



Justified belief, knowledge, and the topology of evidence

Alexandru Baltag¹ · Nick Bezhanishvili¹ · Aybüke Özgün¹  · Sonja Smets¹

Received: 21 January 2022 / Accepted: 4 November 2022 / Published online: 7 December 2022

© The Author(s) 2022

Abstract

We propose a new topological semantics for evidence, evidence-based justifications, belief, and knowledge. Resting on the assumption that an agent's rational belief is based on the available evidence, we try to unveil the concrete relationship between an agent's evidence, belief, and knowledge via a rich formal framework afforded by topologically interpreted modal logics. We prove soundness, completeness, decidability, and the finite model property for the associated logics, and apply this setting to analyze key epistemological issues such as “no false lemma” Gettier examples, misleading defeaters, undefeated justification versus undefeated belief, as well as the defeasibility theories of knowledge.

Keywords Evidence · Justified belief · Knowledge · Epistemic logic · Topological spaces · Completeness · Decidability

1 Introduction

Pioneered by Hintikka (1962), the mainstream approach to epistemic logic is based on the formal ground of relational possible worlds semantics, which provides a relatively simple and flexible way of modeling knowledge and belief. However, this approach is lacking any ingredients to talk about the evidential nature of knowledge or justified belief. One way to correct this is to generalize the standard relational setting to a *topological* one. Indeed, topological spaces emerge naturally as information

Prior work: A shorter version of this paper was published in 2016 in the Proceedings of the 23rd International Workshop on Logic, Language, Information, and Computation (WoLLIC 2016), under the title “Justified Belief and the Topology of Evidence” (Baltag et al., 2016). The current version differs primarily in that it includes the full proofs that were previously omitted (found in Sect. 6), an extended comparison to van Benthem and Pacuit (2011) (Sect. 4), and an extensive discussion on the defeasibility theories of knowledge (Sect. 5). Moreover, both the introduction and conclusion have been substantially revised and extended. (Özgün, 2017, Chapter 5) was developed based on an earlier version of this paper.

✉ Aybüke Özgün
a.ozgun@uva.nl

¹ ILLC, University of Amsterdam, Amsterdam, The Netherlands

structures that can provide a deeper insight into the evidence-based justification of knowledge and belief. For instance, topological notions such as *open*, *closed*, *dense*, and *nowhere dense sets* qualitatively and naturally encode notions such as *measurement/observation*, *closeness*, *smallness*, *largeness*, and *consistency*, all of which will recur with an epistemic interpretation in this work. Moreover, topological spaces come equipped with well-studied basic operators such as the interior and closure operators which—alone or in combination with each other—succinctly interpret different epistemic modalities, providing a better understanding of their axiomatic properties.

In this paper, we propose a topological semantics for various notions of *evidence*, *evidence-based justification*, *belief*, and *knowledge*, and explore the relationship between these epistemic notions. This work builds on the models for evidence, belief, and evidence-management proposed by van Benthem and Pacuit (2011), and van Benthem et al. (2012, 2014), by adopting a *topological* perspective on these notions. The focus is on notions of *belief and knowledge for a rational agent who is in possession of some (possibly false, possibly mutually contradictory) pieces of ‘evidence’*.

A central underlying assumption that we share with van Benthem and Pacuit (2011) is that an agent’s rational beliefs and knowledge are not to be taken as primitive, unjustified concepts, but they are based on, and derived from, a more fundamental notion of ‘evidence’. However, we should stress that this later notion is to be understood here in a wide, inclusive sense: it is not limited to factive evidence, but it may include false or misleading information that nevertheless “looks” to the agent like good enough evidence; also, in addition to acquired evidence (obtained via, e.g., direct observation, measurements, testimony from others, as well as logical inference), this notion includes prior defaults or biases, as well as any type of *a priori* knowledge. As such, we will use the term *basic evidence* to cover all the primitive pieces of (soft, fallible) information available to the subject. On top of these, we will also consider *derived* evidential constructs, that can be built or inferred from the basic ones, and for which we introduce a fine-grained scale of technical terms (*combined evidence*, *arguments* and *justifications*), that will be explained in more detail below. Each of them can be *fallible* or *factive* (true in the actual world), and even when true it can still be *misleading* (in a technical sense, to be formally defined later).

As already mentioned, the notions *justified belief* and *knowledge* we propose are higher-level concepts, definable in terms of the above evidential notions. Before going into details, it may be useful to briefly summarize the relationships between our topological conception and the main positions in Epistemology. First, our setting *is not necessarily committed to an evidentialist epistemology*: while all beliefs and knowledge are justified in terms of the above evidential constructs, we already noted that our ‘evidence’ is mainly a technical term, that may subsume defaults, biases and a priori knowledge. Second, although beliefs and knowledge are derived notions that are justified in terms of evidence, which may suggest some *foundationalist* overtones, our setting *differs in some important respects from the standard foundationalist position*. Our distinction between basic and non-basic ‘evidence’ does *not* lead to a distinction between basic and non-basic beliefs: our “basic” evidence pieces are *not necessarily believed*, indeed it may well happen that *none* of them is believed. Moreover, it may even happen that no (combined) evidence is believed either. Some of our evidential “arguments” (namely, the ones we call “justifications”) will actually be believed, and

other beliefs can be inferred from them. But these evidential justifications are typically not ‘basic’ or primitive in any reasonable sense. Third, we will see that the fundamental feature of our doxastic justifications is their overall consistency with every other available evidence. Our view on justified belief may therefore be said to be essentially *coherentist* in spirit, in that belief is justified if it is entailed by an argument that coheres with the agent’s overall evidential system. On the other hand, unlike in typical coherentism, our theory does not reject the existence of primary, non-inferential forms of evidence such as perceptual evidence or the fact that they can play an important role in our justification system.¹ At the same time, we do not accept such evidence as ‘self-justified’ (or non-inferentially justified) via e.g. perceptual experience: in our setting, only coherence with all other evidence provides doxastic justification. Basic evidence (even perceptual evidence) is not inherently justified, and is not necessarily believed in our framework, unless it coheres with the whole justification system. Fourth, our notion of justified belief fits well with Stalnaker’s view on belief as *subjective certainty* (Stalnaker, 2006): indeed, our notion satisfies Stalnaker’s axiom $B\phi \rightarrow BK\phi$, that equates belief with the “feeling of knowledge”. Fifth, our proposed concept of *knowledge* combines the above-mentioned coherentist view with a strong *reliabilist* flavor: in our setting, knowledge is “correctly justified belief”, where a justification is correct when it doesn’t involve any false evidence or arguments (in addition to not contradicting any other evidence). Such correct justifications provide a reliable process to tracking the truth. This theory of knowledge may at first sound very close to Clark’s “no false lemma” conception (Clark, 1963), but it is subtly different, because our notion of justification is different (requiring coherence with the evidential system). In this sense, our proposal combines features of reliabilism and coherentism. Finally, our topological theory of knowledge can be seen as a sophisticated ‘weakened’ version of *defeasibility theory* (Klein, 1971, 1981; Lehrer, 1990; Lehrer & Paxson, 1969), one that is able to successfully address some of the objections and counterexamples to the defeasibility conception of knowledge, by requiring that the underlying justification remain undefeated by any new non-misleading evidence (though it can be defeated by true but misleading evidence).

It is also important to point out the epistemological issues and conceptions that our proposal does *not* address. Since in this paper we focus on modelling the evidential basis of knowledge and belief, we chose to keep things simple by sticking with the idealized setting of possible-worlds semantics (while only replacing the relational setting with a topological one to deal with evidence). As a consequence, our semantics automatically *enforces closure* of belief and knowledge under logical entailment. In its current form, our topological theory of knowledge is thus incompatible with the epistemological conceptions that deny the closure of knowledge under known entailment², e.g. the sensitivity account (Nozick, 1981), the safety account (Sosa, 1999), the causal accounts (Dretske, 2014, 2016; Goldman, 1967), etc. Another consequence is that our setting runs into the well-known problem of logical omniscience, thus being able to represent only highly idealized reasoners who know/believe all logical and known/believed consequences of what they know/believe. These problems

¹ We thank an anonymous reviewer for bringing this point to our attention.

² We thank an anonymous referee for pointing to us this limitation.

can be fixed. The proposed framework can be easily modified to avoid closure e.g., by requiring belief and knowledge to be exactly one of the evidence pieces (in the spirit of non-monotonic neighbourhood logics, see, e.g., (Chellas, 1980, Chapter 7)) or by employing tools from awareness (Fagin & Halpern, 1987) and topic-sensitive epistemic logics (Berto & Hawke, 2018; Hawke et al., 2020; Özgün & Berto, 2021). See, e.g., Siemers (2021) for a topic-sensitive, hyperintensional version of our proposal where only restricted closure principles for evidence, knowledge, and belief hold. Such a modified variant of our setting can successfully deal with logical non-omniscience, as well as with the genuine cases of non-closure.³ However, all known solutions dilute the simplicity and the extensional-semantic nature of our current topological approach by adding in-built hyperintensional features that come with their own complications. Since in many contexts closure under known entailment does not pose any problems, we choose to present here only the purely semantic core of our proposal, in order to avoid unnecessary complications and to better convey the essence of our topological theory of knowledge in a transparent and simple manner.

1.1 Our proposal in more detail

We will now provide a more detailed overview of the epistemic notions studied in this paper, introduce the modalities we consider, and explain where our work stands in the relevant literature.

As already mentioned, we adopt a possible-worlds semantics, but replace the standard relational setting with a topological one. The *basic pieces of evidence* possessed by an agent are represented simply as nonempty sets of possible worlds. Our topological evidence models will thus come with a designated family of such sets. A *combined evidence* (or just *evidence*, for short) is any nonempty intersection of finitely many pieces of evidence. Note that this notion of evidence is not necessarily factive⁴, since

³ It seems to us that all such genuine cases of non-closure (in which one is really not warranted to believe/know some consequence of a current belief) involve some subtle shift in topic or context from the premise to the conclusion. E.g. Nozick's example of non-closure involves the shift between a day-to-day context or topic in "I have hands" to a wider, more inquisitive or 'philosophical', topic in "I am not a brain in a vat". By making this topic-sensitivity explicit, the hyperintensional version in Siemers (2021) can maintain the closure of belief-knowledge *as long as the same topic is maintained*.

In other, less genuine cases of non-closure (such as the famous Red Barn example), it seems to us that the culprit is a "purist" assumption of a single source of justification for knowledge (be it perceptual evidence, sensitivity, safety, etc). The problem disappears if one admits that logical inference is itself an independent source of knowledge, on a par with the others, and that a full justification of our beliefs may involve a mixture of such sources. E.g., I know that there is a red barn in front of me, because I see a red barn (and my perceptual experience is sensitive to the truth of its claim, since in my country all red barns are genuine); this evidential warrant is not transmitted to the statement "there is a barn in front of me" (since there exist fake barns, so my belief in barns does not track the truth); however, I do know that there is a barn, by a mixed justification that involves both sensitivity and deduction (e.g. seeing a red barn and inferring the existence of a barn). So in these cases, it is intuitive to maintain closure, while giving up the "purist" request for a single source of justification. And indeed, our approach combines evidence and logical inference as sources of knowledge and belief.

⁴ Factive evidence is true in the actual world. In epistemology it is common to reserve the term "evidence" for factive evidence. But we follow here the more liberal usage of this term in (van Benthem & Pacuit, 2011), which agrees with the common understanding in day to day life, e.g. when talking about "uncertain evidence", "fake evidence", "misleading evidence" etc.

Table 1 Evidence modalities and their intended readings

$E_0\varphi$	The agent has a basic (piece of) evidence for φ
$E\varphi$	The agent has a (combined) evidence for φ
$\Box_0\varphi$	The agent has a factive basic (piece of) evidence for φ
$\Box\varphi$	The agent has factive (combined) evidence for φ

the pieces of evidence are possibly false and, moreover, possibly inconsistent with each other. The family of (combined) evidence sets forms a topological basis that generates what we call the *evidential topology*. This is the smallest topology in which all the basic pieces of evidence are open, and it will play an important role in our setting.

For some of these evidential notions, we consider the associated modal operators, e.g. “having (a piece of) basic evidence for a proposition P ” [operator already proposed by van Benthem and Pacuit (2011)], “having (combined) evidence for P ”, “having a (piece of) *factive* evidence for P ” and “having (combined) factive evidence for P ”. Table 1 lists the corresponding evidence modalities together with their intended readings.⁵

In fact, the modality $\Box\varphi$, capturing the concept of “having factive evidence for φ ”, coincides with the topological *interior operator* in the evidential topology. We therefore use the interior semantics of McKinsey and Tarski (1944) to interpret a notion of *factive* evidence. We also show that the two factive variants of evidence-possession operators (\Box_0 and \Box) are more expressive than the non-factive ones (E_0 and E): when interacting with the global modality, the two factive evidence modalities $\Box_0\varphi$ and $\Box\varphi$ can define the non-factive variants $E_0\varphi$ and $E\varphi$, respectively, as well as many other doxastic/epistemic operators (as shown in Proposition 6).

Our semantics for *justification* and *justified belief* is obtained by extending, generalizing, and, to an extent, streamlining the evidence-model framework for belief introduced by van Benthem and Pacuit (2011). The main idea of that setting was that the rational agent tries to form consistent beliefs, by looking at all *strongest finitely-consistent collections of evidence*, and she believes whatever is entailed by all of them.⁶ The consistency of that notion of belief crucially depends on the existence of some such “strongest” evidence, which is of course granted in the finite case (whenver the agent has finitely many pieces of basic evidence) as well as in *some* infinite cases, but it can fail in other cases. As a result, as already noted in (van Benthem et al., 2014), this can lead to *inconsistent beliefs* in the general infinite case, contrary to the spirit of the original proposal.⁷

⁵ The Greek letter φ should be taken as a metavariable for a well-defined sentence in the associated modal language. For the purposes of this introductory section, we need only the intended readings of the listed modal epistemic operators. The recursive definitions of the modal languages employed are given in Sect. 6.

⁶ To be sure, this is still vague since we have not yet specified what a “strongest finitely-consistent collections of evidence” means (we return to formalize these notions in Sect. 3.1.1), however, this much precision should be sufficient to explain the rough idea behind the definition of belief in (van Benthem & Pacuit, 2011) and the notion of justified belief we propose in this paper.

⁷ Another, purely technical drawback of the setting in (van Benthem & Pacuit, 2011) is that the corresponding doxastic logic does not have the finite model property (see van Benthem et al., 2012, Corollary 2.7 or van Benthem et al., 2014, Corollary 1).

Table 2 Belief modalities and their intended readings

$B\varphi$	The agent has justified belief in φ
$B^\varphi\psi$	The agent believes that ψ conditionally on φ

One way to obtain our semantics for evidence-based belief is by, in a sense, weakening the definition from (van Benthem & Pacuit, 2011). According to this revised definition, a proposition P is believed if P is entailed by all *sufficiently strong finitely-consistent collections of evidence*. This notion of belief is equivalent to the one of van Benthem and Pacuit (2011) when the collection of basic pieces of evidence is finite, but the two diverge in the infinite case. Indeed, our semantics *always ensures consistency of belief*, even when the available pieces of evidence are mutually inconsistent, thus fulfilling the project of rationally grounding consistent beliefs on (possibly) inconsistent collections of evidence.

Moreover, our revised definition throws a new light on this notion of belief (even in the case when it is equivalent to the older notion), by connecting it to *topology* and to a notion of *justification*. First, this concept of belief is very natural from a topological perspective: the revised definition is equivalent to saying that P is believed iff it is true in “almost all” epistemically possible states, where “almost all” is interpreted topologically as “all except for a nowhere-dense set”. Second, in order to analyze justified belief, we need some additional evidential notions. An *argument* consists of one or more (combined) evidence sets supporting the same proposition P : in essence, it is a way to provide one or more evidential paths towards a (common) conclusion. A *justification* is an argument that is not contradicted by any other available (combined) evidence; equivalently, a justification is an argument that is not defeated by any other argument (based on the same body of evidence). This is the promised ‘coherentist’ notion of doxastic justification, requiring consistency with all the pieces of evidence possessed by the agent. Our revised definition turns out to be equivalent to requiring that P is believed iff there is some evidence-based justification for P . In this sense, our belief is an *evidentially-justified belief*.

This topological definition of belief can be easily generalized to *conditional* beliefs. Table 2 lists the belief modalities we study in this paper.

Moving on to *knowledge*, there are a number of different notions one may consider. First, there is the ‘infallible’ knowledge, absolutely certain and absolutely indefeasible, akin to van Benthem’s concept of hard information (van Benthem, 2007). This is the standard concept of knowledge used in Computer Science and Game Theory applications, and formalized within the modal epistemic logic S5, based on Kripke models endowed with equivalence relations [(or equivalently, on Aumann’s partitioned models (Aumann, 1999)]. In our single-agent setting, this can be simply defined as the global modality (quantifying universally over all epistemically possible states). For good reasons, most epistemologists do not take this to be a good formalization of our intuitive sense of knowledge. There are very few propositions that can be ‘known’ in this infallible way (apart from logical tautologies, or maybe also things known by introspection). Most facts in science or real life are unknown in this sense. It is therefore more interesting to look at notions of knowledge that are less-than-absolutely-certain, so-called *defeasible knowledge*. As shown by the famous Gettier counterexamples

Table 3 Knowledge modalities and their intended readings

$[\forall]\varphi$	The agent infallibly knows that φ
$K\varphi$	The agent fallibly (or defeasibly) knows that φ

(Gettier, 1963), simply adding factivity to justified belief cannot yield knowledge. True justified belief may be extremely fragile (i.e., it can be too easily lost), and it is consistent with having ‘wrong’ (unreliable) justifications for an accidentally true conclusion.

The notion of defeasible knowledge we propose in this paper is formally defined by saying that “*P* is (fallibly) known iff there is a factive justification for *P*”. Knowledge in this conception is *correctly justified belief*, but with the proviso that the ‘justification’ is taken in the above-mentioned holistic sense (requiring it to be, not only evidence-based, but coherent with every other available evidence). As shown in Sect. 5, this concept of knowledge finds its place in the post-Gettier literature as being stronger than the one characterized by the “no false lemma” of Clark (1963) and weaker than the one described by the defeasibility theory of knowledge championed by Klein (1971, 1981), Lehrer (1990), and Lehrer and Paxson (1969). In our framework, we consider modal operators for both infallible knowledge and defeasible knowledge, but our main focus will be on the latter notion. See Table 3 for the corresponding knowledge modalities and their readings.

Yet another path leading to our proposal in this paper goes via our earlier work (Baltag et al., 2013, 2019b; Özgün, 2013) on a topological semantics for the doxastic-epistemic axioms of Stalnaker (2006). These axioms are meant to capture a notion of fallible knowledge, in close interaction with a notion of *strong belief* defined as *subjective certainty*. The main principle specific to this system was that “believing implies believing that you know”, captured by the axiom $B\varphi \rightarrow BK\varphi$. The topological semantics that was proposed for these concepts in (Baltag et al., 2013, 2019b; Özgün, 2013) was overly restrictive, being limited to the rather unfamiliar class of extremally disconnected and hereditarily extremally disconnected topologies. In the current work, we show that these notions can be interpreted on arbitrary topological spaces without changing their logic. To that end, our definitions of belief and knowledge can be seen as the natural generalizations of the notions in (Baltag et al., 2013, 2019b; Özgün, 2013) to arbitrary topologies.

1.2 Overview of this paper

Section 2 introduces the required topological preliminaries. In Sect. 3, we introduce the evidence models of van Benthem and Pacuit (2011) as well as our *topological* evidence models, and provide semantics for the notions of basic, combined, and factive evidence. We moreover present topological definitions for argument and justification.

