadrightarrow (W, \leq)$ which is continuous and open. It can be proven that given such a map and a valuation V on $(W, \leq)$, if we define $V^f(p) := \{x \in X : fx \in V(p)\}$ it is the case that, for any formula ϕ in the language of S4, $x \vDash \phi$ on (X, τ, V^f) if and only if $fx \vDash \phi$ on $(W, \leq, V)$. Completeness is then a straightforward consequence, for if $\phi \notin$ S4, then there is a model based on a finite rooted preorder $(W, \leq, V)$ refuting ϕ and thus we can refute ϕ on (X, τ, V^f).

In the first section of this chapter we present an adaptation of this result to the dense interior semantics: namely, we prove that $\mathsf{S4.2}_K$ is the logic of any dense-in-itself metrisable space if we read K as according to the semantics defined in 1.3.3. In section 2 we add belief and prove completeness of any such space with respect to Stal. In section 3 we consider fragments of the logic which include the universal modality and show that, despite $\mathbb{R}$ is not a generic model for them, $\mathbb{Q}$ is. In section 4 we provide a sufficient condition for a dense-in-itself metrisable space to be a generic model for these fragments and show that many of the examples provided in the previous chapter are generic models.

2.1 S4.2 as the logic of $\mathbb{R}$

This section is devoted to the proof of this result, our analogue to McKinsey and Tarski's theorem:

Theorem 2.1.1. $\mathsf{S4.2}_K$ *is the logic of any dense-in-itself metrisable space if we read K as dense interior. That is, for any formula in the language of* $\mathsf{S4.2}_K$, $\mathsf{S4.2}_K \vdash \phi$ *if and only if* $(X, \tau) \vDash \phi$ *with the dense interior semantics.*

Before tackling this proof, some preliminaries are needed.

Given a topological space (X, τ) define τ^+ to be the collection of dense open sets in (X, τ):

$$\tau^+ = \{U \in \tau : \mathrm{Cl}\, U = X\} \cup \{\varnothing\}.$$

Lemma 2.1.2. (X, τ^+) *is an extremally disconnected topological space and, for any valuation $V : P \to \mathcal{P}(X)$ and any formula ϕ in the modal language $\mathcal{L}_K$ we have that $\|\phi\|^{(X,\tau,V)}$ under the dense interior semantics coincides with $\|\phi\|^{(X,\tau^+,V)}$ under the interior semantics.*

Proof. First, let us see τ^+ is a topology. X and $\varnothing$ are dense, hence they are in τ^+. Now, if $U, V \in \tau^+$, then $U, V \in \tau$ and they are dense. Due to the density of V, for any nonempty $W \in \tau$ we have that $V \cap W$ is a nonempty open set, and hence $(U \cap V) \cap W = U \cap (V \cap W) \neq \varnothing$. Thus $U \cap V$ is dense in τ and as a consequence $U \cap V \in \tau^+$. Finally, take $\sigma \subseteq \tau^+ \backslash \{\varnothing\}$. For any nonempty $W \in \tau$, we have that $W \cap \bigcup \sigma = \bigcup_{U \in \sigma}(U \cap W)$, which is a union of nonempty sets (by density of each $U \in \sigma$) and hence nonempty. Therefore, $\bigcup \sigma \in \tau^+$.

The fact that (X, τ^+) is extremally disconnected is a consequence of the fact that every open set in τ^+ is dense, hence $\mathrm{Cl}\, U = X \in \tau^+$ for all nonempty $U \in \tau^+$ and of course $\mathrm{Cl}\, \varnothing = \varnothing \in \tau^+$.

The last item is a straightforward subformula induction. The only case that might be more involved is the one with the K operator. Note that, for $A \subseteq X$, we have $\mathrm{Int}_{\tau^+} A = \mathrm{Int}_\tau A$ if $\mathrm{Int}_\tau A$ is dense and $\mathrm{Int}_{\tau^+} A = \varnothing$ otherwise. Hence, $x \in \|K\phi\|^{(X,\tau,V)}$ under the dense interior semantics if and only if $x \in \mathrm{Int}_\tau \|\phi\|^{(X,\tau,V)}$ and $\mathrm{Int}_\tau \|\phi\|^{(X,\tau,V)}$ is dense, if and only if $x \in \mathrm{Int}_{\tau^+} \|\phi\|^{(X,\tau^+,V)}$ if and only if $x \in \|K\phi\|^{(X,\tau^+,V)}$ under the interior semantics. ∎

Corollary 2.1.3. *For any topological space (X, τ), it is the case that $(X, \tau^+) \vDash$ $\mathsf{S4.2}$ under the interior semantics.*

Now, we will be using in this proof the known result that $\mathsf{S4.2}$ is sound and complete with respect to the class of finite rooted frames $(W, \leq)$ in which $\leq$ is a reflexive, transitive and weakly directed relation. Let us make an observation about this class of frames. If a frame is rooted and weakly directed, for every pair of points $x, y \in W$, and given that $x \geq r \leq y$ where r is the root of W, weak directedness grants us the existence of some z such that $x \leq z \geq y$. But this means that, for every pair of points x and y, the set $\uparrow x \cap \uparrow y$ is nonempty, and thus for every pair of nonempty upsets U and V we have that $U \cap V \neq \varnothing$. This means that every nonempty upset is dense in such a frame, and therefore that the topology of upsets $\tau := \mathrm{Up}(W)$ coincides with τ^+. This fact, paired with the previous lemma, immediately gives us this result:

Corollary 2.1.4. *Let $\mathfrak{F} = (W, \leq)$ be a reflexive, transitive and weakly directed rooted frame. Then the dense interior semantics on $(W, \mathrm{Up}(W))$ coincides with the interior semantics on it, which in turn coincides with the standard Kripke semantics on $(W, \leq)$. In other words, in any model based on such a frame*

$$x \vDash K\phi \text{ if and only if } y \vDash \phi \text{ for all } y \geq x.$$

Completeness will follow from this result:

Lemma 2.1.5. *Let (X, τ) be some topological space and $(W, \leq, V)$ a finite, rooted, reflexive, transitive and weakly directed Kripke model. Moreover let*

$$f : (X, \tau^+) \twoheadrightarrow (W, \mathrm{Up}(W))$$

be an onto interior map and define

$$V^f(p) := \{x \in X : fx \in V(p)\}.$$

Then for every $x \in X$ we have that $(X, \tau, V^f), x \vDash \phi$ under the dense interior semantics if and only if $(W, \leq, V), fx \vDash \phi$ under the Kripke semantics.

Proof. We proceed by induction on the structure of ϕ. If ϕ is a propositional variable, the result follows from the construction of V^f. The induction step for the Boolean connectives $\neg$ and $\wedge$ is straightforward.

Now suppose the result holds for ϕ and let $x \vDash K\phi$ for some $x \in X$. This means that there exists some $U \in \tau^+$ with $x \in U \subseteq \|\phi\|^X$. But then $f[U]$ is an open set in $(W, \mathrm{Up}(W))$ which contains fx and, for every $y \in f[U]$, we have $y = fz$ for some $z \in \|\phi\|^X$ and thus, by induction hypothesis, $y \vDash \phi$. Therefore $fx \vDash K\phi$.

Conversely if $fx \vDash K\phi$ we have that there exists an open set $U \in \mathrm{Up}(W)$ such that $fx \in U \subseteq \|\phi\|^W$. But then by continuity $f^{-1}U$ is an open set in τ^+ (thus an open dense set in τ) which contains x and moreover, by induction hypothesis, it is contained in $\|\phi\|^X$, thus $x \vDash K\phi$. ■

So let us take $\mathbb{R}$ with the natural topology and any finite rooted S4.2 model and construct such a map[1]. Completeness will follow. Later we will generalise this result to any dense-in-itself metrisable space.

For the construction, we will use the following fact:

Lemma 2.1.6 (Bezhanishvili et al., 2018)**.** *If $\mathfrak{F} = (W, R)$ is a finite rooted preorder, and (X, σ) is a dense-in-itself metrizable space, there exists a continuous, open and surjective map $f : (X, \sigma) \to (W, \mathrm{Up}_R(W))$.*

Now, let $\mathfrak{F} = (W, \leq)$ be a finite rooted S4.2 frame. Let us find an open, continuous and surjective map $\bar{f} : (\mathbb{R}, \tau^+) \to \mathfrak{F}$.

Note that $\mathfrak{F}$ has a final cluster, i.e., a set $A \subseteq W$ with the property that $w \leq a$ for all $w \in W, a \in A$. Indeed, let $r \in W$ be the root and let $x, y \in W$ be any two maximal elements (which exist, on account that $\mathfrak{F}$ is finite). Since $r \leq x$ and $r \leq y$, by directedness, there is a z such that $x \leq z \geq y$. But by maximality of x and y, we have that $z \leq x$ and $z \leq y$, hence, by transitivity, $x \leq y$ and $y \leq x$: the maximal elements of $\mathfrak{F}$ form a final cluster.

Now, if $\mathfrak{F}$ consists only of this final cluster (i.e. $W = \{w_1, ..., w_n\}$ with $w_i \leq w_j$ for all i, j), its topology of upsets is the set $\{\varnothing, W\}$. In this case we partition $\mathbb{R}$ in n dense sets $\{A_1, ..., A_n\}$ with $\mathrm{Cl}\, A_i = \mathbb{R}$, $A_i \cap A_j = \varnothing$ for $i \neq j$ and $A_1 \cup ... \cup A_n = \mathbb{R}$. Trivially, mapping each $a \in A_i$ to w_i gives us the desired $\bar{f}$.

Now suppose $\mathfrak{F}$ contains more points than just those in its final cluster, Note that, if C is this cluster, $W \backslash C$ with the restriction of $\leq$ forms a finite rooted S4 frame. Call this frame $\mathfrak{F}'$. Let us take a proper subset $A \subseteq \mathbb{R}$ which is closed, nowhere dense[2] and dense-in-itself (for example, the Cantor set[3]). By lemma 2.1.6, we have that there exists an interior map $f : (A, \tau_A) \to \mathfrak{F}'$, where τ_A is the subspace topology, i.e. $\tau_A = \{A \cap U : U \in \tau\}$. Now, if

[1] We wish to thank Guram Bezhanishvili for the idea of this construction.

[2] A *nowhere dense* set is a set with the property that $\mathrm{Int}\,\mathrm{Cl}\, A = \varnothing$. Note that if a set is nowhere dense and closed, its complement is dense.

[3] The Cantor set is the subset of $\mathbb{R}$ that one obtains by taking the closed interval $[0, 1]$ and removing its middle third, only to then do the same to the two remaining intervals and, recursively, removing the middle third of the 2^n intervals that one gets in step n ad

the final cluster of $\mathfrak{F}$ is $C = \{w_1, ..., w_n\}$ let us partition $\mathbb{R}$ in n dense sets $\{A_1, ..., A_n\}$ and set $A'_i = A_i \cap (\mathbb{R}\backslash A)$. Note that A'_i is nonempty for all i and that $\{A'_1, ..., A'_n\}$ constitutes a partition of $\mathbb{R}\backslash A$. Set

$$\bar{f}(a) = \begin{cases} w_i, & a \in A'_i; \\ f(a), & a \in A. \end{cases}$$

Let us see that this $\bar{f}$ is the desired map.

Surjectivity is trivial. For continuity, take an open set U in $\mathfrak{F}$. Note that U contains C (because it is an upset) and that $U\backslash C$ is an upset in $\mathfrak{F}'$. Hence $f^{-1}(U\backslash C)$ is an open set in τ_A, i.e., it equals $V \cap A$ for some $V \in \tau$. Then $\bar{f}^{-1}(U) = \bar{f}^{-1}(C) \cup \bar{f}^{-1}(U\backslash C) = (\mathbb{R}\backslash A) \cup (V \cap A) = (\mathbb{R}\backslash A) \cup V$. This set is open in τ (for it is the union of two opens) and dense (for it contains $\mathbb{R}\backslash A$ which is dense), hence $\bar{f}^{-1}(U) = (\mathbb{R}\backslash A) \cup V \in \tau^+$ and this gives us continuity. For openness, take a dense open set $U \in \tau^+$. It has a nonempty intersection with each of the A'_i, hence $\bar{f}[U \cap (\mathbb{R}\backslash A)] = C$. On the other hand $\bar{f}[U \cap A]$ is an upset in $\mathfrak{F}'$, let us call it V. Hence, $\bar{f}[U] = V \cup C$, which is an upset in $\mathfrak{F}$.

This proof can be very easily extended to an arbitrary dense-in-itself metrizable space, just by making use of the partition lemma:

Lemma 2.1.7 (Partition lemma)**.** *Let X be a dense-in-itself metrizable space and F a nonempty closed discrete subspace of X. Then, for each $n \geq 1$, there is a partition of X $\{G, U_1, ..., U_n\}$ such that G is a dense-in-itself closed nowhere dense subspace of X containing F and each U_i is an open set with the property that there is a discrete subspace F_i of U_i such that* $\mathrm{Cl}\, F_i = F_i \cup G$.

Proof. See Bezhanishvili et al. (2018). ■

Theorem 2.1.8. *Given any dense-in-itself metrisable space and a finite rooted* S4.2 *frame $\mathfrak{F} = (W, R)$, there exists a continuous, open and surjective map $\bar{f} : (X, \tau^+) \twoheadrightarrow (W, \leq)$.*

Proof. Suppose $A = \{w_1, ..., w_n\}$ is the final cluster of $\mathfrak{F}$. If $W = A$, then let $A_1, ..., A_n$ be dense sets partitioning[4] X and map $x \in A_i$ to w_i, as above. Otherwise, take $\{G, U_1, ..., U_n\}$ as in 2.1.7 and consider the continuous, open and surjective map $f : (G, \tau_G) \to \mathfrak{F}'$ which exists as per 2.1.6 and just extend it to $\bar{f} : (X, \tau^+) \to \mathfrak{F}$ mapping each $x \in U_i$ to w_i. This is clearly surjective. It is open, for given $U \in \tau^+$, U has nonempty intersection with each of $U_1, ..., U_n$, hence $\bar{f}[U] = f[U \cap G] \cup A$, which is an upset in $\mathfrak{F}'$ together with the final

infinitum. In other words, the Cantor set is $C = \bigcap_{n\in\omega} C_n$, where $C_0 = [0, 1]$ and

$$C_{n+1} = \{x/3 : x \in C_n\} \cup \{(x + 2)/3 : x \in C_n\}.$$

As a subspace of $\mathbb{R}$, the Cantor set is isomorphic to the Cantor space defined above.

[4]Hewitt (1943) proves that such partition exists in any dense-in-itself metrisable space for $n = 2$. The fact that a dense subspace of X is dense-in-tself and metrisable, plus the fact that a dense subset of a dense subset of X is itself a dense subset of X, allow us to apply it recursively to get the partition for any n.

cluster A, i.e. an upset in $\mathfrak{F}$. It is continuous, for given an upset V in $\mathfrak{F}$, this upset is either $\varnothing$ (in which case $\bar{f}^{-1}(\varnothing) = \varnothing \in \tau^+$) or it equals $V' \cup A$ for some upset V' in $\mathfrak{F}'$. But then

$$\bar{f}^{-1}(V) = f^{-1}(V) \cup \bar{f}^{-1}(A) = f^{-1}(V) \cup (X \backslash G),$$

which, as reasoned above, is a dense open set. ■

We now have enough to prove the main theorem in this section:

Proof of theorem 2.1.1. Soundness follows immediately from 2.1.3. For completeness, take some formula ϕ such that $\phi \notin \mathsf{S4.2}$. Then there exists some finite rooted $\mathsf{S4.2}$ model $(W, \leq, V)$ and some $w \in W$ refuting ϕ. But then we can construct an open, continuous an surjective map $f : (X, \tau) \twoheadrightarrow (W, \leq)$ as above and, by defining $V^f(p) = \{x \in X : fx \in V(p)\}$ and taking some $x \in X$ such that $fx = w$, lemma 2.1.5 gives us that x refutes ϕ. ■

Similar results of soundness and completeness for particular spaces relative to the fragments considered in section 1.3 can be obtained. Let us see some in the following sections.

2.2 Adding belief

The logic Stal introduced in 1.2.2 is the logic of topo-e-models for the belief and knowledge fragment. Recall that the formula $B\phi \leftrightarrow \hat{K}K\phi$ is provable in Stal. In particular, for any formula ϕ in the language $\mathcal{L}_{KB}$, there exists a formula ψ in the language $\mathcal{L}_K$ such that $\vdash_{\mathsf{Stal}} \phi \leftrightarrow \psi$ (indeed, we get ψ by substituting every instance of B in ϕ with $\hat{K}K$).

And thus we have the following:

Theorem 2.2.1. Stal *is sound and complete with respect to any dense-in-itself metrizable space with the dense interior semantics.*

Proof. Soundness follows from the fact that Stal is sound with respect to topo-e-models. For completeness, suppose $\phi \notin \mathsf{Stal}$ and take ψ in the language $\mathcal{L}_K$ such that $\vdash_{\mathsf{Stal}} \phi \leftrightarrow \psi$. Then $\psi \notin \mathsf{S4.2}$, hence there is a valuation on (X, τ^+) making ψ false at some $x \in X$. By soundness and the fact that $\vdash_{\mathsf{Stal}} \phi \leftrightarrow \psi$, we conclude that ϕ is false at x as well. ■

2.3 The global modality $[\forall]$ and the logic of $\mathbb{Q}$

Three fragments including the global modality $[\forall]$ will be considered in the present section: the *knowledge fragment* (the one which includes the K and $[\forall]$ modalities), the *factive evidence fragment* (including $\Box$ and $[\forall]$) and the *evidence fragment* (including $[\forall]$, $\Box$ and $\Box_0$).

