

J.-Y. Beziau, A. Costa-Leite and A. Facchini (eds)
Aspects of Universal Logic, 1-31

The Geometry of Knowledge

Johan van Benthem and Darko Sarenac

Abstract

The most widely used attractive logical account of knowledge uses standard epistemic models, i.e., *graphs* whose edges are indistinguishability relations for agents. In this paper, we discuss more general *topological* models for a multi-agent epistemic language, whose main uses so far have been in reasoning about space. We show that this more geometrical perspective affords greater powers of distinction in the study of common knowledge, defining new collective agents, and merging information for groups of agents.

1 Epistemic logic in its standard guise

1.1 Basic epistemic logic

Epistemic logic is in wide use today as a description of knowledge and ignorance for agents in philosophy [14], computer science [13], [22], game theory [12], and other areas. In this paper, we assume familiarity with the basic language of propositional epistemic logic, interpreted over multi-agent **S4** models whose accessibility relations are reflexive and transitive. Alternative model classes occur, too, such as equivalence relations for each agent in multi-agent **S5**—but our discussion is largely independent from such choices. The key semantic clause about an agent's knowledge of a proposition says that $K_i\phi$ holds at a world x if and only if ϕ is true in all worlds y accessible for i from x . That is, the epistemic knowledge modality is really a modal box $\Box_i\phi$. For technical convenience, we will use the latter notation




Figure 1: In the black central world, 1 does not know if p , while 2 does know that p . In the world to the left, 1 does know that p , so in the central world, 2 does not know if 1 knows that p .

for knowledge in the rest of this paper. The main modern interest in epistemic logic has to do with analyzing iterated knowledge of agents about themselves and what others know, for purposes of communication and interaction. Cf. [4], [9] on systems that combine epistemic logic and dynamic logic to describe information update in groups of agents. A simple example of how the basic logic works is the model in Figure 1.

The universally valid principles in our models are those of multi-agent **S4**. In an epistemic setting, the usual modal axioms get a special flavor. E.g., the iteration axiom $\Box_1\phi \rightarrow \Box_1\Box_1\phi$ now expresses ‘positive introspection’: agents who know something know that they know it. More precisely, we have **S4**-axioms for each separate agent, but no valid further ‘mixing axioms’ for iterated knowledge of agents, such as $\Box_1\Box_2\phi \rightarrow \Box_2\Box_1\phi$. Indeed, the latter implication fails in the above example. For instance, in the world on the left, 1 has no uncertainties, and so 1 knows that 2 knows that p . But 2 does not know there that 1 knows that p , because the latter assertion is false in the central world. Another way of describing the set of valid principles is as a *fusion* $\mathbf{S4} \oplus \mathbf{S4}$ of separate logics **S4** for each agent, a perspective of ‘merging logics’ to which we will return below. In what follows, we shall mostly work with two-agent groups, $G = \{1, 2\}$, since most phenomena of interest can be studied there. Generalizations to finite k -agent cases are straightforward.

1.2 Group knowledge

Perhaps the most interesting topic in an interactive epistemic setting has been the discovery of various notions of what may be called *group knowledge*. Two well-known examples are as follows:

1. $E_G\phi$: every agent in group G knows that ϕ ,
2. $C_G\phi$: ϕ is *common knowledge* in the group G .

The latter notion of group knowledge is much stronger than the former. It has been proposed in the philosophical, economic and linguistic literature as a necessary precondition for coordinated behavior between agents, cf. [16]. The usual semantic definition of common knowledge runs as follows:

$$M, x \models C_{1,2}\phi \text{ iff for all } y \text{ with } x (R_1 \cup R_2)^*y, M, y \models \phi$$

where $x(R_1 \cup R_2)^*y$ if there is a finite sequence of successive steps from either of the two accessibility relations connecting x to y . This relation is the reflexive transitive closure of the union of the relations for both agents. The key valid principles for common knowledge are the following additional axiom and rule:

$$\begin{array}{l} \text{Equilibrium Axiom: } C_{1,2}\phi \leftrightarrow (\phi \wedge (\Box_1 C_{1,2}\phi \wedge \Box_2 C_{1,2}\phi)) \\ \text{Induction Rule: } \frac{\vdash p \rightarrow (\Box_1(q \wedge p) \wedge \Box_2(q \wedge p))}{\vdash p \rightarrow C_{1,2}q} \end{array}$$

This logic is known as $S4_2^C$. It has been shown to be complete and decidable in [13] via a simple variation on similar proofs for propositional dynamic logic.

But there are still further interesting notions of knowledge for a group of agents. A prominent one is so-called *implicit knowledge*, $D_G\phi$, which describes what a group would know if its members decided to merge their information:


$$M, x \models D_{1,2}\phi \text{ iff for all } y \text{ with } xR_1 \cap R_2y, M, y \models \phi$$

where $R_1 \cap R_2$ is the intersection of the accessibility relations for the separate agents. This new notion is technically somewhat different from the earlier two in that, unlike universal and common knowledge, it is not invariant under modal *bisimulations* of epistemic models. It also involves a new

phenomenon of independent epistemic interest: viz. merging the information possessed by different agents. The latter topic will return throughout this paper.

1.3 Agents as epistemic accessibility relations

We can also think of new notions of group knowledge as introducing *new agents*. E.g., C_G defines a new kind of **S4**-agent, since $R_{(1 \cup 2)^*}$ was again a pre-order. Note that $R_1 \cup R_2$ by itself is not a pre-order, so the new ‘agent’ corresponding to the fact that ‘everybody knows’ would have different epistemic properties. In particular, it would lack positive introspection as to what it knows. In contrast, the relation $R_1 \cap R_2$ for D_G is again an **S4**-agent as it stands, since Horn conditions like transitivity and reflexivity are preserved under intersections of relations. So, given a group of individual agents, our logical models suggest new agents. In particular, with two **S4**-agents 1, 2, two additional ones supervene on these, one weaker, one stronger:



All this seems quite rich as an account for epistemic agents. And yet, there are indications that this framework is not yet flexible enough for its tasks.

1.4 Alternative views of common knowledge

Despite the success of the standard epistemic logic framework, there are still doubts about its expressive power and sensitivity. Some recurrent complaints seem endemic to logical approaches as such, like the vexing problem of logical omniscience: agents automatically know all laws of the system. But a more serious concern is the lack of epistemic distinctions in the standard modal setting. Notably, in his well-known critical paper [6], Barwise claimed that a proper analysis of common knowledge must distinguish three different approaches, that we may label

1. countably *infinite iteration* of individual knowledge modalities,
2. the *fixed-point view* of common knowledge as ‘equilibrium’,
3. agents’ having a *shared epistemic situation*.

He then showed how to distinguish all three in a special situation-theoretic framework. As we will see below, however, Barwise’s distinctions make sense in mainstream logic too—provided that we move to a broader topological semantics for the epistemic language involving products of models for individual agents. But before we do that, let us first analyze the reason why standard epistemic logic fails to distinguish the first two options. The third notion of ‘shared understanding’ is somewhat more mysterious, and harder to grasp in a standard relational modal setting. We will have a stab at it in the richer topological models of Section 2.

