

Topological Evidence Logics: Multi-Agent Setting

Alexandru Baltag¹, Nick Bezhanishvili¹, and Saúl Fernández González²

¹ ILLC, Universiteit van Amsterdam

² IRIT, Université de Toulouse

Abstract. We introduce a multi-agent topological semantics for evidence-based belief and knowledge, which extends the *dense interior semantics* developed in [2]. We provide the complete logic of this multi-agent framework together with *generic models* for a fragment of the language. We also define a new notion of group knowledge which differs conceptually from previous approaches.

1 Introduction

A semantic study of epistemic logics, the family of modal logics concerned with what an epistemic agent *believes* or *knows*, has been mostly conducted in the framework of relational structures (Kripke frames) [13]. These are 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.

It is claimed in [13] that the accessibility relation for knowledge must be (minimally) reflexive and transitive. On the syntactic level, this demand translates into the fact that any logic for knowledge based on these frames must contain the axioms of $S4$. This, paired with the fact, famously proven by McKinsey and Tarski [14], that $S4$ is the logic of topological spaces under the *interior semantics* (see [4]), lays the ground for a topological treatment of knowledge. Moreover, McKinsey and Tarski [14] proved that certain *generic* spaces, such as the real line, spaces which intuitively lend themselves to be models for certain situations of knowledge, have $S4$ as their logic.

The semantics outlined in [14] 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 Kripke frames. (See [16] for lengthy discussion on this topic.)

Under this interpretation, knowing a proposition amounts to having evidence for it. This can be an undesirable property, for it constitutes, arguably, an overly simplistic account of what knowledge is. As Gettier [11] argues, there is more to knowledge than ‘true and justified belief’. Depending on the properties one ascribes to knowledge, belief and the relation thereof, one can get different epistemic logics, each with their axioms and rules. For certain applications, one would want, for instance, to operate within a framework in which misleading true evidence can lead to false beliefs. (See [18,16] for more in-depth discussion.)

Inspired by [18] and [6] a new topological semantics was introduced in [2] and explored in depth in [16]. This semantics allows one to talk about knowledge and belief, evidence (both “basic” and “combined”) and a notion of *justification* via the dense-interior operator. Also an epistemic logic complete with respect to the proposed semantics has been given in [2] and [16]. The models for this logic based on the dense-interior semantics in topological spaces are called *topo-e-models*.

In [1] an analogue of the McKinsey-Tarski theorem was proved for the dense-interior semantics: the logic of topological evidence models is sound and complete with respect to any individual topological space (X, τ) which is *dense-in-itself*, *metrizable*, and homeomorphic to the disjoint union $(X, \tau) \cup (X, \tau)$.

The framework defined in [2] is single-agent. In this paper, we introduce a multi-agent topological evidence semantics which generalises the single-agent case and differs substantially from prior approaches. In this sense, we provide several logics of multi-agent models and give some conceptual and theoretical contributions for a notion of group knowledge in this framework.

Outline. In Section 2 we present the (one-agent) notion of topological evidence models introduced in [2] together with some relevant results. In Section 3 we introduce and justify our multi-agent setting, we show how it generalises the single-agent case and we provide the logic for several fragments of the language. In Section 4, we obtain “generic models”, i.e., unique topological spaces whose logic under the semantics previously introduced is exactly the logic of all topological spaces. Section 5 discusses a notion of *group knowledge* in this setting, and gives a sound and complete logic of distributed knowledge. We conclude in Section 6.³

2 Single-agent topological evidence models

The relation between belief and knowledge has historically been one of the main focuses 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 counterexamples of cases of true, justified belief which do not amount to knowledge shattered this paradigm [11].

Stalnaker [18] argues that a relational semantics is insufficient to capture Gettier’s considerations in [11] and, trying to stay close to most of the intuitions of Hintikka in [13], provides an axiomatisation for a system of knowledge and belief in which knowledge is an **S4.2** modality, belief is a **KD45** modality and the following formulas can be proven: $B\phi \leftrightarrow \neg K\neg K\phi$ and $B\phi \leftrightarrow BK\phi$. “Believing p ” is the same as “not knowing you don’t know p ” and belief 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.

A topological semantics in which knowledge is simply the interior modality (i.e., evaluating formulas on a topological space and setting $\|K\phi\| = \text{Int}\|\phi\|$)

³ This paper is based on Saúl Fernández González’s Master’s thesis [10].

proves insufficient to capture these nuances. In [2] a new semantics is introduced, building on the idea of *evidence models* of [6] 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*, etc. 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 “strange” classes of spaces.

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

2.1 The logic of topological evidence models

We briefly present here the framework introduced in [2], see also [16]. Our language is now $\mathcal{L}_{\forall KB\Box\Box_0}$, which includes the modalities K (knowledge), B (belief), $[\forall]$ (infallible knowledge), \Box_0 (basic evidence), \Box (combined evidence).