First let us observe something about the basic evidence fragment. Recall that the logic of this fragment, $\mathsf{Logic}_{\forall\Box}$, consists of $\mathsf{S5}_\forall$ plus $\mathsf{S4}_\Box$ plus the axiom $[\forall]\phi \to \Box\phi$.

This logic is not complete with respect to $\mathbb{R}$. Consider the following formula:

$$[\forall](\Box p \vee \Box\neg p) \to ([\forall]p \vee [\forall]\neg p) \tag{Con}$$

It is the case that (Con) is not derivable in the logic yet it is always true in $\mathbb{R}$. More generally,

Proposition 2.3.1 (Shehtman, 1999)**.** *A topological space* (X, τ) *satisfies* (Con) *if and only if it is connected. Moreover adding* (Con) *to* $\mathsf{Logic}_{\forall\Box}$ *gives us a complete axiomatisation of any dense-in-itself metrisable connected space.*

We could add (Con) to our logic but we have no conceptual use for it. Instead, we will show completeness of this fragment (plus the other two mentioned above which include the global modality) with respect to a dense-in-itself, metrisable yet disconnected space, namely $\mathbb{Q}$.

2.3.1 The knowledge fragment $\mathcal{L}_{\forall K}$

Lemma 2.3.2 (Goranko and Passy, 1992)**.** $\mathsf{Logic}_{\forall K}$ *is sound and complete with respect to finite models of the form* $(X, \leq, V)$ *where* X *is a finite set,* $\leq$ *is a preorder with a final cluster*[5]*, where* K *is read as the Kripke modality for* $\leq$ *and* $[\forall]$ *is read as a universal modality in the model.*

In order to prove completeness it suffices to show the following

Theorem 2.3.3. *For every finite frame* $(W, \leq)$ *where* $\leq$ *is reflexive, transitive and has a final cluster, there exists a surjective dense-interior map*

$$f : (\mathbb{Q}, \tau_{\mathbb{Q}}) \to (W, \mathrm{Up}_{\leq}(W)).$$

For the remainder of this subsection, let $(W, \leq)$ be such a frame. We have the following:

Lemma 2.3.4. *Each finite cofinal preorder is a p-morphic image of a finite disjoint union of finite rooted* $\mathsf{S4.2}$ *frames, via a dense-open and dense-continuous p-morphism.*

Proof. Let $x_1, ..., x_n$ be the minimal elements of M. Now,for $1 \leq i \leq n$ take $M'_i = {\uparrow}x_i \times \{i\}$. Define an order on $M' = M'_1 \cup ... \cup M'_n$ by: $(x, i) \leq (y, j)$ iff $i = j$ and $x \leq y$. $M'_1, ..., M'_n$ are pairwise disjoint finite rooted $\mathsf{S4.2}$ frames (with $A \times \{i\}$ as a final cluster) and $(x, i) \mapsto x$ is a p-morphism from M' onto M. It is easy to see that this mapping is dense-open (for every nonempty open set is dense in M) and dense-continuous (for the preimage of a nonempty M-open set is a M'-open set which contains all the final clusters, and thus is dense). ■

[5] In the sense that there exists $A \subseteq X$ with $x \leq y$ for all $x \in X, y \in A$.

With this:

Proof of Theorem 2.3.3. Let $M_1, ..., M_n$ be a family of pairwise disjoint of finite rooted S4.2 frames whose union $M = M_1 \cup ... \cup M_n$ has $(W, \leq)$ as a p-morphic image.

Take $z_1, ..., z_{n-1} \in \mathbb{R} \backslash \mathbb{Q}$ and consider the intervals $A_1 = (-\infty, z_1)$, $A_n = (z_{n-1}, \infty)$ and $A_i = (z_{i-1}, z_i)$ for $1 < i < n$. Now, each A_i, as a subspace, is homeomorphic to $\mathbb{Q}$ (and thus a dense-in-itself metrisable space). From each $(A_i, \tau|_{A_i})$ we can find a dense-open, dense-continuous and surjective map f_i onto M_i. Then $f = f_1 \cup ... \cup f_n$ is a dense-interior map onto M which, when composed with the p-morphism given by the previous corollary, gives us the desired map. ■

With a very similar proof to that of 2.1.5, we see that, for every $x \in \mathbb{Q}$ and every formula ϕ, $x \vDash \phi$ if and only if $fx \vDash \phi$. Completeness is an immediate consequence of this.

Corollary 2.3.5. $\mathsf{Logic}_{\forall K}$ *is the logic of* $\mathbb{Q}$.

2.3.2 The factive evidence fragment $\mathcal{L}_{\forall\Box}$

Goranko and Passy (1992) show that $\mathsf{Logic}_{\forall\Box}$ is sound and complete with respect to finite relational models of the form $(X, \leq, V)$ where $\leq$ is a preorder.

Thus to prove completeness of this logic with respect to $\mathbb{Q}$ it suffices to find a suitable open and continuous map from $\mathbb{Q}$ onto any such finite frame. And indeed,

Theorem 2.3.6. *Let* $(W, \leq)$ *be any finite preordered frame. Then there exists a open, continuous and surjective map* $f : (\mathbb{Q}, \tau_{\mathbb{Q}}) \to (W, \mathrm{Up}_{\leq}(W))$.

Proof. The proof of this result is very similar to that of 2.3.3, so we will just present a sketch here. If M is a maximal connected subset of W, then M can be written as $B_1 \cup ... \cup B_n$, where each B_i is a rooted upset (indeed, just take $x_1, ..., x_n$ to be the minimal points of M and $B_i = \uparrow x_i$). But we have once again that M is the p-morphic image of a collection of n pairwise disjoint rooted preorders. Indeed, define $M' = \{\langle x, i \rangle : x \in B_i\}$ and $\langle x, i \rangle \leq \langle y, j \rangle$ iff $i = j$ and $x \leq y$, and we have that the map $\langle x, i \rangle \mapsto x$ defines an onto p-morphism. So the whole frame $(W, \leq)$ is the p-morphic image of a finite number of pairwise disjoint finite rooted preorders, $B_1, ..., B_m$, each of which is the image of $\mathbb{Q}$ under an open and continuous map $f_m : \mathbb{Q} \to B_m$. As before, this means that there is such a map from $\mathbb{Q}$ to W. ■

Again, noting that if we define $V^f(p) = \{x \in \mathbb{Q} : fx \in V(p)\}$ we get $x \vDash \phi$ in $(\mathbb{Q}, \tau_{\mathbb{Q}}, V^f)$ if and only if $fx \vDash \phi$ in $(W, \leq, V)$, completeness follows:

Corollary 2.3.7. $\mathsf{Logic}_{\forall\Box}$ *is the logic of* $\mathbb{Q}$, *where* $\Box$ *is read as topological interior and* $[\forall]$ *is read as a global modality.*

2.3.3 Adding basic evidence: the evidence fragment $\mathcal{L}_{\forall\Box\Box_0}$

Let us now account for basic evidence. We take the fragment consisting of the modal operators $\Box$, $[\forall]$ and $\Box_0$. Recall that we interpret formulas of this fragment on topo-e-models (X, τ, E_0, V), where E_0 is a subbasis for (X, τ), like this: $x \in \|\Box_0\phi\|$ if and only if there exists $e \in E_0$ with $x \in e \subseteq \|\phi\|$.

The logic of this fragment is $\mathsf{Logic}_{\forall\Box\Box_0}$, as seen in section 1.3. Baltag et al. (2016) prove that this logic is sound and complete with respect to finite models of the form $(X, \leq, E_0^X, V)$, where $\leq$ is a preorder and E_0^X is a subbasis for $\mathrm{Up}(X)$ with $X \in E_0$.

Suppose that $\phi \notin \mathsf{Logic}_{\forall\Box\Box_0}$. Let us find a subbasis $E_0^{\mathbb{Q}}$ for $(\mathbb{Q}, \tau_{\mathbb{Q}})$, where $\tau_{\mathbb{Q}}$ is the natural topology on $\mathbb{Q}$ and a valuation V such that $(\mathbb{Q}, \tau_{\mathbb{Q}}, E_0^{\mathbb{Q}}, V)$ refutes ϕ.

Now, we know that there exists a model $(X, \leq, E_0^X, V)$ and a point $x \in X$ refuting ϕ. By theorem 2.3.6, we also know that there exists a map $f : \mathbb{Q} \twoheadrightarrow X$ which is surjective, continuous and open. Let us define

$$E_0^{\mathbb{Q}} := \{e \subseteq \mathbb{Q} : f[e] \in E_0^X\}.$$

Let us see that this is a subbasis for $\mathbb{Q}$. First, seeing as $X \in E_0^X$ and $f[\mathbb{Q}] = X$, we have that $\mathbb{Q} \in E_0^{\mathbb{Q}}$, thus $\bigcup E_0^{\mathbb{Q}} = \mathbb{Q}$.

Now, suppose $p \in U \in \tau_{\mathbb{Q}}$. Let us see that there exist $e_1^q, ..., e_n^q \in E_0^{\mathbb{Q}}$ such that $x \in e_1^q \cap ... \cap e_n^q \subseteq U$. Now, we have that $fp \in f[U]$ which is an open set. Since E_0^X is a subbasis for $(X, \leq)$ this means that there exist $e_1^x, ..., e_n^x \in E_0^X$ with $fp \in e_1^x \cap ... \cap e_n^x \subseteq f[U]$. Now set

$$e_i^q := f^{-1}e_i^x \backslash \{y \notin U : fy \in f[U]\}.$$

The fact that $e_i^q \in E_0^{\mathbb{Q}}$ follows from the fact that $f[e_i^q] = e_i^x$. Indeed, if $y \in f[e_i^q]$ then $y \in ff^{-1}e_i^x = e_i^x$ and conversely if $y \in e_i^x$, then either $y \in f[U]$ (in which case $y = fz$ for some $z \in U$ and thus $z \in f^{-1}e_i^x$ and therefore $z \notin \{z' \notin U : fz' \in f[U]\}$, which implies $z \in e_i^q$) or $y \notin f[U]$ (in which case $y = fz$ for some z by surjectivity and $z \notin \{z' \notin U : fz' \in f[U]\}$, thus $z \in e_i^q$). In either case, $y \in f[e_i^q]$.

Finally, note that $e_1^q \cap ... \cap e_n^q \subseteq U$. Indeed, for any $x \in e_1^q \cap ... \cap e_n^q$ we have that $fx \in e_1^x \cap ... \cap e_n^x \subseteq f[U]$, and thus by the definition of the e_i^q's it cannot be the case that $x \notin U$.

So for $p \in U \in \tau_{\mathbb{Q}}$ we have found elements $e_1^q, ...e_n^q \in E_0^{\mathbb{Q}}$ such that $p \in e_1^q \cap ... \cap e_n^q \subseteq U$, and therefore $E_0^{\mathbb{Q}}$ is a subbasis.

We set a valuation on $(\mathbb{Q}, \tau_{\mathbb{Q}}, E_0^{\mathbb{Q}})$ as follows:

$$V^{\mathbb{Q}}(p) = \{x \in \mathbb{Q} : fx \in V(p)\}.$$

Let us show the following result, from which completeness follows.

Lemma 2.3.8. *For any formula ϕ in the language and any $x \in \mathbb{Q}$, we have that $(\mathbb{Q}, \tau_{\mathbb{Q}}, E_0^{\mathbb{Q}}, V^{\mathbb{Q}}), x \vDash \phi$ if and only if $(X, \leq, E_0^X, V), fx \vDash \phi$.*

Proof. This is again an induction on formulas; the only induction step that requires some attention is the one referring to $\Box_0$.

Let $x \vDash \Box_0\psi$. This means that there exists some $e \in E_0^{\mathbb{Q}}$ with $x \in e$ and $y \vDash \psi$ for all $y \in E$. But then $fx \in f(e) \in E_0^X$ and by the induction hypothesis we have $fy \vDash \psi$ for all $fy \in f(e)$ and thus $fx \vDash \Box_0\psi$. Conversely, if $fx \in e' \subseteq \|\psi\|^X$ for some $e' \in E_0^X$, we have that $x \in f^{-1}e' \in E_0^{\mathbb{Q}}$ and $fy \vDash \psi$ for each $y \in f^{-1}e'$ and thus, by induction hypothesis, $y \vDash \psi$. Therefore $x \vDash \Box_0\psi$. ∎

Corollary 2.3.9. $\mathsf{Logic}_{\forall\Box\Box_0}$ *is sound and complete with respect to the class of topo-e-spaces*

$$\{(\mathbb{Q}, \tau_{\mathbb{Q}}, E_0^{\mathbb{Q}}) : E_0^{\mathbb{Q}} \textit{ is a subbasis of } (\mathbb{Q}, \tau_{\mathbb{Q}})\}.$$

To summarise the results in this section:

Theorem 2.3.10. $(\mathbb{Q}, \tau_{\mathbb{Q}})$ *is a generic model for the fragments* $\mathcal{L}_K$, $\mathcal{L}_{\forall\Box}$, $\mathcal{L}_{\forall K}$ *and* $\mathcal{L}_{\forall\Box\Box_0}$.

2.3.4 A condition for generic models

Let us generalise the results in the previous section. While we saw that every dense-in-itself metrisable space is a generic model for $\mathsf{S4.2}$, this result "breaks" when we add the universal modality $[\forall]$. For example, the logic of $\mathbb{R}$ is not $\mathsf{Logic}_{\forall K}$, whereas the logic of $\mathbb{Q}$ is.

One can easily see that the only part in the proof of theorem 2.3.3 which uses a special property of $\mathbb{Q}$ which $\mathbb{R}$ does not have is that in which we partition $\mathbb{Q}$ in n subspaces which are homeomorphic to $\mathbb{Q}$ itself. And it is straightforward that, given a dense-in-itself metrisable space which admits such partition, the proofs in the previous section will work *mutatis mutandis*. Let us then give a necessary and sufficient condition for such a space to have this property.

Definition 2.3.11 (Sum of topological spaces)**.** Given two topological spaces (X, τ) and (Y, σ), their *sum* is the topological space $(X, \tau) \oplus (Y, \sigma)$, whose underlying set is the disjoint union of X and Y, i.e.

$$X \oplus Y = X \times \{1\} \cup Y \times \{2\}$$

and whose topology is

$$\tau \oplus \sigma = \{U \times \{1\} \cup V \times \{2\} : U \in \tau, V \in \sigma\}.$$

We will say that a topological space (X, τ) is *idempotent* whenever (X, τ) is homeomorphic to $(X, \tau) \oplus (X, \tau)$.

The following holds:

Lemma 2.3.12. *A topological space* (X, τ) *is idempotent if and only if it can be partitioned in* n *subspaces homeomorphic to itself for each* $n \geq 1$.

Proof. If (X, τ) admits a partition in two subspaces homeomorphic to itself, since these are disjoint their union (which is X) is the same as their sum, which is homeomorphic to $X \oplus X$.

Conversely, if (X, τ) is idempotent we can reason recursively to find that X is homeomorphic to the sum $X_1 \oplus ... \oplus X_n$ where each X_i is a copy of X. Let $f : X_1 \oplus ... \oplus X_n \to X$ be a homeomorphism. Then $\{f[X_1], ..., f[X_n]\}$ constitutes a partition of X in n subspaces, each of them homeomorphic to X. ■

And thus, we have our general result:

Corollary 2.3.13. *Any dense-in-itself idempotent metrisable space is sound and complete with respect to* $\mathsf{Logic}_{\forall K}$*,* $\mathsf{Logic}_{\forall\Box}$ *and* $\mathsf{Logic}_{\forall\Box\Box_0}$*.*

All the spaces introduced in the preliminaries, except for $\mathbb{R}$ and $\mathcal{T}_2$, are dense-in-itself, metrisable and idempotent spaces. And thus:

Theorem 2.3.14. *The rational numbers* $\mathbb{Q}$*, the irrationals* $\mathbb{I}$*, the Cantor space* $\mathfrak{C}$ *and the Baire space* $\mathfrak{B}$ *are all examples of generic spaces for the fragments* $\mathcal{L}_K$*,* $\mathcal{L}_{\forall\Box}$*,* $\mathcal{L}_{\forall K}$ *and* $\mathcal{L}_{\forall\Box\Box_0}$*.*

2.4 A condition for generic models

Let us generalise the results in the previous section. While we saw that every dense-in-itself metrisable space is a generic model for $\mathsf{S4.2}$, this result "breaks" when we add the universal modality $[\forall]$. For example, the logic of $\mathbb{R}$ is not $\mathsf{Logic}_{\forall K}$, whereas the logic of $\mathbb{Q}$ is.

One can easily see that the only part in the proof of theorem 2.3.3 which uses a special property of $\mathbb{Q}$ which $\mathbb{R}$ does not have is that in which we partition $\mathbb{Q}$ in n subspaces which are homeomorphic to $\mathbb{Q}$ itself. And it is straightforward that, given a dense-in-itself metrisable space which admits such partition, the proofs in the previous section will work *mutatis mutandis*. Let us then give a necessary and sufficient condition for such a space to have this property.

Definition 2.4.1 (Sum of topological spaces)**.** Given two topological spaces (X, τ) and (Y, σ), their *sum* is the topological space $(X, \tau) \oplus (Y, \sigma)$, whose underlying set is the disjoint union of X and Y, i.e.

$$X \oplus Y = X \times \{1\} \cup Y \times \{2\}$$

and whose topology is

$$\tau \oplus \sigma = \{U \times \{1\} \cup V \times \{2\} : U \in \tau, V \in \sigma\}.$$

We will say that a topological space (X, τ) is *idempotent* whenever (X, τ) is homeomorphic to $(X, \tau) \oplus (X, \tau)$.

The following holds:

Lemma 2.4.2. *A topological space (X, τ) is idempotent if and only if it can be partitioned in n subspaces homeomorphic to itself for each $n \geq 1$.*

Proof. If (X, τ) admits a partition in two subspaces homeomorphic to itself, since these are disjoint their union (which is X) is the same as their sum, which is homeomorphic to $X \oplus X$.