1.5 Computing epistemic fixed-points

The above Equilibrium Axiom for the common knowledge operator $C_G\phi$ shows how it may be viewed as defining a fixed-point of an epistemic operator $\lambda X.\phi \wedge \Box_1 X \wedge \Box_2 X$. In conjunction with the Induction Rule, it may even be seen to be a *greatest fixed-point* definable in the standard modal μ -calculus as:

$$C_G\phi := \nu p.\phi \wedge \Box_1 p \wedge \Box_2 p.$$

With a perhaps more familiar modal μ -operator, its existential variant would be defined as a smallest fixed-point

$$\Diamond_G^C\phi := \mu p.\phi \vee \Diamond_1 p \vee \Diamond_2 p.$$

As usual, a greatest fixed-point is defined as the fixed-point of a descending approximation sequence defined over the set of ordinals. We write $[\![\phi]\!]$ for the truth set of ϕ in the relevant model where evaluation takes place:

$$C_{1,2}^0\phi := [\![\phi]\!],$$

$$C_{1,2}^{\kappa+1}\phi := [\phi \wedge \Box_1(C_{1,2}^\kappa\phi) \wedge \Box_2(C_{1,2}^\kappa\phi)],$$

$$C_{1,2}^\lambda\phi := [\bigwedge_{\kappa < \lambda} C_{1,2}^\kappa\phi], \text{ for } \lambda \text{ a limit ordinal.}$$

Finally, we let $C_{1,2}\phi := C_{1,2}^\kappa\phi$ where κ is the least ordinal for which the approximation procedure halts: i.e., $C_{1,2}^{\kappa+1}\phi = C_{1,2}^\kappa\phi$. This approximation procedure must stop at some ordinal because the operator F applied is *monotonic*, a fact which is guaranteed by the positive occurrence of the propositional variable p in the body of F 's definition. As a result, the approximation sequence for a greatest fixed-point operator always descends to subsets, and hence it must stop eventually. In general μ -calculus, reaching this stopping point may take any number of ordinal stages. A standard example is the least-fixed-point formula $\mu p.\Box p$ which computes the so-called ‘well-founded part’ of the binary accessibility relation for the modality. But in certain cases, stabilization is guaranteed to occur by the first infinite stage.

Fact 1.1 *In every relational epistemic model, the approximation procedure for the common knowledge modality stabilizes at $\kappa \leq \omega$.*

This simple behavior is most easily understood by observing that knowledge modalities \Box_i distribute over any infinite conjunction. Thus, $\Box_i(\bigwedge_{n < \omega} C_{1,2}^n\phi)$ is simply $\bigwedge_{n < \omega} \Box_i C_{1,2}^n\phi$ which is equivalent to $\bigwedge_{n < \omega} C_{1,2}^n\phi$. More generally, stabilization for a formula $\nu p.\phi(p)$ is guaranteed by stage ω in any model just in case the syntax defining the monotone approximation operator is constrained as follows [10]. The formula $\phi(p)$ must be a disjunction whose members are constructed using only

1. arbitrary literals $(-)q$,
2. any epistemic formulas that do not contain q at all,
3. conjunctions and universal modalities.

The preceding Fact says that the fixed-point approach to common knowledge and that with countably infinite conjunctions of repeated knowledge modalities are equivalent in the standard setting, as $\nu p.\phi \wedge \Box_1 p \wedge \Box_2 p$ is equivalent to

$$K_{1,2}p := \phi \wedge \Box_1\phi \wedge \Box_2\phi \wedge \Box_1\Box_2\phi \dots$$

This equivalence is often considered a technical convenience. But it may also indicate that our standard models are too weak to make a relevant distinction, and that more general models are needed. As we shall see, these two definitions of common knowledge are different in a *topological* modelling for epistemic logic—and even stronger ones can then be modelled, resembling Barwise’s use of ‘shared situations’.

1.6 Merging Information

Many further interesting issues are raised by a multi-agent epistemic setting. In particular, multi-agent models will often arise by *merging* models for separate agents, or groups of agents, so that common knowledge for the whole group becomes possible at all. One natural way of combining models for two or more agents emphasized in the recent literature on combining modal logics employs *products* of their underlying frames. More precisely,

Definition 1.2 *The product of two frames $\mathcal{F}_1 = (W_1, R_1)$ and $\mathcal{F}_2 = (W_2, R_2)$ is the frame $\mathcal{F}_1 \times \mathcal{F}_2 = (W_1 \times W_2, R_1, R_2)$ with R_1 defined as*

$$(x, y)R_1(z, w) \text{ iff } xR_1z \ \& \ y = w$$

and the relation R_2 defined likewise.

Sometimes one also adds the direct product relation $R_{1,2}$ which requires successor steps in both components. But in the present setting, this is definable as the relational composition of R_1 and R_2 in any order.

This way of combining modal logics is explored in detail in [15]. The separate logics of the component frames are preserved in the product, as is easy to see. But the really interesting question is what happens in the joint language containing both modalities \Box_1 and \Box_2 , which can express interaction between epistemic agents. As it turns out, by a simple argument, product frames automatically validate the following two axioms:

$$\begin{array}{ll} (com) & \Box_1\Box_2p \equiv \Box_2\Box_1p \\ (chr) & \Diamond_1\Box_2p \rightarrow \Box_2\Diamond_1p \end{array}$$

[15] contains much more information on these principles, including general results on when they suffice for axiomatizing the complete logic of frame products over the merge of the component logics. But note that these two

principles were not valid in the general fusion logic $\mathbf{S4} \oplus \mathbf{S4}$ of epistemic agents, as we saw earlier. Figure 1 provided a formal counterexample to *com*. To put such a scenario in words: a student may know that the teacher knows the answer to questions on the test, while the teacher does not know if the student knows the answer. Moreover, if *com* does become valid, common knowledge trivializes, since any finite sequence of knowledge modalities will be equivalent to one of \Box_1, \Box_2 or $\Box_1\Box_2$.

Now there are other notions of merge for epistemic models, and the preceding collapse of common knowledge need not occur with other operations. Often, merging information for single agents or groups of agents is more naturally viewed as an operation on *models*, rather than frames. And in that case, the necessity of obtaining a consistent atomic valuation on pairs of worlds may complicate the above product construction, and thereby block *com* and *chr*. We discuss this issue briefly in Section 2.7. But for our purposes later on with analyzing common knowledge, frame products are important, provided we generalize them, again, to a wider topological setting. In that case, the two undesirable epistemic interaction laws no longer hold, and the above trivialization of common knowledge goes away.

We have now accumulated enough motivation for looking into broader alternative semantics for a multi-agent language, which should be fine-grained enough to distinguish different notions of common knowledge, while being sufficiently robust to still provide a plausible version of epistemic logic. We find this in the following mathematical generalization of relational models.

2 Epistemic Models in Topological Semantics

2.1 From graphs to topological spaces.

One of the major alternatives to relational semantics for modal logics, and historically even the earlier approach, employs *topological* models. Before going into our main epistemic concerns, we present this semantics here with its usual interpretation. Topology is an abstract mathematical theory of space, emphasizing qualitative notions of open environment, closure, boundary, or connectedness.

Definition 2.1 *A topological space \mathcal{X} is a pair (X, τ) where X is a set of*

'points', and the set of 'opens' $\tau \subseteq \wp(X)$ contains X, \emptyset , and is closed under finite intersections and arbitrary unions.