Definition 2.1 (The dense interior semantics). *We interpret sentences on topological evidence models (i.e. tuples (X, τ, E_0, V) where (X, τ, V) is a topological model and E_0 is a subbasis of τ) as follows: $x \in \llbracket K\phi \rrbracket$ iff $x \in \text{Int}\llbracket \phi \rrbracket$ and $\text{Int}\llbracket \phi \rrbracket$ is dense⁴; $x \in \llbracket B\phi \rrbracket$ iff $\text{Int}\llbracket \phi \rrbracket$ is dense; $x \in \llbracket [\forall]\phi \rrbracket$ iff $\llbracket \phi \rrbracket = X$; $x \in \llbracket \Box_0\phi \rrbracket$ iff there is $e \in E_0$ with $x \in e \subseteq \llbracket \phi \rrbracket$; $x \in \llbracket \Box\phi \rrbracket$ iff $x \in \text{Int}\llbracket \phi \rrbracket$. Validity is defined in the standard way.*

We see that “knowing” does not equate “having evidence” in this framework, but it is rather something stronger: in order for the agent to know P , she needs to have a piece of evidence for P which is *dense*, i.e., which has nonempty intersection with (and thus cannot be contradicted by) any other piece of evidence.

Fragments of the logic. The following logics are obtained by considering certain fragments of the language (i.e. certain subsets of the modalities above)⁵.

| | |
|--|---|
| “K-only”, \mathcal{L}_K | S4.2. |
| “Knowledge”, $\mathcal{L}_{\forall K}$ | S5 axioms and rules for $[\forall]$, plus S4.2 for K , plus $[\forall]\phi \rightarrow K\phi$ and $\neg[\forall]\neg K\phi \rightarrow [\forall]\neg K\neg\phi$. |
| “Combined evidence”, $\mathcal{L}_{\forall\Box}$ | S5 for $[\forall]$, S4 for \Box , plus $[\forall]\phi \rightarrow \Box\phi$. |
| “Evidence”, $\mathcal{L}_{\forall\Box\Box_0}$ | S5 for $[\forall]$, S4 for \Box , plus the axioms $\Box_0\phi \rightarrow \Box_0\Box_0\phi$, $[\forall]\phi \rightarrow \Box_0\phi$, $\Box_0\phi \rightarrow \Box\phi$, $(\Box_0\phi \wedge [\forall]\psi) \rightarrow \Box_0(\phi \wedge [\forall]\psi)$. |

We will refer to these logics respectively as S4.2_K , $\text{Logic}_{\forall K}$, $\text{Logic}_{\forall\Box}$ and $\text{Logic}_{\forall\Box\Box_0}$. K and B are definable in the evidence fragments⁶, thus we can think of the logic of $\mathcal{L}_{\forall\Box\Box_0}$ as the “full logic”.

⁴ A set $U \subseteq X$ is dense whenever $\text{Cl}U = X$ or equivalently whenever $U \cap V \neq \emptyset$ for all nonempty open sets V .

⁵ We recall that S4 is the least normal modal logic containing the axioms (T) $\Box\phi \rightarrow \phi$ and (4) $\Box\phi \rightarrow \Box\Box\phi$; that S5 is S4 plus the axiom (5) $\neg\Box\phi \rightarrow \Box\neg\Box\phi$, and that S4.2 is S4 plus the axiom (.2) $\Diamond\Box\phi \rightarrow \Box\Diamond\phi$.

⁶ $K\phi \equiv \Box\phi \wedge [\forall]\Box\Diamond\phi$ and $B\phi \equiv \neg K\neg K\phi$.

3 Going multi-agent

There have been different approaches to a multi-agent logic derived from the framework introduced in [2]. In [17], a two-agent logic with distributed knowledge was defined. However, the semantics of this approach seems to come with some conceptual problems which were discussed in [10]. Another approach, present in [16], 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 section 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. We will limit ourselves to two agents for simplicity in the exposition. Extending these results to any finite number of agents is straightforward.

The problem of density. 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 $x \in U \subseteq \|\phi\|$. However, this approach is highly problematic because it requires the extra assumptions that the same set of worlds is epistemically accessible for both agents, and thus conflates infallible knowledge. This is discussed in more depth in [10]. 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.

3.1 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 3.1. 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. For $x \in X$, let $\Pi_i(x)$ be the unique $\pi \in \Pi_i$ with $x \in \pi$. For $U \subseteq X$, let $\Pi_i[U] = \{\pi \in \Pi_i : \pi \cap U \neq \emptyset\} = \{\Pi_i(x) : x \in U\}$.

A set $U \subseteq X$ is locally dense in $\pi \in \Pi_i$ whenever $\pi \subseteq \text{Cl}_{\tau_i} U$ 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 $\text{Cl}_{\tau_i} U = \bigcup \Pi_i[U]$. Equivalently, U is locally dense if it is locally dense in π for every $\pi \in \Pi_i[U]$.

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

Definition 3.3 (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 $\|\neg\phi\| = X \setminus \|\phi\|$. For $i = 1, 2$ set:

$$\begin{aligned} \mathfrak{M}, x \models K_i \phi \text{ iff } & x \in \text{Int}_{\tau_i} \|\phi\| \\ & \& \text{Int}_{\tau_i} \|\phi\| \text{ is locally dense in } \Pi_i(x). \end{aligned}$$

Consider a topological-partitional model $(X, \tau_1, \tau_2, \Pi_1, \Pi_2, V)$ and set

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

It is straightforward to check that the following holds:

Lemma 3.4. (X, τ_1^*, τ_2^*) is an extremally disconnected bitopological 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.