Conversely, if (X, τ) is idempotent we can reason recursively to find that X is homeomorphic to the sum $X_1 \oplus ... \oplus X_n$ where each X_i is a copy of X. Let $f : X_1 \oplus ... \oplus X_n \to X$ be a homeomorphism. Then $\{f[X_1], ..., f[X_n]\}$ constitutes a partition of X in n subspaces, each of them homeomorphic to X. ■

And thus, we have our general result:

Corollary 2.4.3. *Any dense-in-itself idempotent metrisable space is sound and complete with respect to* $\mathsf{Logic}_{\forall K}$, $\mathsf{Logic}_{\forall\Box}$ *and* $\mathsf{Logic}_{\forall\Box\Box_0}$.

All the spaces introduced in the preliminaries, except for $\mathbb{R}$ and $\mathcal{T}_2$, are dense-in-itself, metrisable and idempotent spaces. And thus:

Theorem 2.4.4. *The rational numbers $\mathbb{Q}$, the irrationals $\mathbb{I}$, the Cantor space $\mathfrak{C}$ and the Baire space $\mathfrak{B}$ are all examples of generic spaces for the fragments $\mathcal{L}_K$, $\mathcal{L}_{KB}$, $\mathcal{L}_{\forall\Box}$, $\mathcal{L}_{\forall K}$ and $\mathcal{L}_{\forall\Box\Box_0}$.*

Chapter 3

Going multi-agent

The present work so far has only accounted for sentences which refer to the epistemic state of a single agent. Given this, introducing several epistemic agents into the framework seems like a very obvious direction in which to steer this ship.

There have been some approaches to a multi-agent logic derived from the framework in Baltag et al. (2016). One of them was the subject of a recent ILLC Master's thesis (Ramírez, 2015), in which a two-agent logic with distributed knowledge was defined. However, the semantics of this approach seems to come with some conceptual problems which will be discussed later in section 5.2.

Another approach is present in Özgün (2017). This approach generalises the one-agent case and is devoid of the aforementioned conceptual issues, yet it uses the semantics of subset space logic: sentences are evaluated at a pair (x, U) where x is a world and U is some neighbourhood of x.

The system introduced in the present chapter and expanded upon in the subsequent ones generalises the one-agent models while maintaining the underlying ideas to the single-agent case, where sentences are evaluated at worlds.

After some discussion in section 1, we present in section 2 the semantics for this framework. We will limit ourselves to two agents for simplicity in the exposition. In the final chapter we will discuss how to easily extend these results to any finite number of agents. Section 3 contains a proof of completeness for the fragment of the language which only accounts for the knowledge of the agents, $\mathcal{L}_{K_1K_2}$. We will show that the logic of this fragment is the fusion logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$. In section 4 we consider other fragments of the logic.

3.1 The problem of density in the two-agent case

A first idea when attempting to incorporate a second epistemic agent would be to simply add a second topology to the single-agent framework and read things in the same way. That is, we could interpret sentences on bitopological spaces (X, τ_1, τ_2) where τ_1 and τ_2 are topologies defined on X, and we say, for $i = 1, 2$, that $x \in K_i\phi$ if and only there is a set $U \in \tau_i$ which is dense in τ_i such that $x \in U \subseteq \|\phi\|$.

There are two flagrant issues with this. One has to do with the interpretation of the density condition in the original framework. For an agent to know p we want her to have a piece of evidence for p ($x \in \text{Int}\,\|p\|$) which is consistent with every other piece of evidence she has ($\text{Int}\,\|p\|$ is dense). But these pieces of evidence have to be contained in some set of worlds which are consistent with the agent's information, and we are equating these sets for both agents. This is made more explicit if we define the density condition in terms of the interior operator $\Box_i$ and the global modality $[\forall]$: the τ_i-interior of $\|p\|$ being dense amounts to $[\forall]\Diamond_i\Box_i p$ being true in the model. And, as discussed in section 2.3, this global modality evaluates in all worlds which the agent considers possible. That is, adopting this semantics requires assuming that the same worlds are compatible with both agents' information, and therefore that their infallible knowledge coincides. This is clearly undesirable.

The second issue is even more salient. Since our original logic for the single-agent framework was $\mathsf{S4.2}_K$ and we do not want to throw into the mix any sort of interaction between the agents, one would expect the logic for the two-agent framework to be the fusion logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$, which simply contains the axioms and rules of $\mathsf{S4.2}$ for each K_i. This is however not the case: consider the formula

$$\phi \equiv \hat{K}_1 K_1 p \to K_2 \hat{K}_1 K_1 p.$$

Now, ϕ is not derivable in $\mathsf{S4.2}_{K_1}+\mathsf{S4.2}_{K_2}$ yet it is valid in bitopological models with the above semantics. Indeed, if $x \vDash \hat{K}_1 K_1 p$ then for every τ_1-open τ_1-dense set U, if $x \in U$ then there exists $y \in U$ such that $y \vDash K_i\phi$, i.e., there exists some τ_1-open τ_1-dense set U_y with $y \in U_y \subseteq \|p\|$. In particular, since $x \in X$ and X is a dense open set, there exists some dense open set $U \subseteq \|p\|$. But then take any $x_0 \in X$ and any τ_1-open τ_1-dense set V including x. We have that $V \cap U \neq \varnothing$ thus there exists some $y \in V$ with $y \in U \cap V \subseteq \|p\|$, whence $y \vDash K_i p$ and thus $x_0 \vDash \hat{K}_i K_i p$. This means that $\|\hat{K}_1 K_1 p\| = X$ and thus $K_2 \hat{K}_1 K_1 p$ holds everywhere. So not only is the logic for these models different from $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$, as one would expect, but, whatever this logic is, it includes a formula expressing that, if one agent considers it to be possible that she herself knows p, then the other agent knows this fact.

Our proposal to eliminate these complications involves making explicit which worlds are compatible with an agent's information at world x. This is done via the use of partitions, and it will be outlined in the next section.

3.2 Topological-partitional models

In order to specify which worlds an agent considers possible, we can define the topologies which encode the evidence of the agents on a common space

X, but we restrict, for each agent and at each world $x \in X$, the set of worlds epistemically accessible to the agent at x. We can still speak about density, but *locally*. A straightforward way to this is through the use of partitions.

Definition 3.2.1. A *topological-partitional model* is a tuple

$$\mathfrak{M} = (X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$$

where V is a valuation, τ_i is a topology defined on X and Π_i is a partition of X with the property that $\Pi_i \subseteq \tau_i$.

The worlds which are compatible with agent i's information at $x \in X$ are now precisely the worlds in the unique cell of the partition Π_i which includes x. The concept of justification comes now in the form of a local notion of density:

Definition 3.2.2. For $x \in X$, let $\Pi_i(x)$ be the unique $\pi \in \Pi_i$ with $x \in \pi$. For $U \subseteq X$, let $\Pi_i[U] = \{\pi \in \Pi_i : \pi \cap U \neq \varnothing\} = \{\Pi_i(x) : x \in U\}$.

A set $U \subseteq X$ is *locally dense in* $\pi \in \Pi_i$ whenever $\pi \subseteq \mathrm{Cl}_{\tau_i} U$ or equivalently when every nonempty open set contained in π has nonempty intersection with U.

We will say that a nonempty set U is *locally dense in* Π_i (or simply *locally dense* if there is no ambiguity) if $\mathrm{Cl}_{\tau_i} U = \bigcup \Pi_i[U]$. Equivalently, U is locally dense if for every $\pi \in \Pi_i[U]$ and every τ_i-open set $V \subseteq \pi$, we have $V \cap U \neq \varnothing$.

With this we can define a semantics for two-agent knowledge:

Definition 3.2.3 (Two-agent locally-dense-interior semantics)**.** Let

$$\mathfrak{M} = (X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$$

be a topological-partitional model and let $x \in X$. As usual, we have $\|p\| = V(p)$, $\|\phi \wedge \psi\| = \|\phi\| \cap \|\psi\|$ and $\|\neg\phi\| = X \backslash \|\phi\|$. For $i = 1, 2$ set:

$$\mathfrak{X}, x \vDash K_i\phi \text{ iff } x \in \mathrm{Int}_{\tau_i} \|\phi\| \\ \& \ \mathrm{Int}_{\tau_i} \|\phi\| \text{is locally dense in } \Pi_i(x).$$

Two remarks about this semantics:

Remark 3.2.4. That this is a generalisation of the one-agent case is a fairly obvious fact. If we have a topological model (X, τ, V) we can turn it in a truth-preserving manner into a topological-partitional model for one agent (X, τ, Π, V) by simply setting $\Pi = \{X\}$. Conversely, if (X, τ, Π, V) is a topological-partitional model for one agent, take $x \in X$ and set $\pi = \Pi(x) \in \Pi$; we can easily see that:

$$(X, \tau, \Pi, V), x \vDash \phi \text{ iff } (\pi, \tau|_\pi, V|_\pi), x \vDash \phi.$$

Remark 3.2.5. It is a routine check that, given a topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2)$ and a nonempty set $U \subseteq X$, the following conditions are equivalent:

1. U is locally dense in Π_i;
2. U is locally dense in π for each $\pi \in \Pi_i[U]$;
3. For each $\pi \in \Pi_i[U]$, $U \cap \pi$ is dense in the subspace topology

$$\tau_i|_\pi := \{V \cap \pi : V \in \tau_i\}.$$

And for $x \in X$ and $\phi \in \mathcal{L}_{K_1 K_2}$ the following conditions are also equivalent:

1. $x \in \|K_i \phi\|$;
2. There exists a Π_i-locally dense open set U such that $x \in U \subseteq \|\phi\|$;
3. There exists an open set $U \in \tau_i$ which is locally dense in $\Pi_i(x)$ such that $x \in U \subseteq \|\phi\|$.
4. There exists an open set $V \subseteq \Pi_i(x)$ which is dense in the subspace topology $\tau_i|_\pi$ such that $x \in V \subseteq \|\phi\|$.

With this in mind, let us see an analogue of lemma 2.1.2: consider a topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$ and set

$$\tau_i^* := \{U \in \tau_i : U \text{ is } \Pi_i\text{-locally dense}\} \cup \{\varnothing\}.$$

The following holds:

Lemma 3.2.6. *(X, τ_i^*) is an extremally disconnected topological space and the locally-dense-interior semantics on $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$ coincides with the interior semantics on $(X, \tau_1^*, \tau_2^*, V)$.*

In particular, given a topological-partitional model $(X, \tau_{1,2}, \Pi_{1,2}, V)$ in which every τ_i-open set is Π_i-locally dense, the locally-dense-interior semantics and the interior semantics coincide.

Proof. The last two statements are immediate consequences of the previous remark, whereas the proof of the first is completely analogous to that of lemma 2.1.2. ■

One last remark before moving on:

Remark 3.2.7. Demanding each element $\pi \in \Pi_i$ to be open might seem like a very strong condition. For example, a connected space such as $\mathbb{R}$ does not admit any such partition other than the trivial one $\Pi_i = \{\mathbb{R}\}$. We could instead do the following:

i. Define topological-partitional models to have arbitrary partitions;

ii. Define $U \subseteq X$ to be locally dense at $\pi \in \Pi_i$ whenever $U \cap \pi$ is dense in the subspace topology $\tau_i|_\pi$;

iii. Set $x \in \|K_i\phi\|$ if and only if there exists $U \in \tau_i$ locally dense in $\Pi_i(x)$ with $x \in U \cap \Pi_i(x) \subseteq \|\phi\|$.

As it turns out, these models can be turned in a truth-preserving manner into topological-partitional models of the kind defined above. Indeed, let $\bar{\tau}_i$ be the topology generated by $\{U \cap \pi : U \in \tau_i, \pi \in \Pi_i\}$. Then clearly $\Pi_i \subseteq \bar{\tau}_i$ and it is a straightforward check that $(X, \tau_i, \Pi_i), x \vDash \phi$ under this semantics if and only if $(X, \bar{\tau}_i, \Pi_i), x \vDash \phi$ under the semantics defined in 3.2.3.

For this reason, we will limit ourselves to the study of models with open partitions.

Let us see an example before moving on:

Example 3.2.8. We have four possible worlds, $X = \{x_{11}, x_{01}, x_{10}, x_{00}\}$ and two agents, Alice and Bob, represented by a and b. Let us consider two propositions, p, which stands for "the house is on fire" and q, standing for "the fire sprinklers went off". Let $V(p) = P = \{x_{11}, x_{10}\}$ and $V(q) = \{x_{11}, x_{01}\}$. The actual world is x_{11}, in which the house is on fire and the fire sprinklers have gone off.

Alice has an app on her phone which lets her know infallibly whether the fire sprinklers have gone off or not. That is, at q-worlds she only considers q-worlds possible, and at $\neg q$-worlds, she only considers $\neg q$-worlds possible. In addition to this, at p-worlds she has fallible evidence that her house is on fire, in the form of a neighbour texting her about the smoke. At $\neg p$-worlds she does not receive this evidence.

At all worlds except x_{01}, Bob installed the fire sprinklers and he (infallibly) knows them to be infallible: if they have gone off, then there is a fire. Thus the only worlds consistent with his information at these three worlds are those in which $q \to p$ holds. He is not as confident when it comes to the fire detector, which lets him know whether a house he has worked on is on fire or not: in p-worlds he has fallible evidence for p and in $\neg p$-worlds he has it for $\neg p$.

Let $\pi_1 = \{x_{11}, x_{01}\}$, $\pi_2 = \{x_{01}, x_{00}\}$, $\pi_3 = \{x_{11}, x_{10}, x_{00}\}$, $\pi_4 = \{x_{01}\}$. Alice's and Bob's partitions are respectively $\Pi_a = \{\pi_1, \pi_2\}$ and $\Pi_b = \{\pi_3, \pi_4\}$. Their topologies τ_a and τ_b are generated respectively by $\{\pi_1, \pi_2, P\}$ and $\{\pi_3, \pi_4, P, X \backslash P\}$. (See figure 3.1.)

At the actual world x_{11}, Alice knows her house is on fire: indeed, $\{x_{11}\}$ is a τ_a-open set, locally dense in π_1 and contained in P, thus $K_a p$ holds. Bob does not know this fact: any τ_b-open set contained in P is not locally dense, because it has empty intersection with the open set $\{x_{00}\}$, thus $\neg K_b p$ holds at x_{11}.

3.2.1 Connected components as equivalence classes

Many of the topological spaces considered in the present thesis come from preorders. As it turns out, preorders come equipped with open partitions in

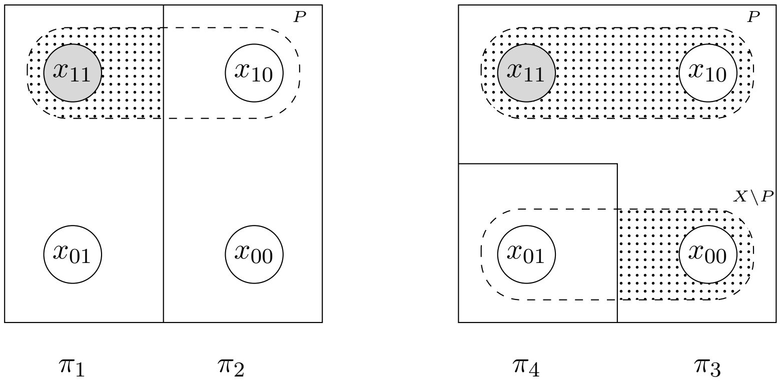

Figure 3.1: The topology and partition of Alice (left) and Bob (right). The dotted areas are the proper open subsets of the cell of each partition which includes the actual world. We can see that x_{11} is in a π_1-locally dense open set but not in a π_3-locally dense one.

the form of their *connected components*.

Definition 3.2.9. Let (X, τ) be a topological space. A set $U \subseteq X$ is said to be *connected* if it does not contain a proper clopen subset.

A *connected component* of (X, τ) is a maximal connected subset of X.

The following result can be found in any topology textbook (see e.g. Munkres, 2000):

Lemma 3.2.10. *The connected components of (X, τ) coincide with the equivalence classes of the equivalence relation: $x \sim y$ if and only if there is a connected subset of X containing x and y.*

The connected components thus give us a straightforward way to define a partition on a topological space. It might however not be the case that the elements of the partition are open sets, thus defining a topological-partitional model this way might not be possible. As the following lemma shows, this is not a problem if the topological space is a preorder (i.e. an Alexandroff space).

Lemma 3.2.11. *Let $(W, \leq)$ be a preordered set. Then:*

i. *The connected components on $(W, \mathrm{Up}(W))$ are open and they coincide with the equivalence classes under the reflexive, transitive and symmetric closure of $\leq$, i.e. the following equivalence relation: $x \sim y$ if and only if there exist $x_0, ..., x_n \in X$ with $x_0 = x, x_n = y$ and $x_k \leq x_{k+1}$ or $x_k \geq x_{k+1}$ for $0 \leq k \leq n-1$.*

ii. *If $(W, \leq)$ is an $\mathsf{S4.2}$ frame (i.e. if $\leq$ is a weakly directed preorder) we have: $x \sim y$ if and only if there exists some $z \in W$ such that $x \leq z \geq y$.*

iii. *If* $(W, \leq)$ *is a forest (i.e. if* $\leq$ *is the reflexive and transitive closure of some relation* $\prec$ *such that every element has at most one* $\prec$*-predecessor), then* $x \sim y$ *if and only if there exists some* $z \in W$ *such that* $x \geq z \leq y$*.*

Proof. (i). Clearly the equivalence class of x is both upward and downward closed, thus every equivalence class is a clopen set. Moreover, it does not contain proper clopen subsets, for if $\varnothing \neq U \subseteq [x]_\sim$ is a clopen set, then take $y \in U$ and $z \in [x]_\sim$. Since there is a path of $\leq$ and $\geq$ from y to z and U is both an upset and a downset, we have that $z \in U$, thus $[x]_\sim$ is connected. And of course it is maximal, for every proper superset of $[x]_\sim$ contains $[x]_\sim$ as a proper clopen subset.