Example 2.2 A typical example is the structure of the rationals with \mathbb{Q} for the set X and the standard metric topology generated by closing the set of bounded open intervals $\{p | q < p < q'\}$ for $p, q, q' \in \mathbb{Q}$ under arbitrary unions. The standard topology on the reals \mathbb{R} is obtained in the same fashion.

The language \mathcal{L} of propositional modal logic is just as before, with a countable set of propositional variables At , and the formulae defined recursively:

$$\phi := p \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond p \mid \Box p$$

On the topological interpretation, Booleans are interpreted as the corresponding set operations, $\Box p$ as the topological interior of the set of points assigned to p , and $\Diamond p$ as the closure of the set assigned to p . More precisely, a topological model $\mathcal{M} = \langle X, \tau, V \rangle$ consists of a topological space $\langle X, \tau \rangle$ with a valuation function $V : At \rightarrow \wp(X)$. The key clauses of the truth definition then read:

$$\begin{aligned} \mathcal{M}, x \models \Box\phi & \text{ iff } (\exists U \in \tau)(x \in U \text{ and } (\forall y \in U)(\mathcal{M}, y \models \phi)), \\ \mathcal{M}, x \models \Diamond\phi & \text{ iff } (\forall U \in \tau)(x \in U \Rightarrow (\exists y \in U)(\mathcal{M}, y \models \phi)). \end{aligned}$$

All topological modalities in this paper satisfy the axioms of the modal logic **S4**, which reflect key properties of the topological interior operation. The interesting epistemic details then lie in the interaction among such modalities. For the moment, we cite two well-known results from [17]:

Theorem 2.3 **S4** is a complete axiomatization of modal \Box interpreted over arbitrary topological spaces.

More striking, and much deeper, is the following result.

Theorem 2.4 **S4** is a complete axiomatization of modal \Box on any metric space that is dense-in-itself.

This theorem shows that **S4** is the complete logic of \mathbb{Q} , \mathbb{R} , \mathbb{Q}^2 , and many other interesting topologies close to our ordinary understanding of space.

Topological semantics generalizes standard modal model theory. A basic example is *bisimulation* for relational models (cf. [2]). Its pervasive invariance properties generalize to topological models.

Definition 2.5 (Topological Bisimulation) *A topo-bisimulation between two topological models $\langle \mathcal{X}, \tau, V \rangle$ and $\langle \mathcal{X}', \tau', V' \rangle$ is a nonempty relation $E \subseteq \mathcal{X} \times \mathcal{X}'$ such that, whenever xEx' , then:*

1. $x \in V(p)$ iff $x' \in V'(p)$, for every proposition letter p ,
2. (forth condition) $x \in U \in \tau$ implies that there is a $U' \in \tau'$, $x' \in U'$ and for every $y' \in U'$ there is a $y \in U$ with yEy' .
3. (back condition) $x' \in U' \in \tau'$ implies that there is a $U \in \tau$, $x \in U$ and for every $y \in U$ there is a $y' \in U'$ with yEy' .

Proposition 2.6 (Invariance for Bisimulation) *Let $\mathcal{M} = \langle \mathcal{X}, \tau, V \rangle$ and $\mathcal{M}' = \langle \mathcal{X}', \tau', V' \rangle$ be models with points x and x' related by some topo-bisimulation. Then, $\mathcal{M}, x \models \phi$ iff $\mathcal{M}', x' \models \phi$ for all modal formulas ϕ .*

The following special case of this result is the topological counterpart of the ‘generated submodels’ in relational semantics. Truth values only depend on what happens in arbitrarily small open neighbourhoods.

Proposition 2.7 (Topological Locality) *Let $\mathcal{X} = \langle X, \tau \rangle$ be a topological space, with $x \in U \in \tau$ and ν some valuation on X . Then, for any formula ϕ , $(\mathcal{X}, \nu), x \models \phi$ iff $(\mathcal{X}|U, \nu|U), x \models \phi$, where $\mathcal{X}|U$ is the topology obtained by taking U as the universe, letting the opens be all sets $U \cap U' \in \tau$, while $\nu|U = \nu(p) \cap U$ for all p .*

Topo-bisimulations are closely related to a more standard topological notion.

Definition 2.8 *Let $\mathcal{X} = (X, \tau_1), \mathcal{Y} = (Y, \tau_2)$ be two topological spaces. A map $f : X \rightarrow Y$ is said to be*

1. open, if the f -image of any open set in τ_1 , is open in τ_2 ,
2. continuous, if the f -inverse image of any open set in τ_2 is open in τ_1 .

It is easy to show that open continuous maps preserve modal theories of topological spaces, just as ‘modal p -morphisms’ preserve theories of relational frames.

This is a good point for stating the general connection between the two classes of models for modal or epistemic languages. Standard relational models can be viewed as a special kind of topological spaces through the following notion.

Definition 2.9 *A topological space \mathcal{X} is Alexandroff if every intersection of open sets of \mathcal{X} is again open.*

Any Alexandroff topology $\mathcal{X} = \langle X, \tau \rangle$ induces a standard relational frame $\langle X, R \rangle$ with a reflexive transitive relation Rxy iff $y \in \bigcap \{U \in \tau \mid x \in U\}$. Conversely, any reflexive transitive relational frame $\langle X, R \rangle$ induces an Alexandroff topology by taking the sets $U_x = \{y \mid Rxy\}$ for each $x \in X$ as a basis for τ . It is easily shown that topological interpretation of modal formulas in a relational model yields the same results as in their associated Alexandroff spaces, and vice versa. In this way, modal logics of relational models describe special sets of topological models. But in general, topological models include settings without a clear relational counterpart. E.g., the standard topologies on \mathbb{Q} and \mathbb{R} are clearly not Alexandroff: any singleton set (a non-open) is the intersection of the open intervals containing it.

There is a recent revival of interest in modal **S4** interpreted over topological spaces, because of its applications to spatial reasoning. [1] and [2] survey the expressive power of **S4** and its extensions for this purpose. We will use a few results from this spatial line later on. But before we cite them, let us make a connection with our major concern of what agents know.

2.2 Topology and information

Dating back to the 1930s, there has also been a more epistemic use of topological models, viz. for *intuitionistic* logic, cf. [20]. In that case, open sets are rather interpreted as ‘pieces of evidence’, e.g., about the location of a point, reflecting the intuitionistic idea of truth-as-provability. We can generalize this idea to epistemic logic, reading the above truth condition for a knowledge modality $\Box_i p$ as saying that there exists a piece of evidence for agent i (viz. an open set in i ’s topology) which validates the proposition p . Alternatively, we could also think of the topology as a collection of

theories or data bases that an agent has at its disposal. [21] contains more abstract versions of this idea. As we will see, one of the side benefits of this information-based interpretation of the epistemic language is that common knowledge arises in a group of agents precisely when they share the same piece of information. But first, we explore the new handle that we get on the issue of merging information structures for different agents.

2.3 Combination of agents in topological products

To deal with epistemic merges, we need some results from recent work on products of topological spaces developed originally in the setting of spatial reasoning in [11].

Products of topological spaces \mathcal{X}, \mathcal{Y} occur quite often, and they support a variety of new topologies. We start with a particularly simple way of ‘lifting’ the two components to *one-dimensional topologies* on the grid space $X \times Y$, which we sometimes visualize as ‘horizontal’ and ‘vertical’ directions in a plane.