In Sect. 4, we introduce our topological semantics for (justified) belief, while comparing our setting to that of van Benthem and Pacuit (2011). We then generalize our semantics of (plain) belief to conditional beliefs.

In Sect. 5, we propose our topological formalization of fallible knowledge, and use it to analyze various issues in the post-Gettier epistemology literature, such as “no false lemma” Gettier examples, stability/defeasibility theories of knowledge, objections based on misleading vs. genuine defeaters, undefeated justification versus undefeated belief, the epistemic role of belief dynamics, etc.

Finally, Sect. 6 presents all our technical results. We completely axiomatize the various resulting logics of evidence, knowledge, and belief, and prove decidability and finite model property results. Our technically most challenging result is the completeness of the richest logic containing the two factive evidence modalities $\Box_0\varphi$ and $\Box\varphi$, as well as the global modality $[\forall]\varphi$. This logic can define all the modal operators mentioned above. While the other proofs are more or less routine, the proof of this result involves a nontrivial combination of known methods (Sect. 6.5).

The paper is organized in such a way that the reader who is interested only in the conceptual contributions can read up to Sect. 6.

2 Topological preliminaries

In this section, we introduce the topological concepts that will be used throughout the paper. We refer to (Dugundji, 1965; Engelking, 1989) for a thorough introduction to topology. The reader who has introductory level knowledge of topology should feel free to skip this section.

Definition 1 (*Topological space*) A *topological space* is a pair (X, τ) , where X is a nonempty set and τ is a family of subsets of X such that $X, \emptyset \in \tau$, and τ is closed under finite intersections and arbitrary unions.

The set X is a *space*; the family τ is called a *topology* on X . The elements of τ are called *open sets* (or *opens*) in the space. If for some $x \in X$ and an open $U \subseteq X$ we have $x \in U$, we say that U is an *open neighborhood* of x . A set $C \subseteq X$ is called a *closed set* if it is the complement of an open set, i.e., it is of the form $X \setminus U$ for some $U \in \tau$. We let $\bar{\tau} = \{X \setminus U \mid U \in \tau\}$ denote the family of all closed sets of (X, τ) .

A point x is called an *interior point* of a set $A \subseteq X$ if there is an open neighbourhood U of x such that $U \subseteq A$. The set of all interior points of A is called the *interior* of A and is denoted by $Int(A)$. Then, for any $A \subseteq X$, $Int(A)$ is an open set and is indeed the largest open subset of A , that is

$$Int(A) = \bigcup \{U \in \tau \mid U \subseteq A\}.$$

Dually, for any $x \in X$, x belongs to the *closure* of A , denoted by $Cl(A)$, if and only if $U \cap A \neq \emptyset$ for each open neighborhood U of x . It is not hard to see that $Cl(A)$ is the smallest closed set containing A , that is

$$Cl(A) = \bigcap \{C \in \bar{\tau} \mid A \subseteq C\},$$

and that $Cl(A) = X \setminus Int(X \setminus A)$ for all $A \subseteq X$. It is well known that the interior Int and the closure Cl operators of a topological space (X, τ) satisfy the following

properties (the so-called Kuratowski axioms) for any $A, B \subseteq X$ (see, e.g., Engelking, 1989, pp. 14-15):

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

A set $A \subseteq X$ is called *dense* in X if $Cl(A) = X$ and it is called *nowhere dense* if $Int(Cl(A)) = \emptyset$. More generally, for any $A, B \subseteq X$, A is called *dense in B* if $B \subseteq Cl(A \cap B)$.

Definition 2 (Topological basis) A family $\mathcal{B} \subseteq \tau$ is called a *basis* for a topological space (X, τ) if every non-empty open subset of X can be written as a union of elements of \mathcal{B} .

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

Given any family $\Sigma = \{A_\alpha \mid \alpha \in I\}$ of subsets of X , there exists a unique, smallest topology $\tau(\Sigma)$ with $\Sigma \subseteq \tau(\Sigma)$ (Dugundji, 1965, Theorem 3.1, p. 65). The family $\tau(\Sigma)$ consists of \emptyset, X , all finite intersections of the A_α , and all arbitrary unions of these finite intersections. Σ is called a *subbasis* for $\tau(\Sigma)$, and $\tau(\Sigma)$ is said to be *generated* by Σ . The set of finite intersections of members of Σ forms a basis for $\tau(\Sigma)$.

Lemma 1 For any two topological space (X, τ) and (X, τ') and $A \subseteq X$, if $\tau \subseteq \tau'$ then $Int_\tau(A) \subseteq Int_{\tau'}(A)$, where Int_τ and $Int_{\tau'}$ are the interior operators of τ and τ' , respectively.

3 Evidence, argument, and justification

In this section, we present the (uniform) evidence models of van Benthem and Pacuit (2011) as well as our *topological* version, and provide the formal semantics of the evidence modalities given in Table 1. In this topological framework, we introduce and study the technical notions of *combined evidence, strongest evidence, strong enough evidence*, (evidence-based) *argument* and *justification*.

3.1 Evidence à la van Benthem and Pacuit

Definition 3 (Evidence models) An *evidence model* is a tuple $\mathfrak{M} = (X, \mathcal{E}_0, V)$, where

- X is a nonempty set of *possible worlds* (or *states*),
- $\mathcal{E}_0 \subseteq \mathcal{P}(X)$ is a family of sets called *basic evidence sets* (or *pieces of evidence*), satisfying $X \in \mathcal{E}_0$ and $\emptyset \notin \mathcal{E}_0$, and

- $V : \text{Prop} \rightarrow \mathcal{P}(X)$ is a valuation function.⁸

The evidence models presented in (van Benthem & Pacuit, 2011; van Benthem et al., 2014) are more general, covering cases in which evidence depends on the actual world, i.e., in which each state may be assigned a different set of neighbourhoods. In this paper, however, we stick with the special case of *uniform models* (given in Definition 3), which corresponds to working with agents who are evidence-introspective (more on this below). Since we never consider the more general case and focus only on the topological extension of their uniform evidence models, we simply use the term *evidence model* exclusively for the uniform evidence models.

Note that evidence models are not necessarily based on topological spaces, i.e., \mathcal{E}_0 is not defined to be a topology (it may not even constitute a topological basis). However, every topology τ constitutes a basic evidence set.⁹ In fact, the family \mathcal{E}_0 is almost an arbitrary nonempty collection of subsets of a given domain, carefully designed to capture certain aspects of the type of evidence that is intended to be formalized. First of all, the subset \mathcal{E}_0 represents the set of evidence the agent has acquired about the actual situation¹⁰ directly via, e.g., testimony, measurement, approximation, computation, or experiment. It is the collection of evidence the agent has gathered so far, and it is all our rational, idealized agent has to form her beliefs and knowledge. The collection of evidence the agent possesses is uniform across the states, i.e., the set of evidence the agent has does not depend on the actual state. This corresponds to working with an evidence-introspective agent, that is, the agent is absolutely sure about what evidence she has and what it does and does not entail.

The two properties of \mathcal{E}_0 , namely, $X \in \mathcal{E}_0$ and $\emptyset \notin \mathcal{E}_0$ impose the following constraints, respectively:

- Tautologies are always evidence, and
- Contradictions never constitute direct evidence.

Unlike the common practice in epistemology, where the term “evidence” is generally reserved for factive evidence, we follow here the more liberal use of the term adopted by van Benthem and Pacuit, that includes *fallible* information coming from a possibly unreliable source: *a piece of evidence in \mathcal{E}_0 does not have to contain the actual state*. This more realistic view on evidence agrees with the common usage in day to day life, e.g. when talking about “uncertain evidence”, “fake evidence”, “misleading evidence”. Moreover, in this setting the pieces of ‘evidence’ may be mutually inconsistent: *the intersection of evidence pieces may be empty*. Indeed, the agent might be collecting evidence from different sources that may or may not be reliable. However, no quantitative measure of reliability or qualitative reliability order is assumed to be given on the elements of \mathcal{E}_0 . Under these assumptions, a rational agent will have to take into account (though not necessarily believe) *every* piece of available evidence,

⁸ Prop is a countable set of propositional variables from which we will recursively define the epistemic languages of interest.

⁹ As an even more special case, we can think of Grove/Lewis Sphere spaces (Grove, 1988; Lewis, 1973). These are topological spaces in which the open sets are nested, i.e. for every $U, U' \in \tau$, we have either $U \subseteq U'$ or $U' \subseteq U$ (see, e.g., Example 1).

¹⁰ Standardly, as in the relational semantics, the actual situation is represented by a state x of X called the *actual state* or the *real world*.

and somehow put these pieces together in a finite and consistent manner. This leads us to the notions of *combined evidence* and *body of evidence*, concepts that will play a crucial role in the formation of consistent beliefs based on fallible evidence.

3.1.1 Bodies of evidence, evidential support, and evidential strength

We call a collection of evidence pieces $F \subseteq \mathcal{E}_0$ *consistent* if $\bigcap F \neq \emptyset$, and *inconsistent* otherwise. To state our definitions, we use the notation $A \subseteq_{fin} B$ to say that A is a *finite subset of B*.

Definition 4 (*(Finite) body of evidence*) Given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$, a *body of evidence* is a nonempty family $F \subseteq \mathcal{E}_0$ of evidence pieces such that every nonempty finite subfamily is consistent. More formally, a nonempty family $F \subseteq \mathcal{E}_0$ is a *body of evidence* if

$$(\forall F' \subseteq_{fin} F)(F' \neq \emptyset \text{ implies } \bigcap F' \neq \emptyset).$$

A *finite body of evidence* $F \subseteq_{fin} \mathcal{E}_0$ is therefore simply a finite set of mutually consistent pieces of evidence, that is, $F \subseteq_{fin} \mathcal{E}_0$ such that $\bigcap F \neq \emptyset$.

Therefore, a body of evidence is simply a collection of evidence pieces that has the finite intersection property, and that represents the agent’s ability of putting evidence pieces together in a *finitely consistent* way.

Given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$, we denote by

$$\mathcal{F} := \{F \subseteq \mathcal{E}_0 \mid (\forall F' \subseteq_{fin} F)(F' \neq \emptyset \text{ implies } \bigcap F' \neq \emptyset)\}$$

the *family of all bodies of evidence* over \mathfrak{M} , and by

$$\mathcal{F}^{fin} := \{F \subseteq_{fin} \mathcal{E}_0 \mid \bigcap F \neq \emptyset\}$$

the *family of all finite bodies of evidence*. Both the interpretation of evidence-based belief of van Benthem and Pacuit (2011) and our proposal for justified belief, as well as the notion of defeasible knowledge we study in a later section crucially rely on the notion of body of evidence. But, in order to be able to talk about these *evidence-based* informational attitudes, we first need to specify what it means for a proposition to be *supported* by a body of evidence.

Remark 1 Throughout Sects. 3–5, we use the following conventions to ease the presentation. Given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$ (or, a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ defined later), we call any subset $P \subseteq X$ a *proposition*. We say a proposition $P \subseteq X$ is true at x if $x \in P$. The Boolean connectives, $\neg, \wedge, \vee, \rightarrow$, on propositions are defined standardly as set operations: for any $P, Q \subseteq X$, we define $\neg P := X \setminus P, P \wedge Q := P \cap Q, P \vee Q := P \cup Q$ and $P \rightarrow Q := (X \setminus P) \cup Q$. Moreover, the Boolean constants \top and \perp are given as $\top := X$ and $\perp := \emptyset$. Following this convention, we define the semantics of the modal operators for evidence, belief, and

knowledge introduced in Tables 1, 2, and 3 as set operators from $\mathcal{P}(X)$ to $\mathcal{P}(X)$ (and for the binary modality of conditional belief, from $\mathcal{P}(X) \times \mathcal{P}(X)$ to $\mathcal{P}(X)$). These set operators give rise to the interpretations of the corresponding modalities of the full language \mathcal{L} (given in Sect. 6) in a standard way.

Definition 5 (*Evidential support*) Given an evidence model $\mathfrak{M}=(X, \mathcal{E}_0, V)$ and a proposition $P \subseteq X$, a body of evidence F supports P if P is true in every state satisfying all the evidence in F , i.e., if $\bigcap F \subseteq P$.

It is easy to see that a body of evidence F is inconsistent iff it supports every proposition (since $\emptyset \subseteq P$, for all P). The *strength order* between bodies of evidence is given by inclusion: $F \subseteq F'$ means that F' is at least as strong as F . Note that stronger bodies of evidence support more propositions: if $F \subseteq F'$ then every proposition supported by F is also supported by F' . A body of evidence is *maximal* (or, *strongest*) if it is a maximal element of the poset (\mathcal{F}, \subseteq) , i.e., if it is not a proper subset of any other such body. We denote by

$$Max_{\subseteq}\mathcal{F}:=\{F \in \mathcal{F} \mid (\forall F' \in \mathcal{F})(F \subseteq F' \Rightarrow F = F')\}$$

the *family of all maximal bodies of evidence* of a given evidence model. By Zorn's Lemma, every body of evidence can be strengthened to a maximal body of evidence, i.e.,

$$\forall F \in \mathcal{F} \exists F' \in Max_{\subseteq}\mathcal{F} (F \subseteq F').$$

Therefore, in particular, every evidence model has at least one maximal body of evidence, that is, $Max_{\subseteq}\mathcal{F} \neq \emptyset$.

In fact, for *finite* bodies of evidence, the notions of evidential support and strength can be represented in a more concise way via the notion of combined evidence, which, to anticipate further, is represented by basic open sets of the evidential topology generated from \mathcal{E}_0 (see Sect. 3.2).

3.1.2 Combined evidence and evidential basis

Definition 6 (*Combined evidence*) Given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$, a *combined evidence* (or just *evidence*, for short) is any nonempty intersection of finitely many basic evidence pieces. In other words, a nonempty subset $e \subseteq X$ is a combined evidence if $e = \bigcap F$, for some $F \in \mathcal{F}^{fin}$.

A combined evidence therefore is just a repackaging of a finite body of evidence in terms of its intersection. We denote by

$$\mathcal{E}:=\{\bigcap F \mid F \in \mathcal{F}^{fin}\}$$

the *family of all (combined) evidence*, which in fact constitutes a topological basis over X . We will return to the topological versions of evidence models in Sect. 3.2.

The definitions *evidential support* and *strength* are adapted for the elements of \mathcal{E} in an obvious way. A (combined) evidence $e \in \mathcal{E}$ *supports* a proposition $P \subseteq X$ if $e \subseteq P$. In this case, we also say that e is *evidence for* P . The natural strength order between combined evidence sets therefore is given by the reverse inclusion: $e \supseteq e'$ means that e' is at least as strong as e . This is both to fit with the strength order on bodies of evidence (since $F \subseteq F'$ implies $\bigcap F \supseteq \bigcap F'$), and to ensure that stronger evidence supports more propositions (since, if $e \supseteq e'$, then every proposition supported by e is supported by e').

Recall that \mathcal{E}_0 represents the collection of evidence pieces that are directly observed by the agent. The elements of the derived set \mathcal{E} therefore serve as indirect evidence which is obtained by combining finitely many pieces of direct evidence together in a consistent way. This does not mean that all of this evidence is necessarily true. We say that some (basic or combined) evidence $e \in \mathcal{E}$ is *factive evidence* at state $x \in X$ whenever it is true at x , i.e., if $x \in e$. Similarly, a body of evidence F is *factive* if all the pieces of evidence in F are *factive*, i.e., if $x \in \bigcap F$.

Having presented the primary semantic concepts used in the representation of (basic and combined) evidence, we proceed with our topological setting.

3.2 Evidence in topological evidence models

For any nonempty set X and any family Σ of subsets of X , we can construct a topology on this domain by simply closing Σ under finite intersections and arbitrary unions (recall the definition of subbasis given in Sect. 2). Therefore, every evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$ can be associated with an *evidential topology* that is generated by the set of basic evidence pieces \mathcal{E}_0 , or equivalently, by the family of all combined evidence \mathcal{E} . In this section, we introduce our topological evidence models, and provide topological formalizations of our notions of *argument* and *justification*. We moreover give the semantics for the modalities $E_0\varphi$ and $E\varphi$ denoting possession of basic and combined evidence, respectively, as well as for their *factive* versions $\Box_0\varphi$ and $\Box\varphi$.

Definition 7 (*Topological evidence model*) A *topological evidence model* (or, in short, a *topo-e-model*) is a tuple $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, where (X, \mathcal{E}_0, V) is an evidence model and $\tau = \tau_{\mathcal{E}}$ is the topology generated by the family of combined evidence \mathcal{E} (or equivalently, by the family of basic evidence sets \mathcal{E}_0), which is called the *evidential topology*.

The families \mathcal{E}_0 and \mathcal{E} obviously generate the same topology: \mathcal{E} is the closure of \mathcal{E}_0 under nonempty finite intersections. We denote the evidential topology by $\tau_{\mathcal{E}}$ only because the family \mathcal{E} of combined evidence forms a *basis* of this topology (and we omit the subscript \mathcal{E} when it is contextually clear). Since any family $\mathcal{E}_0 \subseteq \mathcal{P}(X)$ generates a topology over X , topo-e-models are just another way to present the evidence models described in Definition 3. We use this special terminology to stress our focus on the induced topological structures, and to avoid ambiguities, since our definition of belief in topo-e-models will be different from the notion of belief in evidence models defined in (van Benthem & Pacuit, 2011).




Fig. 1 From \mathcal{E}_0 to $\tau\mathcal{E}$; from direct evidence to argument

Arguments Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ and a proposition $P \subseteq X$, an *argument for P* is a union $U = \bigcup \mathcal{E}'$ of some nonempty family of (combined) evidence $\mathcal{E}' \subseteq \mathcal{E}$, each separately supporting P (i.e., $e \subseteq P$ for all $e \in \mathcal{E}'$, or equivalently, $U \subseteq P$). Epistemologically, an argument for P provides *multiple evidential paths* $e \in \mathcal{E}'$ to support the common conclusion P . Topologically, an argument for P is the same as a *nonempty open subset of P*: a set of states U is an argument for P iff $U \in \tau$ and $U \subseteq P$. Therefore, the open $Int(P)$ forms the *weakest (most general) argument for P*, since it is the largest open subset of P . (See Fig. 1 for the construction of $\tau\mathcal{E}$ from \mathcal{E}_0 and the notions corresponding to their elements.)

Justifications A *justification for P* is an argument U for P that is consistent with every (combined) evidence (i.e., $U \cap e \neq \emptyset$ for all $e \in \mathcal{E}$, that is, $U \cap U' \neq \emptyset$ for all $U' \in \tau \setminus \{\emptyset\}$). Justifications are thus defined to be arguments that are *undefeated* (i.e., whose negations are not supported) by any available evidence or any other argument based on the evidence. Topologically, a justification for P is just a *dense open subset of P*: a set of states U is a justification for P iff $U \in \tau$ such that $U \subseteq P$ and $Cl(U) = X$. As for evidence, an argument or a justification U for P is said to be *factive* (or “correct”) if it is true in the actual world x , i.e., if $x \in U$.

The fact that arguments are open in the generated topology encodes the principle that *any argument should be evidence-based*: whenever an argument is correct, then it is supported by some factive evidence. To anticipate further: in our setting, justifications will form the basis of *belief*, while correct justifications will form the basis of *fallible (defeasible) knowledge*. But before moving to justified belief and fallible knowledge, we introduce a stronger, irrevocable form of knowledge that is captured by the global modality.