(ii). From right to left it is trivial. From left to right, let us take a path $\{x_0 = x, x_1, ..., x_n = y\}$ such that $x_k \leq x_{k+1}$ or $x_{k+1} \leq x_k$ for all $0 \leq k \leq n-1$. By transitivity, we can assume that $\leq$ and $\geq$ alternate, i.e., for each $1 \leq k \leq n-1$ we have that either $x_{k-1} \leq x_k \geq x_{k+1}$ or $x_{k-1} \geq x_k \leq x_{k+1}$. Now, set $y_0 = x_0$ and $y_n = x_n$. For each $1 \leq k \leq n-1$, if $y_{k-1} \leq x_k \geq x_{k+1}$, then set $y_k = x_k$ and, if $y_{k-1} \geq x_k \leq x_{k+1}$ we have that, since y_{k-1} and x_{k+1} are both successors of x_k and $\leq$ is a directed relation, there must exist some $t \in X$ such that $y_{k-1} \leq t \geq x_{k+1}$. In this case set $y_k = t$. What we get at the end of this process is a path of the form $x = y_0 \leq ... \leq y_k \geq ... \geq y_n = y$, and by setting $z = y_k$ we obtain the desired result.

(iii). If $\leq$ is the reflexive and transitive closure of $\prec$, we can also define $\sim$ on W in terms of $\prec$: $x \sim y$ if there are $x_0 = x, x_1, ..., x_n = y$ with $x_k \prec x_{k+1}$ or $x_k \succ x_{k+1}$ for all k. The proof goes exactly as in the previous item, noting that $x_{k-1} \prec x_k \succ x_{k+1}$ implies $x_{k-1} = x_{k+1}$. ■

Note that item (ii) entails that each upset in a directed preorder is $\sim$-locally dense. Indeed, take x and y in the same equivalence class. Item (ii) gives us that $\uparrow x \cap \uparrow y \neq \varnothing$, thus every pair of nonempty upsets contained in the same connected component has nonempty intersection.

This fact plus the last item in lemma 3.2.6 have an immediate consequence:

Corollary 3.2.12. *Let* $(X, \leq_1, \leq_2, \sim_1, \sim_2, V)$ *be a model in which each* $\leq_i$ *is a weakly directed preorder and* $\sim_i$ *is the equivalence relation given by:* $x \sim_i y$ *if and only if there exists* $z \in X$ *such that* $x \leq_i z \geq_i y$*. Then the locally-dense-interior semantics on this model coincide with the Kripke semantics on* $(X, \leq_1, \leq_2, V)$*.*

3.3 The fusion logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$

The logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ is the logic that simply includes the axioms and rules of S4.2 for both modalities K_1 and K_2. As advanced in the introduction to this chapter, we have the following:

Theorem 3.3.1. $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ *is the logic of topological-partitional models for two agents.*

Proof. Checking soundness is completely analogous to the one-agent case (see Baltag et al., 2016).

Completeness (plus the finite model property of $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$) follow from corollary 3.2.12 and the following fact:

Lemma 3.3.2. $\mathsf{S4.2} + \mathsf{S4.2}$ *is sound and complete with respect to the class of finite frames of the form* $(W, \leq_1, \leq_2)$ *where each* $\leq_i$ *is a directed preorder.*

Now take $\phi \notin \mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$. Then ϕ is refuted on some birelational weakly directed preorder $(W, \leq_1, \leq_2, V)$ and therefore, as per corollary 3.2.12, it is refuted on the topological-partitional model $(W, \mathrm{Up}_{\leq_1}(W), \mathrm{Up}_{\leq_2}(W), \Pi_1, \Pi_2, V)$, where Π_i is the set of $\leq_i$-connected components. ■

3.4 Other fragments

Let us now consider other fragments of the logic. For this we add to our language the *infallible knowledge modalities* $[\forall]_i$, the *evidence modalities* $\Box_i$, and the *belief modalities* B_i, for $i = 1, 2$, and their respective duals $[\exists]_i$, $\Diamond_i$ and $\hat{B}_i$. We read these on topological-paritional models $(X, \tau_{1,2}, \Pi_{1,2}, V)$ as follows:

$$x \in \|[\forall]_i\phi\| \text{ iff } \Pi_i(x) \subseteq \|\phi\|;$$
$$x \in \|\Box_i\phi\| \text{ iff } x \in \mathrm{Int}_{\tau_i} \|\phi\|;$$
$$x \in \|B_i\phi\| \text{ iff } \mathrm{Int}_{\tau_i} \|\phi\| \text{ is locally dense in } \Pi_i(x).$$

Analogously to the one-agent case, we can check that the following equalities hold:

$$\|K_i\phi\| = \|\Box_i\phi \wedge [\forall]_i\Diamond_i\Box_i\phi\|$$
$$\|B_i\phi\| = \|\hat{K}_iK_i\phi\|$$

We can also tweak our models to be able to talk about *basic evidence*:

Definition 3.4.1. A *topological-partitional evidence model* (topo-part-e model for short) is a tuple

$$\mathfrak{M} = (X, \tau_1, \tau_2, \Pi_1, \Pi_2, E_1^0, E_2^0, V)$$

where $(X, \tau_{1,2}, \Pi_{1,2}, V)$ is a topological-partitional model and each E_i^0 is a subbasis for τ_i.

We define *basic evidence modalities* $\Box_i^0$ (with dual $\Diamond_i^0$) on topo-part-e models by:

$$x \in \|\Box_i^0\phi\| \text{ iff there exists } e \in E_i^0 \text{ with } x \in e \subseteq \|\phi\|.$$

Much like in the one-agent framework, we are interested in looking at fragments of this logic. We will focus on the *knowledge fragment* $\mathcal{L}_{K_i\forall_i}$, the *knowledge-belief fragment* $\mathcal{L}_{K_iB_i}$, and the *factive evidence fragment* $\mathcal{L}_{\Box_i\forall_i}$.

3.4.1 The factive evidence fragment $\mathcal{L}_{\Box_i\forall_i}$

The logic for this fragment is $\mathsf{Logic}_{\Box_i\forall_i}$, which is the least logic which includes

- the axioms and rules of $\mathsf{S4}$ for $\Box_i$;
- the axioms and rules of $\mathsf{S5}$ for $[\forall_i]$;
- the axiom $[\forall_i]\phi \to \Box_i\phi$ for $i = 1, 2$.

Soundness for topological-partitional models is a rather simple check: the $\mathsf{S4}$ rules for the topological interior hold, for $\mathrm{Int}\, P \subseteq P \cap \mathrm{Int}\,\mathrm{Int}\, P$ and so do the $\mathsf{S5}$ rules for $[\forall]_i$, which are defined via equivalence relations. The fact that each equivalence class is open takes care of the axiom $[\forall]_i\phi \to \Box_i\phi$.

For completeness, we can use the Sahlqvist completeness theorem (see Blackburn et al., 2001) and note that the axioms of $\mathsf{Logic}_{\Box_i\forall_i}$ are Sahlqvist formulas and thus canonical and the canonical Kripke model for this logic is of the shape $(X, \leq_1, \leq_2, \sim_1, \sim_2)$, where each $\leq_i$ is a preorder (as per the $\mathsf{S4}$ axioms) and each $\sim_i$ constitutes an equivalence relation (as per the $\mathsf{S5}$ axioms). Moreover, the axiom $[\forall_i]\phi \to \Box_i\phi$ grants us that $x \leq_i y$ implies $x \sim_i y$ and thus that the $\sim_i$-equivalence classes are $\leq_i$-open sets. In other words, this canonical model is a topological-partitional model.

Therefore if $\phi \notin \mathsf{Logic}_{\Box_i\forall_i}$, then ϕ will be refuted in the canonical model, whence we have a topological-partitional model refuting it. And thus, we have completeness.

Theorem 3.4.2. $\mathsf{Logic}_{\Box_i\forall_i}$ *is sound and complete with respect to topological-partitional models.* ■

3.4.2 The knowledge fragment $\mathcal{L}_{K_i\forall_i}$

The logic of the fragment with all the knowledge modalities, $K_1, K_2, [\forall]_1$ and $[\forall]_2$ is $\mathsf{Logic}_{K_i\forall_i}$, the least logic including the axioms and rules of $\mathsf{S4}$ for each K_i, $\mathsf{S5}$ for each $[\forall]_i$ plus the following axioms for $i = 1, 2$:

$$\begin{array}{ll}(\mathrm{A}) & [\forall]_i\phi \to K_i\phi;\\ (\mathrm{B}) & [\exists]_i K_i\phi \to [\forall]_i\hat{K}_i\phi.\end{array}$$

Note that the .2 axiom for K_i is derivable from (A) and (B).

Soundness is a routine check, whereas for completeness we can again resort to Sahlqvist. The canonical model is of the shape $(X, \leq_1, \leq_2, \sim_1, \sim_2)$ where each $\leq_i$ is a weakly directed preorder and each $\sim_i$ is an equivalence relation. Moreover the Sahlqvist first order correspondent of axiom (A) gives us that $x \leq_i y$ implies $x \sim_i y$ and axiom (B) tells us that, if $x \sim_i y$, then there exists some z such that $x \leq_i z \geq_i y$. These two facts, together with item (ii) of lemma 3.2.11, imply that the $\sim_i$-equivalence classes are exactly the $\leq_i$-connected components. And thus the Kripke semantics on this model coincide with the locally-dense-interior semantics on the topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2)$ where $\tau_i = \text{Up}\leq_i(X)$ and Π_i are the $\leq_i$-connected components. Completeness follows.

Theorem 3.4.3. $\mathsf{Logic}_{K_i\forall_i}$ *is sound and complete with respect to topological-partitional models.* ■

3.4.3 The knowledge-belief fragment $\mathcal{L}_{K_iB_i}$

The logic of the knowledge-belief fragment is $\mathsf{Stal}_1 + \mathsf{Stal}_2$, i.e. the logic that consists of adding together the axioms and rules of Stal (introduced in subsection 1.2.2) for K_1 and B_1 and for K_2 and B_2. Explicitly, the $\mathsf{S4}$ axioms and rules for K_i plus the axioms, for $i = 1, 2$, in table 3.1.

(PI$_i$)	$B_i\phi \to K_iB_i\phi$;
(NI$_i$)	$\neg B_i\phi \to K_i\neg B_i\phi$;
(KB$_i$)	$K_i\phi \to B_i\phi$;
(CB$_i$)	$B_i\phi \to \neg B_i\neg\phi$;
(FB$_i$)	$B_i\phi \to B_iK_i\phi$.

Table 3.1: Extra axioms of $\mathsf{Stal}_1 + \mathsf{Stal}_2$.

Now, it is easy to check that these axioms are sound. For completeness, let us note that, in an analogous manner to the one-agent case (subsection 1.2.2) we have:

i. $\mathsf{Stal}_1 + \mathsf{Stal}_2 \vdash B_i\phi \leftrightarrow \hat{K}_iK_i\phi$ for $i = 1, 2$;

ii. $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ is contained in $\mathsf{Stal}_1 + \mathsf{Stal}_2$.

Thus if a formula ϕ in the language $\mathcal{L}_{K_iB_i}$ is not provable in $\mathsf{Stal}_1 + \mathsf{Stal}_2$, we can rewrite it as per (i.) into a formula in the language $\mathcal{L}_{K_i}$ which, by (ii.), is not provable in $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$. By completeness of the latter, there is a topological-partitional countermodel for ϕ, and completeness of $\mathsf{Stal}_1 + \mathsf{Stal}_2$ follows.

Chapter 4

Generic models for two agents

So far so good. We have a multi-agent logic for the class of topological-partitional models. However, we can get something else. In the spirit of chapter 2, we want to find a *generic model* for the two-agent logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$, that is, a single topological-partitional model whose logic is precisely $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

In the first section we find such a space in the quaternary tree $\mathcal{T}_{2,2}$, drawing on a similar result by van Benthem, Bezhanishvili, ten Cate, and Sarenac (2006) for $\mathsf{S4}+\mathsf{S4}$ with the Kripke semantics, finding a suitable partition and showing completeness. In section 2 we will account for the fact that there exists a surjective interior map onto the quaternary tree from $\mathbb{Q}\times\mathbb{Q}$ with two particular topologies and show that there exists a partition on $\mathbb{Q}\times\mathbb{Q}$ which makes it into a generic model.

4.1 The quaternary tree $\mathcal{T}_{2,2}$

The quaternary tree $\mathcal{T}_{2,2}$ is a full infinite tree with two relations R_1 and R_2 such that each node of the tree has exactly four successors, two of them being R_1-successors and the other two being R_2-successors, as it appears in figure 4.1.

By setting T to be the set of points of $\mathcal{T}_{2,2}$ and $\leq_i$ to be the reflexive and transitive closure of R_i for $i = 1, 2$, we can see $\mathcal{T}_{2,2} = (T, \leq_1, \leq_2)$ as a birelational preordered frame. As it turns out, the logic of this frame (under the Kripke semantics) is $\mathsf{S4}+\mathsf{S4}$.

4.1.1 Completeness of $\mathsf{S4}+\mathsf{S4}$: defining an interior map

The following result is proven in van Benthem et al. (2006):

Theorem 4.1.1. *$\mathcal{T}_{2,2}$ is sound and complete with respect to $\mathsf{S4}+\mathsf{S4}$.*

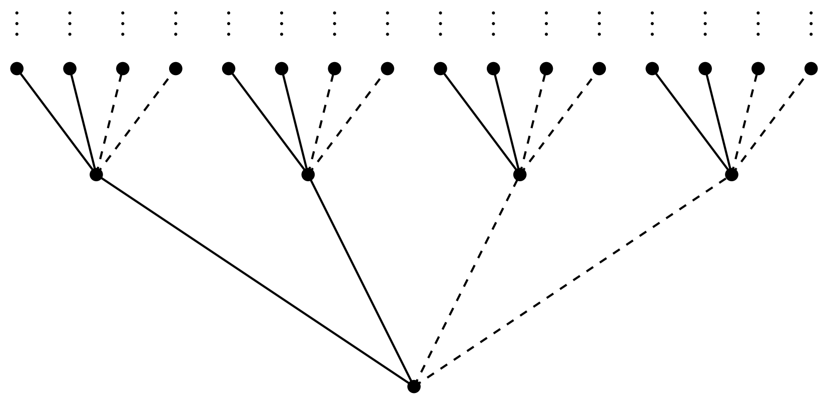

Figure 4.1: The quaternary tree $\mathcal{T}_{2,2}$. R_1 and R_2 are represented respectively by the continuous and the dashed lines.

The proof of this result goes along the lines of the proofs in chapter 2, by building a p-morphism from $\mathcal{T}_{2,2}$ onto any rooted finite birelational preordered model $(W, \leq_1, \leq_2, V)$. We will use this p-morphism in the next subsection. Let us sketch its construction here.

Let $w \in W$ and let $v_1, ..., v_n$ be an enumeration of w's $\leq_i$-successors. The *finite* $\leq_i$ *plug* for w is a binary branching tree whose root is labelled by w and has two successors, the left one labelled by w and the right one being a leaf labelled by v_1. The left successor is itself related to two points, the right one being a leaf labelled by v_2 and the left one being labelled by w and having itself two successors, etcetera. Repeating this process n times until we run out of successors gives this finite $\leq_i$ plug, drawn in in figure 4.2.

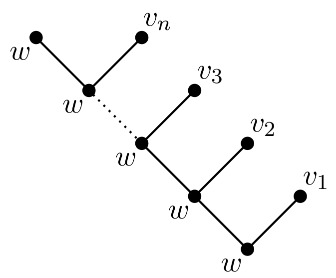

Figure 4.2: Finite $\leq_i$ plug for $w \in W$.

Now we do the following construction: we start with a single node, which we label with the root r of W. Then on that node we "plug" (in the sense that we "make sprout" from it) the $\leq_1$ finite plug for r and the $\leq_2$ finite plug for r (see fig. 4.3). In each successive step, for each node, if it is labelled by w, we make sprout from it the $\leq_1$ and $\leq_2$ finite plugs for w if we have not done it in a previous step.

What we obtain in the limit by repeating this process recursively is pre-

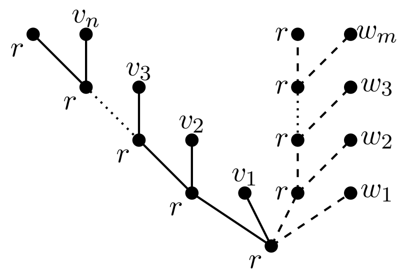

Figure 4.3: First step of the construction, with $v_1, ..., v_n$ being the $\leq_1$-successors of r and $w_1, ..., w_m$ being its $\leq_2$-successors.

cisely the quaternary tree $\mathcal{T}_{2,2}$ with each of its nodes labelled by some element in W. We thus have a (surjective) map

$$\mathsf{label} : \mathcal{T}_{2,2} \twoheadrightarrow W$$

assigning to each point its label. Two remarks about this map:

- If a node $x \in \mathcal{T}_{2,2}$ is labelled by w, every node which x sees via $\leq_i$ is by construction labelled by some $\leq_i$-successor of w. This takes care of the forth condition of the p-morphism:

$$x \leq_i y \text{ implies } \mathsf{label}\, x \leq_i \mathsf{label}\, y.$$

- If a point $w \in W$ has v as a $\leq_i$-successor, any node labelled by w in the tree has a successor labelled by v. This takes care of the back condition of the p-morphism:

$$\mathsf{label}\, x \leq_i v \text{ implies } \exists y (x \leq_i y \,\&\, \mathsf{label}\, y = v).$$

Thus, label is a surjective map which is continuous and open in both topologies. If we now define a valuation on $\mathcal{T}_{2,2}$ by

$$V^{\mathcal{T}_{2,2}}(p) = \{x \in \mathcal{T}_{2,2} : \mathsf{label}\, x \in V(p)\},$$

we can refute in $\mathcal{T}_{2,2}$ any formula which is refuted in W. Completeness of $\mathsf{S4} + \mathsf{S4}$ follows.