One last remark before proceeding with the main results: at first glance demanding each element $\pi \in \Pi_i$ to be open may seem as 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 $(X, \tau_i, \Pi_i), x \models \phi$ under this semantics if and only if $(X, \bar{\tau}_i, \Pi_i), x \models \phi$ under the semantics in Def. 3.3.

For this reason, we will limit ourselves to the study of models with open partitions. Let us now look at an example:

Example 3.5. 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 and q . Let $V(p) = P = \{x_{11}, x_{10}\}$ and $V(q) = \{x_{11}, x_{01}\}$. The actual world is x_{11} , in which p and q hold.

At q -worlds Alice 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 p . At $\neg p$ -worlds she does not receive this evidence.

The only worlds consistent with Bob's information are those in which $q \rightarrow p$ holds. Moreover, in p -worlds he has fallible evidence for p and in $\neg p$ -worlds he has it for $\neg p$.




Fig. 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 contained in P but not in a π_3 -locally dense one.

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 $\{\pi_3, \pi_4, P, X \setminus P\}$. (See Fig. 1.)

At the actual world x_{11} , Alice knows p yet Bob does not: indeed, $\{x_{11}\}$ is a τ_a -open set, locally dense in π_1 and contained in P , thus $K_a p$ holds. And 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} .

Certain topological spaces come equipped with open partitions, in the form of their *connected components*.

Definition 3.6. 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. [15]):

Lemma 3.7. *The connected components of (X, τ) coincide with the equivalence classes of the relation: $x \sim y$ if and only if there is a connected subset of X containing x and y .*

The following lemma, whose proof is straightforward, shows that the connected components of an Alexandroff space are always open:

Lemma 3.8. *Let (W, \leq) be a preordered set. Then:*

- i. *The connected components on $(W, \text{Up}(W))$ 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 $x_0, \dots, x_n \in X$ 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 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 $x \leq z \leq y$.*
- 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). Let us see that $[x]_{\sim}$ is clopen and connected. Clearly it is both upward and downward closed. Moreover, if $\emptyset \neq U \subseteq [x]_{\sim}$ is a clopen set, 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.

(ii). Take a path $(x_0 = x, x_1, \dots, x_n = y)$ such that $x_k \leq x_{k+1}$ or $x_{k+1} \leq x_k$ for all $0 \leq k \leq n-1$, and note that $x_{k-1} \geq x_k \leq x_{k+1}$ implies that there exists a certain x'_k such that $x_{k-1} \leq x'_k \geq x_{k+1}$. Applying this successively we reach a chain $x = x'_0 \leq \dots \leq x'_k \geq \dots \geq x'_n = y$.

(iii). Similar to (ii.), noting that $x_{k-1} \prec x_k \succ x_{k+1}$ implies $x_{k-1} = x_{k+1}$.

For $x \in W$, we shall denote $\uparrow x := \{z \in W : x \leq z\}$. 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 $\uparrow x \cap \uparrow y \neq \emptyset$, thus every pair of nonempty upsets contained in the same connected component has nonempty intersection.

This fact plus the last item in Lemma 3.4 have an immediate consequence:

Corollary 3.9. *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) .*

As an immediate consequence of this, plus the fact that S4.2 $_{K_1}$ + S4.2 $_{K_2}$ is the logic of frames (W, \leq_1, \leq_2) where each \leq_i is a weakly directed preorder, we have:

Theorem 3.10. *S4.2 $_{K_1}$ + S4.2 $_{K_2}$ is the logic of topological-partitional models for two agents.*

3.2 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 interpret these on topological-partitional models $(X, \tau_{1,2}, \Pi_{1,2}, V)$ as follows:

$$\begin{aligned} x \in \|\forall_i \phi\| &\text{ iff } \Pi_i(x) \subseteq \|\phi\|; \\ x \in \|\Box_i \phi\| &\text{ iff } x \in \text{Int}_{\tau_i} \|\phi\|; \\ x \in \|\hat{B}_i \phi\| &\text{ iff } \text{Int}_{\tau_i} \|\phi\| \text{ is locally dense in } \Pi_i(x). \end{aligned}$$

Analogously to the one-agent case, we can check that the following equalities hold: $\|K_i \phi\| = \|\Box_i \phi \wedge [\forall]_i \Diamond_i \Box_i \phi\|$; $\|B_i \phi\| = \|\hat{K}_i K_i \phi\|$.

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* $\mathcal{L}_{K_i B_i}$, and the *factive evidence fragment* $\mathcal{L}_{\Box_i \forall_i}$.

The factive evidence fragment $\mathcal{L}_{\Box_i \forall_i}$. The logic for this fragment is $\text{Logic}_{\Box_i \forall_i}$, which is the least normal modal logic which includes

- the axioms and rules of S4 for \Box_i ;
- the axioms and rules of S5 for $[\forall]_i$;
- the axiom $[\forall]_i \phi \rightarrow \Box_i \phi$ for $i = 1, 2$.