Definition 2.10 *Let $\mathcal{X} = \langle X, \eta \rangle$ and $\mathcal{Y} = \langle Y, \theta \rangle$ be two topological spaces. Suppose $A \subseteq X \times Y$. We say that A is horizontally open (H-open) if for any $(x, y) \in A$ there exists $U \in \eta$ such that $x \in U$ and $U \times \{y\} \subseteq A$. Similarly, we say that A is vertically open (V-open) if for any $(x, y) \in A$ there exists $V \in \theta$ such that $y \in V$ and $\{x\} \times V \subseteq A$. If A is both H- and V-open, then we call it HV-open. Dual closed sets are defined as usual.*

We can now interpret the modal operators \Box_1 and \Box_2 of the combined language $\mathcal{L}_{\Box_1, \Box_2}$ in product models $\langle X \times Y, \tau_1, \tau_2 \rangle$ with some arbitrary valuation for proposition letters. The two key clauses will read as follows:

$$(x, y) \models \Box_1 \phi \text{ iff } (\exists U \in \eta)(x \in U \ \& \ \forall u \in U : (u, y) \models \phi)$$

$$(x, y) \models \Box_2 \phi \text{ iff } (\exists V \in \theta)(y \in V \ \& \ \forall v \in V : (x, v) \models \phi)$$

In order to visualize this semantics, it helps to think of ‘grids’ of ordered pairs where one topology runs along horizontal lines, and the other along vertical ones. Next, we say that a formula ϕ of the language $\mathcal{L}_{\Box_1, \Box_2}$ is *valid* at (x, y) in a product space $X \times Y$ if for every valuation on that space $(x, y) \models \phi$. The following proposition then tells us that the structural theories of component topologies (or agents’ knowledge) ‘lift’ to the product




Figure 2: In A, the valuation $\nu(p) = (\bigcup_{x \in (-1,0)} \{x\} \times (x, -x)) \cup (\{0\} \times (-1, 1)) \cup (\bigcup_{x \in (0,1)} \{x\} \times (-x, x))$ falsifies *com* at $(0, 0)$. In B, $\nu'(p) = \bigcup \{ \{\frac{1}{n}\} \times (-\frac{1}{n}, \frac{1}{n}) : n \in \mathbb{N} \}$ falsifies *chr* at $(0, 0)$.

space without any additions. Unlike the case of products of relational frames in Section 1.6, topological product does not automatically enforce new interaction principles between agents.

Proposition 2.11 *A formula ϕ constructed from atoms, Booleans and the modal operator \Box_1 is valid at a point $(x, y) \in \langle X \times Y, \tau_1, \tau_2 \rangle$ iff ϕ is valid at x in \mathcal{X} . The same is true for the language with \Box_2 only, by taking the right projection.*

This was a result for the separate sublanguages of the agents. Moving to the joint language $\mathcal{L}_{\Box_1 \Box_2}$, it can be shown that the earlier product interaction principles *chr* and *com* fail on topological products. Figure 2 shows graphically how these failures occur for suitable valuations ν, ν' on the two-dimensional real plane:

$$(\mathbb{R} \times \mathbb{R}, \nu), (0, 0) \not\models \Diamond_V \Box_H p \rightarrow \Box_H \Diamond_V p$$

$$(\mathbb{R} \times \mathbb{R}, \nu'), (0, 0) \not\models \Box_H \Box_V p \rightarrow \Box_V \Box_H p$$

Next, we turn to matters of complete axiomatization. The following result from [11] says that topological products perform the most minimal merge of modal logics, without interactive side-effects for modalities.

Theorem 2.12 *The fusion logic $S4 \oplus S4$ is complete with respect to products of arbitrary topological spaces.*

As in the single-agent case, one can prove stronger results for particular structures, and in fact, we have the following:

Theorem 2.13 *$S4 \oplus S4$ is complete with respect to $\mathbb{Q} \times \mathbb{Q}$.*

A detailed proof of this result can be found in [11]. For later reference, we give a sketch here.

The first major observation to be made is that $S4 \oplus S4$ is complete for the infinite quaternary tree $T_{2,2}$, using a standard modal unravelling procedure for countable relational models. To transfer modal counter-examples from that tree to topological products, we need to make a second step, showing that $T_{2,2}$ is the image of an *HV*-open subset of the ‘rational plane’ $\mathbb{Q} \times \mathbb{Q}$ under some *HV*-continuous and *HV*-open map. Such a map is constructed in stages via the following procedure, which is easily visualized. Let $T_{2,2}^n$ be the nodes of $T_{2,2}$ of *R*-depth *n*. Now, iteratively label a sequence of growing subsets of $\mathbb{Q} \times \mathbb{Q}$ with nodes of $T_{2,2}$, as follows:

Stg 0: Label $(0, 0)$ with the root *r* of the tree $T_{2,2}$.

Stg 1: Label $(-1, 0)$ with the immediate left R_1 -successor, and $(1, 0)$ with the immediate right R_1 -successor of *r*; also label $(0, -1)$ with the immediate left R_2 -successor, and $(0, 1)$ with the immediate right R_2 -successor of *r*. Call these four points *environmental points at the distance $\frac{1}{3^0}$* .

Stg *n*: The environmental points labelled at Stage *n* – 1 are at the distance no smaller than $\frac{1}{3^{n-1}}$. Now for each of labelled points we create four environmental points at the distance $\frac{1}{3^n}$ —two at the vertical distance $\frac{1}{3^n}$ and two at the horizontal distance $\frac{1}{3^n}$ —and label them with respective immediate R_1 - and R_2 -successors in the tree.

This procedure labels a subset of $\mathbb{Q} \times \mathbb{Q}$ which can be contracted, modulo isomorphism, to an *HV*-open subset of $\mathbb{Q} \times \mathbb{Q}$. Moreover, there is an obvious map *f* taking labelled points in this set to nodes in the tree $T_{2,2}$.

A straightforward verification shows that this map is both HV -continuous and HV -open. Obviously, we can copy any valuation on the tree to one on $\mathbb{Q} \times \mathbb{Q}$ backward along the map f . Thus, if some modal formula is refuted in the root of the tree under some valuation, we get a topo-bisimulation with a model whose domain is a HV -open subset of the rational plane. By the above Locality Lemma 2.7, this counter-example can be lifted to the whole model $\mathbb{Q} \times \mathbb{Q}$, which is what we wanted.

Thus the fusion $\mathbf{S4} \oplus \mathbf{S4}$ is the logic of two epistemic agents combined into one framework using topological products, without any dramatic interaction enforced as in the case of products of relational frames. This result gives us the technical means to analyze different versions of common knowledge in a concrete setting of merged multi-agent models.