4.1.2 Completeness of $\mathcal{T}_{2,2}$ with respect to $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$

Let us now bring this to our realm. We want to think of $\mathcal{T}_{2,2}$ as a topological-partitional model. For this, we turn to its connected components.

As per item (iii) of lemma 3.2.11, we know that the connected components are given by the equivalence relation: $x \sim_i y$ if and only if there exists a z such that $x \geq_i z \leq_i y$. Note that for each $x \in \mathcal{T}_{2,2}$ and $i = 1, 2$, the set of $\leq_i$-predecessors of x forms a finite chain (and in particular, there is a least predecessor x_0 of x, which does not have any $\leq_i$ predecessors other than itself). These two facts give us the following characterisation:

Lemma 4.1.2. *The $\leq_i$-connected components of $\mathcal{T}_{2,2}$ are exactly the upsets of the form $\uparrow_i x_0$, where x_0 does not have any $\leq_i$-predecessors other than itself.*

Now, let $(W, \leq_1, \leq_2, V)$ be a finite model whose underlying frame is a rooted birelational weakly directed preorder. We can define a map $\mathsf{label} : \mathcal{T}_{2,2} \twoheadrightarrow W$ and a valuation $V^{\mathcal{T}_{2,2}}$ as above. Let σ_i be the topology of $\leq_i$-upsets of W and $\equiv_i$ be the equivalence relation determining the connected components. Recall that $\mathfrak{W} = (W, \sigma_{1,2}, \equiv_{1,2}, V)$ is a topological-partitional model in which every σ_i-open set is $\equiv_i$-locally dense. Moreover, we have:

Lemma 4.1.3. *For $x \in \mathcal{T}_{2,2}$, $w \in W$ and $i = 1, 2$, let $[x]_{\sim_i}$ and $[w]_{\equiv_i}$ be the respective equivalence classes (i.e. the respective connected components containing x and w). The following holds:*

i. *For any $x \in \mathcal{T}_{2,2}$, $\mathsf{label}[x]_{\sim_i} \subseteq [\mathsf{label}\, x]_{\equiv_i}$.*

ii. *Let $x_0 \in \mathcal{T}_{2,2}$ and let U be a (locally dense) σ_i-open set such that $\mathsf{label}\, x_0 \in U \subseteq [\mathsf{label}\, x_0]_{\equiv_i}$. Then*

$$U' := \bigcup \{\uparrow_i x : x \sim_i x_0 \,\&\, \mathsf{label}\, x \in U\}$$

is a locally dense upset such that $x_0 \in U' \subseteq [x_0]_{\sim_i}$.

Proof. (i). Set $y \sim_i x$. Then there is some z such that $y \geq_i z \leq_i x$ and thus, since the map label preserves order, we have that $\mathsf{label}\, y \geq_i \mathsf{label}\, z \leq_i \mathsf{label}\, x$ and thus $\mathsf{label}\, y \equiv_i \mathsf{label}\, x$.

(ii). U' is an upset because it is a union of upsets and $x_0 \in U' \subseteq [x_0]_{\sim_i}$ by construction. Let us see that it is locally dense. Take some $z \in \mathcal{T}_{2,2}$ such that $\uparrow_i z \subseteq [x_0]_{\sim_i}$. Now, $\mathsf{label}(\uparrow_i z)$ is an open set (by opennes of label) and $\mathsf{label}(\uparrow_i z) \subseteq \mathsf{label}[x_0]_{\sim_i} \subseteq [\mathsf{label}\, x_0]_{\equiv_i}$. By local density of U there exists some $a \in U \cap \mathsf{label}(\uparrow_i z)$. That is, for some $z' \geq_i z$ we have $\mathsf{label}\, z' = a$ and $\mathsf{label}\, z' \in U$, thus by construction $z' \in \uparrow_i z \cap U'$ and thus $\uparrow_i z \cap U' \neq \varnothing$. ∎

As a consequence:

Proposition 4.1.4. *For any $x \in \mathcal{T}_{2,2}$ and any formula ϕ in the language, $\mathcal{T}_{2,2}, x \vDash \phi$ if and only if $\mathfrak{W}, \mathsf{label}\, x \vDash \phi$.*

Proof. This is once again an induction on the structure of formulas in which the only involved case is the induction step corresponding to the K_i modalities.

Suppose $x \vDash K_i \phi$. Then there exists some locally dense open set U with $x \in U \subseteq [x]_{\sim_i}$ such that $y \vDash \phi$ for all $y \in U$. But then

$$\mathsf{label}\, x \in \mathsf{label}\, U \subseteq \mathsf{label}[x]_{\sim_i} \subseteq [\mathsf{label}\, x]_{\equiv_i},$$

this last inclusion given by (i) of the previous lemma, and $\mathsf{label}\, U$ is a locally dense open set in W: it is open because label is an open map and it is locally dense because every open set in W is locally dense. Moreover, for every $\mathsf{label}\, y \in \mathsf{label}\, U$ we have by induction hypothesis that $\mathsf{label}\, y \vDash \phi$. Thus $\mathsf{label}\, x \vDash K_i\phi$.

Conversely, suppose $\mathsf{label}\, x \vDash K_i\phi$. Then there exists a (locally dense) σ_i-open set U with $\mathsf{label}\, x \in U \subseteq [\mathsf{label}\, x]_{\equiv_i}$ such that $w \vDash \phi$ for all $w \in U$. But then by part (ii) of the previous lemma

$$U' := \bigcup\{\uparrow_i z : z \sim_i x \,\&\, \mathsf{label}\, z \in U\}$$

is a locally dense upset such that $x \in U' \subseteq [x]_{\sim_i}$. Now take $y \in U'$. We have that $y \geq_i z$ for some $z \in [x]_{\sim_i}$ with $\mathsf{label}\, z \in U$. But since label is order preserving we have that $\mathsf{label}\, y \geq_i \mathsf{label}\, z$ and thus $\mathsf{label}\, y \in U$, which means that $\mathsf{label}\, y \vDash \phi$ and thus, by induction hypothesis, $y \vDash \phi$. This means that $U' \subseteq \|\phi\|^{\mathcal{T}_{2,2}}$ and thus $x \vDash K_i\phi$. ■

Completeness is now an immediate consequence.

Corollary 4.1.5. *The quaternary tree* $(\mathcal{T}_{2,2}, \leq_1, \leq_2, \sim_1, \sim_2)$ *is a generic model for* $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

4.2 Completeness for $\mathbb{Q} \times \mathbb{Q}$

In the present section we will show that it is possible to define two topologies and two equivalence relations on the product space $\mathbb{Q} \times \mathbb{Q}$ which make it into a generic topological-partitional space for $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

These topologies will be the vertical and horizontal topologies, which can be defined on a product $X \times Y$ and, in a way, "lift" the topologies of the components.

Definition 4.2.1. Let (X, τ) and (Y, σ) be two topological spaces. A set $U \subseteq X \times Y$ is said to be *horizontally open* if, for every $(x, y) \in U$, there exists $V \in \tau$ such that $x \in V$ and $V \times \{y\} \subseteq U$. The *horizontal topology* τ_H is the topology defined on $X \times Y$ by the horizontally open sets or, equivalently, the topology on $X \times Y$ generated by the basis

$$\mathcal{B}_H = \{U \times \{y\} : U \in \tau, y \in Y\}.$$

Similarly, the *vertical topology* τ_V is the topology on $X \times Y$ generated by the basis

$$\mathcal{B}_V = \{\{x\} \times V : x \in X, V \in \sigma\}.$$

In particular, if we take both components to be $\mathbb{Q}$ with the natural topology, we obtain our bitopological space $(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V)$. An important result about this space is the following:

Theorem 4.2.2 (van Benthem et al., 2006)**.** $\mathsf{S4} + \mathsf{S4}$ *is the logic of* $(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V)$ *under the interior semantics.*

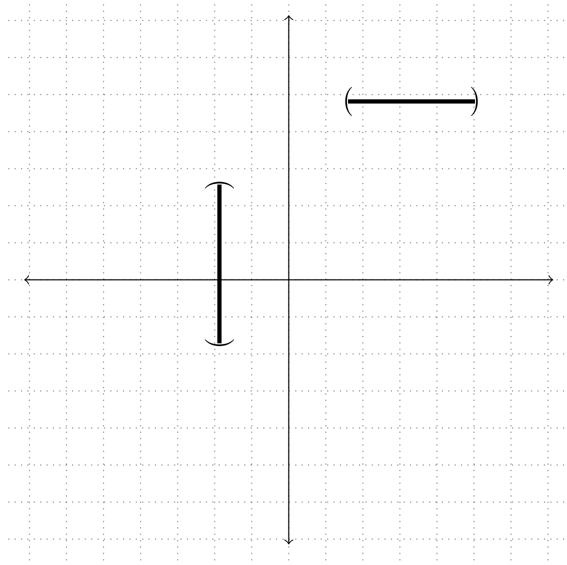

Figure 4.4: $\mathbb{Q} \times \mathbb{Q}$ with a vertically open set and a horizontally open set

Now the idea is to be able to define some partition on $\mathbb{Q} \times \mathbb{Q}$ which will give us the desired completeness result. Note that we cannot shelter ourselves in the connected components this time, for the connected components in $(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V)$ are the singletons, which are not even open sets. The main aim of this section is to show that such a partition exists.

Now, let (X, τ_1, τ_2) be a bitopological space and

$$\mathfrak{Y} = (Y, \sigma_1, \sigma_2, \sim_1, \sim_2, V)$$

be a topological partitional model. Moreover, let

$$f : (X, \tau_1, \tau_2) \twoheadrightarrow (Y, \sigma_1, \sigma_2)$$

be a surjective map which is open and continuous in both topologies. Define two equivalence relations $\equiv_1$ and $\equiv_2$ on X by:

$$x \equiv_i y \text{ if and only if } fx \sim_i fy.$$

Define a valuation on X by $V^f(p) = \{x \in X : fx \in V(p)\}$. The following holds:

Proposition 4.2.3. $\mathfrak{X} = (X, \tau_1, \tau_2, \equiv_1, \equiv_2, V^f)$ *is a topological evidence model and, for every formula* ϕ *in the language and every* $x \in X$ *we have that* $\mathfrak{X}, x \vDash \phi$ *if and only if* $\mathfrak{Y}, fx \vDash \phi$.

Proof. Checking that $\mathfrak{X}$ is a topological partitional model amounts to checking that each equivalence class is an open set. Let $[x]_{\equiv_i}$ be the equivalence class under $\equiv_i$ of some $x \in X$. Note that, by definition of $\equiv_i$, the image of this class coincides with the equivalence class of fx, i.e. $f[x]_{\equiv_i} = [fx]_{\sim_i}$. Indeed, if $fx \sim_i fy$ then $y \in [x]_{\equiv_i}$ and thus $fy \in f[x]_{\equiv_i}$ and conversely if $y \in f[x]_{\equiv_i}$ then $y = fx'$ for some $x' \equiv_i x$ which means that $y = fx' \sim_i fx$. Now, $[fx]_{\sim_i}$ is an equivalence class and thus an open set and, since f is continuous, $f^{-1}f[x]_{\equiv_i}$ is also an open set. So it suffices to show that $f^{-1}f[x]_{\equiv_i} = [x]_{\equiv_i}$. And indeed, if $z \in f^{-1}f[x]_{\equiv_i}$ then $fz \in f[x]_{\equiv_i} = [fx]_{\sim_i}$ which means that $fz \sim_i fx$ and thus $z \equiv_i x$.

The second result is an induction on formulas. For the propositional variables and the induction steps corresponding to the Boolean connectives the result is straightforward. Now suppose that for some ϕ it is the case that, for all x, $\mathfrak{X}, x \vDash \phi$ if and only if $\mathfrak{Y}, fx \vDash \phi$, and let $\mathfrak{X}, x \vDash K_i\phi$. This means that there exists some open set $U \in \tau_i$ such that $x \in U \subseteq \|\phi\|^{\mathfrak{X}}$ and U is locally dense in $[x]_{\equiv_i}$, i.e., for every nonempty open set $V \subseteq [x]_{\equiv_i}$, it is the case that $U \cap V \neq \varnothing$. But then we have that $fx \in f[U]$, $f[U]$ is an open set (by openness of f) which is contained in $f\|\phi\|^{\mathfrak{X}}$ (and thus, by induction hypothesis, in $\|\phi\|^{\mathfrak{Y}}$) and $f[U]$ is locally dense in $[fx]_{\sim_i}$. Indeed, suppose V is an open set contained in $[fx]_{\sim_i}$. then $f^{-1}V$ is an open set contained in $f^{-1}[fx]_{\sim_i} = [x]_{\equiv_i}$ which implies that there exists some $z \in f^{-1}V \cap U$ and thus some $fz \in V \cap f[U]$. Conversely, suppose $\mathfrak{Y}, fx \vDash K_i\phi$. There is an open set $U \subseteq \|\phi\|^{\mathfrak{Y}}$ which includes fx and which is locally dense on $[fx]_{\sim_i}$. Then $f^{-1}U$ is an open set including x which is contained in $f^{-1}\|\phi\|^{\mathfrak{Y}} = \|\phi\|^{\mathfrak{X}}$ and moreover it is locally dense on $[x]_{\equiv_i}$: indeed, if V is an open set contained in $[x]_{\equiv_i}$, then fV is an open set contained in $[fx]_{\sim_i}$ and thus there exists some $y \in fV \cap [fx]_{\sim_i}$. But then $y = fz$ for some $z \in V$ and $z \in V \cap f^{-1}[fx]_{\sim_i} = V \cap [x]_{\equiv_i}$, whence $\mathfrak{X}, x \vDash K_i\phi$. ■

Now, the proof of theorem 2.3.3 in van Benthem et al. (2006) uses completeness of $\mathcal{T}_{2,2}$ with respect to $\mathsf{S4} + \mathsf{S4}$ and shows that there exists an onto map $f : \mathbb{Q} \times \mathbb{Q} \to \mathcal{T}_{2,2}$, open and continuous in both τ_H and τ_V. The previous proposition plus this fact grants us the existence of a partition which makes $\mathbb{Q} \times \mathbb{Q}$ a generic model for $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

Corollary 4.2.4. *Let $f : \mathbb{Q} \times \mathbb{Q} \to \mathcal{T}_{2,2}$ be some onto interior map. Define $(x, y) \equiv^f_H (x', y')$ if and only if $f(x, y)$ and $f(x', y')$ belong to the same $\leq_1$-connected component and $(x, y) \equiv^f_V (x', y')$ if and only if their images are in the same $\leq_2$ connected component. Then $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$ is sound and complete with respect to the topological-partitional space*

$$(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V, \equiv^f_H, \equiv^f_V).$$

Chapter 5

Distributed and common knowledge

We have so far a multi-agent framework whose logic simply combines the axioms of the single-agent logic for each of the agents.

Of course, one would like to go beyond this. The reason why having multiple agents in the same framework is desirable goes further than the possibility of modelling what each of them individually knows: we want to be able to consider concepts like what the group knows, what they would come to know after exchanging information, what they know each other to know, and know each other to know each other to know, etcetera.

In the present chapter we consider the notions of *distributed* and *common* knowledge applied to this framework. The first section presents both concepts and comments on their relational semantics, which will inspire the topological semantics introduced later. In section 2, we discuss earlier approaches to distributed knowledge in the dense-interior framework to subsequently motivate and present our own, complete with an axiomatisation. Section 3 is a first approach to a notion of common knowledge in this setting.

5.1 Knowledge of the group in a relational setting

A very basic way to account for the knowledge of a group is via a notion of "everybody knows that". That is, via a modal $E\phi$ which holds whenever all the epistemic agents know that ϕ. $E\phi$ is veridical (in the sense that ϕ holds when $E\phi$ does) but it is not introspective: it would be undesirable if all agents knowing ϕ were to entail that all agents know they all know it. It may well be the case that one agent is not aware of some other knowing ϕ.

This notion seems rather weak and moreover our current framework already has the expressive power to account for it. Indeed, we can just have $E\phi$ as an abbreviation for $K_1\phi \wedge K_2\phi$. A notion directly related to this is that of *common knowledge*. One of the most standard ways[1] to think of common

[1]For other ways, see Barwise (1988). This will be discussed below.

knowledge is, loosely speaking, as an infinite conjunction

$$C\phi = \bigwedge_{n\in\omega} E\phi,$$

where $E^0\phi = \phi$, $E^{n+1}\phi = EE^n\phi$. In other words, ϕ is common knowledge if it is true and everybody knows that it is true and everybody knows that everybody knows, etcetera. As it is put in Fagin, Halpern, Moses, and Vardi (2004), we can read $C\phi$ as "every fool knows that ϕ".

A much weaker notion is that of *distributed knowledge*. We can think of it as whatever the group knows implicitly, or whatever would become known if all the agents were to share their information. Not only does the group know ϕ if one agent in the group knows it, but the group also knows things that no individual agent knows yet can be derived from the information of several agents. For example, if agent 1 knows p to be the case, and agent 2 knows $p \to q$ to be the case, then together they know q, even if individually no one does.

In relational semantics, if $\mathcal{A}$ is a finite group of agents and, for each $a \in \mathcal{A}$, K_a is the Kripke modality corresponding to some relation R_a, then we can think of D as the Kripke modality corresponding to the relation $\bigcap_{a\in\mathcal{A}} R_a$ and C as the one corresponding to $(\bigcup_{a\in\mathcal{A}} R_a)^*$, the reflexive and transitive closure of $\bigcup_{a\in\mathcal{A}} R_a$. (See van Ditmarsch, van der Hoek, and Kooi, 2007.)