Infallible knowledge: possessing hard information We use $[\forall]$ for the so-called *global* modality, which associates to every proposition $P \subseteq X$, some other proposition $[\forall]P$, given by putting:

$$[\forall]P = \begin{cases} X & \text{if } P = X \\ \emptyset & \text{otherwise.} \end{cases}$$

In other words, $[\forall]P$ is true (at any state) iff P is true at all states. In this setting, $[\forall]P$ is interpreted as “absolutely certain, *infallible knowledge*”, defined as truth in all the worlds that are consistent with the agent’s information.¹¹ This is a limit notion

¹¹ In a multi-agent model, some worlds might be consistent with one agent’s information, while being ruled out by another agent’s information. Therefore, in a multi-agent setting, $[\forall_i]$ will only quantify over

capturing a very strong form of knowledge encompassing all epistemic possibilities. It is *irrevocable*, i.e., it cannot be lost or weakened by any information gathered later. In this respect, $[\forall]P$ could be best described as *possession of hard information*. Its dual $[\exists]P := \neg[\forall]\neg P$ expresses the fact that P is consistent with (all) the agent’s hard information.¹²

The notion of infallible knowledge $[\forall]\varphi$ is not very widely applicable, and the thesis that all knowledge is infallible has been harshly criticized by many epistemologists (see, e.g., Hintikka, 1962). However, having the global modality as an operator in our framework is useful for both conceptual and technical reasons: while it formalizes the intuitive notion of hard evidence, and it distinguishes it from “softer” types of information such as fallible knowledge, the global modality adds to the expressive power of modal languages. In particular, when combined with the evidential modalities $\Box_0\varphi$ and $\Box\varphi$ introduced below, it will allow us to define as abbreviations all the other epistemic and doxastic operators considered in this paper (see Proposition 6).

Having basic evidence for a proposition For every proposition $P \subseteq X$, we can define another proposition E_0P by putting:¹³

$$E_0P = \begin{cases} X & \text{if } \exists e \in \mathcal{E}_0 (e \subseteq P) \\ \emptyset & \text{otherwise.} \end{cases}$$

The modal sentence E_0P therefore captures possession of basic (direct) evidence for the proposition P , thus reads as “the agent *has basic evidence for P*”. In other words, E_0P states that P is supported by some basic piece of evidence. Additionally, we introduce a *factive* version of this proposition, \Box_0P , that is read as “the agent has *factive basic evidence for P*”, and is given by

$$\Box_0P = \{x \in X \mid \exists e \in \mathcal{E}_0 (x \in e \text{ and } e \subseteq P)\}.$$

Having (combined) evidence for a proposition The above notions of evidence possession based on having basic evidence for a propositions can be generalized to having (combined) evidence for a proposition. This way, we obtain two other evidence operators: EP , meaning that “the agent *has (combined) evidence for P*”, and $\Box P$, meaning that “the agent *has factive (combined) evidence for P*”. More precisely, EP and $\Box P$ are given as follows:

$$EP = \begin{cases} X & \text{if } \exists e \in \mathcal{E} (e \subseteq P) \\ \emptyset & \text{otherwise} \end{cases}$$

all the states in agent i ’s current information cell (according to a partition Π_i of the state space reflecting agent i ’s hard information).

¹² We ask the reader not to confuse \forall and $[\forall]$: while we use the former to abbreviate “for all” in the metalanguage, the latter is the global modality operating on propositions. Similarly for \exists and $[\exists]$: the former abbreviates “there exists” in the metalanguage and the latter is the existential modality defined as $\neg[\forall]\neg$.

¹³ Van Benthem and Pacuit (2011) denote this by $\Box P$ and it is denoted by $[E]P$ in (van Benthem et al., 2014). We use E_0P for this notion, since we reserve the notation EP for having *combined* evidence for P , and $\Box P$ for having *combined factive* evidence for P .

$$\Box P = \{x \in X \mid \exists e \in \mathcal{E} (x \in e \text{ and } e \subseteq P)\}.$$

Since \mathcal{E} is a basis of the evidential topology $\tau_{\mathcal{E}}$, we have that *the agent has evidence for a proposition P iff she has an argument for P* . So $\Box P$ can also be interpreted as “having an argument for P ”. Similarly, $\Box P$ can be interpreted as “having a *correct* argument for P ”. Moreover, \Box operator for having combined factive evidence coincides with the topological interior operator:

$$Int(P) = \Box P,$$

where Int is the interior operator of the evidential topology $\tau_{\mathcal{E}}$.

4 Justified belief

In this section, we propose a topological semantics for a notion of evidence-based justified belief. One way to do this is by modifying the belief definition proposed by van Benthem and Pacuit (2011) based on evidence models, so we start by recapitulate their proposal. While the two definitions are equivalent on evidence models carrying a *finite* collection of evidence pieces \mathcal{E}_0 , our notion is better behaved in general, since it is *always consistent*, and in fact it *satisfies the axioms of the standard doxastic logic KD45 on all topo-e-models*. We then provide several equivalent characterizations of the proposed notion of belief, in particular one in terms of evidential justification and others in purely topological terms. We also generalize this setting to *conditional* beliefs.

4.1 Belief à la van Benthem and Pacuit

Given an evidence model, van Benthem and Pacuit (2011) define belief by putting, for any proposition P :

P is believed iff every maximal (i.e., strongest) body of evidence supports P .

We denote this notion by Bel ¹⁴. More formally, given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$ and a proposition $P \subseteq X$,

$$Bel P \text{ holds (at any state) iff } (\forall F \in Max_{\subseteq \mathcal{F}})(\bigcap F \subseteq P).$$

¹⁴ In the finite case and many other (but not all) cases, this definition is equivalent to treating plausibility models as a special case of evidence models where the plausibility relation is given by the *evidential* plausibility order $\sqsubseteq_{\mathcal{E}}$ defined as

$$x \sqsubseteq_{\mathcal{E}} y \text{ iff } (\forall e \in \mathcal{E}_0)(x \in e \text{ implies } y \in e) \text{ (equivalently, } (\forall e \in \mathcal{E})(x \in e \text{ implies } y \in e)),$$

and applying the standard semantics of belief on plausibility models as “truth in all the most plausible states”. The relation between evidence models and plausibility models, as well as the connection between the notions of belief defined on these structures are subtle. We refer to (van Benthem and Pacuit, 2011, Sect. 5) and (van Benthem et al., 2014, Sect. 3) for details.




Fig. 2 $\mathfrak{M} = (\mathbb{N}, \mathcal{E}_0, V)$

However, as can be seen directly from the above definition, Bel is inconsistent on evidence models whose every maximal body of evidence is inconsistent.

Example 1 Consider the evidence model $\mathfrak{M} = (\mathbb{N}, \mathcal{E}_0, V)$, where the state space is the set \mathbb{N} of natural numbers, $V(p) = \emptyset$, and the basic evidence family is $\mathcal{E}_0 = \{[n, \infty) \mid n \in \mathbb{N}\}$ (see Fig. 2). The only maximal body of evidence in \mathcal{E}_0 is \mathcal{E}_0 itself. However, $\bigcap \mathcal{E}_0 = \emptyset$. So $Bel \perp$ holds in \mathfrak{M} .

This phenomenon happens only in (some cases of) *infinite* models, so it is *not* due to the inherent mutual inconsistency of the available evidence. At a high level, the source of the problem seems to be the tension between the way the agent combines her evidence pieces and the way she forms her beliefs based her evidence: while she puts her evidence pieces together in a *finitely* consistent way, having consistent beliefs requires possibly infinite collections to have nonempty intersections. More precisely, even though it is guaranteed by definition that every finite subfamily of a maximal body of evidence is consistent, the whole maximal body of evidence may actually be inconsistent. Therefore, in order to avoid this problem, we could instead focus on *maximal finite bodies* of evidence as blocks of evidence forming beliefs: these are, by definition, guaranteed to be always consistent. However, this solution inevitably restricts the class of evidence models we can work with, simply because an infinite evidence model might not have any maximal finite body of evidence. To illustrate this, we can think of the evidence model presented in Example 1: the set of basic evidence \mathcal{E}_0 is the only maximal body of evidence in $(\mathbb{N}, \mathcal{E}_0, V)$, and it is infinite. Therefore, in order to eventually be able to provide a belief logic of all evidence models that formalizes a notion of consistent belief, further adjustments in the definition of Bel are warranted. To this end, we propose to “weaken” the above definition, by focusing on the *finite bodies of evidence that are “strong enough”* (instead of the “strongest” such bodies).

4.2 Justified belief: our proposal

It seems to us that the intended goal (only partially fulfilled) of the above-mentioned definition of belief was to ensure that the agents are able to form consistent beliefs based on the (possibly false and possibly mutually contradictory) available evidence. We think this to be a natural requirement for *idealized rational* agents, and so we consider doxastic inconsistency to be a bug, not a feature, of the above framework. Hence, we now propose a notion that produces in a natural way—with no need for further restrictions—only consistent beliefs, and that also agrees with the one in (van Benthem & Pacuit, 2011) in the finite case (and other cases specified below).

The intuition behind our proposal is that a proposition P is believed iff *it is supported by all “sufficiently strong” evidence*. We therefore say that P is believed, and write BP , iff every finite body of evidence can be strengthened to some finite body of evidence which




Fig. 3 $\mathfrak{M} = (\mathbb{N} \cup \{\spadesuit\}, \mathcal{E}_0, V)$

supports P . More formally, given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$ and a proposition $P \subseteq X$,

$$BP \text{ holds (at any state) iff } \forall F \in \mathcal{F}^{fin} \exists F' \in \mathcal{F}^{fin} (F \subseteq F' \text{ and } \bigcap F' \subseteq P).$$

The notion of belief B (like *Bel*) is a “global” notion, which depends only on the agent’s evidence, not on the actual world, so it is either true in all possible worlds, or false in all possible worlds. We therefore have

$$BP := \begin{cases} X & \text{if } \forall F \in \mathcal{F}^{fin} \exists F' \in \mathcal{F}^{fin} (F \subseteq F' \text{ and } \bigcap F' \subseteq P) \\ \emptyset & \text{otherwise.} \end{cases}$$

This reflects the assumption that beliefs are internal (and fully transparent) to the agent (Baltag et al., 2008).

It is easy to see that, unlike *Bel*, our notion of belief B is *always consistent* (i.e., $B\perp = B\emptyset = \emptyset$), since no finite body of evidence has an empty intersection. Moreover, it satisfies the axioms of the standard doxastic logic KD45 (see Sect. 6.3). As shown in Example 2, our notion of belief B and *Bel* are in general incompatible (even in cases when *Bel* is consistent). On the other hand, these two notions coincide on a restricted class of evidence models (see Proposition 1).

Example 2 The models below show that B and *Bel* are in general not comparable. More precisely, the first model illustrates that BP does not imply *Bel* P and the second model shows that *Bel* P does not imply BP even when *Bel* is consistent.

Consider the evidence model $\mathfrak{M} = (\mathbb{N} \cup \{\spadesuit\}, \mathcal{E}_0, V)$, where \mathbb{N} is the set of natural numbers, $V(p) = \emptyset$, and the set of basic evidence is $\mathcal{E}_0 = \{e_i \mid i \in \mathbb{N}\} \cup \{n \mid n \in \mathbb{N}\}$ where $e_i = [i, \infty) \cup \{\spadesuit\}$ (see Fig. 3).

We then have that

$$Max_{\subseteq}(\mathcal{F}) = \{\{e_i \mid i \in \mathbb{N}\}\} \cup \{\{e_i \mid i \leq n\} \cup \{m\} \mid n, m \in \mathbb{N} \text{ with } m \geq n\}.$$

Therefore, for any $F \in Max_{\subseteq}(\mathcal{F})$, we have

$$\bigcap F = \begin{cases} \{\spadesuit\} & \text{if } F = \{e_i \mid i \in \mathbb{N}\}, \\ \{m\} & \text{if } F = \{e_i \mid i \leq n\} \cup \{m\} \text{ with } m \geq n. \end{cases}$$

We thus obtain that $\bigcup_{F \in Max_{\subseteq}(\mathcal{F})} \bigcap F = \mathbb{N} \cup \{\spadesuit\}$. This means that $Bel(\mathbb{N} \cup \{\spadesuit\}) = Bel\top$ holds in \mathfrak{M} , and moreover, $\mathbb{N} \cup \{\spadesuit\}$ is the only proposition that is believed according to the belief definition of van Benthem and Pacuit (2011). Thus, in particular, $Bel(\mathbb{N}) = \emptyset$,




Fig. 4 $\mathfrak{M}' = (\mathbb{N} \cup \{\spadesuit\}, \mathcal{E}'_0, V)$

hence, $Bel(\mathbb{N})$ does not hold in \mathfrak{M} (i.e., no state in $\mathbb{N} \cup \{\spadesuit\}$ makes $Bel(\mathbb{N})$ true). On the other hand, we have $F \in \mathcal{F}^{fin}$ iff $F = \{e_i \mid i \in I\}$, or $F = \{e_i \mid i \in I\} \cup \{m\}$ for some $I \subseteq_{fin} \mathbb{N}$ and $m \geq \max(I)$, where $\max(I)$ is the greatest natural number in I . Therefore, for every $F \in \mathcal{F}^{fin}$, we have

$$\bigcap F = \begin{cases} [\max(I), \infty) \cup \{\spadesuit\} & \text{if } F = \{e_i \mid i \in I\}, \\ \{m\} & \text{if } F = \{e_i \mid i \in I\} \cup \{m\} \text{ for } m \geq \max(I). \end{cases}$$

This implies that, any finite body F of the form $\{e_i \mid i \in I\} \cup \{m\}$ already supports \mathbb{N} . Moreover, if $F = \{e_i \mid i \in I\}$, there exists a stronger finite body F' of the form $F' = \{e_i \mid i \in I\} \cup \{m\}$ for some $m \geq \max(I)$ that supports \mathbb{N} . We therefore have that $B(\mathbb{N})$ holds in \mathfrak{M} . Hence, in general, BP does not imply $BelP$.

Now consider the evidence model $\mathfrak{M}' = (\mathbb{N} \cup \{\spadesuit\}, \mathcal{E}'_0, V)$ based on the same domain as \mathfrak{M} , and where $V(p) = \emptyset$ and the basic evidence family $\mathcal{E}'_0 = \{[n, \infty) \cup \{\spadesuit\} \mid n \in \mathbb{N}\}$ (see Fig. 4). The only maximal body of evidence in \mathcal{E}'_0 is \mathcal{E}'_0 itself, and $\bigcap \mathcal{E}'_0 = \{\spadesuit\}$. Therefore, we have $\neg Bel \perp$ true in \mathfrak{M}' , i.e., Bel is consistent in \mathfrak{M}' . Moreover, in particular, $Bel\{\spadesuit\}$ is true in \mathfrak{M} . On the other hand, for all finite bodies $F \in \mathcal{F}^{fin}$, we have $\{\spadesuit\} \subsetneq \bigcap F$, implying that $\neg B\{\spadesuit\}$ is true in \mathfrak{M}' . Therefore, even when Bel is consistent, $BelP$ does not imply BP .

There are special cases where Bel and B do coincide. First of all, B coincides with Bel on the evidence models with finite basic evidence sets \mathcal{E}_0 . More generally, Bel and B coincide on all *maximally compact* evidence models: the ones in which every body of evidence is equivalent to a finite body of evidence. More formally, an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$ is called *maximally compact* if it satisfies the property

$$\forall F \in \mathcal{F} \exists F' \in \mathcal{F}^{fin} (\bigcap F = \bigcap F') \tag{MC}$$

Proposition 1 *For all maximally compact evidence models $\mathfrak{M}=(X, \mathcal{E}_0, V)$ and $P \subseteq X$, we have $BelP = BP$.*

Proof Let $\mathfrak{M} = (X, \mathcal{E}_0, V)$ be a maximally compact evidence model and $P \subseteq X$.

(\subseteq) Suppose $BelP$ holds in \mathfrak{M} , i.e., suppose that for all $F \in \text{Max}_{\subseteq} \mathcal{F}$, we have $\bigcap F \subseteq P$. Now let $F' \in \mathcal{F}^{fin}$. By Zorn’s Lemma, F' can be extended to a maximal body of evidence $F'' \in \mathcal{F}$. Note that, since F'' extends F' , i.e., $F' \subseteq F''$, we have $\bigcap F'' \subseteq \bigcap F'$. Since \mathfrak{M} is maximally compact, there is $F_0 \in \mathcal{F}^{fin}$ such that $\bigcap F'' = \bigcap F_0$. Now consider the family of evidence $F_0 \cup F'$. Since $\bigcap F_0 = \bigcap F'' \subseteq \bigcap F'$, we have $\bigcap (F_0 \cup F') = \bigcap F_0 \cap \bigcap F' = \bigcap F_0 \neq \emptyset$. Therefore, the family of evidence $F_0 \cup F'$ is a finite body of evidence, i.e., $F_0 \cup F' \in \mathcal{F}^{fin}$. Obviously, $F_0 \cup F'$ extends F' , i.e., $F' \subseteq F_0 \cup F'$. Moreover, since $BelP$ holds in \mathfrak{M} , we have that $\bigcap F'' \subseteq P$. We then obtain $\bigcap (F_0 \cup F') = \bigcap F_0 = \bigcap F'' \subseteq P$. We have therefore proven that the finite body of evidence $F_0 \cup F'$

extends F' and it entails P . As F' has been chosen arbitrarily from \mathcal{F}^{fin} , we conclude that BP holds in \mathfrak{M} .

(\supseteq) Suppose BP holds in \mathfrak{M} , i.e., suppose that for all $F \in \mathcal{F}^{fin}$, there exists $F' \in \mathcal{F}^{fin}$ such that $F \subseteq F'$ and $\bigcap F' \subseteq P$. Let $F'' \in \text{Max}_{\subseteq} \mathcal{F}$. Then, since \mathfrak{M} is maximally compact, there exists $F_0 \in \mathcal{F}^{fin}$ such that $\bigcap F'' = \bigcap F_0$. Moreover, since BP holds in \mathfrak{M} , there exists $F_1 \in \mathcal{F}^{fin}$ such that $F_0 \subseteq F_1$ and $\bigcap F_1 \subseteq P$. Besides, since $\bigcap F_1 \subseteq \bigcap F_0 = \bigcap F''$ and F'' is maximal, we in fact have $F_1 \subseteq F''$ (otherwise, there exists $e \in \mathcal{E}_0$ such that $e \in F_1$ but $e \notin F''$. Therefore, as $\bigcap F_1 \subseteq \bigcap F''$, we would have $\bigcap F_1 \subseteq \bigcap (F'' \cup \{e\})$, and thus $\bigcap (F'' \cup \{e\}) \neq \emptyset$, contradicting maximality of F'' .) Therefore, $\bigcap F'' \subseteq \bigcap F_1$, and thus, $\bigcap F_1 = \bigcap F''$. Then, together with $\bigcap F_1 \subseteq P$, we obtain $\bigcap F'' \subseteq P$. As F'' has been chosen arbitrarily from $\text{Max}_{\subseteq} \mathcal{F}$, we conclude that $BelP$ holds in \mathfrak{M} . \square

Another important feature of our belief definition is that B is a purely topological notion, as stated in the following proposition, which, in turn, constitutes a justification for our use of topo-e-models rather than working with only evidence models.