2.4 Common knowledge in product spaces

The earlier definitions of common knowledge still make sense in topological models. For instance, countably infinite iteration of all finite sequences of alternating knowledge modalities for the individual agents 1, 2 is as before:

$$K_{1,2}p := \bigwedge_n K_{1,2}^n p,$$

with $K_{1,2}^n p$ defined inductively as follows:

$$\begin{aligned} K_{1,2}^0 p &:= p \\ K_{1,2}^{n+1} p &:= \Box_1(K_{1,2}^n p) \wedge \Box_2(K_{1,2}^n p) \end{aligned}$$

And the same is true for the fixed-point definition

$$C_{1,2}\phi := \nu p. \phi \wedge \Box_1 p \wedge \Box_2 p,$$

provided we make the appropriate adjustments in computing fixed points. In particular, the monotone operations generated by formulas positive in p now work a bit differently from before. In relational models, the operator \Box_i applied to a set X yielded $\Box_i(X) = \{y \mid \forall x (R_i y x \rightarrow x \in X)\}$, making the modality a bounded universal quantifier. In topological semantics,

however, the relevant operator is

$$\Box_i(X) = \{y \mid \exists U \in \tau_i \ \& \ \forall x(x \in U \rightarrow x \in X)\}$$

This reads a modality as an existential quantifier over open sets followed by a universal quantifier over elements of those sets. This two-quantifier combination complicates matters when approximating greatest or smallest fixed-points. Indeed, the definitions of common knowledge by fixed-points and by countably infinite iteration will now diverge. Here is a first indication why this may happen. The topological semantics validates the finitary logic **S4**, but it diverges from the relational validities in its infinitary behaviour.

Fact 2.14 *Topological interior does not distribute over infinite conjunctions:*

$$\Box_i \bigwedge_n p_n \text{ is not always equivalent to } \bigwedge_n \Box_i p_n$$

Take the standard topology on \mathbb{Q} . Define a valuation ν with, for all n , $\nu(p_n) = (-\frac{1}{n}, \frac{1}{n})$. Note that the intersection of these open sets is the singleton 0. Then $\bigwedge_n \Box_i p_n$ is true at 0, whereas $\Box_i \bigwedge_n p_n$ is not true anywhere. This result, though suggestive, is not yet a proof that the two definitions of common knowledge diverge. To do that, we will show that given a set p , the operator $K_{1,2}p$ does not always define a horizontally and vertically open set. Since the fixed-point version of $C_{1,2}p$ is always open in both these senses, the two cannot be the same.

We construct the relevant example by choosing a countable sequence of points in the rational plane $\mathbb{Q} \times \mathbb{Q}$ horizontally converging to the origin $(0, 0)$. The first point in the sequence makes $\Box_1 p$ true but not $\Box_2 \Box_1 p$, the second $\Box_1 \Box_2 p, \Box_2 \Box_1 p$ but not $\Box_2 \Box_1 \Box_2 p$, etc. This is possible by Theorem 2.12 for the logic of \Box_1, \Box_2 : no finite iteration level of knowledge implies the next in the fusion logic **S4** \oplus **S4**, and hence situations as described must exist in suitable models over $\mathbb{Q} \times \mathbb{Q}$. In particular, at each point of the sequence, $K_{1,2}$ will be false, and hence $\Box_1 K_{1,2} p$ is false at the origin $(0, 0)$. It then remains to show that $K_{1,2} p$ itself does hold at $(0, 0)$, but this will happen because of a well-chosen total valuation $\nu(p)$ for p on $\mathbb{Q} \times \mathbb{Q}$. To make this work, we make a number of more precise observations— while also slightly changing the formulas involved:

Theorem 2.15 $K_{1,2}p \rightarrow \Box_1 K_{1,2}p$ is not valid on topological product spaces.

Let ψ_n be the formula $\Box_1(K_{1,2}^n p) \rightarrow \Box_2(K_{1,2}^n p)$.

Fact 2.16 (a) For all n , ψ_n is not a theorem of the fusion logic **S4** \oplus **S4**.

(b) There is a model M_n on $\mathbb{Q} \times \mathbb{Q}$ such that $M_n, (0, 0) \not\models \Box_2(K_{1,2}^n p)$, and for all $q \in \mathbb{Q}$, $M_n, (q, 0) \models K_{1,2}^n p$.

As for (a), one can easily construct finite fusion frames invalidating any given principle ψ_n .

(b) Since **S4** \oplus **S4** is complete for $\mathbb{Q} \times \mathbb{Q}$, by (a) there is a model M'_n such that $M'_n, (0, 0) \not\models \psi_n$, that is,

$$M'_n, (0, 0) \models \Box_1(K_{1,2}^n p)$$

as well as

$$M'_n, (0, 0) \not\models \Box_2(K_{1,2}^n p).$$

It follows that there is an open interval $((-q, 0), (q, 0))$ and every $(q', 0)$ in this interval satisfies $K_{1,2}^n p$. By Locality (Proposition 2.7), in $(-q, q) \times \mathbb{Q}$ with the valuation from M'_n restricted to this space it is still true that $\Box_2(K_{1,2}^n p)$ fails at $(0, 0)$ and that $K_{1,2}^n p$ holds at each point $(q', 0)$. But $(-q, q) \times \mathbb{Q}$ is homeomorphic to $\mathbb{Q} \times \mathbb{Q}$ itself, and hence the valuation of M'_n transfers to $\mathbb{Q} \times \mathbb{Q}$ via the homeomorphism.

Fact 2.17 There is a sequence of positive irrational numbers converging to 0 such that for any two adjacent numbers r, r' in the sequence, the distance $r - r'$ is a rational number.

Take for instance $\sqrt{2}, \sqrt{2}-1, \sqrt{2}-1.4, \sqrt{2}-1.41$, etc. Next, for each rational interval, we form squares S_1, S_2, \dots of decreasing sizes over these intervals bounded by the separating irrationals [see Figure 3]. In the above example, the first square would be $(\sqrt{2}, \sqrt{2}-1) \times (-\frac{1}{2}, \frac{1}{2})$, the second $(\sqrt{2}-1, \sqrt{2}-1.4) \times (-0.2, 0.2)$, etc. Each of these squares is still homeomorphic to the rational plane $\mathbb{Q} \times \mathbb{Q}$ with some valuation for the proposition letter p .

Now, we create a new big model M over $\mathbb{Q} \times \mathbb{Q}$ as follows. In the sequence of squares S_n , we embed the earlier counter-examples M_n into S_n in such a way that its horizontal axis becomes the horizontal axis of the square S_n . This ensures that $K_{1,2}^n p$ holds everywhere on S_n 's X -axis while $\Box_2(K_{1,2}^n p)$ fails somewhere on it. Outside of the squares, we put every point of the total rational plane in $V(p)$. Now we can prove the earlier informal assertion.




Figure 3:

Claim 2.18 (a) $M, (0,0) \models K_{1,2}p$
 (b) $M, (0,0) \not\models \Box_1 K_{1,2}p$.

(a) We will prove that for all n , $K_{1,2}^n p$ holds at $(0,0)$. The proof is by induction. First note that any point on the y axis or to the left of it (except $(0,0)$) sits in an open circle interior in which p is true everywhere. Inside such a circle, these points evidently satisfy all formulas $K_{1,2}^n p$, and hence by Locality again, they also satisfy all these formulas in the whole model M .

Now we consider the origin $(0,0)$. The base step is simple: $K_{1,2}^0 p$ is true by the definition of $\nu(p)$. Next consider the inductive step $K_{1,2}^n p \Rightarrow K_{1,2}^{n+1} p$, where $K_{1,2}^{n+1} p$ is $\Box_1(K_{1,2}^n p) \wedge \Box_2(K_{1,2}^n p)$. We show that the two conjuncts hold separately. To see that $\Box_2(K_{1,2}^n p)$ holds at $(0,0)$ we need an open set $((0,y), (0,-y))$ with $K_{1,2}^n p$ true at each point in this set. Evidently, this formula holds at $(0,0)$ itself by the inductive hypothesis. And it holds at any other point on the Y axis by the preceding observation about open p -circles.