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

For completeness, we can use the Sahlqvist completeness theorem (see [9]) and note that the axioms of $\text{Logic}_{\Box_i \forall_i}$ 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 (due to the S4 axioms) and each \sim_i constitutes an equivalence relation (due to the S5 axioms). Moreover, the axiom $[\forall]_i \phi \rightarrow \Box_i \phi$ 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 $\phi \notin \text{Logic}_{\Box_i \forall_i}$, then ϕ will be refuted in the canonical model, whence we have a topological-partitional model refuting it. And thus, we have completeness. \square

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 $\text{Logic}_{K_i \forall_i}$, the least logic including the axioms and rules of S4 for each K_i , S5 for each $[\forall]_i$ plus the following axioms for $i = 1, 2$:

- (A) $[\forall]_i \phi \rightarrow K_i \phi$;
- (B) $[\exists]_i K_i \phi \rightarrow [\forall]_i \hat{K}_i \phi$.

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 the Sahlqvist theorem. 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 3.8, 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. \square

The knowledge-belief fragment $\mathcal{L}_{K_i B_i}$. The logic of the knowledge-belief fragment is $\text{Stal}_1 + \text{Stal}_2$ the least normal modal logic including the S4 axioms and rules for K_i plus the following axioms, for $i = 1, 2$:

$$\begin{array}{ll} (\text{PI}_i) & B_i \phi \rightarrow K_i B_i \phi; & (\text{NI}_i) & \neg B_i \phi \rightarrow K_i \neg B_i \phi; \\ (\text{KB}_i) & K_i \phi \rightarrow B_i \phi; & (\text{CB}_i) & B_i \phi \rightarrow \neg B_i \neg \phi; \\ (\text{FB}_i) & B_i \phi \rightarrow B_i K_i \phi. \end{array}$$

We have that $\text{S4.2}_{K_1} + \text{S4.2}_{K_2} \cup \{B_i \phi \leftrightarrow \hat{K}_i K_i \phi : \phi \in \mathcal{L}_{K_i B_i}\} \subseteq \text{Stal}_1 + \text{Stal}_2$ and thus, if a formula ϕ in the language $\mathcal{L}_{K_i B_i}$ is not provable in $\text{Stal}_1 + \text{Stal}_2$, we can rewrite it as per into a formula in the language \mathcal{L}_{K_i} which is not provable in $\text{S4.2}_{K_1} + \text{S4.2}_{K_2}$. By completeness of the latter, there is a topological-partitional countermodel for ϕ , and completeness of $\text{Stal}_1 + \text{Stal}_2$ follows. \square

4 Generic models for two agents

In their famous paper [14], McKinsey and Tarski prove that S4 is not only the logic of topological spaces when one considers the interior semantics (i.e. when one reads $\|K\phi\| = \text{Int } \|\phi\|$), but that there are single topological spaces, such as the real line \mathbb{R} or the rationals \mathbb{Q} , whose logic is precisely S4. In [1], the authors of this paper have been concerned with finding *generic models* such as these for the logic of single-agent topo-e-models. In this section we provide two examples of generic models for the multi-agent logic, i.e., two topological-partitional spaces whose logic is precisely $\text{S4.2}_{K_1} + \text{S4.2}_{K_2}$.

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 2.

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.




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

It is proven in [5] that the logic of this frame under the usual Kripke semantics is $\mathbf{S4} + \mathbf{S4}$. This result is a corollary of the following proposition which we will use in our proof:

Proposition 4.1 ([5]). *Given a finite frame $\mathfrak{F} = (W, \preceq_1, \preceq_2)$, where \preceq_1 and \preceq_2 are both preorders, there exists a p -morphism from $\mathcal{T}_{2,2}$ onto \mathfrak{F} , i.e., a surjective map $\mathfrak{p} : \mathcal{T}_{2,2} \rightarrow \mathfrak{F}$ such that, for $i = 1, 2$, (i) $x \leq_i y$ implies $(\mathfrak{p}x) \preceq_i (\mathfrak{p}y)$, and (ii) $(\mathfrak{p}x) \preceq_i v$ implies there exists $y \in T$ such that $x \leq_i y$ and $\mathfrak{p}y = v$.*

Completeness of $\mathcal{T}_{2,2}$ with respect to $\mathbf{S4.2}_{K_1} + \mathbf{S4.2}_{K_2}$. 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 3.8, 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 4.2. *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 $\mathfrak{p} : \mathcal{T}_{2,2} \rightarrow W$ 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 4.3. *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). Then the following holds:*

- i. For any $x \in \mathcal{T}_{2,2}$, $\mathfrak{p}[x]_{\sim_i} \subseteq [\mathfrak{p}x]_{\equiv_i}$.

- ii. Let $x_0 \in \mathcal{T}_{2,2}$ and let U be a (locally dense) σ_i -open set such that $\mathbf{p}x_0 \in U \subseteq [\mathbf{p}x_0]_{\equiv_i}$. Then $U' := \bigcup\{\uparrow_i x : x \sim_i x_0 \ \& \ \mathbf{p}x \in U\}$ 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 \mathbf{p} preserves order, we have that $\mathbf{p}y \geq_i \mathbf{p}z \leq_i \mathbf{p}x$ and thus $\mathbf{p}y \equiv_i \mathbf{p}x$.