Proposition 2 *In every topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, the following are equivalent, for any proposition $P \subseteq X$:*

1. BP holds (at any state)
(i.e., $\forall F \in \mathcal{F}^{fin} \exists F' \in \mathcal{F}^{fin} (F \subseteq F' \text{ and } \bigcap F' \subseteq P)$);
2. every evidence can be strengthened to some evidence supporting P
(i.e., $\forall e \in \mathcal{E} \exists e' \in \mathcal{E} (e' \subseteq e \cap P)$);
3. every argument (for anything) can be strengthened to an argument for P (i.e., $\forall U \in \tau \setminus \{\emptyset\} \exists U' \in \tau \setminus \{\emptyset\} (U' \subseteq U \cap P)$);
4. there is a justification for P , i.e., there is some argument for P which is consistent with any available evidence
(i.e., $\exists U \in \tau (U \subseteq P \text{ and } \forall e \in \mathcal{E} (U \cap e \neq \emptyset))$);
5. P includes some dense open set
(i.e., $\exists U \in \tau (U \subseteq P \text{ and } Cl(U) = X)$);
6. $Int(P)$ is dense in τ (i.e., $Cl(Int(P)) = X$), or equivalently, $X \setminus P$ is nowhere dense
(i.e., $Int(Cl(X \setminus P)) = \emptyset$);
7. $[\forall] \diamond \square P$ holds (at any state) (i.e., $[\forall] \diamond \square P = X$), or equivalently, $[\forall] \diamond \square P \neq \emptyset$.

Proof The equivalence of (1), (2) and (3) is easy, and follows directly from the definitions of combined evidence and argument. The equivalence of (5) and (6) is also straightforward (recall that $Int(P)$ is the largest open contained in P). The equivalence of (4) and (5) simply follows from the definitions of arguments and dense sets. For the equivalence of (6) and (7), recall that $[\forall]$ is the global modality, \square is interior, and \diamond is closure. For the equivalence of (3) and (4):

(3) \Rightarrow (4): Suppose that (3) holds and consider the open set $Int(P)$. We will show that $Int(P)$ is a justification for P , i.e., $Int(P) \cap e \neq \emptyset$ for all $e \in \mathcal{E}$. Let $e \in \mathcal{E}$. By (3), since $e \in \mathcal{E} \subseteq \tau \setminus \{\emptyset\}$, there exists $U_0 \in \tau \setminus \{\emptyset\}$ such that $U_0 \subseteq e \cap P$. We then have $Int(U_0) \subseteq Int(e \cap P) = Int(e) \cap Int(P)$. Therefore, since U_0 and e are open sets, we obtain $U_0 \subseteq e \cap Int(P)$. As $U_0 \neq \emptyset$, we conclude that $e \cap Int(P) \neq \emptyset$.

(4) \Rightarrow (3): Suppose that (4) holds, i.e., suppose that there exists $U_0 \in \tau$ such that (a) $U_0 \subseteq P$ and (b) $U_0 \cap e \neq \emptyset$ for all $e \in \mathcal{E}$. Let $U \in \tau$ with $U \neq \emptyset$. Now consider the open set $U \cap U_0$. Since \mathcal{E} is a basis of τ , there exists $e \in \mathcal{E}$ such that $e \subseteq U$. Therefore, by (b), the intersection $U \cap U_0 \neq \emptyset$, thus, $U \cap U_0 \in \tau \setminus \{\emptyset\}$. By (a), we also have $U \cap U_0 \subseteq U \cap P$. \square

Proposition 2 deserves a closer look. First, it describes the topological properties of our notion of belief. Second, it states that our belief is the same as “justified” belief, but more specifically one whose justification is an evidence-based argument that is consistent with every available evidence. The equivalence of (1.), (2.), and (3.) shows that we can define BP in equivalent ways by using only basic evidence pieces (i.e., the elements of \mathcal{E}_0), or by using only combined evidence (i.e., the elements of \mathcal{E}), or by using only the open sets of the generated evidential topology τ . Proposition 2.4. proves that our definition of belief indeed gives us a conception of *evidentially justified belief*. The requirement that any justification of a believed proposition must be *open* in the evidential topology means that the justification is ultimately based on the available evidence; while the requirement that the justification is *dense* (in the same topology) means that all the agent’s beliefs must be consistent with every piece of evidence. Therefore, believed propositions, according to our definition, are those for which there is some evidential justification that is consistent with every available (basic or combined) evidence. Moreover, whenever a proposition P is believed, there exists a *weakest (most general) justification* for P , namely the open set $Int(P)$. Items (5.) and (6.) provide topological reformulations of the above items. In particular, Proposition 2.6. shows that our proposal is very natural from a topological perspective: it is equivalent to saying that P is believed iff the complement of P is nowhere dense. Since nowhere dense sets are one of the topological concepts of “small” or “negligible” sets, this amounts to believing propositions iff they are true in *almost all* epistemically-possible worlds, where “almost all” spelled out topologically as “everywhere but a nowhere dense part of the model”. Finally, Proposition 2.7. tells us that belief is definable in terms of the operators $[\forall]$ and \square .

4.3 Conditional belief on Topo-e-models

The belief semantics given in Sect. 4.2 can be generalized to *conditional beliefs* $B^Q P$ by relativizing the plain belief definition BP to the given condition Q . The current setting requires a careful treatment of the aforementioned relativization (as recognized already in van Benthem and Pacuit 2011) since some of the agent’s evidence might be inconsistent with the condition Q . While evaluating beliefs under the assumption that the given condition Q is true, one should focus only on the evidence that is consistent with Q by neglecting the evidence pieces that are disjoint with Q . Therefore, in order to define conditional beliefs, we need a relativized version of the notion of consistent (bodies of) evidence.

Given an evidence model $\mathfrak{M} = (X, \mathcal{E}_0, V)$, for any subsets $Q, A \subseteq X$, we say that A is *Q-consistent* iff $Q \cap A \neq \emptyset$. Moreover, a body of evidence F is called *Q-consistent* iff $\bigcap F \cap Q \neq \emptyset$. We can then define conditional beliefs based on these notions of conditional consistency. We say that P is *believed given Q*, and write $B^Q P$, iff every finite *Q-consistent body of evidence can be strengthened to some finite Q-consistent body of evidence supporting the proposition $Q \rightarrow P$* .

An analogue of Proposition 2 providing different characterizations can also be proven for conditional belief:

Proposition 3 *In every topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, the following are equivalent, for any two propositions $P, Q \subseteq X$ with $Q \neq \emptyset$:*

1. $B^Q P$ holds (at any state)
(i.e., $\forall F \in \mathcal{F}^{fin}(\bigcap F \cap Q \neq \emptyset \Rightarrow \exists F' \in \mathcal{F}^{fin}(\bigcap F' \cap Q \neq \emptyset \text{ and } \bigcap F' \subseteq \bigcap F \cap (Q \rightarrow P)))$);
2. every Q -consistent evidence can be strengthened to some Q -consistent evidence supporting $Q \rightarrow P$
(i.e., $\forall e \in \mathcal{E}(e \cap Q \neq \emptyset \Rightarrow \exists e' \in \mathcal{E}(e' \cap Q \neq \emptyset \text{ and } e' \subseteq e \cap (Q \rightarrow P)))$);
3. every Q -consistent argument can be strengthened to a Q -consistent argument for $Q \rightarrow P$
(i.e., $\forall U \in \tau(U \cap Q \neq \emptyset \Rightarrow \exists U' \in \tau(U' \cap Q \neq \emptyset \text{ and } U' \subseteq U \cap (Q \rightarrow P)))$);
4. there is some Q -consistent argument for $Q \rightarrow P$ whose intersection with any Q -consistent evidence is Q -consistent
(i.e., $\exists U \in \tau(U \cap Q \neq \emptyset \text{ and } U \subseteq Q \rightarrow P \text{ and } \forall e \in \mathcal{E}(e \cap Q \neq \emptyset \Rightarrow (U \cap e) \cap Q \neq \emptyset))$);
5. $Q \rightarrow P$ includes some Q -consistent open set which is dense in Q
(i.e., $\exists U \in \tau(U \cap Q \neq \emptyset \text{ and } U \subseteq Q \rightarrow P \text{ and } Q \subseteq Cl(U \cap Q))$);
6. $Int(Q \rightarrow P)$ is dense in Q
(i.e., $Q \subseteq Cl(Q \cap Int(Q \rightarrow P))$);
7. $[\forall](Q \rightarrow \diamond(Q \wedge \square(Q \rightarrow P)))$ holds (at any state) (i.e., $[\forall](Q \rightarrow \diamond(Q \wedge \square(Q \rightarrow P))) = X$), or equivalently, $[\forall](Q \rightarrow \diamond(Q \wedge \square(Q \rightarrow P))) \neq \emptyset$.

Proof The equivalence of (1), (2), (3) is easy and directly follows from the semantics of $B^Q P$, and the definitions of Q -consistent evidence and Q -consistent argument. For the equivalence between (5) and (6), consider the weakest argument $Int(Q \rightarrow P)$ for $Q \rightarrow P$ as the relevant open set. And, for the equivalence of (6) and (7), recall that $[\forall]$ is the universal quantifier, \square is interior, and \diamond is closure. We here show only the equivalence of (3) and (4), and between (4) and (5) in details.

(3) \Rightarrow (4): Suppose that (3) holds and consider the weakest argument $Int(Q \rightarrow P)$ for $Q \rightarrow P$. Since $X \in \mathcal{E}$ and X is Q -consistent, by (3), there exists a stronger $U \in \tau$ such that $U \cap Q \neq \emptyset$ and $U \subseteq Q \rightarrow P$. Since $Int(Q \rightarrow P)$ is the largest open with $Int(Q \rightarrow P) \subseteq Q \rightarrow P$, we obtain $U \subseteq Int(Q \rightarrow P) \subseteq Q \rightarrow P$ for any such U , therefore, $Int(Q \rightarrow P)$ is also Q -consistent. Let $e \in \mathcal{E}$ be such that $e \cap Q \neq \emptyset$. Therefore, since $\mathcal{E} \subseteq \tau$, by (3), there exists $U' \in \tau$ such that $U' \cap Q \neq \emptyset$ and $U' \subseteq e \cap (Q \rightarrow P)$. By the previous argument, we know that $U' \subseteq Int(Q \rightarrow P)$, thus, $U' \subseteq e \cap Int(Q \rightarrow P) \neq \emptyset$. And, since U' is Q -consistent, the result follows.

(4) \Rightarrow (3): Suppose that (4) holds, i.e., suppose that there is $U_0 \in \tau$ such that (a) $U_0 \cap Q \neq \emptyset$, (b) $U_0 \subseteq Q \rightarrow P$ and (c) for all $e \in \mathcal{E}$ with $e \cap Q \neq \emptyset$, we have $(U_0 \cap e) \cap Q \neq \emptyset$. Let $U \in \tau$ be such that $U \cap Q \neq \emptyset$ and consider the open set $U \cap U_0$. Since $U \cap Q \neq \emptyset$ and \mathcal{E} is a basis for τ , there exists $e_0 \in \mathcal{E}$ such that $e_0 \subseteq U$ and $e_0 \cap Q \neq \emptyset$. Therefore, by (c), we have that $(U_0 \cap e_0) \cap Q \neq \emptyset$, thus, the open set $U_0 \cap e_0$ is Q -consistent. Moreover, since $U_0 \subseteq Q \rightarrow P$ and $e_0 \subseteq U$, we obtain $U_0 \cap e_0 \subseteq U \cap (Q \rightarrow P)$.

(4) \Leftrightarrow (5): For the left-to-right direction, suppose (4) holds as in the above case, and toward showing $Q \subseteq Cl(U_0 \cap Q)$, let $x \in Q$ and $e \in \mathcal{E}$ such that $x \in e$. Therefore, e is Q -consistent, i.e., $e \cap Q \neq \emptyset$. Then, by (4), we obtain $(U_0 \cap e) \cap Q \neq \emptyset$, implying that $x \in Cl(U_0 \cap Q)$. For the right-to-left direction, suppose (5) holds with U_0 the witness and let $e \in \mathcal{E}$ be such that $e \cap Q \neq \emptyset$. This means that there is $y \in e \cap Q$, thus, $y \in Q$. Then, by (5), $y \in Cl(U_0 \cap Q)$. Therefore, as $y \in e \in \mathcal{E}$, we conclude $(U_0 \cap Q) \cap e \neq \emptyset$. \square

5 Knowledge

As already mentioned, the notion of *infallible knowledge*—represented by the global modality $[\forall]$ introduced in Sect. 3.2—has a very limited scope: there are very few things we could know in this strong sense, maybe, say, only logical-mathematical tautologies. We now proceed to define a weaker and thus more widely applicable notion of knowledge, which better approximates the common usage of the word.

More concretely, the concept of (fallible) knowledge we propose is based on *factive justifications*. Formally, given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, we set

$$KP = \{x \in X \mid \exists U \in \tau (x \in U \subseteq P \text{ and } Cl(U) = X)\}.$$

In other words: *KP holds at a world x* iff *P includes a dense open neighborhood of x*. Similarly to the cases for belief and conditional beliefs (recall Propositions 2 and 3), we can provide several equivalent definitions of *KP* on topo-e-models as follow.

Proposition 4 *Let $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ be a topo-e-model and $x \in X$ be the actual world. The following are equivalent for all $P \subseteq X$:*

1. *KP holds at x in \mathfrak{M}*
(i.e., $\exists U \in \tau (x \in U \subseteq P \text{ and } Cl(U) = X)$);
2. *there is some factive justification for P at x, i.e., there is some factive argument for P at x which is consistent with any available evidence*
(i.e., $\exists U \in \tau (x \in U \subseteq P \text{ and } \forall e \in \mathcal{E}(U \cap e \neq \emptyset))$);
3. *Int(P) contains the actual state and is dense in τ*
(i.e., $x \in Int(P)$ and $Cl(Int(P)) = X$);
4. $\Box P \wedge BP$ holds at x.

Proof The proof is similar to the proof of Proposition 2. For the equivalence of (1.) and (2.), recall that \mathcal{E} constitutes a basis for τ . The equivalence of (2.) and (3.) is also straightforward (recall that $Int(P)$ is the largest open set contained in P). For the equivalence of (3.) and (4.), see Proposition 2.6. and recall that \Box is interpreted as the interior operator. \square

Therefore, as the equivalence of items 1. and 2. of Proposition 4 shows, our proposal equates “knowledge” with *correctly justified belief*: belief based on *true* justifications. We will see that our ‘coherentist’ notion of justification makes this notion subtly different from the influential “no false lemma” account of knowledge. But first, we should note that our proposal does *not* simply boil down to “justified true belief”. This would clearly be vulnerable to Gettier-type counterexamples (Gettier, 1963). To better explain the distinction, we illustrate in the example below the proposed semantics for justified belief and knowledge, as well as the connection between the two notions.

Example 3 Consider the topo-e-model $\mathfrak{M} = ([0, 1], \mathcal{E}_0, \tau, V)$, where $\mathcal{E}_0 = \{(a, b) \cap [0, 1] \mid a, b \in \mathbb{R}, a < b\}$ and $V(p) = \emptyset$. The generated topology τ is the standard topology of open intervals restricted to $[0, 1]$. Let $P = [0, 1] \setminus \{\frac{1}{n} \mid n \in \mathbb{N}\}$ be the proposition stating that “the actual state is *not* of the form $\frac{1}{n}$, for any $n \in \mathbb{N}$ ” (see Fig. 5). Since the complement $\neg P = [0, 1] \setminus P = \{\frac{1}{n} \mid n \in \mathbb{N}\}$ is nowhere dense (i.e., $Int(Cl(\neg P)) = Int(\neg P) = \emptyset$), the agent *believes P*, and e.g. $U = \bigcup_{n \geq 1} (\frac{1}{n+1}, \frac{1}{n})$ is a justification for P , that is, U is a




Fig. 5 $([0, 1], \tau)$

dense open subset of P . This belief is *true at world* $0 \in P$. But this *true belief is not known* at 0: no justification for P is true at 0, since P does not include any open neighborhood of 0, so $0 \notin \text{Int}(P)$ and hence $0 \notin KP$. This shows that $KP \neq P \wedge BP$. Moreover, P is known in all the *other* states $x \in P \setminus \{0\}$, since

$$\forall x \in P \setminus \{0\} \exists \epsilon > 0 (x \in (x - \epsilon, x + \epsilon) \subseteq P),$$

therefore $x \in \text{Int}(P)$.

A brief note on Stalnaker’s epistemic-doxastic system (Stalnaker, 2006): it is easy to see that K together with justified belief B satisfies Stalnaker’s Full Belief principle $BP = BKP$ (see Table 5 for the complete list of his axioms). These operators in fact satisfy all the axioms and rules of the system Stalnaker’s logic of knowledge and belief on all topo-e-models, thus, on *all topological spaces*, not only on the restricted class of extremally disconnected spaces. We prove the soundness and completeness of Stalnaker’s system with respect to all topo-e-models in Sect. 6.4.

One interesting property of this weaker type of knowledge is it being *defeasible* in the light of new information, even when the new information is true. In contrast, the usual assumption in epistemic logic is that *knowledge acquisition is monotonic*. As a result, logicians typically assume that knowledge is *irrevocable*: once acquired, it cannot be defeated by any further evidence gathered later. In our setting, the only irrevocable knowledge is the absolutely certain one (true in all epistemically-possible worlds), captured by the operator $[\forall]$. Clearly, K is *not* irrevocable.

5.1 Knowledge is defeasible

Gettier (1963)—with his famous counterexamples against the account of knowledge as justified true belief—triggered an extensive discussion in epistemology that is concerned with understanding what knowledge is, and in particular, with identifying the exact properties and conditions that render a piece of justified true belief knowledge. Epistemologists have made various proposals such as, among others, the *no false lemma* (Clark, 1963), the *defeasibility analysis of knowledge* (Klein, 1971, 1981; Lehrer, 1990; Lehrer & Paxson, 1969), the *sensitivity account* (Nozick, 1981), the *safety account* (Sosa, 1999), and the *contextualist account* (DeRose, 2009).¹⁵ While there is still very little agreement as to which proposal gives a satisfactory solution to the Gettier challenge, the extent of the post-Gettier literature at the very least shows that the relation between justified belief and knowledge is very delicate, and it is not an easy task, if possible, to identify a unique notion of knowledge that can deal with all kinds of intuitive counterexamples. However, as Rott states, one can accept that all these proposals “capture important intuitions that can

¹⁵ For an overview of responses to the Gettier challenge and a detailed discussion, we refer the reader to (Ichikawa & Steup, 2013; Rott, 2004).

in some way or other be regarded as relevant to the question whether or not a given belief constitutes a piece of knowledge” (Rott, 2004, p. 469).

In this section, we argue that the conception of knowledge captured by our modality K is stronger than Clark’s “no false lemma” (Clark, 1963), and very close to (though subtly different from) the so-called defeasibility theory of knowledge held by Klein (1971, 1981), Lehrer (1990), and Lehrer and Paxson (1969). But providing an extensive philosophical comparison with all the aforementioned theories of knowledge is way beyond the scope of this paper, so we leave this task for future work.

Clark’s influential “no false lemma” proposal requires a *correct* “justification”—one that doesn’t use any falsehood—for a piece of belief to constitute knowledge (Clark, 1963). While this may sound very similar to our definition of knowledge K , our proposal imposes a stronger implicit requirement than Clark’s, since our concept of justification requires consistency with all the available (combined) evidence. In our terminology, Clark only requires a *factive argument* for P . So Clark’s approach is local, assessing a knowledge claim based only on the truth of the evidence pieces (and the correctness of the inferences) that are used to justify it. In contrast, our proposed notion of knowledge inherits the ‘holistic’ character of our proposed concept of belief: to count as justifications, evidential arguments first need to be checked against all (the other arguments that can be constructed from the agent’s) current evidence. So a knowledge claim is assessed by checking both the truth of the underlying argument and its consistency with *all* of the agent’s acceptance system.