Let us put aside common knowledge for a moment and focus on what this relational definition of distributed knowledge means. $D\phi$ holds at a world w if ϕ holds at every world which is reachable by every agent. As van Ditmarsch et al. (2007) put it,

> The idea being that if one agent considers t a possibility, given s, but another does not, the latter could 'inform' the first that he need not consider t.

Now, how do we cash this out topologically?

5.2 Distributed knowledge

We have seen that, in a relational setting, distributed knowledge is the modality of the intersection of the relations. Topologically, if we read knowledge as interior and limit ourselves to Alexandroff topologies, the interior semantics in a bitopological Alexandroff space (X, τ_1, τ_2) correspond to the Kripke semantics on some preorder $(X, \leq_1, \leq_2)$ such that $\tau_i = \mathrm{Up}\leq_i(X)$. If distributed knowledge is the modality corresponding to the preorder $\leq_1 \cap \leq_2$, what is its topological counterpart? The answer resides in the *join topology*.

Definition 5.2.1. Given two topologies τ_1 and τ_2, the *join of* τ_1 *and* τ_2 is the least topology containing both τ_1 and τ_2. More explicitly,

$$\tau_1 \vee \tau_2 := \{U \cap V : U \in \tau_1, V \in \tau_2\}.$$

It is routine to check that (i). $\tau_1 \vee \tau_2$ is a topology, (ii). it contains τ_1 and τ_2 and (iii). any topology containing τ_1 and τ_2 is a superset of it. And so is the following lemma:

Lemma 5.2.2. *If $\leq_1$ and $\leq_2$ are perorders on X,*

$$\mathrm{Up}(\leq_1 \cap \leq_2) = \mathrm{Up}(\leq_1) \vee \mathrm{Up}(\leq_2).$$

Distributed knowledge as the interior in the join topology makes sense if we look at knowledge from an evidential perspective: under this view, the group knows p whenever agent 1 and agent 2 have each a piece of evidence which, when combined, result in a piece of evidence for p.

Although our framework is not as simple, we will be using the join topology to talk about distributed knowledge in this section.

5.2.1 A problematic approach

So what exactly amounts to distributed knowledge in our framework? A very direct way to translate the ideas in the previous section would be this: we say that $D\phi$ holds at w whenever agent 1 and agent 2 have each a piece of evidence which, when put together, constitute a justification for w.

Now, we know that a *justification* is a piece of evidence which cannot be contradicted by any other potential evidence. In the multi-agent framework presented, this potential contradictory evidence is limited to the set of worlds that the agent considers compatible with the world of evaluation x. That is, agent i knows ϕ whenever she has an evidence for ϕ which cannot be trumped by any piece of evidence contained in $\Pi_i(x)$. In this spirit, we can narrow our scope of what amounts to a justification to the set of worlds that both agents consider compatible with their information at x, namely $\Pi_1(x) \cap \Pi_2(x)$.

Formally, $x \in \|D\phi\|$ if and only if there exist $U_1 \in \tau_1, U_2 \in \tau_2$ such that $x \in U_1 \cap U_2 \subseteq \|\phi\|$ and $\mathrm{Cl}_{\tau_1 \vee \tau_2}(U_1 \cap U_2) \supseteq \Pi_1(x) \cap \Pi_2(x)$. Note that this notion corresponds to the interior in the topology $(\tau_1 \vee \tau_2)^*$ of locally dense $\tau_1 \vee \tau_2$-open sets, where "local density" is measured in terms of the partition $\Pi = \{\pi_1 \cap \pi_2 : \pi_i \in \Pi_i\}$.

The trouble with this intuitive approach is that distributed knowledge, if defined like this, does not satisfy the axioms that one would come to expect. For example, it might be the case that something which is known by one agent is not known by the group.

Example 5.2.3 (Example 5.5.3 from Özgün, 2017)**.** Let $X = \{x_{10}, x_{11}, x_{01}\}$, $P = \{x_{10}, x_{11}\}$, $Q = \{x_{11}, x_{01}\}$, and let x_{10} be the actual world. Let τ_m be the topology generated by $\{P, Q\}$, i.e.

$$\tau_m = \{X, \varnothing, P, Q, \{x_{11}\}\}.$$

Let τ_d be the topology $\{X, \varnothing, \{x_{10}, x_{01}\}\}$. Finally, let $\Pi_m = \Pi_d = \{X\}$ and set $V(p) = P$, $V(q) = Q$.

We have that $(X, \tau_m, \tau_d, \Pi_m, \Pi_d, V)$ is a topological partitional model in which $K_m p$ holds at the actual world x_{10}. Indeed, P is a τ_m-dense open set such that $x_{10} \in P \subseteq \|p\|$. However, if we read distributed knowledge in terms of the join topology, Dp does not hold: the join topology is

$$\tau_d \vee \tau_m = \{X, \varnothing, P, Q, \{x_{10}, x_{01}\}, \{x_{10}\}, \{x_{01}\}, \{x_{11}\}\}$$

and no piece of evidence contained in P is dense in this topology, for it will have empty intersection with $\{x_{01}\}$.

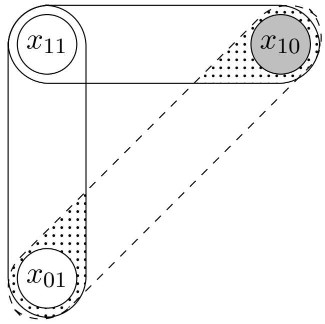

Figure 5.1: The continuous line represents the evidence of agent m, whereas the dashed line represents the evidence of agent d. The singleton sets $\{x_{01}\}$ and $\{x_{10}\}$ become open in the distributed topology.

This is of course an issue and two attempts at solving it are presented in the remainder of this subsection.

Ramírez (2015) acknowledges the problem and proposes the following semantics: models are now of the shape

$$(X, \tau_1, \tau_2, \tau, \Pi_1, \Pi_2, V),$$

where $(X, \tau_{1,2}, \Pi_{1,2}, V)$ is a topological-partitional model in our sense and τ is a topology which contains both τ_1 and τ_2. The language includes the modalities K_1, K_2 and D and modal sentences are read as follows:

$x \in \|K_i\phi\|$ iff there is some $U \in \tau_i$ such that $x \in U \subseteq \|\phi\|$ and $\mathrm{Cl}_{\tau|_{\Pi_i(x)}} U = \Pi_i(x)$;

$x \in \|D\phi\|$ iff there exist $U_1 \in \tau_1$ and $U_2 \in \tau_2$ such that $x \in U_1 \cap U_2 \subseteq \|\phi\|$ and $\mathrm{Cl}_{\tau|_{\Pi_1(x)\cap\Pi_2(x)}}(U \cap V) = \Pi_1(x) \cap \Pi_2(x)$.

Ramírez (2015) then provides a complete logic for these models, which includes axioms that one might desire such as $K_i\phi \to D\phi$ and $\hat{K}_i K_i \phi \to K_i \hat{D}\phi$.

There seem to be two issues with this. The first is formal: it does not seem very clear how this semantics generalises the single-agent case. The second one is conceptual: by considering closure in the τ topology, which contains the join $\tau_1 \vee \tau_2$, we are effectively demanding that each agent's evidence for a proposition is compared to not only the evidence she has, but whatever evidence the other agent has which happens to be contained in her set of compatible worlds. It seems as if what each agent knows is not independent from the other agent's knowledge, and thus it is not clear how one would, within this framework, remove distributed knowledge from the mix in order to simply consider a model for two-agent private knowledge while maintaining this "distributed" topology τ.

Another way to confront the fact that individual knowledge does not amount to distributed knowledge is to simply preserve the intuitive semantics presented in the previous subsection and deny it is an issue at all. Precisely one of the strengths of the definition of knowledge developed in Baltag et al. (2016) is its *defeasibility*, meaning that something which is known can stop being known when the agent is presented with evidence, even if factual, which trumps the evidence that made her come to know it. Özgün (2017) talks about *misleading evidence*, which is (loosely speaking) a piece of true, factual evidence that produces some fake evidence when added to the information the agent has.

The following informal example makes this clear:

Example 5.2.4. Morwenna is an epistemic agent who wonders about the attendance to a certain party which is going on while she writes her thesis. All the evidence she has is that two people, Jonathan and Hana, were planning on attending this party. Both people informed her of this fact individually. In particular, she has evidence that Hana is attending the party and this does not contradict any other evidence she has. As it turns out, Hana is indeed attending the party and, by our account, Morwenna *knows* this fact: she has a correctly justified belief in it.

Dean is another epistemic agent who has evidence that Hana and Jonathan have to take care of an injured rabbit. In fact he has evidence that, if either of them is at the party, the other one will not be, as one has to stay home looking for the animal.

Now, if Dean and Morwenna were to put their evidence together, then Morwenna's evidence that Hanna is at the party is no longer a *justification*, since it's defeated by one of Dean's pieces of evidence. And thus, even if Morwenna knows this fact, the group $\{\text{Dean}, \text{Morwenna}\}$ does not.

(Note that example 5.2.3 above is the formalisation of this story.)

This was the subject of an essay for the course Dynamic Epistemic Logic by Ethan Lewis and myself. In it, we gave a simpler version of the multi-agent semantics which included interior operators for each agent, a "distributed evidence" modality $\Box_D$ and global modalities for the partitions Π_1, Π_2 and $\Pi = \{\pi_1 \cap \pi_2 : \pi_i \in \Pi_i\}$. The knowledge modalities were defined from this and a complete logic was provided. However, the logic of this fragment does not seem to contain any axioms relating K_i and D, which is rather strange.

The idea of distributed knowledge discussed so far (and this is indirectly present in Ramírez (2015) and the aforementioned essay) reflects what the group could come to know if they put their evidence together and acted, in a way, as a collective agent. This is more an account of *implicit evidence* of the group rather than its *implicit knowledge*.

Is this the account of distributed knowledge we want? Halpern and Moses (1992) seem to have a different idea:

> [I]t is also often desirable to be able to reason about the knowledge that is distributed in the group, i.e., what someone who could combine the knowledge of all of the agents in the group would know. Thus, for example, if Alice knows ϕ and Bob knows $\phi \Rightarrow \psi$, then the knowledge of ψ is distributed among them, even though it might be the case that neither of them individually knows ψ. (...) [D]istributed knowledge corresponds to what a (fictitious) 'wise man' (one that knows exactly what each individual agent knows) would know.

The desired interpretation of 'distributed knowledge' here is that of a 'wise man' who has the information of what each agent knows, as opposed to what evidence they have, and whose knowledge stems from what they actually know. Thus, from this lens, one might want to keep misleading evidence out of the equation, and consider that a proposition known by one agent is known by this hypothetical wise man and thus constitutes distributed knowledge of the group.

There seem to be good reasons to stick to a notion of distributed knowledge which disregards the idea of 'putting evidence together' and which is based solely on the knowledge of the agents, whose logic would contain axioms like $K_1\phi \to D\phi$. In the next subsection we present a way to have such a notion.

5.2.2 Our proposal: the semantics

We again have a language with two modal operators K_1 and K_2 for the knowledge of each agent plus an operator D for distributed knowledge.

Definition 5.2.5 (Semantics for distributed knowledge)**.** Let

$$\mathfrak{X} = (X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$$

be a topological-partitional model. We read $\|p\|$, $\|\phi \wedge \psi\|$ $\|\neg\phi\|$ and $\|K_i\phi\|$ as usual, using the locally-dense-interior semantics of definition 3.2.3, whereas D is read as follows:

$$x \in \|D\phi\| \text{ iff there exist } U_1 \in \tau_1,\, U_2 \in \tau_2 \text{ such that } U_i \text{ is } \Pi_i\text{-locally dense and } x \in U_1 \cap U_2 \subseteq \|\phi\|.$$

Conceptually, what we are demanding here is that each agent has a justification which, when combined together, entail ϕ. While the semantics defined in the previous subsection amounted to reading distributed knowledge as the interior in the topology $(\tau_1 \vee \tau_2)^*$, what we are doing here is reading it as interior in $\tau_1^* \vee \tau_2^*$.

5.2.3 The logic of distributed knowledge

Let Logic_{K_iD} be the least set of formulas containing:

- The $\mathsf{S4.2}$ axioms and rules for K_1 and for K_2;
- The $\mathsf{S4}$ axioms and rules for D;
- The axioms $K_i\phi \to D\phi$ for $i = 1, 2$.

Theorem 5.2.6. Logic_{K_iD} *is sound and complete with respect to topological-partitional models.*

We will dedicate the present subsection to showing this fact.

Soundness. That every topological-partitional model satisfies the $\mathsf{S4.2}$ axioms for K_i can be proven exactly as in section 3.2. That D satisfies the $\mathsf{S4}$ axioms is a consequence of D being read as $\mathrm{Int}_{\tau_1^* \vee \tau_2^*}$. And for the two extra axioms, if $x \vDash K_i\phi$, then there exists $U_i \in \tau_i^*$ with $x \in U_i \subseteq \|\phi\|$. Let $j \neq i$ and, by taking $U_j = X$, which is a Π_j-locally dense τ_j-open set, we get $x \in U_i \cap U_j \subseteq \|\phi\|$ and thus $x \vDash D\phi$.

Completeness. We will use maximal consistent sets. A *maximal consistent set* is a set T of formulas in the language which is consistent (i.e. there are no $\phi_1, ..., \phi_n$ in T such that $\phi_1 \wedge ... \wedge \phi_n \to \bot$ is derivable in the logic) and maximally so (i.e. no proper superset of T is consistent).

The following is true for any maximal consistent set T (see e.g. Blackburn et al., 2001):

i. T is closed under logical equivalence.

ii. For each formula ϕ in the language, either $\phi \in T$ or $\neg\phi \in T$.

iii. $\phi \wedge \psi \in T$ if and only if $\phi \in T$ and $\psi \in T$.

iv. $\phi \vee \psi \in T$ if and only if $\phi \in T$ or $\psi \in T$.

Moreover,

Lemma 5.2.7 (Lindenbaum's lemma)**.** *Given a consistent set of formulas* Γ *in the language,* Γ *can be extended to a maximal consistent set.*

Let X be the set of maximal consistent sets over the language. We define R_i and R_D on X as follows: given $T, S \in X$,

$$TR_iS \text{ iff } K_i\phi \in T \text{ implies } \phi \in S \text{ for all } \phi \text{ in the language;}$$
$$TR_DS \text{ iff } D\phi \in T \text{ implies } \phi \in S \text{ for all } \phi \text{ in the language.}$$

Note that $R_D \subseteq R_i$ for $i = 1, 2$. Indeed, if TR_DS and $K_i\phi \in T$, then $D\phi \in T$ as per the axiom $K_i\phi \to D\phi$ and thus $\phi \in S$.

A *labelled path over X* is a path

$$\alpha = T_0 \xrightarrow{i_1} T_1 \xrightarrow{i_2} ... \xrightarrow{i_n} T_n,$$

where $T_0, ..., T_n \in X$ and $i_1, ..., i_n \in \{R_1, R_2, R_D\}$. Given $S \in X$ and a path $\alpha = T_0 \xrightarrow{i_1} T_1 \xrightarrow{i_2} ... \xrightarrow{i_n} T_n$, we define

$$\mathsf{last}\,\alpha := T_n \text{ and } \alpha \xrightarrow{i} S := T_0 \xrightarrow{i_1} T_1 \xrightarrow{i_2} ... \xrightarrow{i_n} T_n \xrightarrow{i} S.$$

Now, let $\mathcal{T}$ be the smallest set of labelled paths over X such that:

i. $T_0 \in \mathcal{T}$;

ii. For $i = 1, 2$, if $\alpha \in \mathcal{T}$ and $(\mathsf{last}\,\alpha)R_iT$, then $\alpha \xrightarrow{R_i} T \in \mathcal{T}$;

iii. If $\alpha \in \mathcal{T}$ and $(\mathsf{last}\,\alpha)R_DT$, then $\alpha \xrightarrow{R_D} T \in \mathcal{T}$.

For $i = 1, 2, D$ we define the following relations on $\mathcal{T}$: $\alpha \prec_i \beta$ if and only if $\alpha = T_0 \xrightarrow{i_1} ... \xrightarrow{i_n} T_n$ and $\beta = T_0 \xrightarrow{i_1} ... \xrightarrow{i_n} T_n \xrightarrow{R_i} S$ for some $T_0, ..., T_n, S \in X$. In other words,

$$\prec_i = \{\langle \alpha, \alpha \xrightarrow{R_i} S \rangle : \langle \mathsf{last}\,\alpha, S \rangle \in R_i\}.$$

We have thus given $\mathcal{T}$ the structure of a forest. Indeed, every $\alpha \in \mathcal{T}$ has at most one predecessor under $\prec_1 \cup \prec_2 \cup \prec_D$. Now let us define three preorders on $\mathcal{T}$: let $\leq_1$ be the reflexive and transitive closure of $\prec_1 \cup \prec_D$, that is,

$$\alpha \leq_1 \beta \text{ iff there exist } \alpha_0 = \alpha, \alpha_1, ..., \alpha_n = \beta$$
$$\text{with } \alpha_k \prec_1 \alpha_{k+1} \text{ or } \alpha_k \prec_D \alpha_{k+1}.$$

Similarly we define $\leq_2$ to be the reflexive and transitive closure of $\prec_2 \cup \prec_D$ and $\leq_D$ to be the reflexive and transitive closure of $\prec_D$. Note that by construction $\leq_D = \leq_1 \cap \leq_2$.