(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, $\mathbf{p}(\uparrow_i z)$ is an open set (by openness of \mathbf{p}) and $\mathbf{p}(\uparrow_i z) \subseteq \mathbf{p}[x_0]_{\sim_i} \subseteq [\mathbf{p}x_0]_{\equiv_i}$. By local density of U there exists some $a \in U \cap \mathbf{p}(\uparrow_i z)$. That is, for some $z' \geq_i z$ we have $\mathbf{p}z' = a$ and $\mathbf{p}z' \in U$, thus by construction $z' \in \uparrow_i z \cap U'$ and thus $\uparrow_i z \cap U' \neq \emptyset$.

As a consequence:

Proposition 4.4. *For any $x \in \mathcal{T}_{2,2}$ and any formula ϕ in the language, $\mathcal{T}_{2,2}, x \models \phi$ if and only if $\mathfrak{W}, \mathbf{p}x \models \phi$.*

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

$$\mathbf{p}x \in \mathbf{p}U \subseteq \mathbf{p}[x]_{\sim_i} \subseteq [\mathbf{p}x]_{\equiv_i},$$

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

Conversely, suppose $\mathbf{p}x \models K_i \phi$. Then there exists a (locally dense) σ_i -open set U with $\mathbf{p}x \in U \subseteq [\mathbf{p}x]_{\equiv_i}$ such that $w \models \phi$ for all $w \in U$. But then by part (ii) of the previous lemma $U' := \bigcup\{\uparrow_i z : z \sim_i x \ \& \ \mathbf{p}z \in U\}$ 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 $\mathbf{p}z \in U$. But since \mathbf{p} is order preserving we have that $\mathbf{p}y \geq_i \mathbf{p}z$ and thus $\mathbf{p}y \in U$, which means that $\mathbf{p}y \models \phi$ and thus, by induction hypothesis, $y \models \phi$. This means that $U' \subseteq \|\phi\|^{\mathcal{T}_{2,2}}$ and thus $x \models K_i \phi$.

Completeness is now an immediate consequence.

Corollary 4.5. *S4.2 $_{K_1}$ + S4.2 $_{K_2}$ is sound and complete with respect to the quaternary tree $(\mathcal{T}_{2,2}, \leq_1, \leq_2, \sim_1, \sim_2)$.*

The product $\mathbb{Q} \times \mathbb{Q}$. Let us now 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 S4.2 $_{K_1}$ + 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 4.6. Let (X, τ) and (Y, σ) be two topological spaces. The horizontal and vertical topologies, τ_H and τ_V , are the topologies on $X \times Y$ generated, respectively, by the bases

$$\mathcal{B}_H = \{U \times \{y\} : U \in \tau, y \in Y\} \text{ and } \mathcal{B}_V = \{\{x\} \times V : x \in X, V \in \sigma\}.$$

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 4.7 ([5]). $S4 + S4$ is the logic of $(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V)$ under the interior semantics.

Now we shall show there exists a 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.

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

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

be a surjective map which is open and continuous in both topologies. We shall call this an *onto interior map*. 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 4.8. $\mathfrak{X} = (X, \tau_1, \tau_2, \equiv_1, \equiv_2, V^f)$ is a topological evidence model and, for every formula ϕ in the language and every $x \in X$ we have that $\mathfrak{X}, x \models \phi$ if and only if $\mathfrak{Y}, fx \models \phi$.

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 the image of this class coincides with the equivalence class of fx , i.e. $f[x]_{\equiv_i} = [fx]_{\sim_i}$. Indeed, $fy \in [fx]_{\sim_i}$ implies $y \in [x]_{\equiv_i}$, and thus $fy \in f[x]_{\equiv_i}$; conversely $y \in f[x]_{\equiv_i}$ implies $y = fx'$ for some $x' \equiv_i x$ and thus $y = fx' \sim_i fx$. Now, $[fx]_{\sim_i}$ is an equivalence class and thus an open set and, since f is continuous, $f^{-1}[fx]_{\sim_i}$ is also an open set. So it suffices to show that $f^{-1}[fx]_{\sim_i} = [x]_{\equiv_i}$. And indeed, if $z \in f^{-1}[fx]_{\sim_i}$ then $fz \in [fx]_{\sim_i} = [fx]_{\sim_i}$ 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 , $\mathfrak{X}, x \models \phi$ if and only if $\mathfrak{Y}, fx \models \phi$, and let $\mathfrak{X}, x \models K_i \phi$. 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 $U \cap V \neq \emptyset$. But then we have that $fx \in f[U]$, the set $f[U]$ is open (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 $\mathfrak{Y}, fx \models K_i\phi$. 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 $f[V]$ is an open set contained in $[fx]_{\sim_i}$ and thus there exists some $y \in f[V] \cap [fx]_{\sim_i}$. 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 $\mathfrak{X}, x \models K_i\phi$.

It is proven in [5] that there exists an onto map $f : \mathbb{Q} \times \mathbb{Q} \rightarrow \mathcal{T}_{2,2}$, 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 $\text{S4.2}_{K_1} + \text{S4.2}_{K_2}$.