On the other hand, the defeasibility theory of knowledge, roughly speaking, defines knowledge as *undefeated justified belief*: justified belief that cannot be defeated by any *factive* evidence that might be gathered later (though it may be defeated by false evidence). In its simplest version, called by Rott (2004) *stable belief theory* or *stability theory of knowledge*, it says that *the agent knows P if and only if*

1. P is true
2. she believes that P , and
3. her belief in P cannot be defeated by new *factive* information.

In other words, given a true proposition P , the agent knows P iff the belief in P is *stable* for true information. The stable belief theory has been challenged for being too weak to characterize knowledge: the agent may keep her (true) belief stable, while continuously adopting newer justifications. Each of these justifications is wrong and can be defeated, but the belief itself remains undefeated. A more developed version of defeasibility theory, as held by Lehrer and others, insists that, in order to know P , not only the belief in P has to stay stable, but also its justification (i.e. what we call here “an argument for P ”) should be undefeated. More precisely, according to this strong version of defeasibility theory, *the agent knows P if and only if*

1. P is true
2. she believes that P ,
3. her belief in P cannot be defeated by new *factive* information, and
4. her ‘justification’ (=argument, in our sense) is undefeated by new *factive* information.

In this sense, for the agent to know P there must exist an argument for P that is believed conditional on every true evidence. Clearly, this implies that the belief in P is stable, but the converse fails. As already observed, the problem is that, when confronted with various new pieces of evidence, the agent might keep switching between different justifications (for believing P), thus, she may keep believing in P conditional on any such new true

evidence without actually having any good, robust justification (i.e., one that remains itself undefeated by all true evidence) (see Example 5). To have knowledge, we thus need a *stable justification*.¹⁶

However, the above interpretations of both the stability and the defeasibility theory were also attacked as being *too strong*: if we allow as potential defeaters *all* factive propositions (i.e. all sets of worlds P containing the actual world), then there are intuitive examples showing that knowledge KP can be defeated (Klein, 1980, 1981). Here is such an example discussed by Klein (1981), a leading proponent of the defeasibility theory. Loretta filled in her federal taxes, following very carefully all the required procedures on the forms, doing all the calculations and double checking everything. Based on this evidence, she correctly believes that she owes \$500, and she seems perfectly justified to believe this. So it seems obvious that she *knows* this. But suppose now that, being aware of her own fallibility, she asks her accountant to check her return. The accountant finds no errors (when there are in fact some errors in her calculation, yet not affecting the correct result that she owes \$500), and so he sends her his reply reading “Your return contains *no* errors”; but he inadvertently leaves out the word “no”. If Loretta would learn the true fact that the accountant’s letter actually reads “Your return contains errors”, she would lose her true belief that she owed \$500. So it seems that there exist defeaters that are true but “misleading”. We formalize this counterexample in Example 4 and show that our knowledge K is neither stable nor indefeasible. In order to make the formalization more succinct, we first introduce an operation of *evidence addition* and some notation.

Definition 8 (*Evidence added topo-e-model*) Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ and a nonempty $P \subseteq X$, we can define a P -added topo-e-model \mathfrak{M}^{+P} as $\mathfrak{M}^{+P} = (X, \mathcal{E}_0^{+P}, \tau^{+P}, V)$, where $\mathcal{E}_0^{+P} = \mathcal{E}_0 \cup \{P\}$ and τ^{+P} is the topology generated by \mathcal{E}_0^{+P} .

It is easy to see that \mathfrak{M}^{+P} is a topo-e-model, since $\emptyset \notin \mathcal{E}_0^{+P}$ and $X \in \mathcal{E}_0^{+P}$, and τ^{+P} is the evidential topology generated by \mathcal{E}_0^{+P} . Moreover, the set of combined evidence \mathcal{E}^{+P} of \mathfrak{M}^{+P} can be described as

$$\mathcal{E}^{+P} = \mathcal{E} \cup \{e \cap P \mid e \in \mathcal{E} \text{ with } e \cap P \neq \emptyset\},$$

which clearly constitutes a basis for τ^{+P} .

Example 4 Consider the model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, where $X = \{x_1, x_2, x_3\}$, $V(p) = \emptyset$, $\mathcal{E}_0 = \{X, O_1, O_2\}$, $O_1 = \{x_1, x_2\}$, $O_2 = \{x_2, x_3\}$ (see Fig. 6). The resulting set of combined evidence is $\mathcal{E} = \{X, O_1, O_2, \{x_2\}\}$. Assume that the actual world is x_1 . Then O_1 is *known*, since $x_1 \in \text{Int}(O_1) = O_1$ and $\text{Cl}(O_1) = X$. Now consider the model $\mathfrak{M}^{+O_3} = (X, \mathcal{E}_0^{+O_3}, \tau^{+O_3}, V)$ obtained by adding the new evidence $O_3 = \{x_1, x_3\}$ (as in Definition 8). We have $\mathcal{E}_0^{+O_3} = \{X, O_1, O_2, O_3\}$, so $\mathcal{E}^{+O_3} = \{X, O_1, O_2, O_3, \{x_1\}, \{x_2\}, \{x_3\}\}$. Note that the new evidence is *true* ($x_1 \in O_3$). However, O_1 is *not even believed* in \mathfrak{M}^{+O_3} anymore, since $O_1 \cap \{x_3\} = \emptyset$, so O_1 is no longer dense in τ^{+O_3} . Therefore, O_1 is *no longer known* after the true evidence O_3 was added.

¹⁶ Lehrer uses the metaphor of an *Ultra-Justification Game* (Lehrer, 1990), according to which ‘knowledge’ is based on arguments that survive a game between the Believer and an omniscient truth-telling Critic, who tries to defeat the argument by using both the Believer’s current “justification system” and *any new true evidence* (see, Fiutek, 2013, Sect. 5.2 for a formalization of Lehrer’s ultra-justification game).




Fig. 6 From \mathfrak{M} to \mathfrak{M}^{+O_3}

Klein’s story corresponds to taking O_1 to represent Loretta’s direct evidence (based on careful calculations) that she owes \$500, O_2 to represent her prior evidence (either based on past experience, or just being one of Loretta’s default assumptions) that the accountant doesn’t make mistakes in his replies to her, and O_3 the potential new evidence provided by the letter. In conclusion, our notion of knowledge is incompatible with the above-mentioned strong interpretations of both stability and defeasibility theory, thus confirming the objections raised against them.

Klein’s solution is that one should exclude such *misleading* defeaters, which may “unfairly” defeat a good justification. But how can we distinguish them from genuine defeaters? Klein’s diagnosis, in Foley’s more succinct formulation, is that “a defeater is misleading if it justifies a falsehood in the process of defeating the justification for the target belief” (Foley, 2012, p. 96). In the example, the falsehood is that the accountant had discovered errors in Loretta’s tax return. It seems that the new evidence O_3 (the existence of the letter as actually written) supports this falsehood, but how? According to us, it is the combination $O_2 \cap O_3$ of the new (true) evidence O_3 with the old (false) evidence O_2 that supports the new falsehood: the true fact (about the letter saying what it says) entails a falsehood *only* if it is taken in conjunction with Loretta’s prior evidence (or blind trust) that the accountant cannot make mistakes. So intuitively, *misleading defeaters are the ones which may lead to new false conclusions when combined with some of the old evidence*.


5.2 Misleading evidence and weakly indefeasible knowledge

We proceed now to formalize the distinction between misleading and genuine (i.e., non-misleading) defeaters. Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, a state $x \in X$ and a proposition $Q \subseteq X$,

- Q is *misleading* at $x \in X$ with respect to \mathcal{E} if evidence-addition with Q produces some false new evidence;

equivalently, and more formally, if there is some $e \in \mathcal{E}^{+Q} \setminus \mathcal{E}$ such that $x \notin e$, i.e., if there is some $e \in \mathcal{E}$ such that $x \notin (e \cap Q)$ and $(e \cap Q) \notin \mathcal{E} \cup \{\emptyset\}$. A proposition $Q \subseteq X$ is called *nonmisleading* if Q is not misleading. It is easy to see that *old evidence* $e \in \mathcal{E}$ is by definition nonmisleading with respect to \mathcal{E} (i.e., each $e \in \mathcal{E}$ is nonmisleading with respect to \mathcal{E}), and *new nonmisleading* evidence must be true (i.e., if $Q \subseteq X$ is nonmisleading at x and $Q \notin \mathcal{E}$, then $x \in Q$).

Fig. 7 $\mathfrak{M} = (X, \mathcal{E}_0, V)$: the continuous ellipses represent the currently available pieces of evidence, while the dashed ones represent the other nonmisleading propositions



We are now in the position to formulate precisely the weakened versions of both stability and defeasibility theories that we are looking for. The weak stability theory will stipulate that *the agent knows P if and only if*

1. P is true
2. she believes that P,
3. her belief in P cannot be defeated by any *nonmisleading* evidence.

On the other hand, the weak defeasibility theory requires that there exists some justification (argument) for P that is undefeated by every nonmisleading proposition. More precisely, the weak defeasibility theory strengthens the above described weak stability theory by the following “stable justification” clause:

4. her belief in its justification is undefeated by any *nonmisleading* evidence.

Finally, we also provide a third formulation, which one might call *epistemic coherence theory*, saying that *P is known iff there exists some justification (argument) for P which is consistent with every nonmisleading proposition*. While our proposed notion of knowledge is stronger than the one described by the weak stability theory, as illustrated by Example 5, it coincides with the ones defined by the weak defeasibility and epistemic coherence theories (see Proposition 5). In particular, the following counterexample shows that weak stability is (only a necessary, but) not a sufficient condition for knowledge K :

Example 5 Consider the model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, where $X = \{x_0, x_1, x_2\}$, $V(p) = \emptyset$, $\mathcal{E}_0 = \{X, O_1, O_2\}$ with $O_1 = \{x_1\}$, $O_2 = \{x_1, x_2\}$ (see Fig. 7). The resulting set of combined evidence is $\mathcal{E} = \mathcal{E}_0$. Assume that the actual world is x_0 and let $P = \{x_0, x_1\}$. Then, P is believed in \mathfrak{M} (since its interior $Int(P) = \{x_1\}$ is dense in τ) but it is *not known* (since $x_0 \notin Int(P) = \{x_1\}$). However, we can show that *P is believed in \mathfrak{M}^{+Q} for any nonmisleading Q at x_0* . For this, note that the family of nonmisleading propositions (at x_0) is $\mathcal{E} \cup \{P, \{x_0\}\} = \{X, O_1, O_2, P, \{x_0\}\}$. It is easy to see that for each set Q in this family, BP holds in \mathfrak{M}^{+Q} .

One should stress that our counterexample agrees with the position taken by most proponents of the defeasibility theory: stability of (justified) belief is not enough for knowledge. Intuitively, what happens in the above example is that, although the agent continues to believe P given any nonmisleading evidence, her justification keeps changing. For example, while the only justification for believing P in \mathfrak{M} is O_1 , the evidence O_1 is no longer dense in model $\mathfrak{M}^{+\{x_0\}}$, therefore, cannot constitute a justification for P in $\mathfrak{M}^{+\{x_0\}}$. On the other hand, another argument in $\mathfrak{M}^{+\{x_0\}}$, namely $\{x_0, x_1\}$ forms a justification for P in $\mathfrak{M}^{+\{x_0\}}$, thus P is still believed in $\mathfrak{M}^{+\{x_0\}}$, but, based on a different justification. Therefore, there is *no* uniform justification for P that works for every nonmisleading evidence Q.

The next result shows that *our notion of knowledge exactly matches the weakened version of defeasibility theory*, as well as the *epistemic coherence formulation*:

Proposition 5 Let $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ be a topo-e-model and $x \in X$ be the actual world. The following are equivalent for all $P \subseteq X$:

1. KP holds at x in \mathfrak{M} .
2. There is an argument (justification) for P that cannot be defeated by any nonmisleading proposition; i.e. $\exists U \in \tau \setminus \{\emptyset\}$ such that $U \subseteq P$ and BU holds in \mathfrak{M}^{+Q} for all nonmisleading $Q \subseteq X$ (at x with respect to \mathcal{E}).
3. There is an argument (justification) for P that is consistent with every nonmisleading proposition; i.e. $\exists U \in \tau \setminus \{\emptyset\}$ such that $U \subseteq P$ and $U \cap Q \neq \emptyset$ for all nonmisleading $Q \subseteq X$ (at x with respect to \mathcal{E}).

Proof (1) \Rightarrow (2): Suppose $x \in KP$. This means, by Proposition 4.3., that $x \in Int(P)$ and $Cl(Int(P)) = X$. Now consider the argument $Int(P)$. Obviously $Int(P) \in \tau \setminus \{\emptyset\}$ and $Int(P) \subseteq P$. Let Q be a nonmisleading proposition at x with respect to \mathcal{E} , and Cl^{+Q} and Int^{+Q} denote the closure and the interior operators of τ^{+Q} , respectively. We only need to show that $Int^{+Q}(Int(P))$ is dense in (X, τ^{+Q}) , i.e., that for all $e \in \mathcal{E}^{+Q}$, we have $e \cap Int^{+Q}(Int(P)) \neq \emptyset$. Let $e \in \mathcal{E}^{+Q}$. Then, by the definition of \mathcal{E}^{+Q} , we have two cases: (1) $e \in \mathcal{E}$, or (2) $e \notin \mathcal{E}$ but $e = e' \cap Q$ for some $e' \in \mathcal{E}$. Since Q is nonmisleading, the latter case entails that $x \in e$. If $e \in \mathcal{E}$, we have $e \cap Int^{+Q}(Int(P)) \neq \emptyset$ since $Int(P) \subseteq Int^{+Q}(Int(P))$ (by Lemma 1) and $Int(P)$ is dense in (X, τ) . If $e \notin \mathcal{E}$ and $e = e' \cap Q$ for some $e' \in \mathcal{E}$ with $x \in e$, we obtain $x \in e \cap Int^{+Q}(Int(P))$ since $x \in Int(P) \subseteq Int^{+Q}(Int(P))$, thus, $e \cap Int^{+Q}(Int(P)) \neq \emptyset$. Therefore, $Int^{+Q}(Int(P))$ is dense in (X, τ^{+Q}) , i.e., $B(Int(P))$ holds in \mathfrak{M}^{+Q} .

(2) \Rightarrow (3): Suppose that (2) holds, i.e., there is a $U \in \tau \setminus \{\emptyset\}$ such that $U \subseteq P$ and $Cl^{+Q}(Int^{+Q}(U)) = X$ for all nonmisleading $Q \subseteq X$ (at x with respect to \mathcal{E}). Let Q be nonmisleading at x with respect to \mathcal{E} . Since $Cl^{+Q}(Int^{+Q}(U)) = X$, we have that $e \cap Int^{+Q}(U) \neq \emptyset$ for all $e \in \mathcal{E}^{+Q}$. As Q is nonmisleading at x , we in particular have $\emptyset \neq Q = Q \cap X \in \mathcal{E}^{+Q}$ (by the definition of \mathcal{E}^{+Q} and the fact that $X \in \mathcal{E}$). Hence, it follows from (2) that $Q \cap Int^{+Q}(U) \neq \emptyset$. Since $Int^{+Q}(U) \subseteq U$, we obtain $U \cap Q \neq \emptyset$.

(3) \Rightarrow (1): Assume that $U \in \tau \setminus \{\emptyset\}$ is such that $U \subseteq P$ and $U \cap Q \neq \emptyset$ holds for all nonmisleading Q (at x with respect to \mathcal{E}). Clearly, this implies that U is consistent with all $e \in \mathcal{E}$, i.e., that $e \cap U \neq \emptyset$ (since available evidence is by definition nonmisleading), so U is a justification for P (i.e., $X = Cl(U) = Cl(Int(P))$). So, to show that KP holds at x , it is enough to show that $x \in Int(P)$. For this, take the proposition $Q = \{x\}$, which obviously is nonmisleading at x , hence by (3) we must have $U \cap \{x\} \neq \emptyset$, i.e. $x \in U$. Then, $x \in U \in \tau$ and $U \subseteq P$ give us $x \in Int(P)$, as desired. \square

6 Logics for evidence, justified belief and knowledge

This section constitutes the technical heart of the paper and is devoted to our results concerning soundness, completeness, decidability, and the finite model property for several logics of evidence, belief, and knowledge (Sects. 6.3–6.5). In order to keep this section self-contained and fix some notation, we first recapitulate, in a concise way, the formal syntax and semantics for the notions presented in the previous sections.

6.1 Logics for evidence, justified belief, and knowledge

Syntax The full language \mathcal{L} of evidence, belief, and knowledge we consider is defined recursively by the grammar

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid E_0\varphi \mid E\varphi \mid \Box_0\varphi \mid \Box\varphi \mid B\varphi \mid B^\varphi\varphi \mid K\varphi \mid [\forall]\varphi$$

where $p \in \text{Prop}$. We employ the usual abbreviations for propositional connectives \top , \perp , \vee , \rightarrow , \leftrightarrow , and for the dual modalities \hat{B} , \hat{K} , \hat{E} etc. except that some of them have special abbreviations: $[\exists]\varphi := \neg[\forall]\neg\varphi$ and $\diamond\varphi := \neg\Box\neg\varphi$. We will follow the usual rules for the elimination of the parentheses. Several fragments of the language \mathcal{L} is of particular interest: \mathcal{L}_B the fragment having the belief modality B as the only modality; \mathcal{L}_K having only the knowledge modality K ; and some bimodal fragments such as \mathcal{L}_{KB} having only operators K and B ; $\mathcal{L}_{[\forall]K}$ having only operators $[\forall]$ and K ; and the trimodal fragment $\mathcal{L}_{[\forall]\Box_0\Box}$ having only the modalities $[\forall]$, \Box_0 , and \Box .

Semantics We interpret the language \mathcal{L} on topo-e-models in an obvious way, following the definitions of the corresponding operators provided in previous sections.

Definition 9 (Topo-e-Semantics for \mathcal{L}) Given a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, we extend the valuation map V to an interpretation map $\llbracket \cdot \rrbracket^{\mathfrak{M}} : \mathcal{L} \rightarrow \mathcal{P}(X)$ recursively as follows:¹⁷

$$\begin{aligned} \llbracket p \rrbracket^{\mathfrak{M}} &= V(p) \\ \llbracket \neg\varphi \rrbracket^{\mathfrak{M}} &= X \setminus \llbracket \varphi \rrbracket^{\mathfrak{M}} \\ \llbracket \varphi \wedge \psi \rrbracket^{\mathfrak{M}} &= \llbracket \varphi \rrbracket^{\mathfrak{M}} \cap \llbracket \psi \rrbracket^{\mathfrak{M}} \\ \llbracket E_0\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists e \in \mathcal{E}_0 (e \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}})\} \\ \llbracket E\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists e \in \mathcal{E} (e \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}})\} \\ \llbracket \Box_0\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists e \in \mathcal{E}_0 (x \in e \text{ and } e \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}})\} \\ \llbracket \Box\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists U \in \tau (x \in U \text{ and } U \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}})\} \\ \llbracket B\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists U \in \tau (U \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}} \text{ and } Cl(U) = X)\} \\ \llbracket B^\varphi\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists U \in \tau (\emptyset \neq U \cap \llbracket \theta \rrbracket^{\mathfrak{M}} \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}} \text{ and } Cl(U \cap \llbracket \theta \rrbracket^{\mathfrak{M}}) \supseteq \llbracket \theta \rrbracket^{\mathfrak{M}})\} \\ \llbracket K\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \exists U \in \tau (x \in U \subseteq \llbracket \varphi \rrbracket^{\mathfrak{M}} \text{ and } Cl(U) = X)\} \\ \llbracket [\forall]\varphi \rrbracket^{\mathfrak{M}} &= \{x \in X \mid \llbracket \varphi \rrbracket^{\mathfrak{M}} = X\} \end{aligned}$$

We omit the superscript \mathfrak{M} when the model is contextually clear. Given a $\Gamma \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$, we say that φ is a *logical consequence* of Γ , denoted by $\Gamma \models \varphi$, iff for all topo-e-models $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ and all $x \in X$: if $x \in \llbracket \psi \rrbracket$ for all $\psi \in \Gamma$, then $x \in \llbracket \varphi \rrbracket$. As a special case, *validity*, $\models \varphi$, is truth at all worlds of all topo-e-models. φ is called *invalid*, denoted by $\not\models \varphi$, if it is not a validity, that is, if there is a topo-e-model $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ and a possible world $x \in X$ such that $x \notin \llbracket \varphi \rrbracket$. We say that a formula φ is *valid in a topo-e-model* $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$, denoted by $\mathfrak{M} \models \varphi$, if $\llbracket \varphi \rrbracket = X$. *Soundness* and *completeness* with respect to topo-e-models are defined standardly (see, e.g., Blackburn et al., 2001, Chapter 4.1).