Next we show that $\Box_1(K_{1,2}^n p)$ holds at $(0,0)$. This time we need an interval of the form $((-y,0), (x,0))$ with $K_{1,2}^n p$ true at every point in the interval. Here, points in $((y,0), (0,0))$ are covered by the observation about open p -circles again, and the origin itself by the inductive hypothesis. Then, looking toward the right, by the construction of the squares S_n , we know that $K_{1,2}^n p$ holds everywhere at the horizontal axis of S_n , and the same

obviously remains true for S_m with $m > n$. Thus, for the desired right end-point $(x, 0)$ we can take any point on the horizontal axis of the square S_n . Since every point in $((0, 0), (x, 0))$ is in some S_m for $m \geq n$, we have the desired interval, and hence $\Box_1(K_{1,2}^n p)$ is true at the origin. In this connection, the idea behind our ‘gluing’ the squares at irrationals was that inside $\mathbb{Q} \times \mathbb{Q}$, there are then no boundary points to consider.

(b) To see that $\Box_1 K_{1,2} p$ fails at $(0, 0)$, we observe that in any horizontal open interval I around $(0, 0)$ there is a point where $K_{1,2} p$ fails. Note that for some n , the horizontal axis of S_n is a subset of I , by our construction of ever smaller squares S_n , and hence there is a point inside our interval where $\Box_2(K_{1,2}^n p)$ fails, and hence also $K_{1,2} p$, as desired.

Corollary 2.19 $K_{1,2} p$ is not equivalent to $C_{1,2} p$ in topological models.

Corollary 2.20 Stabilization of the fixed-point version of $C_{1,2} X$ may occur later than ordinal stage ω .

Thus, the topological setting achieves a natural separation between the first two definitions of common knowledge that Barwise distinguished. Moreover, our method raises further issues. First, it is rather ‘logicky’, and one might want a concrete independently motivated set of points in the rational plane for which the separation occurs. Also, it would be of interest to determine the exact (countable) ordinals at which epistemic fixed-point definitions do stabilize in this model.

This still leaves Barwise’s third account of common knowledge in terms of ‘shared situations’. We shall return to this matter in Section 2.6.

2.5 Complete logic of common knowledge on topo-products

Now what is the basic logic of the greatest fixed-point common knowledge modality $C_{1,2}$ on topological models? Perhaps surprisingly, the general answer is: ‘the same as that for relational $\mathbf{S4}$ -models’. The reason is that the usual system $\mathbf{S4}_2^C$ already has principles for common knowledge that are satisfied by the fixed-point definition. Moreover, that system is complete w.r.t. relational models [13], and the latter are Alexandroff topological models at the same time. More interesting is what happens in our topological product models. In fact, the logic does not change here either, but this time, the argument takes a little more thought.

Theorem 2.21 $S4_2^C$ is complete for products of arbitrary topologies. In fact it is even the complete logic of $\mathbb{Q} \times \mathbb{Q}$.

The completeness argument runs along the lines of the earlier one for the language without common knowledge: this is why we sketched the main proof steps for Theorem 2.12 in some detail. By the usual completeness proof with respect to relational models, any non-theorem of $S4_2^C$ fails on some finite rooted modal model. Next, such a model can be unravelled via a bisimulation into the double-binary branching tree $T_{2,2}$ with an appropriate valuation. Now we do the labelling construction described in the proof of Theorem 2.12. In the end, this procedure produced a topo-bisimulation between the given model on $T_{2,2}$ and some model on the rational plane $\mathbb{Q} \times \mathbb{Q}$. Now the only thing we need to observe is that topo-bisimulations do not just preserve truth values of ordinary modal formulas. They also evidently preserve truth values of formulas in any modal language allowing *infinite* conjunctions and disjunctions of formulas. And, the latter observation gives us exactly what we need to transfer counterexamples to formulas in the epistemic language with common knowledge viewed as a fixed-point operator.

Fact 2.22 *Topological bisimulations preserve arbitrary fixed-point formulas.*

In any given model M , any modal fixed-point formula ϕ is equivalent to some modal formula $\phi(\alpha)$ which has no fixed-point operators any more, but which uses infinite conjunctions and disjunctions up to a size determined by the ordinal α to ‘unwind’ approximation sequences. What this α depends on is the size of the model M . Moreover, it does not matter if we unwind up to any higher ordinal. Now, suppose that some fixed-point formula ϕ is true at M, s , and E is a bisimulation connecting s to t in a model N, t . Let α^* be the maximum of the unwinding ordinals for ϕ in the two models M, N . Then $\phi(\alpha^*)$ is true at s in M , and therefore also true at t in N . It follows that the original fixed-point formula ϕ is true in N, t .

Even so, given the difference between $C_{1,2}\phi$ and $K_{1,2}\phi$ that we have now found, a new completeness question arises, yet to be solved:

Question:

What is the complete logic of $K_{1,2}\phi$?

Given all this emphasis on geometrical models like the rational plane, can we really claim that they are also epistemically relevant? Our discussion only shows their use as visualizations of abstract distinctions. Whether there is any deeper *informational* meaning to $\mathbb{Q} \times \mathbb{Q}$ still remains to be seen.

In the remainder of this paper, we discuss some further aspects of the topological semantics for knowledge, analogous to those raised in Section 1.

2.6 More on epistemic agents as topologies

In relational semantics, agents were really just accessibility relations. Likewise, in our topological models, agents are topologies! As was explained in Section 2.2, what the agent knows in a world of some model is what holds there according to the box modality of its topology. Let us now draw some comparisons with the situation in Section 1.3., where two agents 1, 2 generated at least two further ‘introspective collective agents’, one being their supremum $R_{(1 \cup 2)^*}$ leading to common knowledge, and the other their infimum $R_1 \cap R_2$ leading to ‘implicit knowledge’ for the group. The topological semantics gives us interesting counterparts to these operations.

Remark.

Introspection principles If we are less strict in our logic, without requiring positive introspection, then many further options arise, just as with relational models. If we are more strict, as in relational **S5**-models with negative introspection, then we must only use topologies that do satisfy the axiom $\phi \rightarrow \Box \Diamond \phi$. It is easy to see that, on T_0 spaces in which all singletons are closed, imposing this principle makes the topology discrete, trivializing the epistemic logic. But then, even a weak separation axiom like T_0 is not plausible epistemically. On general spaces, $\phi \rightarrow \Box \Diamond \phi$ corresponds to the property that every set is a subset of the interior of its closure. Unpacked further this says that:

$$\forall x, \exists U \in \tau : x \in U \ \& \ \forall y \in U, y \in V \in \tau : x \in V$$

This means the space is a union of open sets whose points have the same open neighbourhoods – which is a topological counterpart of relational **S5**

models.