Corollary 4.9. *Let $f : \mathbb{Q} \times \mathbb{Q} \rightarrow \mathcal{T}_{2,2}$ be some onto interior map. Define $(x, y) \equiv_i^f (x', y')$ iff $f(x, y)$ and $f(x', y')$ belong to the same \leq_i -connected component for $i = 1, 2$. Then $\text{S4.2}_{K_1} + \text{S4.2}_{K_2}$ is sound and complete with respect to*

$$(\mathbb{Q} \times \mathbb{Q}, \tau_H, \tau_V, \equiv_1^f, \equiv_2^f).$$

This existence result is not in itself very satisfactory, and it leads to the immediate question: what do these partitions look like? One could easily think of suitable candidates, (such as

$$\equiv_1 = \{\mathbb{Q} \times (\pi + k, \pi + k + 1) : k \in \mathbb{Z}\} \ \& \ \equiv_2 = \{(\pi + k, \pi + k + 1) \times \mathbb{Q} : k \in \mathbb{Z}\},$$

for instance) but none of the ‘obvious’ candidates seem to give us completeness (in the example above, $[\forall]_1\phi \rightarrow \Box_2\phi$ holds everywhere, while not being a theorem of the logic). This problem is open for future work.

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.

In the present section we consider the notions of *distributed* and *common* knowledge applied to this framework.

5.1 Distributed knowledge

We can think of *distributed knowledge* 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 $p \rightarrow q$ 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 $\bigcap_{a \in \mathcal{A}} R_a$.

Let us remark something here: given two preorders \leq_1 and \leq_2 defined on a set X , let τ_i be the topology of \leq_i -upwards closed sets for $i = 1, 2$. The Kripke semantics on (X, \leq_1, \leq_2) correspond with the interior semantics on (X, τ_1, τ_2) , and the collection of upwards-closed sets of the relation $\leq_1 \cap \leq_2$ is precisely the *join topology* $\tau_1 \vee \tau_2$, i.e., the least topology containing $\tau_1 \cup \tau_2$, or, equivalently, the topology generated by $\{U_1 \cap U_2 : U_i \in \tau_i\}$. We will be using join topologies in our approach.

A problematic approach. What exactly amounts to distributed knowledge in our framework? A very direct way to translate the ideas presented so far 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 ϕ (i.e., a *locally dense* piece of evidence).

This approach, while intuitive, has two issues. On the one hand, it might be the case that an agent has a piece of evidence for ϕ which is dense in her topology (i.e., she knows ϕ) yet, when the evidence of both agents is put together, the corresponding evidence is no longer locally dense in the partition of the join topology (i.e., the group does not know ϕ).⁷ Obviously, this is undesirable.

On the other hand, this notion 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*. According to [12],

It 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. Thus, from this lens, one might want to keep misleading evidence out of the equation, and consider that this hypothetical ‘wise man’ forms his knowledge based not on what the agents have evidence for, but rather on what the agents *actually know*.

⁷ In [10] this is discussed in more depth and an example is provided.

On this account, instead of each agent having a piece of *evidence* that, when combined together, constitute a *justification* for ϕ , we would want for each to have a justification which combine into a piece of evidence for ϕ . (For a more in-depth argument, see [10]).

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 $K_1\phi \rightarrow D\phi$. In the following we present a way to have such a notion.

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 5.1 (Semantics for D). Let $\mathfrak{X} = (X, \tau_{1,2}, \Pi_{1,2}, V)$ be a topological-partitional model. We read $\|p\|$, $\|\phi \wedge \psi\|$, $\|\neg\phi\|$ and $\|K_i\phi\|$ as in Def. 3.3, and:

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

While the problematic semantics outlined above 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 $\tau_1^* \vee \tau_2^*$.

The logic of distributed knowledge. Let $\text{Logic}_{K_i D}$ be the least set of formulas containing:

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

Theorem 5.2. $\text{Logic}_{K_i D}$ is sound and complete with respect to topological-partitional models.

We will dedicate the rest of this subsection to showing this fact.

Soundness. That every topological-partitional model satisfies the S4.2 axioms for K_i can be proven exactly as in section 3.1. That D satisfies the S4 axioms is a consequence of D being read as $\text{Int}_{\tau_1^* \vee \tau_2^*}$. And for the two extra axioms, if $x \models K_i\phi$, 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 $x \models D\phi$.

Completeness. 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 $TR_D S$ and $K_i \phi \in T$, then $D\phi \in T$ as per the axiom $K_i \phi \rightarrow D\phi$ and thus $\phi \in S$.

A *labelled path over X* is a path

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

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

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

Now, let \mathcal{T} be the smallest set of labelled paths over X such that: (i.) The path T_0 (of length 0) belongs to \mathcal{T} ; (ii.) For $i = 1, 2$, if $\alpha \in \mathcal{T}$ and $(\text{last } \alpha)R_i T$, then $\alpha \xrightarrow{R_i} T \in \mathcal{T}$; (iii.) If $\alpha \in \mathcal{T}$ and $(\text{last } \alpha)R_D T$, then $\alpha \xrightarrow{R_D} T \in \mathcal{T}$.

For $i = 1, 2, D$ we define: $\alpha \prec_i \beta$ if and only if $\alpha = \beta \xrightarrow{R_i} S$ for some $S \in X$.