¹⁷ We remind the reader not to confuse \exists and $[\exists]$: while we use the former to abbreviate “there exists” in the metalanguage, the latter is the existential modality defined as $\neg[\forall]\neg$ in the object language \mathcal{L} . Similarly for \forall and $[\forall]$: the former abbreviates “for all” in the metalanguage and the latter is the global modality in the object language \mathcal{L} .

It is not hard to see that the above defined semantics for the modalities of \mathcal{L} corresponds exactly to the semantic operators given in Sects. 3–5: e.g. $\llbracket [\forall]\varphi \rrbracket = [\forall]\llbracket \varphi \rrbracket$, $\llbracket [\Box]\varphi \rrbracket = \Box\llbracket \varphi \rrbracket = \text{Int}(\llbracket \varphi \rrbracket)$, etc. Moreover, while all modalities except for E_0 and \Box_0 capture topological properties of topo-e-models, i.e., they can be interpreted directly in (X, τ, V) , a *topo-model*, the expressivity of the full language goes beyond the purely topological properties: the meaning of E_0 and \Box_0 does not only depend on the evidential topology, but also depends on the basic evidence set \mathcal{E}_0 . From the point of expressivity, the most important fragment of \mathcal{L} is the trimodal language $\mathcal{L}_{[\forall]\Box_0\Box}$ since it is equally expressive as the full language \mathcal{L} with respect to the topo-e-models:

Proposition 6 *The following equivalences are valid:*

1. $B\varphi \leftrightarrow [\forall]\Diamond\Box\varphi$
2. $E\varphi \leftrightarrow [\exists]\Box\varphi$
3. $E_0\varphi \leftrightarrow [\exists]\Box_0\varphi$
4. $K\varphi \leftrightarrow \Box\varphi \wedge [\forall]\Diamond\Box\varphi$
5. $B^\theta\varphi \leftrightarrow [\forall](\theta \rightarrow \Diamond(\theta \wedge \Box(\theta \rightarrow \varphi)))$

Proof The proof follows easily from the semantics clauses of the modalities given in Definition 9. □

Therefore, all the other modalities of \mathcal{L} can be defined in $\mathcal{L}_{[\forall]\Box_0\Box}$. For this reason, instead of focusing on the full language \mathcal{L} , we present soundness, completeness, and decidability results for the *factive evidence fragment* $\mathcal{L}_{[\forall]\Box_0\Box}$: its importance comes from its expressive power. We moreover provide sound and complete axiomatizations for the pure doxastic fragment \mathcal{L}_B , the pure epistemic fragments \mathcal{L}_K and $\mathcal{L}_{[\forall]K}$, and finally for the epistemic-doxastic fragment \mathcal{L}_{KB} . As the semantics of $[\forall]$, B , and K can be defined only based on the evidential topology (without referring to \mathcal{E}_0), we will state the corresponding soundness and completeness results simply with respect to topo-models¹⁸. Notions of *validity in a topo-model*, *validity*, *soundness*, and *completeness wrt topo-models* are defined standardly, similarly to those for topo-e-models. For $\mathcal{L}_{[\forall]\Box_0\Box}$, we need the complete structure of the topo-e-models as the semantics of \Box_0 depends on the basic evidence set \mathcal{E}_0 and cannot be recovered purely topologically.

Before moving on to the technical results, we briefly recall the following standard terminology of Hilbert-style axiom systems and set some notation. Given a logic L defined by a finitary Hilbert-style axiom system, an *L-derivation/proof* is a finite sequence of formulas such that each element of the sequence is either an axiom of L , or obtained from the previous formulas in the sequence by one of the inference rules. A formula φ is called *L-provable*, or, equivalently, a *theorem* of L , if it is the last formula of some L -proof. In this case, we write $\vdash_L \varphi$ (or, equivalently, $\varphi \in L$). For any set of formulas Γ and any formula φ , we write $\Gamma \vdash_L \varphi$ if there exist finitely many formulas $\varphi_1, \dots, \varphi_n \in \Gamma$ such that $\vdash_L \varphi_1 \wedge \dots \wedge \varphi_n \rightarrow \varphi$. We say that Γ is *L-consistent* if $\Gamma \not\vdash_L \perp$, and *L-inconsistent* otherwise. A formula φ is *consistent with* Γ if $\Gamma \cup \{\varphi\}$ is L -consistent (or, equivalently, if $\Gamma \not\vdash_L \neg\varphi$). Finally, a set of formulas Γ is *maximally consistent* if it is L -consistent and any set of formulas properly containing Γ is L -inconsistent, i.e. Γ cannot be extended to another L -consistent set. We drop mention of the logic L when it is clear from the context.

For the axiomatizations of the well-known normal unimodal logics, we refer to (Blackburn et al., 2001, Chapter 4). Here we use the standard naming conventions for these logics

¹⁸ A topo-model is a tuple (X, τ, V) , where (X, τ) is a topological space and $V : \text{Prop} \rightarrow \mathcal{P}(X)$ is a standardly defined valuation function. The semantic clauses of the language \mathcal{L} minus the components E_0 and \Box_0 in topo-models are as given in Definition 9.

and the relevant axioms, and add the axiomatized operator as a subscript to their names: e.g., the normal modal logic KD45 for B is denoted by $KD45_B$, S5 for $[\forall]$ by $S5_{[\forall]}$, and axiom $(.2)$ for K by $(.2_K)$. We provide the axioms of our multi-modal logics in tables in the relevant sections. Finally, $L+(\varphi)$ denotes the smallest modal logic containing L and φ , i.e., $L+(\varphi)$ is the smallest set of formulas (in the corresponding language) that contains L and φ , and is closed under the inference rules of L .

6.2 The knowledge fragments \mathcal{L}_K and $\mathcal{L}_{[\forall]K}$: S4.2 $_K$ and $L_{[\forall]K}$

In this section, we focus on the two knowledge fragments \mathcal{L}_K and $\mathcal{L}_{[\forall]K}$, and provide sound and complete axiomatizations for the associated logics. While the fragment having only the modality K leads to the familiar system S4.2 $_K$, the full knowledge fragment having both K and $[\forall]$ gives us the axiomatization $L_{[\forall]K}$ presented below, in Table 4.

The soundness results, as usual, are shown by proving that all axioms are validities and that all derivation rules preserve validities. These proofs are elementary for most axioms and derivation rules, we here show only the relatively trickier cases. For completeness of both S4.2 $_K$ and $L_{[\forall]K}$, we rely on their completeness wrt Kripke models and the connection between their Kripke models and topo-models explained in Sect. 6.2.1.

6.2.1 Soundness and completeness of S4.2 $_K$

The relatively harder case in the soundness proof of S4.2 $_K$ is the normality axiom (K_K) for the knowledge modality K , whose validity follows from the following lemma and the fact that the interior operator commutes with finite intersections (see Sect. 2).

Lemma 2 *Given a topological space (X, τ) and any two subsets $U_1, U_2 \subseteq X$, if U_1 is open dense and U_2 is dense, then $U_1 \cap U_2$ is dense.*

Proof Let (X, τ) be a topological space and $U_1, U_2 \subseteq X$. Suppose U_1 is an open dense and U_2 is a dense set in (X, τ) . Since U_1 is open and dense we have that $W \cap U_1$ is open and non-empty for any non-empty open set W . Thus, since U_2 is dense, we also have that $(W \cap U_1) \cap U_2 \neq \emptyset$. Therefore, $W \cap (U_1 \cap U_2) \neq \emptyset$ for any nonempty $W \in \tau$, i.e., $U_1 \cap U_2$ is dense as well. \square

For completeness, we rely on the completeness of S4.2 $_K$ wrt its Kripke models and their connection to topological models.

Connection between S4.2-frames and topological spaces Let (X, R) be a *transitive* Kripke frame. A nonempty subset $C \subseteq X$ is called *cluster* if (1) for each $x, y \in C$ we have xRy , and (2) there is no $D \subseteq X$ such that $C \subsetneq D$ and D satisfies (1). A point $x \in X$ is called a *maximal point* if there is no $y \in X$ such that xRy and $\neg(yRx)$. We call a cluster a *final cluster* if all its points are maximal. It is not hard to see that for any final cluster C of (X, R) and any $x \in C$, we have $R(x) = C$. A transitive Kripke frame (X, R) is called *cofinal* if it has a unique final cluster C such that for each $x \in X$ and $y \in C$ we have xRy .

Lemma 3 *S4.2 $_K$ is sound and complete with respect to the class of reflexive and transitive cofinal frames.*

Proof See, e.g., (Chagrov and Zakharyashev, 1997, Chapter 5). \square

As well-known, given a reflexive and transitive Kripke frame (X, R) , we can construct an Alexandroff space¹⁹ (X, τ_R) by defining τ_R to be the set of all upsets²⁰ of (X, R) (see, e.g., van Benthem and Bezhanishvili, 2007, Sect. 2).

Lemma 4 *For every reflexive transitive cofinal frame (X, R) and nonempty $U \in \tau_R$, we have $Cl(U) = X$ in (X, τ_R) .*

Proof Let (X, R) be a reflexive and transitive cofinal frame and let $C \subseteq X$ denote its final cluster. By construction, $C \in \tau_R$ and moreover $C \subseteq U$, for all nonempty $U \in \tau_R$. Therefore, for every nonempty $U, V \in \tau_R$, we have $V \cap U \supseteq C \neq \emptyset$. Hence, $Cl(U) = X$ for any nonempty $U \in \tau_R$.

Given a reflexive and transitive Kripke model $\mathcal{M} = (X, R, V)$, let $B(\mathcal{M}) = (X, \tau_R, V)$ denote the corresponding topo-model (where τ_R is the set of all upsets of (X, R)). For any formula φ in the relevant object language, $\|\varphi\|^{\mathcal{M}}$ denotes the set of worlds in $\mathcal{M} = (X, R, V)$ that make φ true with respect to the standard Kripke semantics (where $[\forall]$ is interpreted as the global modality).

Proposition 7 *For every reflexive and transitive cofinal Kripke model $\mathcal{M} = (X, R, V)$ and all $\varphi \in \mathcal{L}_{[\forall]K}$,*

$$\|\varphi\|^{\mathcal{M}} = \llbracket \varphi \rrbracket^{B(\mathcal{M})},$$

where $B(\mathcal{M}) = (X, \tau_R, V)$.

Proof The proof follows by subformula induction on φ ; cases for the propositional variables, the Boolean connectives and the modality $[\forall]$ are elementary. So assume inductively that the result holds for ψ ; we must show that it holds also for $\varphi := K\psi$. Let $\mathcal{M} = (X, R, V)$ be a reflexive and transitive cofinal Kripke model and $x \in X$.

(\subseteq) Suppose $x \in \|K\psi\|^{\mathcal{M}}$. This implies that $x \in R(x) \subseteq \|\psi\|^{\mathcal{M}}$. By induction hypothesis, we obtain $R(x) \subseteq \llbracket \psi \rrbracket^{B(\mathcal{M})}$. Since $x \in R(x) \in \tau_R$, we have $x \in Int(\llbracket \psi \rrbracket^{B(\mathcal{M})})$. Then, by Lemma 4, $Cl(Int(\llbracket \psi \rrbracket^{B(\mathcal{M})})) = X$. Therefore, $x \in \llbracket K\psi \rrbracket^{B(\mathcal{M})}$.

(\supseteq) Suppose $x \in \llbracket K\psi \rrbracket^{B(\mathcal{M})}$. This means, by the topological semantics of K , that $x \in Int(\llbracket \psi \rrbracket^{B(\mathcal{M})})$ and that $Cl(Int(\llbracket \psi \rrbracket^{B(\mathcal{M})})) = X$. Then, by induction hypothesis, $x \in Int(\|\psi\|^{\mathcal{M}})$ and $Cl(Int(\|\psi\|^{\mathcal{M}})) = X$. The former implies that there is an open set $U \in \tau_R$ such that $x \in U \subseteq \|\psi\|^{\mathcal{M}}$. In particular, since $R(x)$ is the smallest open neighbourhood of x , we obtain $R(x) \subseteq \|\psi\|^{\mathcal{M}}$. Therefore, $x \in \|K\psi\|^{\mathcal{M}}$. \square

Theorem 1 *S4.2_K is sound and complete with respect to the class of all topo-models.*

Proof For completeness, let $\varphi \in \mathcal{L}_K$ such that $\varphi \notin S4.2_K$. Then, by Lemma 3, there exists a Kripke model $\mathcal{M} = (X, R, V)$ based on the reflexive and transitive cofinal frame (X, R) such that $\|\varphi\|^{\mathcal{M}} \neq X$. Thus, by Proposition 7, we have $\llbracket \varphi \rrbracket^{B(\mathcal{M})} \neq X$, where $B(\mathcal{M}) = (X, \tau_R, V)$ is the corresponding topo-model. \square

6.2.2 Soundness and completeness of $\mathcal{L}_{[\forall]K}$

The full knowledge fragment $\mathcal{L}_{[\forall]K}$ having both K and $[\forall]$ yields the axiomatic system $\mathcal{L}_{[\forall]K}$ given in Table 4.

¹⁹ A topological space (X, τ) is an Alexandroff space if τ is closed under arbitrary intersections, i.e., $\bigcap \mathcal{A} \in \tau$ for all $\mathcal{A} \subseteq \tau$.

²⁰ $A \subseteq X$ is called an upset of (X, R) if for each $x, y \in X, xRy$ and $x \in A$ imply $y \in A$.

Table 4 The axiomatization of $L_{[\forall]K}$

(CPL)	All classical propositional tautologies and (MP)
(S5 $_{[\forall]}$)	All S5 axioms and rules for the modality $[\forall]$
(S4 $_K$)	All S4 axioms and rules for the modality K
(Ax1)	$[\forall]\varphi \rightarrow K\varphi$
(Ax2)	$[\exists]K\varphi \rightarrow [\forall]\hat{K}\varphi$

Theorem 2 $L_{[\forall]K}$ is sound and complete with respect to the class of all topo-models.

Proof Soundness is easy to see, we here only prove that the axiom $([\exists]K\varphi \rightarrow [\forall]\hat{K}\varphi)$ is valid in all topo-models. Let $\mathfrak{M} = (X, \tau, V)$ be a topo-model, $\varphi \in \mathcal{L}_{[\forall]K}$, and $x \in X$ such that $x \in \llbracket [\exists]K\varphi \rrbracket$. This means that there exist $y \in X$ such that $y \in \text{Int}(\llbracket \varphi \rrbracket)$ and $Cl(\text{Int}(\llbracket \varphi \rrbracket)) = X$. Note that for any $z \in X$,

$$z \in \llbracket \hat{K}\varphi \rrbracket \text{ iff } z \notin \text{Int}(\llbracket \neg\varphi \rrbracket) \text{ or } Cl(\text{Int}(\llbracket \neg\varphi \rrbracket)) \neq X,$$

(see Proposition 4-3.). Therefore, in order to show $\llbracket \hat{K}\varphi \rrbracket = X$, it suffices to show that $Cl(\text{Int}(\llbracket \neg\varphi \rrbracket)) \neq X$. Since $y \in \text{Int}(\llbracket \varphi \rrbracket)$, we know that $\text{Int}(Cl(\llbracket \varphi \rrbracket)) \neq \emptyset$ (as $\text{Int}(\llbracket \varphi \rrbracket) \subseteq \text{Int}(Cl(\llbracket \varphi \rrbracket))$). Hence, $Cl(\text{Int}(\llbracket \neg\varphi \rrbracket)) \neq X$. We therefore obtain that $\llbracket \hat{K}\varphi \rrbracket = X$, hence, $[\forall]\hat{K}\varphi$ holds everywhere in \mathfrak{M} .

For completeness, we use a well-known Kripke completeness result for the logic obtained by extending S4.2 $_K$ with the universal modality $[\forall]$. More precisely, it has been shown in (Goranko & Passy, 1992) that the modal system $L_{[\forall]K}^0 := S5_{[\forall]} + S4.2_K + ([\forall]\varphi \rightarrow K\varphi)$, simply obtained by replacing (Ax2) in Table 4 by the axiom $(.2_K) := \hat{K}K\varphi \rightarrow K\hat{K}\varphi$, is complete with respect to the class of reflexive and transitive cofinal Kripke frames when K is interpreted as the standard Kripke modality and $[\forall]$ as the global modality. It is not hard to see that the axiom $(.2_K)$ is derivable in $L_{[\forall]K}$ (by using Ax1 and Ax2 in Table 4), hence, $L_{[\forall]K}$ is stronger than $L_{[\forall]K}^0$, i.e., that $L_{[\forall]K}^0 \subseteq L_{[\forall]K}$. Let $\varphi \in \mathcal{L}_{[\forall]K}$ such that $\varphi \notin L_{[\forall]K}$. Thus, $\varphi \notin L_{[\forall]K}^0$. Then, by the relational completeness of $L_{[\forall]K}^0$, there exists a reflexive and transitive cofinal Kripke model $\mathcal{M} = (X, R, V)$ such that $\|\varphi\|^{\mathcal{M}} \neq X$. Then, by Proposition 7, we obtain $\llbracket \varphi \rrbracket^{B(\mathcal{M})} \neq X$, where $B(\mathcal{M}) = (X, \tau_R, V)$. \square

6.3 The belief fragment \mathcal{L}_B : KD45 $_B$


In this section, we prove that the logic of belief on all topo-models is the standard belief system KD45 $_B$, and it moreover has the finite model property with respect to the class of topo-models.

6.3.1 Soundness of KD45 $_B$

Proposition 8 KD45 $_B$ is sound with respect to the class of all topo-models.