Now let us see what the $\leq_1$- and $\leq_2$-connected components look like. By part (iii) of lemma 3.2.11, we know that the connected components of the topology of upsets of $\leq_i$ ($i = 1, 2$) are given by the equivalence relation:

$$\alpha \sim_i \beta \text{ iff there exists } \gamma \text{ such that } \alpha \geq_i \gamma \leq_i \beta.$$

But then we have a path $\gamma = \gamma_0 \prec_{i_1} \gamma_1 \prec_{i_2} ... \prec_{i_n} \gamma_n = \alpha$ with $i_1, ..., i_n \in \{i, D\}$, thus for every $k = 0, ..., n-1$ we have

$$(\mathsf{last}\,\gamma_k) R_i (\mathsf{last}\,\gamma_{k+1}) \text{ or } (\mathsf{last}\,\gamma_k) R_D (\mathsf{last}\,\gamma_{k+1}).$$

Now, since $R_D \subseteq R_i$ and R_i is transitive, this gives us $(\mathsf{last}\,\gamma) R_i (\mathsf{last}\,\alpha)$. Similarly, we get that $(\mathsf{last}\,\gamma) R_i (\mathsf{last}\,\beta)$. Therefore we have the following result:

Lemma 5.2.8. *If α and β belong to the same $\leq_i$-connected component on $\mathcal{T}$, then $\mathsf{last}\,\alpha$ and $\mathsf{last}\,\beta$ belong to the same R_i-connected component in X.*

Moreover, there is an alternative characterisation of the connected components, similar to that in lemma 4.1.2, which we will find useful:

Lemma 5.2.9. *The $\leq_i$-connected components correspond to upsets of the form $\uparrow_i \alpha_0$, where α_0 has no $\leq_i$-predecessors other than itself.*

Proof. Clearly, if $\beta, \gamma \in \uparrow_i \alpha_0$ then $\beta \sim_i \gamma$ and conversely, since the set of $\leq_i$-predecessors of any point in this tree-like structure forms a finite chain, let α_0 be the least $\leq_i$-predecessor of some α. Then for every $\beta \in [\alpha]_\sim$ we have that there exists some γ with $\alpha \geq_i \gamma \leq_i \beta$ and thus $\alpha_0 \leq_i \gamma \leq_i \beta$. ■

We have given $\mathcal{T}$ the structure of a topological-partitional space and by defining $V^{\mathcal{T}}(p) = \{\alpha \in \mathcal{T} : p \in \mathsf{last}\,\alpha\}$ we have a topological-partitional model and we can prove the following:

Lemma 5.2.10 (Truth lemma)**.** *For every $\alpha \in \mathcal{T}$ and ϕ in the language, $\alpha \vDash \phi$ if and only if $\phi \in \mathsf{last}\,\alpha$.*

Proof. This is again an induction on formulas in which the base case for the propositional variables follows from the definition of $V^{\mathcal{T}}$ and the induction steps for the Boolean connectives are routine.

Now, suppose the result holds for ϕ and $K_i\phi \in \mathsf{last}\,\alpha$. We need to define a locally dense open set U_i such that $\alpha \in U_i \subseteq [\alpha]_{\sim_i}$ and with the property that, for every $\beta \in U_i$, $\phi \in \mathsf{last}\,\beta$, which will give us, by induction hypothesis, that $U_i \subseteq \|\phi\|$. By lemma 5.2.9, we have that $[\alpha]_{\sim_i} = \uparrow_i \alpha_i$ for some $\alpha_i \in \mathcal{T}$. In other words, every $\beta \in [\alpha]_{\sim_i}$ is of the form

$$\beta = \alpha_i \xrightarrow{R_i \text{ or } R_D} T_i \xrightarrow{R_i \text{ or } R_D} ... \xrightarrow{R_i \text{ or } R_D} T_n.$$

Let us now partition $[\alpha]_{\sim_i}$ in two sets:

$$\begin{aligned} V_D &:= \{\beta \in [\alpha]_{\sim_i} : \alpha_i \leq_D \beta\}; \\ V_i &:= \{\beta \in [\alpha]_{\sim_i} : \alpha_i \leq_i \beta \,\&\, \alpha_i \not\leq_D \beta\}. \end{aligned}$$

Note that the elements in V_D are of the form

$$\beta = \alpha_i \xrightarrow{R_D} T_1 \xrightarrow{R_D} ... \xrightarrow{R_D} T_n,$$

the elements in V_i are of the form

$$\beta = \alpha_i \xrightarrow{r_1} T_1 \xrightarrow{r_2} ... \xrightarrow{r_n} T_n \text{ with } r_k \in \{R_i, R_D\} \text{ and at least one } r_k = R_i,$$

and each element in $[\alpha]_{\sim_i}$ is in exactly one of V_i, V_D. Let us define U_i as follows:

$$U_i := \{\beta \in V_D : (\mathsf{last}\,\alpha)R_D(\mathsf{last}\,\beta)\} \cup \{\gamma \in V_i : (\mathsf{last}\,\alpha)R_i(\mathsf{last}\,\gamma)\}.$$

The following holds:

i. $\alpha \in U_i$ by construction.

ii. *U_i is an upset.* Take any $\beta \in U_i$. If $\beta \prec_i \gamma$ then $\gamma = \beta \xrightarrow{R_i} S$ for some $S \in X$ and we clearly have $\gamma \in V_i$ and $(\mathsf{last}\,\alpha)R_i(\mathsf{last}\,\beta)R_iS$, thus $(\mathsf{last}\,\alpha)R_iS$. If $\beta \prec_D \gamma$ then $\beta = \gamma \xrightarrow{R_D} S$ and, if $\beta \in V_D$ we then have that $\gamma \in V_D$ and $(\mathsf{last}\,\alpha)R_D(\mathsf{last}\,\beta)R_DS$ (thus $(\mathsf{last}\,\alpha)R_D(\mathsf{last}\,\gamma)$) whereas if $\beta \in V_i$ we have that $\gamma \in V_i$ and similarly (given that $R_D \subseteq R_i$), $(\mathsf{last}\,\alpha)R_iS$. In any case $\gamma \in U_i$.

iii. *U_i is locally dense.* Take any $\beta \in [\alpha]_{\sim_i}$. By lemma 5.2.8, we have that $\mathsf{last}\,\beta$ and $\mathsf{last}\,\alpha$ are in the same R_i-connected component and, since R_i is an $\mathsf{S4.2}$ relation, part (ii) of lemma 3.2.11 gives us that there exists some $S \in X$ with $(\mathsf{last}\,\alpha)R_iS$ and $(\mathsf{last}\,\beta)R_iS$ and thus we have that $\beta \xrightarrow{R_i} S \in V_i$ with $(\mathsf{last}\,\alpha)R_i\,\mathsf{last}(\beta \xrightarrow{R_i} S)$ so $\beta \xrightarrow{R_i} S \in U_i \cap {\uparrow_i}\beta$.

iv. *For every $\beta \in U_i$, we have $\phi \in \mathsf{last}\,\beta$.* Indeed, given that $K_i\phi \in \mathsf{last}\,\alpha$ and that $(\mathsf{last}\,\alpha)R_i(\mathsf{last}\,\beta)$, we have that $\phi \in \mathsf{last}\,\beta$.

Thus $\alpha \vDash K_i\phi$, as we intended to prove.

Conversely, if $\alpha \vDash K_i\phi$, there exists some locally dense open set U_i with $\alpha \in U_i \subseteq [\alpha]_{\sim_i} \cap \|\phi\|$. But then if $\mathsf{last}\,\alpha R_iS$ we have that $\alpha \prec_i \alpha \xrightarrow{R_i} S$, thus since U_i is an upset we have $\alpha \xrightarrow{R_i} S \in U_i$, which means $\alpha \xrightarrow{R_i} S \in \|\phi\|$ and by induction hypothesis $\phi \in S$. Hence we have that every R_i-successor of $\mathsf{last}\,\alpha$ contains ϕ, which gives $K_i\phi \in \mathsf{last}\,\alpha$.

Now suppose $D\phi \in \mathsf{last}\,\alpha$. Define U_1 and U_2 as above. They are locally dense open sets contained respectively in $[\alpha]_{\sim_1}$ and $[\alpha]_{\sim_2}$. Moreover, $\alpha \in U_1 \cap U_2$ by construction. We simply need to see that $U_1 \cap U_2 \subseteq \|\phi\|$. First let us note the following: if $\beta \in [\alpha]_{\sim_1} \cap [\alpha]_{\sim_2} = {\uparrow_1}\alpha_1 \cap {\uparrow_2}\alpha_2$, then β is simultaneously of the form

$$\beta = \alpha_1 \xrightarrow{R_1 \text{ or } R_D} T_1 \xrightarrow{R_1 \text{ or } R_D} ... \xrightarrow{R_1 \text{ or } R_D} T_n$$

and of the form

$$\beta = \alpha_2 \xrightarrow{R_2 \text{ or } R_D} S_1 \xrightarrow{R_2 \text{ or } R_D} ... \xrightarrow{R_2 \text{ or } R_D} S_m.$$

The only way for both this things to be true is if β is of the form

$$\beta = \alpha_2 \xrightarrow{R_D} S_1 \xrightarrow{R_D} ... \xrightarrow{R_D} S_m$$

and α_2 is of the form

$$\alpha_2 = \alpha_1 \xrightarrow{R_1 \text{ or } R_D} T_1 \xrightarrow{R_1 \text{ or } R_D} ... \xrightarrow{R_1} T_n$$

or vice versa. Let us assume the former without loss of generality. In particular we have that, if $\beta \in U_1 \cap U_2$, then $\beta \in V_D^{[2]}$ and hence $(\mathsf{last}\,\alpha)R_D(\mathsf{last}\,\beta)$. And since $D\phi \in \mathsf{last}\,\alpha$, this entails that $\phi \in \mathsf{last}\,\beta$ and thus that $\beta \in \|\phi\|$. We have thus proven that $\alpha \vDash D\phi$.

For the converse, if $\alpha \vDash D\phi$ then $\alpha \in U_1 \cap U_2 \subseteq \|\phi\|$ for some $\leq_i$-locally dense $U_i \subseteq [\alpha]_{\sim_i}$. But then if $(\mathsf{last}\,\alpha)R_D S$ we have $\alpha \leq_D \alpha \xrightarrow{R_D} S$ and since $\leq_D = \leq_1 \cap \leq_2$ and U_1 and U_2 are respectively a $\leq_1$ and a $\leq_2$-upset, we have that $\alpha \xrightarrow{R_D} S \in U_1 \cap U_2$ and thus $\alpha \xrightarrow{R_D} S \vDash \phi$ which by induction hypothesis gives $\phi \in S$. This entails $D\phi \in \mathsf{last}\,\alpha$. ■

Completeness follows from this: if $\phi \notin \mathsf{Logic}_{K_iD}$, then $\{\neg\phi\}$ is consistent and can be extended as per Lindenbaum's lemma to some maximal consintent set $T_0 \in X$. We then unravel the tree around T_0 as discussed above and we have ourselves a topological-partitional model rooted in $\alpha = T_0$ with $\alpha \nvDash \phi$ as per the truth lemma.

5.3 Common knowledge

In the context of epistemic logic, one can think of *common knowledge* as that which "every fool knows". This informal definition can be formally cashed out in several intuitive ways when one is modelling an epistemic situation. Barwise (1988) presents a comparison of three approaches to common knowledge:

(1) *The iterate approach.* A fact ϕ is common knowledge for a group of agents when ϕ is true, all agents know that it is true, all agents know that all agents know that it is true, etc. If we have two relevant agents, a and b, and $E\psi$ is short for $K_a\psi \wedge K_b\psi$, then $C\phi$ corresponds to the infinite conjunction

$$\phi \wedge E\phi \wedge EE\phi \wedge EEE\phi \wedge ...$$

(2) *The fixed-point approach.* This is an approach in which common knowledge refers back to itself. The idea here is that, if ϕ is the proposition which expresses "it is common knowledge for agents a and b that p", then ϕ is equivalent to "a and b know (p and ϕ)".

(3) *The shared environment approach.* According to this approach, a and b have common knowledge of p whenever there exists a "situation" s such that the following are facts of s: (i) p; (ii) a knows s; (iii) b knows s.

Barwise (1988) goes on to argue that, despite the fact that the first approach is the most "orthodox account" of common knowledge in the field of logic, being the approach considered in such influential works as Lewis (1969) and Halpern and Moses (1990), and despite the fact that early literature considered this approach equivalent to the fixed point one, (1) and (2) offer in fact distinct accounts and the fixed point approach provides "the right theoretical analysis of the pretheoretic notion of common knowledge".

One thing to note is that (1) and (2) are in fact equivalent approaches if one is modelling an epistemic situation within a relational framework. If K_i is the Kripke modality corresponding to some relation R_i, then we will read C in both cases as the Kripke modality corresponding to the reflexive and transitive closure of the union of their relations, $R_C = (R_1 \cup R_2)^*$. That is, $C\phi$ will be true at x just in case ϕ is true at any world that can be reached from each via a path of R_1's and R_2's. However, as shown in van Bethem and Sarenac (2004) this equivalence disappears once we are working in a topological setting. If one is working topologically, one has to make a choice.

Our proposal amounts to reading the common knowledge modality C as the interior in the intersection topology $\tau_1^* \cap \tau_2^*$. More explicitly:

Definition 5.3.1 (Semantics for common knowledge)**.** Let

$$\mathfrak{X} = (X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$$

be a topological-partitional model. We read $\|p\|$, $\|\phi \wedge \psi\|$, $\|\neg\phi\|$ and $\|K_i\phi\|$ as in definition 3.2.3 and:

$$\mathfrak{X}, x \models C\phi \text{ iff } \begin{array}{l}\text{there exists } U \in \tau_1 \cap \tau_2 \text{ locally dense in } \Pi_1 \text{ and in } \Pi_2 \\ \text{such that } x \in U \subseteq \|\phi\|.\end{array}$$

This amounts to the following: there is common knowledge of ϕ at x whenever there exists a common factive justification for ϕ.

Let us point out two things about this approach. The first one is that it matches rather well with the notion of distributed knowledge given in the previous section. If D and C correspond, respectively, to the meet and join of the corresponding accessibility relations on Kripke model, in this setting they are made to correspond to the join and meet, respectively, of the τ_i^* topologies. Moreover, much like distributed knowledge, this notion of common knowledge corresponds directly with the relational definition when we are dealing with a topological-partitional model stemming from two $\mathsf{S4.2}$ relations: if R_1 and R_2 are $\mathsf{S4.2}$, τ_i is the topology of R_i-upsets and Π_i is the set of R_i-connected components, then $\tau_1^* \cap \tau_2^*$ contains exactly the upsets of $(R_1 \cup R_2)^*$.

The second observation is that, in the spirit of Barwise (1988), this definition is precisely the fixed point account of common knowledge. As pointed out in van Bethem and Sarenac (2004) and expanded upon in Bezhanishvili and van der Hoek (2014), the fixed point approach can be expressed in the notation of mu-calculus as

$$C\phi = \nu p(\phi \wedge Ep),$$

where p is a propositional variable which does not appear in ϕ. We read

$$\|\nu p\psi\| = \bigcup\{U \in \mathcal{P}(X) : U \subseteq \|\psi\|^{V_p^U}\},$$

where V_p^U is the valuation assigning U to p and $V(q)$ to $q \neq p$.

In particular,

$$\|C\phi\| = \bigcup\{U \in \mathcal{P}(X) : U \subseteq \|\phi \wedge Ep\|^{V_p^U}\}.$$

Now, $\|\phi \wedge Ep\|^{V_p^U} = \|\phi\|^{V_p^U} \cap \|K_1 p\|^{V_p^U} \cap \|K_2 p\|^{V_p^U}$. Since p does not appear in ϕ, we have that $\|\phi\|^{V_p^U} = \|\phi\|$. On the other hand, $\|K_i p\|^{V_p^U} = \mathrm{Int}_{\tau_i^*} \|p\|^{V_p^U} = \mathrm{Int}_{\tau_i^*} U$.

Thus a set U is contained in $\|\phi \wedge Ep\|^{V_p^U}$ if and only if U is a τ_1^*-open and τ_2^*-open subset of $\|\phi\|$. And thus,

$$\|C\phi\| = \bigcup\{U \in \mathcal{P}(X) : U \in \tau_1^* \cap \tau_2^* \,\&\, U \subseteq \|\phi\|\} = \mathrm{Int}_{\tau_1^* \cap \tau_2^*} \|\phi\|,$$

which is precisely our account of common knowledge.

Some theorems in the logic of topological-partitional models with common knowledge are the following:

i. The S4.2 axioms for K_i;

ii. the S4 axioms for C;

iii. the *fixed point axiom* $C\phi \rightarrow E(C\phi \wedge \phi)$;

iv. the *induction axiom* $C(\phi \rightarrow E\phi) \rightarrow (E\phi \rightarrow C\phi)$.

Proposition 5.3.2 (Soundness)**.** *All the theorems above are valid on topological-partitional models with the semantics of definition 5.3.1.*

Proof. That i., ii. and iii. hold for topological-partitional models is a very straightforward check. Item iv. is more involved. It amounts to checking that, on any such model, and for any $P \subseteq X$,

$$C(\neg P \vee (K_1 P \cap K_2 P)) \cap P \subseteq CP.$$

Now, let $x \in C(\neg P \vee (K_1 P \cap K_2 P))$. By the semantics of 5.3.1 this means that there exists some $U \in \tau_1^* \cap \tau_2^*$ such that

$$x \in U \subseteq \neg P \cup (\mathrm{Int}_{\tau_1^*} P \cap \mathrm{Int}_{\tau_2^*} P).$$

Call $V := U \cap \mathrm{Int}_{\tau_1^*}$. Now, V is a τ_1^*-open set. Note that $V \subseteq U \cap \mathrm{Int}_{\tau_2^*}$ and $U \cap \mathrm{Int}_{\tau_2^*} \subseteq V$ and thus V is also a τ_2^*-open set. Moreover, V includes x and it is contained in P. Thus there exits some $V \in \tau_1^* \cap \tau_2^*$ with $x \in V \subseteq P$, hence $x \in CP$. ■

Whether the preceding list of formulas constitutes a complete axiomatisation of the logic of common knowledge for topological-partitional models is a question that remains open.