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} : for $i = 1, 2$, let \leq_i be the reflexive and transitive closure of $\prec_i \cup \prec_D$ 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 3.8, we know that the connected components of the topology of upsets of \leq_i ($i = 1, 2$) are given by the equivalence relation: $\alpha \sim_i \beta$ iff there exists γ such that $\alpha \geq_i \gamma \leq_i \beta$. The definition of \leq_i plus the fact that $R_D \subseteq R_i$ entail that $(\text{last } \gamma)R_i(\text{last } \alpha)$ and $(\text{last } \gamma)R_i(\text{last } \beta)$. Therefore we have the following result:

Lemma 5.3. *If α and β belong to the same \leq_i -connected component on \mathcal{T} , then $\text{last } \alpha$ and $\text{last } \beta$ 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 4.2, which we will find useful:

Lemma 5.4. *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.*

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

Lemma 5.5 (Truth lemma). *For every $\alpha \in \mathcal{T}$ and ϕ in the language, $\alpha \models \phi$ if and only if $\phi \in \text{last } \alpha$.*

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 $K_i \phi \in \text{last } \alpha$. 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 $\beta \in U_i$, $\phi \in \text{last } \beta$, which will give us, by induction hypothesis, that $U_i \subseteq \|\phi\|$. By lemma 5.4, 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

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

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

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

Note that the elements in $V_D^{[i]}$ are of the form $\beta = \alpha_i \xrightarrow{R_D} T_1 \xrightarrow{R_D} \dots \xrightarrow{R_D} T_n$, and the elements in V_i are of the form

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

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

$$U_i := \{\beta \in V_D^{[i]} : (\text{last } \alpha)R_D(\text{last } \beta)\} \cup \{\gamma \in V_i : (\text{last } \alpha)R_i(\text{last } \gamma)\}.$$

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 $(\text{last } \alpha)R_i(\text{last } \beta)R_i S$, thus $(\text{last } \alpha)R_i S$. If $\beta \prec_D \gamma$ then $\beta = \gamma \xrightarrow{R_D} S$ and, if $\beta \in V_D^{[i]}$ we then have that $\gamma \in V_D^{[i]}$ and $(\text{last } \alpha)R_D(\text{last } \beta)R_D S$ (thus $(\text{last } \alpha)R_D(\text{last } \gamma)$) whereas if $\beta \in V_i$ we have that $\gamma \in V_i$ and similarly (given that $R_D \subseteq R_i$), $(\text{last } \alpha)R_i S$. In any case $\gamma \in U_i$.
- iii. U_i is locally dense. Take any $\beta \in [\alpha]_{\sim_i}$. By lemma 5.3, we have that $\text{last } \beta$ and $\text{last } \alpha$ are in the same R_i -connected component and, since R_i is an S4.2 relation, part (ii) of Lemma 3.8 gives us that there exists some $S \in X$ with $(\text{last } \alpha)R_i S$ and $(\text{last } \beta)R_i S$ and thus we have $\beta \xrightarrow{R_i} S \in U_i \cap \uparrow_i \beta$.
- iv. $\phi \in \text{last } \beta$ for every $\beta \in U_i$ (given that $K_i \phi \in \text{last } \alpha$ and $(\text{last } \alpha)R_i(\text{last } \beta)$).

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 $\alpha \in U_i \subseteq [\alpha]_{\sim_i} \cap \|\phi\|$. Since U_i is an upset, if $(\text{last } \alpha)R_i S$, we have $\alpha \xrightarrow{R_i} S \in U_i$, which means $\alpha \xrightarrow{R_i} S \in \|\phi\|$ and by induction hypothesis $\phi \in S$. Every R_i -successor of $\text{last } \alpha$ includes ϕ , which gives $K_i \phi \in \text{last } \alpha$.

Now suppose $D\phi \in \text{last } \alpha$. 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

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

and of the form

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

These can only be true at the same time if β is of the form

$$\beta = \alpha_j \xrightarrow{R_D} S_1 \dots \xrightarrow{R_D} S_m \text{ with } \alpha_j = \alpha_i \xrightarrow{R_i \text{ or } R_D} T_1 \dots \xrightarrow{R_i \text{ or } R_D} T_k$$

for $i \neq j \in \{1, 2\}$. Let us assume w.l.o.g. that $i = 1, j = 2$. In particular we have that, if $\beta \in U_1 \cap U_2$, then $\beta \in V_D^{[2]}$ and hence $(\text{last } \alpha)R_D(\text{last } \beta)$. Since $D\phi \in \text{last } \alpha$, this entails that $\phi \in \text{last } \beta$ and thus that $\beta \in \|\phi\|$, whence $\alpha \models D\phi$.

For the converse, if $\alpha \models D\phi$ 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 $(\text{last } \alpha)R_D S$ 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 $\alpha \xrightarrow{R_D} S \models \phi$ which by induction hypothesis gives $\phi \in S$. This entails $D\phi \in \text{last } \alpha$.

Completeness follows from this: if $\phi \notin \text{Logic}_{K_i D}$, then $\{\neg\phi\}$ is consistent and can be extended as per Lindenbaum's lemma to some maximal consistent 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 $\alpha \not\models \phi$ as per the truth lemma.

5.2 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. [3] compares the following 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 $E\phi$ is an abbreviation of $K_1\phi \wedge K_2\phi$, then

$$C\phi \equiv \phi \wedge E\phi \wedge EE\phi \wedge EEE\phi \wedge \dots$$

- (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] goes on to argue that, 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”.