Proof The soundness, as usual, is shown by proving that all axioms are validities and that all derivation rules preserve validities. The cases for the axioms (4 $_B$) and (5 $_B$) and the inference rules are elementary, whereas the validity of (K $_B$) in the class of all topological

Fig. 8 An example of a brush and of a pin, where the top ellipses illustrate the final clusters and an arrow relates the state it started from to every element in the cluster



spaces follows from Lemma 2 as follows. Let $\mathfrak{M} = (X, \mathcal{E}_0, \tau, V)$ and $\varphi, \psi \in \mathcal{L}_B$. We need to show that $\llbracket B(\varphi \wedge \psi) \leftrightarrow B\varphi \wedge B\psi \rrbracket = X$, i.e., that $\llbracket B(\varphi \wedge \psi) \rrbracket = \llbracket B\varphi \wedge B\psi \rrbracket$. Let $x \in B(\varphi \wedge \psi)$. This implies, by the semantics of B that $\llbracket B(\varphi \wedge \psi) \rrbracket = X$, i.e., $Cl(Int(\llbracket \varphi \wedge \psi \rrbracket)) = X$. We therefore obtain that $X = Cl(Int(\llbracket \varphi \wedge \psi \rrbracket)) = Cl(Int(\llbracket \varphi \rrbracket) \cap Int(\llbracket \psi \rrbracket)) \subseteq Cl(Int(\llbracket \varphi \rrbracket)) \cap Cl(Int(\llbracket \psi \rrbracket)) = \llbracket B\varphi \wedge B\psi \rrbracket$. For the other direction, suppose $x \in \llbracket B\varphi \wedge B\psi \rrbracket$. We therefore have $x \in \llbracket B\varphi \rrbracket$ and $x \in \llbracket B\psi \rrbracket$. Then, by the semantics of B , we obtain that $Cl(Int(\llbracket \varphi \rrbracket)) = X$ and $Cl(Int(\llbracket \psi \rrbracket)) = X$. This means that both $Int(\llbracket \varphi \rrbracket)$ and $Int(\llbracket \psi \rrbracket)$ are dense in (X, τ) . Hence, by Lemma 2, we obtain $Cl(Int(\llbracket \varphi \rrbracket) \cap Int(\llbracket \psi \rrbracket)) = X$. Similarly to the argument above, we then have $X = Cl(Int(\llbracket \varphi \rrbracket) \cap Int(\llbracket \psi \rrbracket)) = Cl(Int(\llbracket \varphi \wedge \psi \rrbracket)) = \llbracket B(\varphi \wedge \psi) \rrbracket$. \square

6.3.2 Completeness of $KD45_B$

For completeness, we use the following connection between the $KD45$ -Kripke frames and topological spaces.

Connection between $KD45$ -frames and topological spaces. Recall that $KD45$ -frames are serial, transitive and Euclidean Kripke frames. Since truth of modal formulas with respect to the standard relational semantics is preserved under taking generated submodels (see, e.g., Blackburn et al., 2001, Proposition 2.6), we can use the following simplified relational structures as Kripke frames of $KD45_B$.

Definition 10 (*Brush/pin*)

- A relational frame (X, R) is called a *brush* if there exists a nonempty subset $C \subseteq X$ such that $R = X \times C$;
- A brush is called a *pin* if $|X \setminus C| = 1$.

Clearly, if such a C exists, it is unique and it is the *final cluster* of the brush. It is easy to see that every brush is serial, transitive, and Euclidean (see Fig. 8). For the proof of the following lemma see, e.g., (Chagrov and Zakharyashev, 1997, Chapter 5) and (Blackburn et al., 2001, Chapters 2, 4).

Lemma 5 $KD45_B$ is a sound and complete with respect to the class of brushes, and with respect to the class of pins. In fact, $KD45_B$ is sound and complete with respect to the class of finite pins.

We can build a topological space from a given pin. For any frame (X, R) , let R^+ denote the reflexive closure of R , defined as

$$R^+ = R \cup \{(x, x) \mid x \in X\}.$$




Fig. 9 From pins to topological spaces

Given a pin (X, R) , the set $\tau_{R^+} = \{R^+(x) \mid x \in X\}$ constitutes a topology on X . In fact, in this special case of pins, we have $\tau_{R^+} = \{X, C, \emptyset\}$ where C is the final unique cluster of (X, R) . Therefore, it is easy to see that (X, τ_{R^+}) is a topological space.²¹ In fact, (X, τ_{R^+}) is a generalized Sierpiński space where C does not have to be a singleton (see Fig. 9).

This construction leads to a natural correspondence between pins and topological spaces for the language \mathcal{L}_B . In particular, for any Kripke model $\mathcal{M} = (W, R, V)$ based on a pin, we set $I(\mathcal{M}) = (X, \tau_{R^+}, V)$. Moreover, any two such models \mathcal{M} and $I(\mathcal{M})$ make the same formulas of \mathcal{L}_B true at the same states, as shown in Proposition 9.

Proposition 9 For all $\varphi \in \mathcal{L}_B$ and any Kripke model $\mathcal{M} = (W, R, V)$ based on a pin,

$$\|\varphi\|^{\mathcal{M}} = \llbracket \varphi \rrbracket^{I(\mathcal{M})}.$$

Proof The proof follows by subformula induction on φ ; cases for the propositional variables and the Boolean connectives are elementary. So assume inductively that the result holds for ψ ; we must show that it holds also for $\varphi := B\psi$. Observe that, given a Kripke model $\mathcal{M} = (W, R, V)$ based on a pin (X, R) , we have

$$\|B\psi\|^{\mathcal{M}} = \begin{cases} X & \text{if } \|\psi\|^{\mathcal{M}} \supseteq C \\ \emptyset & \text{otherwise} \end{cases} \quad \text{and} \quad \llbracket B\psi \rrbracket^{I(\mathcal{M})} = \begin{cases} X & \text{if } \llbracket \psi \rrbracket^{I(\mathcal{M})} \supseteq C \\ \emptyset & \text{otherwise} \end{cases}$$

where C is the final cluster of (X, R) . By induction hypothesis, we have $\llbracket \psi \rrbracket^{I(\mathcal{M})} = \|\psi\|^{\mathcal{M}}$, therefore, $\llbracket B\psi \rrbracket^{I(\mathcal{M})} = \|B\psi\|^{\mathcal{M}}$. □

Theorem 3 $KD45_B$ is sound and complete with respect to the class of all topo-models. Moreover, $KD45_B$ has the finite model property.

Proof Soundness is given in Proposition 8. For completeness, let $\varphi \in \mathcal{L}_B$ such that $\varphi \notin KD45_B$. Then, by Lemma 5, there exists a finite pin $\mathcal{M} = (X, R, V)$ with $\|\varphi\|^{\mathcal{M}} \neq X$. Thus, by Proposition 9, we have that $\llbracket \varphi \rrbracket^{I(\mathcal{M})} \neq X$, where $I(\mathcal{M}) = (X, \tau_{R^+}, V)$ is the corresponding topo-model. Since $I(\mathcal{M}) = (X, \tau_{R^+}, V)$ is finite, we have also shown that $KD45_B$ has the finite model property. □

²¹ τ_{R^+} is in fact an Alexandroff hereditarily extremally disconnected space. However, these extra properties are not of interest in this paper.

Table 5 Stalnaker’s knowledge and belief logic *Stal*

(CPL)	All class. prop. taut. and (MP)	
(S4 _K)	All S4 axioms and rules for <i>K</i>	
(D _B)	$B\varphi \rightarrow \neg B\neg\varphi$	Consistency of belief
(sPI)	$B\varphi \rightarrow K B\varphi$	Strong positive introspection
(sNI)	$\neg B\varphi \rightarrow K\neg B\varphi$	Strong negative introspection
(KB)	$K\varphi \rightarrow B\varphi$	Knowledge implies belief
(FB)	$B\varphi \rightarrow BK\varphi$	Full belief

6.4 The knowledge-belief fragment \mathcal{L}_{KB} : *Stal* revisited

In this section, we show that Stalnaker’s system *Stal* of knowledge and belief, given in Table 5, is sound and complete with respect to the class of all topo-models under our proposed semantics for knowledge and belief.²² As noted in the introduction, in previous work Baltag et al. (2013, 2019b), we provided a topological completeness result for this system for the restricted class of extremally disconnected spaces. Therefore, we here show that the topological semantics presented in this paper generalizes the one proposed in Baltag et al. (2013, 2019b) for Stalnaker’s combined system *Stal*.

Theorem 4 *Stal* is sound and complete with respect to the class of all topo-models.

Proof For soundness, we here only show the validity of the axiom (FB): the validity proofs of the other axioms are either trivial or follow from the previous results. Let $\mathfrak{M} = (X, \tau, V)$ be a topo-model, $\varphi \in \mathcal{L}_{KB}$ and $x \in X$. Suppose $x \in \llbracket B\varphi \rrbracket$. Hence, $\llbracket B\varphi \rrbracket \neq \emptyset$. This implies, by the semantics of *B*, that $\llbracket B\varphi \rrbracket = Cl(Int(\llbracket \varphi \rrbracket)) = X$. Recall that $x \in \llbracket K\varphi \rrbracket$ iff $x \in Int(\llbracket \varphi \rrbracket)$ and $Cl(Int(\llbracket \varphi \rrbracket)) = X$. By the assumption, we already know that $Cl(Int(\llbracket \varphi \rrbracket)) = X$. Thus, in this particular case, $\llbracket K\varphi \rrbracket = Int(\llbracket \varphi \rrbracket)$. Therefore, $X = Cl(Int(\llbracket \varphi \rrbracket)) = Cl(Int(Int(\llbracket \varphi \rrbracket))) = Cl(Int(\llbracket K\varphi \rrbracket))$ implying that $BK\varphi$ holds everywhere in \mathfrak{M} .

For completeness, we follow a similar method as in the proof of Theorem 2. Let $\varphi \in \mathcal{L}_{KB}$ such that $\varphi \notin \text{Stal}$. Then, since $\vdash_{\text{Stal}} B\varphi \leftrightarrow \hat{K}K\varphi$, there exists a $\psi \in \mathcal{L}_K$ such that $\vdash_{\text{Stal}} \varphi \leftrightarrow \psi$ (this is obtained by replacing every occurrence of *B* in φ by $\hat{K}K$). Therefore, $\psi \notin \text{Stal}$. Moreover, since $S4.2_K \subseteq \text{Stal}$, we obtain $\psi \notin S4.2_K$. Then, by Theorem 1, there exists a topo-model $\mathfrak{M} = (X, \tau, V)$ such that $\llbracket \psi \rrbracket \neq X$. Since *Stal* is sound with respect to all topo-models and $\vdash_{\text{Stal}} \varphi \leftrightarrow \psi$, we conclude $\llbracket \varphi \rrbracket \neq X$. □

6.5 The factive evidence fragment $\mathcal{L}_{\llbracket \vee \rrbracket \square_0 \square}$: $\text{Log}_{\llbracket \vee \rrbracket \square_0 \square}$

The logic $\text{Log}_{\llbracket \vee \rrbracket \square_0 \square}$ of factive evidence is given by the axiom schemas and inference rules in Table 6 over the language $\mathcal{L}_{\llbracket \vee \rrbracket \square_0 \square}$.

This section presents the proof of Theorem 5. Strong completeness and strong finite model property are defined standardly (see, e.g., Blackburn et al., 2001, Definition 4.10-Proposition 4.12 and Definition 6.6, respectively).

²² What justifies the properties of knowledge and belief stated in *Stal* may be debatable, though not in the scope of this paper. We refer to (Bjorndahl & Özgün, 2019) for a topological-based reformulation of Stalnaker’s system.

Table 6 The axiomatization of $\text{Log}_{\forall\Box\Box_0}$

(CPL)	All classical propositional tautologies and (MP)
(S5 $_{[\forall]}$)	All S5 axioms and rules for the modality $[\forall]$
(S4 $_{\Box}$)	All S4 axioms and rules for the modality \Box
(4 $_{\Box_0}$)	$\Box_0\varphi \rightarrow \Box_0\Box_0\varphi$
Universality (U)	$[\forall]\varphi \rightarrow \Box_0\varphi$
Factive Evidence (FE)	$\Box_0\varphi \rightarrow \Box\varphi$
Pullout ^a	$(\Box_0\varphi \wedge [\forall]\psi) \rightarrow \Box_0(\varphi \wedge [\forall]\psi)$
Monotonicity rule for \Box_0	from $\varphi \rightarrow \psi$, infer $\Box_0\varphi \rightarrow \Box_0\psi$

^a This axiom originates in (van Benthem et al., 2012, 2014), where it is stated as an equivalence rather than an implication. But the converse is provable in our system from the Monotonicity rule for \Box_0 , (FE), and S4 $_{\Box}$

Theorem 5 *The logic $\text{Log}_{\forall\Box\Box_0}$ of factive evidence is sound and strongly complete with respect to the class of all topo-models. Moreover, it has the strong finite model property, therefore, it is decidable.*

The proof of Theorem 5 is technically the most challenging result of this paper. The key difficulty consists in guaranteeing that the natural topology for which \Box acts as interior operator is exactly the topology generated by the neighborhood family associated to \Box_0 . Though the main steps of the proof may look familiar, involving known methods (a canonical quasi-model construction, a filtration argument, and then making multiple copies of the worlds to yield a finite model with the right properties), addressing the above-mentioned difficulty requires a non-standard application of these methods, as well as a number of additional notions and results, and a careful treatment of each of the steps. The plan of the proof is as follows. Since the soundness proof is straightforward, we here focus on completeness and the finite model property (then decidability follows immediately). We first prove strong completeness of $\text{Log}_{\forall\Box\Box_0}$ with respect to a *canonical quasi-model*. We then continue with proving the strong finite quasi-model property for $\text{Log}_{\forall\Box\Box_0}$ via a filtration argument. In the last step, we prove that every finite quasi-model is equivalent to a finite Alexandroff quasi-model by making multiple copies of the worlds in order to put the model in the right shape. As Alexandroff quasi-models are modally equivalent to Alexandroff topo-e-models (Proposition 11), the result follows.

6.5.1 Quasi-model construction

A *quasi-model* is a tuple $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$, where (X, \mathcal{E}_0, V) is an evidence model and \leq is a preorder such that every $e \in \mathcal{E}_0$ is an upset of (X, \leq) (see footnote 20 to recall the definition of an upset). Given a preordered set (X, \leq) , the set $Up_{\leq}(X)$ denotes the *set of all upsets of (X, \leq)* . We denote by $\uparrow x = \{y \in X \mid x \leq y\}$ the *upset generated by x* . We use the same notations as for topo-e-models, for example, \mathcal{E} for the closure of \mathcal{E}_0 under nonempty finite intersections, and $\tau_{\mathcal{E}}$ for the topology generated by \mathcal{E} .

The *semantics* for the language $\mathcal{L}_{[\forall]\Box\Box}$ on quasi-models is defined the same way as on topo-e-models (see Definition 9), *except* that for \Box we (do *not* use the topology, but instead we) use the standard Kripke semantics based on the relation \leq . More precisely, the

semantics for the modalities $[\forall]$, \Box_0 , and \Box are given by the following clauses:

$$\begin{aligned} \llbracket [\forall]\varphi \rrbracket^{\mathcal{M}} &= \{x \in X \mid \llbracket \varphi \rrbracket^{\mathcal{M}} = X\} \\ \llbracket \Box_0\varphi \rrbracket^{\mathcal{M}} &= \{x \in X \mid \exists e \in \mathcal{E}_0 (x \in e \text{ and } e \subseteq \llbracket \varphi \rrbracket^{\mathcal{M}})\} \\ \llbracket \Box\varphi \rrbracket^{\mathcal{M}} &= \{x \in X \mid \forall y \in X (x \leq y \text{ implies } y \in \llbracket \varphi \rrbracket^{\mathcal{M}})\} \end{aligned}$$

We again omit the superscripts for the model when it is clear from the context.

A quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ is called *Alexandroff* if the topology $\tau_{\mathcal{E}}$ is Alexandroff and \leq is the specialization preorder $\sqsubseteq_{\mathcal{E}}$ on X , where

$$x \sqsubseteq_{\mathcal{E}} y \text{ iff } x \in Cl(\{y\}) \text{ (i.e., } \forall U \in \tau_{\mathcal{E}} (x \in U \text{ implies } y \in U)).$$

Proposition 10 *For every quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ the following are equivalent:*

1. \mathcal{M} is Alexandroff;
2. $\tau_{\mathcal{E}} = Up_{\leq}(X)$;
3. for every $x \in X$, $\uparrow x$ is in $\tau_{\mathcal{E}}$.

Proof (1) \Rightarrow (3): Suppose \mathcal{M} is Alexandroff, i.e., $\tau_{\mathcal{E}}$ is Alexandroff and $\leq = \sqsubseteq_{\mathcal{E}}$. Let $x \in X$. Then we have: $\uparrow x = \{y \in X \mid x \leq y\} = \{y \in X \mid x \sqsubseteq_{\mathcal{E}} y\} = \{y \in X \mid \forall U \in \tau_{\mathcal{E}} (x \in U \Rightarrow y \in U)\} = \bigcap \{U \in \tau_{\mathcal{E}} \mid x \in U\}$. Since $\tau_{\mathcal{E}}$ is an Alexandroff space, we have $\bigcap \{U \in \tau_{\mathcal{E}} \mid x \in U\} \in \tau_{\mathcal{E}}$, and hence $\uparrow x = \bigcap \{U \in \tau_{\mathcal{E}} \mid x \in U\} \in \tau_{\mathcal{E}}$.

(3) \Rightarrow (2): It is easy to see that $\tau_{\mathcal{E}} \subseteq Up_{\leq}(X)$ (since $\tau_{\mathcal{E}}$ is generated by \mathcal{E}_0 and every element of \mathcal{E}_0 is an upset of (X, \leq)). Now let $A \in Up_{\leq}(X)$. Since A is an upset of (X, \leq) , we have $A = \bigcup \{\uparrow x \mid x \in A\}$. Then, by (3) (and $\tau_{\mathcal{E}}$ being closed under arbitrary unions), we obtain $A \in \tau_{\mathcal{E}}$.

(2) \Rightarrow (1): Suppose $\tau_{\mathcal{E}} = Up_{\leq}(X)$ and let $\mathcal{A} \subseteq \tau_{\mathcal{E}}$. By (2), every $U \in \mathcal{A}$ is an upset wrt of (X, \leq) , hence, $\bigcap \mathcal{A}$ is an upset as well. Therefore, by (2), $\bigcap \mathcal{A} \in \tau_{\mathcal{E}}$. This proves that $\tau_{\mathcal{E}}$ is Alexandroff. (2) also implies that $\uparrow x$ is the least open neighbourhood of x in $\tau_{\mathcal{E}}$, i.e., $\uparrow x \subseteq U$, for all U such that $x \in U \in \tau_{\mathcal{E}}$. Therefore, \leq is included in $\sqsubseteq_{\mathcal{E}}$. For the other direction, suppose $x \sqsubseteq_{\mathcal{E}} y$. This implies, in particular, that $y \in \uparrow x$ (since $x \in \uparrow x \in \tau_{\mathcal{E}}$), i.e., $x \leq y$. □

There is a natural one-to-one correspondence between Alexandroff quasi-models and Alexandroff topo-e-models, given by putting, for any Alexandroff quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$, $B(\mathcal{M}) = (X, \mathcal{E}_0, \tau_{\mathcal{E}}, V)$. Moreover, \mathcal{M} and $B(\mathcal{M})$ satisfy the same formulas of $\mathcal{L}_{[\forall]\Box_0\Box}$ at the same points, as shown in Proposition 11 below.

Proposition 11 *For all $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box}$ and every Alexandroff quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$, we have*

$$\llbracket \varphi \rrbracket^{\mathcal{M}} = \llbracket \varphi \rrbracket^{B(\mathcal{M})}.$$

Proof The proof follows by subformula induction on φ ; cases for the propositional variables, the Boolean connectives, and the modalities $[\forall]$ and \Box_0 are trivial as the semantics for these cases are defined exactly the same way in both structures. For the modality \Box , recall that it is interpreted as the interior operator of the topology $\tau_{\mathcal{E}}$ and use Proposition 10. □

Therefore, as stated by Proposition 11, Alexandroff quasi-models provide just another presentation of Alexandroff topo-e-models with respect to the language $\mathcal{L}_{[\forall]\Box_0\Box}$.