Our favorite setting for studying new collective agents are the product models that we used so far. We start with a simple but perhaps surprising observation. Common knowledge as a greatest fixed-point corresponds to taking the following very natural operation on the given topologies for the individual agents. Consider the *intersection* $\tau_{1 \cap 2}$ of the earlier topologies τ_1 and τ_2 on a product space. It is easy to see that this is again a topology: all closure conditions are satisfied. Now we observe the following connection:


Fact 2.23 $\forall M \forall x, M, x \models C_{1,2}\phi$ iff $M, x \models [1 \cap 2]\phi$

We will show that the truth sets $\llbracket C_{1,2}\phi \rrbracket$ and $\llbracket [1 \cap 2]\phi \rrbracket$ are identical in all models. First, $\llbracket C_{1,2}\phi \rrbracket \in \tau_i$ for $i \in \{1, 2\}$ since the truth set is a fixed-point of $\nu p. \phi \wedge \Box_1 p \wedge \Box_2 p$. But then $\llbracket C_{1,2}\phi \rrbracket \in \tau_{1 \cap 2}$ by the definition, and so $\llbracket C_{1,2}\phi \rrbracket \subseteq \llbracket [1 \cap 2]\phi \rrbracket$. Next, $\llbracket [1 \cap 2]\phi \rrbracket$ satisfies $\llbracket \Box_i [1 \cap 2]\phi \rrbracket = \llbracket [1 \cap 2]\phi \rrbracket$ for $i \in \{1, 2\}$. Hence $\llbracket [1 \cap 2]\phi \rrbracket$ is a fixed-point. Since $\llbracket C_{1,2}\phi \rrbracket$ is the greatest fixed-point, $\llbracket [1 \cap 2]\phi \rrbracket \subseteq \llbracket C_{1,2}\phi \rrbracket$.

It is worth observing that this argument holds in general, for any two given topologies on some space, not just the vertical and horizontal ones in products. In fact, intersection of topologies is the counterpart, under the model-to-topology transformation sketched earlier, of taking the reflexive transitive closure of given accessibility relations.

Thus, we also expect a topological counterpart for the earlier operation of *relational intersection*, which modelled implicit group knowledge D_G . This should be the union of two topologies, and then closing off in the minimal way that produces a topology again. The result is the *sum topology* $\tau_1 + \tau_2$ which takes all pairwise intersections of opens of the two topologies as a basis. The latter topology need not always be of great interest. E.g., on our recurrent topo-product $\mathbb{Q} \times \mathbb{Q}$, it will just be the discrete topology, making every point an open. From an informational perspective, this means that merging the information that we get about points in the horizontal and vertical directions fixes their position uniquely.

The result of all this is again an inclusion diagram:



Let us now return to the three distinctions made in [6]. So far, we have separated the countably infinite conjunction view from the greatest fixed-point view of common knowledge. What about the third view of having a ‘shared situation’? In some ways, using the intersection topology seems to model this. Its opens are precisely those information pieces that are accepted by both agents. But if that is the case, then we have not separated the second and third notions. Fact 2.23 tells us precisely that the two amount to the same thing. But topological product models have further resources! In particular, so far, we have not discussed what topologists would call the real *product topology* τ on spaces $X \times Y$. This topology is defined by letting the sets $U \times V$ form a basis, where U is open in \mathcal{X} and V is open in \mathcal{Y} . An example is the natural metric topology on the plane $\mathbb{Q} \times \mathbb{Q}$, used briefly in the argument for Claim 2.18, with open circles around points as neighbourhoods. The agent corresponding to this new group concept τ only accepts very strong collective evidence for any proposition. Here are two relevant results from [11]:

Theorem 2.24 *The epistemic box modality for the true product topology is not definable in the language of the separate modalities \Box_1, \Box_2 , even when we add fixed-point operators.*

Theorem 2.25 *The complete logic including the true product topology is the smallest normal modal logic in the language of three modalities \Box, \Box_1, \Box_2 that contains (i) the S4 axioms for \Box_1, \Box_2 and \Box , (ii) $\Box p \rightarrow \Box_1 p$ and $\Box p \rightarrow \Box_2 p$.*

Thus, we have found an even stronger notion of common knowledge that might be said to model Barwise’s third stage. Nevertheless, there are some difficulties with this identification. For instance, unlike the preceding two operations of intersection and union closure, true product topology has no general definition on arbitrary models for our language, as it exploits the

product structure essentially. This makes it rather specialized, and this same fact is also reflected in the poverty of the complete logic given above. Nevertheless, there are also interesting logical aspects to this situation. In contrast with the *sequential quantification* embodied in the greatest fixed-point reading of common knowledge, the true product modality reads more like a *branching quantifier* as defined in [7]. We do not know what to make epistemically of this tantalizing analogy at this stage.

2.7 Operations that are safe for topo-bisimulation

To illustrate the preceding notions of knowledge and agency a bit further, we add a brief digression on simulations between topological models.

In relational semantics for modal languages, most natural operations $f(R_1, R_2)$ have the property of being *safe for bisimulation*, that is,

- any given bisimulation between two models w.r.t. the relations R_1, R_2 is also a bisimulation for the relation $f(R_1, R_2)$.

This says that the new operation stays at the same level of model structure as the old. The regular operations of composition, union, and iteration on binary relations are all safe in this sense, while a typical non-safe operation is *intersection*. Safety is a natural extension of invariance for static formulas to dynamic transition relations ([10] has a complete characterization of all first-order definable safe operations). Safety constrains the repertoire of definable transition relations within one given model. In general process theories, new relations can also be constructed out of old while forming a new model at the same time, as happens with products for concurrent processes in Process Algebra. In that setting, safety for operations generalizes to *respect for bisimulation*, e.g., if we let \cong signify bisimulation:

- if $M \cong M'$ and $N \cong N'$, then $f(M, N) \cong f(M', N')$.

Most natural product operations show respect for bisimulation. As a check on our new notions, we can also look at operations on topologies in the same way, substituting the above topological bisimulations for the usual relational ones.

Of the repertoire of regular operations, only a small part matters in our perspective. When working only with reflexive transitive relations, composition and union by themselves do not qualify as operations, and we




Figure 4:

need to take $*$ -closures. And for reflexive-transitive $R_1, R_2, (R_1 \cup R_2)^*$ and $(R_1; R_2)^*$ yield even the same relation. The topological counterpart for the latter operation was *intersection of topologies* $\tau_1 \cap \tau_2$, as noted above. Fact 2.23 expressed the observation that the modality for this is the same as the common knowledge fixed-point modality for the modal operators $[\tau_1], [\tau_2]$. The latter is invariant for topological bisimulations by earlier observations. Indeed we have the following

Fact 2.26 *Intersection of topologies is safe for topological bisimulation.*

Let E be a relation between topological models M, N which is a topological bisimulation for their two separate topologies, as in Figure 4.

For a start, let sEt , and $s \in U$ with U in $\tau_1 \cap \tau_2$. Since E is a bisimulation w.r.t. τ_1 , there is a τ_1 -open set V in M' such that every point $v \in V$ is E -related to some point u in U . Likewise, there is a τ_2 -open set W in M' such that every point $v \in W$ is E -related to some point u in U . Now, it may be tempting to take the intersection of V and W at t for the required matching neighbourhood of U , but this need not be open in either topology. Instead, we consider every E -link between points u in U and points v in the union $V \cup W$. Using the bisimulation properties again, there are again both τ_1 and τ_2 -open neighbourhoods for all such points u , which satisfy the backward zigzag condition toward U . Continuing this procedure countably many times, the union of all these successively produced subsets of M' is both τ_1 - and τ_2 open, and moreover, it still satisfies the correct backward zigzag condition w.r.t. the original open neighbourhood U of s in M . The argument in the opposite direction is similar.