Chapter 6

Discussion

The work contained in this Master's thesis is a dive into several aspects of the dense interior semantics defined on topological evidence models, furthering the results in Baltag et al. (2016).

Whenever it was possible, we found particular spaces whose logic is precisely the epistemic logic introduced in the aforementioned paper. These spaces provide a "natural" or (as we referred to it) "generic" model in which to study epistemic concepts such as knowledge, belief, evidence and justification, which interact with each other following the considerations in Stalnaker (2006). We showed that any dense-in-itself metrisable space, such ar $\mathbb{R}$, is a generic model for the knowledge-only fragment of this logic and that some of these spaces, such as $\mathbb{Q}$, are a generic model for several other fragments of the language of topological evidence logic.

We then introduced a second agent in topo-e-models, obtaining a multi-agent semantics for these models, along with a brief conceptual and theoretical study of notions of "group knowledge" for this group of agents. We showed how this semantics generalises the single agent case and we provided a complete logic for our two-agent models. Moreover, mirroring the single-agent case, we found generic spaces with respect to which the logic is sound and complete: the quaternary tree $\mathcal{T}_{2,2}$ and the rational plane $\mathbb{Q} \times \mathbb{Q}$.

In the process of writing this thesis some questions remained unanswered, and other questions (which are out of the scope of the present work but nonetheless quite interesting) sprung.

In the first section of this chapter we show that the results in this thesis are easily generalisable to any finite number of agents. In the second section we will comment on some of the loose ends and pose some questions which could constitute relevant lines of future research.

6.1 Generalising to any number of agents

The results included in the "multi-agent" part of this thesis have been limited, for the sake of simplicity, to the two-agent case. It is however very easy to generalise these results to any finite number of agents n. Let us outline this generalisation:

Our topological partitional models for n agents are now of the form $(X, \tau_1, ..., \tau_n, \Pi_1, ..., \Pi_n, V)$, where X is a set, V is a valuation, each τ_i is a topology on X and each Π_i is a τ_i-open partition of X.

Note that sections 3.3 and 3.4, which study the logics of different fragments, at no point require that $n = 2$ in their proofs. Thus, the logic of the different fragments of these n-agent topological partitional models is exactly the logic of the respective fragment for the two-agent case, where i now ranges over $\{1, ..., n\}$. For instance, the logic of the $K_1, ..., K_n$ fragments is the fusion logic $\mathsf{S4.2}_{K_1} + ... + \mathsf{S4.2}_{K_n}$, i.e., the axioms and rules of $\mathsf{S4.2}$ for each operator.

The results of chapter 5 work similarly for n agents, using when necessary the join and meet topologies of $\tau_1, ..., \tau_n$. Here, $\tau_1 \vee ... \vee \tau_n$ is the topology $\{U_1 \cap ... \cap U_n : U_i \in \tau_i\}$, and we read D as interior in $\tau_1^* \vee ... \vee \tau_n^*$. In particular, the logic of distributed knowledge has the same axioms and rules as $\mathsf{Logic}_{K_i D}$, as defined in subsection 5.2.3, with $1 \leq i \leq n$.

The only part in the multi-agent discussion in which we used the fact that $n = 2$ is chapter 4, when talking about the spaces $\mathcal{T}_{2,2}$ and $\mathbb{Q} \times \mathbb{Q}$. We can generalise the quaternary tree to n agents via the $2 \times n$-ary tree $\mathcal{T}_{2\times n}$, i.e., the infinite branching tree with n relations $R_1, ..., R_n$ such that every node has exactly $2n$ successors: a left R_i-successor and a right R_i-successor for each $1 \leq i \leq n$. Now, arguing analogously to section 4.1 we can unravel any finite frame with n $\mathsf{S4.2}$ relations into $\mathcal{T}_{2\times n}$, and this construction comes along with an onto p-morphism which gives us completeness of $\mathcal{T}_{2\times n}$ with respect to $\mathsf{S4.2}_{K_1} + ... + \mathsf{S4.2}_{K_n}$. For more details about $\mathcal{T}_{2\times n}$ and this unravelling, see appendix A in Sarenac (2006).

For the other space, instead of taking $\mathbb{Q} \times \mathbb{Q}$ with the vertical and horizontal topology, we consider $\mathbb{Q}^n$ and define n topologies $\tau_1, ..., \tau_n$ on it, where τ_k is the topology generated by

$$\mathcal{B}_k = \{\{x_1\} \times ... \times \{x_{k-1}\} \times U_k \times \{x_{k+1}\} \times ... \times \{x_n\} : U_k \in \tau_{\mathbb{Q}} \,\&\, x_i \in \mathbb{Q} \text{ for } i \neq k\}.$$

We can define and open subspace of (something homeomorphic to) $\mathbb{Q}^n$ by $Y = \bigcup_{k\in\omega} Y_k$, where $Y_0 = \{(0, ..., 0)\}$ and

$$Y_{k+1} = Y_k \cup \{(x_1, ..., x_i \pm 1/3^k, ..., x_n) : (x_1, ..., x_i, ..., x_n) \in Y_k \,\&\, 1 \leq i \leq n\}.$$

Now, if we are working in one dimension (that is to say, $Y_0 = \{0\}$ and $Y_{k+1} = Y_k \cup \{x \pm 1/3^k : x \in Y_k\}$), this space Y is homeomorphic to $\mathbb{Q}$. Let us call this one-dimensional space (X, τ). We can see that, working again on n dimensions, Y is an open subspace of X^n. We can define an interior map $f : Y \to \mathcal{T}_{2\times n}$ as follows:

We map $(0, ..., 0)$ to the root r of the tree. Now, suppose we have mapped all points in Y_k to some nodes in the tree and take $x \in Y_{k+1} \backslash Y_k$. Then either $x = (x_1, ..., x_i - 1/3^k, ..., x_n)$ for some $(x_1, ..., x_n) \in Y_k$, in which case we map x to the left R_i-successor of $f(x_1, ..., x_n)$, or $x = (x_1, ..., x_i + 1/3^k, ..., x_n)$, in which case we map it to its right R_i-successor.

An argument similar to the proof in van Benthem and Bezhanishvili (2007) for the case $n = 2$ shows that this is indeed a surjective continuous and open map thus completeness (for a certain equivalence relation) follows.

6.2 Open questions

Completeness of $\mathsf{Logic}_{\forall\Box\Box_0}$ with respect to $\mathbb{Q}$ with a particular subbasis

While in chapter 2 we showed several of the logics in the introduction to be sound and complete with respect to singleton classes of models, we failed to provide a single topo-e-model for the fragment involving the basic evidence modality. Instead, we showed that the corresponding logic is sound and complete with respect to the class of topological evidence models based on $(\mathbb{Q}, \tau_{\mathbb{Q}})$ with arbitrary subbases.

There were attempts in the writing of this thesis to find one particular subbasis for $\mathbb{Q}$ that would give us a topological evidence model whose logic is precisely $\mathsf{Logic}_{\forall\Box\Box_0}$. This needs to be a subbasis which is not a basis (for otherwise $\Box\phi \leftrightarrow \Box_0\phi$ would be a theorem of the logic). One obvious candidate is perhaps the most paradigmatic case of subbasis-which-is-not-a-basis, namely

$$\mathcal{S} = \{(a, \infty), (-\infty, b) : a, b \in \mathbb{Q}\}.$$

This does not work. To show why, let us find a formula that is consistent in the logic yet it cannot be satisfied by any model based on $\mathbb{Q}$ with the aforementioned subbasis. Consider the following formula, with three propositional variables p_1, p_2, p_3:

$$\gamma = \bigwedge_{i=1,2,3} (\Box_0 p_i \wedge [\exists]\Box_0\neg p_i) \bigwedge_{i\neq j\in\{1,2,3\}} [\exists](\Box_0 p_i \wedge \neg\Box_0 p_j).$$

First of all, note that, in any topo-e-model, $\|\Box_0\phi\|$ is a union of elements in the subbasis. Indeed, for all $x \in \|\Box_0\phi\|$, there is an e_x in the subbasis with $x \in e_x \subseteq \|\phi\|$. But then $\|\Box_0\phi\| = \bigcup_{x\in\|\Box_0\phi\|} e_x$. In particular, in the topology of $\mathbb{Q}$ with the subbasis E_0 as defined above, we have that $\|\Box_0\phi\|^{\mathbb{Q}}$ is always of the form

$$\|\Box_0\phi\|^{\mathbb{Q}} = (-\infty, a) \cup (b, \infty)$$

for some $a, b \in \mathbb{R} \cup \{-\infty, \infty\}$ (here, we call $(-\infty, -\infty) = (\infty, \infty) = \varnothing$ and $(-\infty, \infty) = \mathbb{Q}$).

Moreover, if the set $\|\Box_0\phi \wedge [\exists]\Box_0\neg\phi\|^{\mathbb{Q}}$ is nonempty, then $\|\Box_0\phi\|^{\mathbb{Q}}$ is of the form (a, ∞) of or the form $(-\infty, a)$ for some $a \in \mathbb{R}$. Indeed, suppose x satisfies $[\exists]\Box_0\neg\phi$. Then there exists a y and an $e \in E_0$ with $y \in e \subseteq \|\neg\phi\|$. Let us assume without loss of generality that $e = (b', \infty)$ for some b'. Then $\|\Box_0\phi\|$ cannot be of the form $(-\infty, a) \cup (b, \infty)$ with $b \neq \infty$, for otherwise take $z > b, b'$ and we have $z \vDash \phi \wedge \neg\phi$. So $\|\Box_0\phi\| = (-\infty, a)$. Similarly, if $e = (-\infty, a')$, then $\|\Box_0\phi\| = (b, \infty)$ for some b.

With all this:

Lemma 6.2.1. *γ is not satisfied in any model based in $\mathbb{Q}$ with the aforementioned subbasis.*

Proof. Suppose γ is satisfied in such a model. By the previous observation, the first conjunct gives that $\|\Box_0 p_i\|$ needs to be of the form (a, ∞) or $(-\infty, a)$ for some $a \in \mathbb{R} \cup \{\infty, -\infty\}$. By the second conjunct, the sets $\|\Box_0 p_i\|$ and $\|\Box_0 p_j\|$ need to be incomparable for $i \neq j$. But of course, at least two of the sets $\|\Box_0 p_i\|$ have to be of the same form (either $(-\infty, a_i)$ and $(-\infty, a_j)$ or (a_i, ∞) and (a_j, ∞)), hence obviously it cannot be the case that the three sets are incomparable: contradiction. ■

Lemma 6.2.2. *γ is consistent in the logic.*

Proof. We use the fact (see Baltag et al., 2016) that the logic is complete with respect to *quasi-models* of the form $(X, \leq, E_0, V)$, where $\leq$ is a preorder and E_0 is a collection of $\leq$-upsets. $[\forall]$ is read globally, $\Box$ is read as the Kripke modality for $\leq$ and we read $x \in \|\Box_0 \phi\|$ if and only if there is some $e \in E_0$ with $x \in e \subseteq \|\phi\|$.

Let $(X, \leq)$ be the following poset:

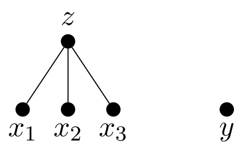

and call $e_i = \{x_i, z\}$ for $i = 1, 2, 3$. Let $E_0 = \{e_1, e_2, e_3, \{y\}, X\}$ and $V(p_i) = e_i$ for $i = 1, 2, 3$. It is clear that $(X, \leq, E_0, V)$ is a quasi-model and that $z \vDash \Box_0 pi$, $x_i \vDash \Box_0 p_i \wedge \neg\Box_0 p_j$ and $y \vDash \Box_0 \neg p_i$.

Thus $z \vDash \gamma$ and γ is therefore consistent in the logic. ■

Since every model based on $\mathbb{Q}$ with E_0 as a subbasis makes $\neg\gamma$ true yet $\nvdash_{\mathsf{Logic}_{\forall\Box\Box_0}} \neg\gamma$, incompleteness follows.

My (unfounded) conjecture here is that no particular subbasis will give us completeness. Proving this result if this conjecture is true, or otherwise finding such a subbasis if false, would be an interesting line of future work.

Common knowledge

As we mentioned in section 5.3, whether the list of theorems provided constitutes a complete axiomatisation of the logic of common knowledge is an open question.

Strong completeness

Kremer (2013) shows that $\mathsf{S4}$ is strongly complete with respect to any dense-intself metrisable space with the interior semantics.

Is $\mathsf{S4.2}$ strongly complete with respect to $\mathbb{R}$ with the dense interior semantics? What about the other logics and their respective generic models?

Other generic models for the two-agent logic

We have provided two generic models for the logic $\mathsf{S4.2}_{K_1} + \mathsf{S4.2}_{K_2}$.

What are some other generic models for this logic? One would speculate that, for instance, the binary Cantor space (i.e. the topological space of branches through $\mathcal{T}_{2,2}$) should be such a model, generalising the single-agent case.

Proving this would be an interesting line of research, and so would be finding generic models for some of the other multi-agent logics provided in section 3.4, as well as for the logic of distributed knowledge in section 5.2.

Dynamic topo-e-models

As discussed in the preliminaries, one of the advantages of the framework of topo-e-models over other Stalnakerian approaches is that they behave well dynamically. Baltag et al. (2016) consider several dynamic extensions of the language $\mathcal{L}_{\forall\Box\Box_0}$, including modalities for public announcements, evidence addition, evidence upgrade and feasible evidence combination, and give complete logics for each of these extensions.

This thesis has been exclusively concerned with "static" models. Combining these dynamic results with the results in this thesis is a very obvious and very interesting line of future research. In particular, finding the multi-agent logic with public announcements, accounting for a notion of private announcements or even finding generic models for any of these extensions are very relevant (and probably quite fruitful) lines of research.

Bibliography

Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. The topology of belief, belief revision and defeasible knowledge. In *International Workshop on Logic, Rationality and Interaction*, pages 27–40. Springer, 2013.

Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. Justified belief and the topology of evidence. In *International Workshop on Logic, Language, Information, and Computation*, pages 83–103. Springer, 2016.

Jon Barwise. Three views of common knowledge. In *Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 365–379, 1988.

Guram Bezhanishvili, Nick Bezhanishvili, Joel Lucero-Bryan, and Jan van Mill. A new proof of the McKinsey–Tarski theorem. *Studia Logica*, Mar 2018.

Nick Bezhanishvili and Wiebe van der Hoek. Structures for epistemic logic. In *Johan van Benthem on logic and information dynamics*, pages 339–380. Springer, 2014.

Patrick Blackburn, Maarten De Rijke, and Yde Venema. *Modal Logic*, volume 53. Cambridge University Press, 2001.

Ryszard Engelking. *General topology*. Sigma series in pure mathematics. Heldermann Verlag, 1989. ISBN 9783885380061.

Ronald Fagin, Joseph Y Halpern, and Moshe Y Vardi. What can machines know?: On the properties of knowledge in distributed systems. *Journal of the ACM (JACM)*, 39(2):328–376, 1992.

Ronald Fagin, Joseph Y Halpern, Yoram Moses, and Moshe Vardi. *Reasoning about knowledge*. MIT press, 2004.

Edmund L Gettier. Is justified true belief knowledge? *Analysis*, 23(6):121–123, 1963.

Valentin Goranko and Solomon Passy. Using the universal modality: gains and questions. *Journal of Logic and Computation*, 2(1):5–30, 1992.

Joseph Y Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM (JACM)*, 37(3):549–587, 1990.

Joseph Y Halpern and Yoram Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence*, 54(3):319–379, 1992.

Edwin Hewitt. A problem of set-theoretic topology. *Duke Mathematical Journal*, 10(2):309–333, 1943.

Jaakko Hintikka. *Knowledge and belief: an introduction to the logic of the two notions*. Contemporary philosophy. Cornell University Press, 1962.

Kevin T Kelly. *The logic of reliable inquiry*. OUP USA, 1996.

Philip Kremer. Strong completeness of s4 for any dense-in-itself metric space. *The Review of Symbolic Logic*, 6(3):545–570, 2013.

David Lewis. *Convention: A philosophical study*. Harvard University Press, 1969.

John Charles Chenoweth McKinsey and Alfred Tarski. The algebra of topology. *Annals of mathematics*, pages 141–191, 1944.

James R Munkres. *Topology*. Prentice Hall, 2000.

Aybüke Özgün. *Evidence in Epistemic Logic: A Topological Perspective*. PhD thesis, University of Amsterdam, ILLC, 2017.

Aldo Ramírez. Topological models for group knowledge and belief. Master's thesis, University of Amsterdam, 2015.

Helena Rasiowa and Roman Sikorski. *The Mathematics of Metamathematics*. Institut Mathematyczny, Polskiej Akademii Nauk: Monographie Mathematyczne. PWN-Polish Scientific Publishers, 1970.

Darko Sarenac. *Products of topological modal logics*. PhD thesis, Stanford University, 2006.

Valentin Shehtman. "Everywhere" and "Here". *Journal of Applied Non-Classical Logics*, 9(2-3):369–379, 1999.

Robert Stalnaker. On logics of knowledge and belief. *Philosophical studies*, 128(1):169–199, 2006.

Chris Steinsvold. *Topological models of belief logics*. City University of New York, 2006.

Johan van Benthem and Guram Bezhanishvili. Modal logics of space. In *Handbook of spatial logics*, pages 217–298. Springer, 2007.

Johan van Benthem and Eric Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1-3):61, 2011.

Johan van Benthem, Guram Bezhanishvili, Balder ten Cate, and Darko Sarenac. Multimodal logics of products of topologies. *Studia Logica*, 84 (3):369–392, 2006.

Johan van Bethem and Darko Sarenac. The geometry of knowledge. *Travaux de logique*, 17:1–31, 2004.

Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic epistemic logic*. Springer Science & Business Media, 2007.

Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, 2002.