Moreover, while (1) and (2) are equivalent in relational semantics, as shown in [7] 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 5.6 (Common knowledge semantics). Let $\mathfrak{X} = (X, \tau_{1,2}, \Pi_{1,2}, V)$ be a topological-partitional model. We read

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

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

Much like our account of 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 S4.2 relations: if R_1 and R_2 are 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 $(R_1 \cup R_2)^*$.

Another observation is that, in the spirit of [3], this definition is precisely the fixed point account of common knowledge. As pointed out in [7] and expanded in [8], 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, $\|C\phi\| = \bigcup \{U \in \mathcal{P}(X) : U \subseteq \|\phi \wedge Ep\|^{V_p^U}\}$. It is straightforward to check that this last set equals

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

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 5.7 (Soundness). *All the theorems above are valid on topological-partitional models with the semantics of definition 5.6.*

Proof. That i., ii. and iii. hold for topological-partitional models is a 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 5.6 this means that there exists some $U \in \tau_1^* \cap \tau_2^*$ such that

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

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

6 Conclusions and future work

This paper presents a multi-agent generalisation for the dense interior semantics defined on topological evidence models, furthering the results in [2].

This was achieved by introducing a second epistemic agent and a partition-based semantics. We showed how this semantics generalises the single agent case and we provided a complete logic for our two-agent models. Moreover, ‘generic spaces’ were provided with respect to which the logic is sound and complete: the quaternary tree $\mathcal{T}_{2,2}$ and the rational plane $\mathbb{Q} \times \mathbb{Q}$. Along with this, a brief conceptual and theoretical study of notions of “group knowledge” for this group of agents was developed.

Some questions remain unanswered (and some potentially interesting results were out of the scope of this investigation). Among these are the following:

- We have proven (Cor. 4.9) that there exist partitions \equiv_1 and \equiv_2 making $\mathbb{Q} \times \mathbb{Q}$ a model for the logic $\text{S4.2}_{K_1} + \text{S4.2}_{K_2}$. What would be an example of such partitions?
- Are $\mathcal{T}_{2,2}$ and $\mathbb{Q} \times \mathbb{Q}$ generic models for any (all) of the fragments of the language considered in Section 3.2? For the distributed knowledge logic defined in the last section?
- Can we more broadly characterize a class of topological - partitional spaces which are generic for the logic? For example it is shown in [1] that, for the one-agent case, any topological space which is dense-in-itself, metrizable and idempotent is a generic model for the logic. Is a similar result true in the multi-agent setting?
- Does the list of theorems presented in Section 5.2 constitute a complete axiomatization of the logic of common knowledge?

Acknowledgements. We wish to thank Guram Bezhanishvili for very fruitful discussions on these topics. Special thanks go to the reviewers for their very valuable comments, which helped us improve the paper.

References

1. Baltag, A., Bezhanishvili, N., Fernández González, S.: The McKinsey-Tarski theorem for topological evidence logics. In: International Workshop on Logic, Language, Information, and Computation. pp. 177–194. Springer (2019)
2. Baltag, A., Bezhanishvili, N., Özgün, A., Smets, S.: Justified belief and the topology of evidence. In: International Workshop on Logic, Language, Information, and Computation. pp. 83–103. Springer (2016)
3. Barwise, J.: Three views of common knowledge. In: Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge. pp. 365–379 (1988)
4. van Benthem, J., Bezhanishvili, G.: Modal logics of space. In: Handbook of spatial logics, pp. 217–298. Springer (2007)
5. van Benthem, J., Bezhanishvili, G., ten Cate, B., Sarenac, D.: Multimodal logics of products of topologies. *Studia Logica* **84**(3), 369–392 (2006)
6. van Benthem, J., Pacuit, E.: Dynamic logics of evidence-based beliefs. *Studia Logica* **99**(1-3), 61 (2011)
7. van Bethem, J., Sarenac, D.: The geometry of knowledge. *Travaux de logique* **17**, 1–31 (2004)
8. Bezhanishvili, N., van der Hoek, W.: Structures for epistemic logic. In: Johan van Benthem on logic and information dynamics, pp. 339–380. Springer (2014)
9. Blackburn, P., De Rijke, M., Venema, Y.: *Modal Logic*, vol. 53. Cambridge University Press (2001)
10. Fernández González, S.: *Generic Models for Topological Evidence Logics*. Master’s thesis, University of Amsterdam (2018)
11. Gettier, E.L.: Is justified true belief knowledge? *Analysis* **23**(6), 121–123 (1963)
12. Halpern, J.Y., Moses, Y.: A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence* **54**(3), 319–379 (1992)
13. Hintikka, J.: *Knowledge and belief: an introduction to the logic of the two notions*. Contemporary philosophy, Cornell University Press (1962)
14. McKinsey, J.C.C., Tarski, A.: The algebra of topology. *Annals of mathematics* pp. 141–191 (1944)
15. Munkres, J.R.: *Topology*. Prentice Hall (2000)
16. Özgün, A.: *Evidence in Epistemic Logic: A Topological Perspective*. Ph.D. thesis, University of Amsterdam, ILLC (2017)
17. Ramírez, A.: *Topological Models for Group Knowledge and Belief*. Master’s thesis, University of Amsterdam (2015)
18. Stalnaker, R.: On logics of knowledge and belief. *Philosophical studies* **128**(1), 169–199 (2006)