Figure 5:

This result may sound strange because intersection of binary relations led to non-invariance for bisimulation. But the topological counterpart of this operation was the sum topology $\tau_1 + \tau_2$ defined above, and its behaviour is indeed unsafe.

Fact 2.27 *Taking the sum of topologies is not safe for topological bisimulation.*

The counterexample is the same as for the relational case. Consider the two three-point models of Figure 5, with their topologies plus a binary relation E between their points as indicated.

Note that the sum topology on the left-hand side has the singleton set $\{s\}$ as an open, whereas the sum topology on the right has only the whole two-element space for a non-empty open. Also, the relation E is a bisimulation for both topologies τ_1 and τ_2 . Next, consider the link sEt , with the open subset $\{s\}$ on the left. The only matching open set on the right can be $\{s', v\}$, but this fails to satisfy the backward zigzag condition, as sEv does not hold.

Finally, more general operations may produce new topologies over combined spaces. Our characteristic example was topological product as in Definition 2.10.

Fact 2.28 *Topological products $\tau_1 \times \tau_2$ respect topological bisimulation.*

Let E_1 be a bisimulation w.r.t. τ_1 between models M, M' , and likewise E_2 a bisimulation w.r.t. τ_2 between models N, N' . Now define a bisimulation E between $M \times N, M' \times N'$ by setting:

$$(s, t)E(s', t') \text{ iff } sE_1s' \text{ and } tE_2t'.$$

Given Definition 2.10, it is completely straightforward to check that E is a bisimulation w.r.t both topologies on the product.

In contrast to this, taking a product of two topological spaces with the true product topology τ introduced a little while ago does not respect topological bisimulation. The reason is the earlier fact that the true product modality \square is not invariant for topological bisimulations w.r.t. the two component topologies.

2.8 Merging information revisited

Finally, we make a few comments on the issue of merging epistemic situations. We have shown that products of topological spaces are a natural setting for combining knowledge by different agents, and for distinguishing various forms of knowledge in the group of all agents. But as in Section 1, there is a broader question behind this. Our topological products are just one way of merging information models. The general subject of merging epistemic models goes far beyond the scope of this paper (cf. [8] for more on this topic). We only make one general point here which seems relevant to our move from relational semantics to topological models.

In general, we need to *specify* what we want to happen with existing knowledge and ignorance of agents when merging their information. Suppose we are given two epistemic models M for group G_1 and N for G_2 , where G_1, G_2 overlap. In that case, we may want to require that the intersection group does not learn anything new in the ‘merge model’ $M * N$, at least w.r.t. formulas in its old language. This situation is reminiscent of the process of *amalgamation* of relational models in semantic proofs of the interpolation theorem for the basic modal language (cf. [3] for an elementary exposition). Such proofs often start with a $G_1 \cap G_2$ *bisimulation* between models M, s and N, t , which serves as an initial connection between the two different settings. The relevant merge $M * N$ then turns out to be a *submodel of the full product* $M \times N$, viz. just those pairs which stand in that bisimulation. One then shows that the projections from pairs to the original models M, N are bisimulations for the separate languages. Hence, formulas in the intersection of the two languages retain one unambiguous truth value: the one they had before under the bisimulation. In the case

of interpolation theorems for shared modalities, this amalgamation construction has to be complicated, but the point remains the same. General merging of models for groups of agents may presuppose some initial connection, and its effects on modal formulas can be prescribed to some extent. In particular, we need not accept all pairs in a product as members of a merge model. Once we do this, the connection between topological models and relational models becomes more complicated, as we could also try to get the results of this paper with sub-product constructions on relational models. We refer to [18] for details.

3 Conclusion

Topological semantics for epistemic logic is a natural extension of the usual relational modelling. It provides distinctions that can be used to differentiate between various notions of common knowledge, and define various sorts of collective agents. Also, using product spaces, topological semantics suggests 'low-interaction' merges for epistemic models for separate groups of agents. Thus, we believe that there are good reasons for further development of this currently still marginal perspective.

References

- [1] M. Aiello, *Spatial reasoning: theory and practice*, ILLC Dissertation Series, no. 2002-02, University of Amsterdam, (2002).
- [2] M. Aiello, J. van Benthem, and G. Bezhanishvili, *Reasoning about space: the modal way*, Journal of Logic and Computation, pages 1–32, (2002).
- [3] H. Andréka, J. van Benthem and I. Németi, *Modal logics and bounded fragments of predicate logic*, Journal of Philosophical Logic 27:3, pages 217-274, (1998).
- [4] A. Baltag, L.S. Moss, and S. Solecki, *The logic of public announcements, common knowledge and private suspicions*, Proceedings TARK, 43-56, Morgan Kaufmann Publishers, Los Altos (1998).
- [5] J. Barwise, *On branching quantifiers in English*, Journal of Philosophical Logic 8, pages 47-80, (1979).
- [6] J. Barwise, *Three views of common knowledge*. Proceedings of TARK, pages 365-379, Morgan Kaufmann Publishers, Los Altos (1988).
- [7] J. Barwise, and R. Cooper. *Generalized quantifiers and natural language*. Linguistics and Philosophy 4, pages 159-219, (1981).
- [8] J. van Benthem. *Two logical concepts of information*, manuscript, ILLC, University of Amsterdam (2004).
- [9] J. van Benthem. *One is a lonely number*, ILLC Tech Report 2002-27, (2002).
- [10] J. van Benthem. *Exploring logical dynamics*, CSLI publications, Stanford (1997).

- [11] J. van Benthem, G. Bezhanishvili, B. ten Cate, and D. Sarenac, *Modal logics for products of topologies*, to appear, (2004).
- [12] K.G. Binmore. *Game theory and the social contract*. MIT Press, Cambridge, (1994).
- [13] R. Fagin, J.H. Halpern, Y. Moses, M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, (1994).
- [14] J. Hintikka. *The logic of epistemology and the epistemology of logic: selected essays*. Kluwer Academic Press, Boston, (1989).
- [15] D.M. Gabbay, A. Kurucz, F. Wolter and M. Zakharyashev. *Many-dimensional modal logics: theory and applications*. Studies in Logic and the Foundations of Mathematics, Volume 148. Elsevier, (2003).
- [16] D.K. Lewis. *Convention: a philosophical study*. Harvard University Press, Cambridge, (1969).
- [17] J. C. C. McKinsey and Alfred Tarski, *The algebra of topology*, Ann. of Math. (2) **45**, pages 141–191, (1944).
- [18] D. Sarenac. *Modal logic and topological products*, Ph.D Thesis, Stanford University. Forthcoming, (2005).
- [19] V.B. Shehtman, *Two-dimensional modal logics*, Mathematical notices of the USSR Academy of Sciences **23**, pages 417–424, (1978). (Translated from the Russian.)
- [20] A.S. Troelstra, and D. van Dalen. *Constructivism in Mathematics*, Volumes 1 and 2. North-Holland, Amsterdam, (1988).
- [21] S. Vickers. *Topology via logic*. New York, Cambridge University Press, (1989).
- [22] M.J. Wooldridge. *An introduction to multiagent systems*. J. Wiley, New York, (2002).

Johan van Benthem
Institute of Logic, Language and Information
University of Amsterdam
The Netherlands

johan@science.uva.nl
and
Department of Philosophy
Stanford University
USA
johan@csl.stanford.edu

Darko Sarenac
Department of Philosophy
Stanford University
USA
sarenac@stanford.edu