Having introduced the auxiliary notions and facts, we are ready to prove Theorem 5. This proof goes through *three steps*:

1. strong completeness for quasi-models;
2. strong finite quasi-model property; and
3. every finite quasi-model is modally equivalent to a finite Alexandroff quasi-model (hence, to a topo-e-model).

Step 1: strong completeness for quasi-models The proof follows via a canonical quasi-model construction.

Lemma 6 (Lindenbaum’s Lemma) *Every $\text{Log}_{\forall\Box_0\Box}$ -consistent set can be extended to a maximally consistent one.*

Let us now fix a consistent set of sentence $\Phi_0 \subseteq \mathcal{L}_{[\forall]\Box_0\Box}$. Our goal is to construct a quasi-model for Φ_0 . By Lemma 6, there exists a maximally consistent set T_0 such that $\Phi_0 \subseteq T_0$. For any two maximally consistent sets T and S of $\text{Log}_{\forall\Box_0\Box}$, we put:

$$T \sim S \text{ iff for all } \varphi \in \mathcal{L}_{[\forall]\Box_0\Box} : ([\forall]\varphi \in T \text{ implies } \varphi \in S),$$

$$T \leq S \text{ iff for all } \varphi \in \mathcal{L}_{[\forall]\Box_0\Box} : (\Box\varphi \in T \text{ implies } \varphi \in S).$$

Since $[\forall]$ is an S5 modality, \sim is an equivalence relation. Similarly, as \Box is an S4 modality, \leq is a preorder. Moreover, since $\vdash [\forall]\varphi \rightarrow \Box\varphi$ (by axioms (U) and (FE) in Table 6), we obtain that \leq is included in \sim , i.e., $\leq \subseteq \sim$.

Definition 11 (*Canonical quasi-model for T_0*) The canonical quasi model for T_0 is defined as $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$, where

- $X = \{T \subseteq \mathcal{L}_{[\forall]\Box_0\Box} \mid T \text{ is a maximally consistent set with } T \sim T_0\}$;
- $\mathcal{E}_0 = \{\widehat{\Box_0\varphi} \mid \varphi \in \mathcal{L}_{[\forall]\Box_0\Box} \text{ with } [\exists]\Box_0\varphi \in T_0\}$, where $\widehat{\theta} := \{T \in X \mid \theta \in T\}$ for any $\theta \in \mathcal{L}_{[\forall]\Box_0\Box}$;
- \leq is the restriction of the above preorder \leq to X ; and
- $V(p) = \hat{p}$.

In the following, variables T, S, \dots range over X .

Lemma 7 $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ is a quasi-model.

Proof In order to show that \mathcal{M} is a quasi model, we need to show that (1) $X \in \mathcal{E}_0$ and $\emptyset \notin \mathcal{E}_0$, (2) \leq is a preorder, and (3) $\mathcal{E}_0 \subseteq Up_{\leq}(X)$. Note that (2) follows from the fact that \Box is an S4 modality.

(1): Since $\vdash \Box_0\top$ (by $\text{Nec}_{[\forall]}$ and axiom (U) in Table 6), we have $\widehat{\Box_0\top} = X$. Moreover, by axiom (T $_{[\forall]}$), we obtain $[\exists]\Box_0\top \in T_0$, hence, $\widehat{\Box_0\top} = X \in \mathcal{E}_0$. And, obviously, $\emptyset \notin \mathcal{E}_0$.

(3): Let $e \in \mathcal{E}_0$. By the definition of \mathcal{E}_0 , we have $e = \widehat{\Box_0\varphi}$ for some $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box}$ such that $[\exists]\Box_0\varphi \in T_0$. Now suppose $T, S \in X$ with $T \in \widehat{\Box_0\varphi}$ (i.e., $\Box_0\varphi \in T$) and $T \leq S$. Note that $\vdash \Box_0\varphi \rightarrow \Box\Box_0\varphi$ (by axioms (4 $_{\Box_0}$) and (FE)). Therefore, $\Box\Box_0\varphi \in T$. Since $T \leq S$, we then obtain $\Box_0\varphi \in S$, i.e., $S \in \widehat{\Box_0\varphi}$. Thus, as S has been chosen arbitrarily, we conclude that $e \in Up_{\leq}(X)$. □

Lemma 8 (Existence Lemma for $[\forall]$) *For every $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box}$, $[\exists]\widehat{\varphi} \neq \emptyset$ iff $\widehat{\varphi} \neq \emptyset$.*

Proof (\Rightarrow) Suppose $[\exists]\widehat{\varphi} \neq \emptyset$, i.e., there is $T \in X$ such that $T \in [\exists]\widehat{\varphi}$. This means $[\exists]\varphi \in T$. This implies that the set $\Gamma := \{[\forall]\psi \in \mathcal{L}_{[\forall]\Box_0\Box} \mid [\forall]\psi \in T\} \cup \{\varphi\}$ is consistent. Otherwise, there exist finitely many sentences $[\forall]\psi_1, \dots, [\forall]\psi_n \in T$ such that $\vdash ([\forall]\psi_1 \wedge \dots \wedge [\forall]\psi_n) \rightarrow \neg\varphi$. But then, since $[\forall]$ is an S5-modality, we obtain that $\vdash ([\forall]\psi_1 \wedge \dots \wedge [\forall]\psi_n) \rightarrow [\forall]\neg\varphi$. Hence, as $[\forall]\psi_1 \wedge \dots \wedge [\forall]\psi_n \in T$, we get $[\forall]\neg\varphi \in T$, which combined with $[\exists]\varphi \in T$, implies that T is inconsistent, contradicting T being consistent. Therefore, given that Γ is consistent, by Lindenbaum’s Lemma (Lemma 6), there exists some maximally consistent set S such that $\Gamma \subseteq S$. It is easy to see that this implies $\varphi \in S$ and $S \sim T \sim T_0$ (i.e., $S \in X$). Therefore, $S \in \widehat{\varphi}$ implying that $\widehat{\varphi} \neq \emptyset$.

(\Leftarrow) Suppose $\widehat{\varphi} \neq \emptyset$, i.e., there is $T \in X$ such that $T \in \widehat{\varphi}$. Then, since $\varphi \rightarrow [\exists]\varphi \in T$ (by axiom $(T_{[\forall]})$), we obtain $[\exists]\varphi \in T$, implying that $[\exists]\widehat{\varphi} \neq \emptyset$. \square

Lemma 9 (Existence Lemma for \Box) *For every $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box}$ and $T \in X$, $T \in \widehat{\Diamond\varphi}$ iff there is $S \in \widehat{\varphi}$ such that $T \leq S$.*

Proof (\Rightarrow) Assume $T \in \widehat{\Diamond\varphi}$, that is, $\Diamond\varphi \in T$. This implies that the set $\Gamma := \{\Box\psi \in \mathcal{L}_{[\forall]\Box_0\Box} \mid \Box\psi \in T\} \cup \{\varphi\}$ is consistent. Otherwise there exist finitely many sentences $\Box\psi_1, \dots, \Box\psi_n \in T$ such that $\vdash (\Box\psi_1 \wedge \dots \wedge \Box\psi_n) \rightarrow \neg\varphi$. But then, since \Box is an S4-modality, we obtain that $\vdash (\Box\psi_1 \wedge \dots \wedge \Box\psi_n) \rightarrow \Box\neg\varphi$. Hence, as $\Box\psi_1 \wedge \dots \wedge \Box\psi_n \in T$, we get $\Box\neg\varphi \in T$, which combined with $\Diamond\varphi \in T$, implies that T is inconsistent, contradicting T being consistent. Therefore, given that Γ is consistent, by Lindenbaum’s Lemma (Lemma 6), there exists some maximally consistent set S such that $\Gamma \subseteq S$. It is easy to see that this implies $\varphi \in S$ and $T \leq S$. Since \leq is included in \sim , we also obtain $S \sim T \sim T_0$, i.e., $S \in X$. Therefore, $S \in \widehat{\varphi}$.

(\Leftarrow) Suppose there is $S \in \widehat{\varphi}$ such that $T \leq S$. Then, by definition of \leq , $\Diamond\varphi \in T$, i.e., $T \in \widehat{\Diamond\varphi}$. \square

Lemma 10 (Existence Lemma for \Box_0) *For every $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box}$ and $T \in X$, $T \in \widehat{\Box_0\varphi}$ iff there exist $e \in \mathcal{E}_0$ such that $T \in e$ and $e \subseteq \widehat{\varphi}$.*

Proof (\Rightarrow) Suppose $T \in \widehat{\Box_0\varphi}$, i.e. $\Box_0\varphi \in T$. Since $T \sim T_0$, we get $[\exists]\Box_0\varphi \in T_0$. This means $\widehat{\Box_0\varphi} \in \mathcal{E}_0$. Taking $e := \widehat{\Box_0\varphi}$, we get $e \in \mathcal{E}_0$ and $T \in e$. Moreover, since $\vdash \Box_0\varphi \rightarrow \varphi$, we obtain $e = \widehat{\Box_0\varphi} \subseteq \widehat{\varphi}$.

(\Leftarrow) Suppose there is $e \in \mathcal{E}_0$ such that $T \in e$ and $e \subseteq \widehat{\varphi}$. Then, by the definition of \mathcal{E}_0 , we obtain that $e = \widehat{\Box_0\theta}$ for some θ such that $[\exists]\Box_0\theta \in T_0$. Therefore, $T \in e = \widehat{\Box_0\theta} \subseteq \widehat{\varphi}$. This implies that the set $\Gamma := \{\Box_0\theta\} \cup \{[\forall]\psi \in \mathcal{L}_{[\forall]\Box_0\Box} : [\forall]\psi \in T\} \cup \{\neg\varphi\}$ is inconsistent. Otherwise, by Lindenbaum’s Lemma (Lemma 6), there exists a $S \in X$ such that $\Box_0\theta \in S$ and $\neg\varphi \in S$. The former means that $S \in \widehat{\Box_0\theta}$ and the latter means (since S is maximal) that $S \notin \widehat{\varphi}$. Thus, $S \in \widehat{\Box_0\theta} \setminus \widehat{\varphi}$, contradicting the assumption $\widehat{\Box_0\theta} \subseteq \widehat{\varphi}$. Therefore, given that Γ is inconsistent, there exists a finite set $\{[\forall]\psi_1, \dots, [\forall]\psi_n\} \subseteq \Gamma$ such that $\vdash \bigwedge_{i \leq n} [\forall]\psi_i \rightarrow (\Box_0\theta \rightarrow \varphi)$. Since $[\forall]$ is a normal modality and T is maximal,

$\bigwedge_{i \leq n} [\forall]\psi_i = [\forall]\gamma$ for some $[\forall]\gamma \in T$. We then have

1. $\vdash [\forall]\gamma \rightarrow (\Box_0\theta)$ *CPL*
2. $\vdash ([\forall]\gamma \wedge \Box_0\theta) \rightarrow \varphi$ *CPL*
3. $\vdash \Box_0([\forall]\gamma \wedge \Box_0\theta) \rightarrow \Box_0\varphi$ *Monotonicity of \Box_0*
4. $\vdash \Box_0\Box_0([\forall]\gamma \wedge \theta) \rightarrow \Box_0\varphi$ *Pullout axiom, right – to – left*
5. $\vdash \Box_0([\forall]\gamma \wedge \theta) \rightarrow \Box_0\varphi$ *since $\vdash \Box_0\varphi \leftrightarrow \Box_0\Box_0\varphi$*
6. $\vdash ([\forall]\gamma \wedge \Box_0\theta) \rightarrow \Box_0\varphi$ *Pullout axiom*

Therefore, since $[\forall]\gamma, \Box_0\theta \in T$, and T is maximal, we obtain $\Box_0\varphi \in T$, i.e., $T \in \widehat{\Box_0\varphi}$. □

Lemma 11 (Truth Lemma) *For every formula $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box_0}$, we have*

$$\|\varphi\|^{\mathcal{M}} = \widehat{\varphi}.$$

Proof The proof follows standardly by subformula induction on φ , where the inductive step for each modality uses the corresponding Existence Lemma, as usual. □

Proposition 12 *$\text{Log}_{\forall\Box_0\Box_0}$ is sound and strongly complete for quasi-models.*

Proof Let Φ_0 be a $\text{Log}_{\forall\Box_0\Box_0}$ -consistent set of formulas. Then, by Lindenbaum’s Lemma (Lemma 6), Φ_0 can be extended to a maximally consistent set T_0 . We can then construct a canonical quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ for T_0 as in Definition 11, and by Lemma 11, obtain that $T_0 \in \|\varphi\|^{\mathcal{M}}$ for all $\varphi \in \Phi_0$. □

Step 2: strong finite quasi-model property In this section, we prove that the logic $\text{Log}_{\forall\Box_0\Box_0}$ has the strong finite quasi-model property. We do so via a filtration argument using the canonical model described in Definition 11.

Let φ_0 be a $\text{Log}_{\forall\Box_0\Box_0}$ -consistent formula. By Lemma 6, there exist a maximally consistent set T_0 such that $\varphi_0 \in T_0$. Consider the canonical quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ for T_0 (as given in Definition 11). We will use two facts about this model:

1. $\|\varphi\|^{\mathcal{M}} = \widehat{\varphi}$, for all $\varphi \in \mathcal{L}_{[\forall]\Box_0\Box_0}$; and
2. $\mathcal{E}_0 = \{\widehat{\Box_0\varphi} \mid [\exists]\Box_0\varphi \in T_0\} = \{\|\Box_0\varphi\|^{\mathcal{M}} \mid [\exists]\Box_0\varphi \in T_0\}$.

Closure conditions for Σ : Let Σ be a *finite* set such that: (1) $\varphi_0 \in \Sigma$; (2) Σ is closed under subformulas; (3) if $\Box_0\varphi \in \Sigma$ then $\Box_0\Box_0\varphi \in \Sigma$; (4) Σ is closed under single negations; (5) $\Box_0\top \in \Sigma$. For $T, S \in X$, put

$$T \equiv_{\Sigma} S \text{ iff for all } \psi \in \Sigma (T \in \|\psi\|^{\mathcal{M}} \text{ iff } S \in \|\psi\|^{\mathcal{M}}),$$

and denote by $|T| := \{S \in X \mid T \equiv_{\Sigma} S\}$ the equivalence class of T modulo \equiv_{Σ} . Also, put $X^f = \{|T| \mid T \in X\}$, and more generally put $e^f = \{|T| \mid T \in e\}$ for every $e \in \mathcal{E}_0$. We now define a *filtration* $\mathcal{M}^f = (X^f, \mathcal{E}_0^f, \leq^f, V^f)$ of \mathcal{M} through Σ , where

- $X^f = \{|T| \mid T \in X\}$;
- $|T| \leq^f |S|$ iff for all $\Box_0\psi \in \Sigma$ ($T \in \|\Box_0\psi\|^{\mathcal{M}}$ implies $S \in \|\Box_0\psi\|^{\mathcal{M}}$);
- $\mathcal{E}_0^f = \{e^f \mid e = \widehat{\Box_0\psi} = \|\Box_0\psi\|^{\mathcal{M}} \in \mathcal{E}_0 \text{ for some } \psi \text{ such that } \Box_0\psi \in \Sigma\}$;

- $V^f(p) = \{|T| : T \in V(p)\}$ for all $p \in \Sigma$, and $V^f(p) = \emptyset$ otherwise.

Lemma 12 \mathcal{M}^f is a finite quasi-model (of size bounded by a computable function of φ_0).

Proof Since Σ is finite, there are only finitely many equivalence classes modulo \equiv_Σ . Therefore, X^f is finite. In fact, X^f has at most $2^{|\Sigma|}$ states. It is obvious that \leq^f is a preorder. Moreover, since $X = \|\Box_0 \top\|^{\mathcal{M}}$ and $\Box_0 \top \in \Sigma$, we have $X^f \in \mathcal{E}_0^f$. Also, since $e \neq \emptyset$ for all $e \in \mathcal{E}_0$, we have each $e^f \in \mathcal{E}_0^f$ nonempty. So we only have to prove that the evidence sets e^f are upsets of (X^f, \leq^f) . For this, let $e^f \in \mathcal{E}_0^f$, $|T|, |S| \in X^f$ such that $|T| \in e^f$ and $|T| \leq^f |S|$. We need to show that $|S| \in e^f$. By the definition of \mathcal{E}_0^f , we know that $e = \widehat{\Box_0 \psi} = \|\Box_0 \psi\|^{\mathcal{M}}$ for some $\Box_0 \psi \in \Sigma$. From $|T| \in e^f$, it follows that there is some $T' \equiv_\Sigma T$ such that $T' \in e = \|\Box_0 \psi\|^{\mathcal{M}}$, and since $\Box_0 \psi \in \Sigma$, we have $T \in \|\Box_0 \psi\|^{\mathcal{M}}$. Therefore, since $\vdash \Box_0 \psi \rightarrow \Box_0 \psi$ (this is easy to see from axioms (4_{\Box_0}) and (FE) stated in Table 6), we have $T \in \|\Box_0 \psi\|^{\mathcal{M}}$. But $\Box_0 \psi \in \Sigma$ (by the closure assumptions on Σ), so $|T| \leq^f |S|$ gives us $S \in \|\Box_0 \psi\|^{\mathcal{M}}$. By the axiom (T_{\Box_0}) , we obtain $S \in \widehat{\Box_0 \psi} = e$, hence $|S| \in e^f$. \square

Lemma 13 (Filtration Lemma) For every formula $\varphi \in \Sigma$, we have $\|\varphi\|^{\mathcal{M}^f} = \{|T| \mid T \in \|\varphi\|^{\mathcal{M}}\}$.

Proof The proof follows by subformula induction on $\varphi \in \Sigma$; cases for the propositional variables, the Boolean connectives, and the modalities $[\forall]\psi$ and $\Box\psi$ are treated as usual (in the last case using the filtration property of \leq^f that: if $T \leq S$ then $|T| \leq^f |S|$). We only prove here the inductive case for $\varphi := \Box_0 \psi$:

(\Rightarrow) Let $|T| \in \|\Box_0 \psi\|^{\mathcal{M}^f}$. This means that there exists some $e^f \in \mathcal{E}_0^f$ s.t. $|T| \in e^f \subseteq \|\psi\|^{\mathcal{M}^f}$. By the definition of \mathcal{E}_0^f , there exists some $\chi \in \mathcal{L}_{[\forall]\Box_0}$ such that $\Box_0 \chi \in \Sigma$ and $e = \widehat{\Box_0 \chi} = \|\Box_0 \chi\|^{\mathcal{M}} \in \mathcal{E}_0$. From $|T| \in e^f$, it follows that there is some $T' \equiv_\Sigma T$ such that $T' \in e = \|\Box_0 \chi\|^{\mathcal{M}}$, and since $\Box_0 \chi \in \Sigma$, we have $T \in \|\Box_0 \chi\|^{\mathcal{M}} = e$. Now let $S \in e$ be any element of e . Then, by the definition of e^f and the assumption that $e^f \subseteq \|\psi\|^{\mathcal{M}^f}$, we obtain $|S| \in e^f \subseteq \|\psi\|^{\mathcal{M}^f}$. So, $|S| \in \|\psi\|^{\mathcal{M}^f}$. Therefore, by the induction hypothesis, $S \in \|\psi\|^{\mathcal{M}}$, hence, $e \subseteq \|\psi\|^{\mathcal{M}}$. Thus, we have found an evidence set $e \in \mathcal{E}_0$ such that $T \in e \subseteq \|\psi\|^{\mathcal{M}}$, i.e., shown that $T \in \|\Box_0 \psi\|^{\mathcal{M}}$.

(\Leftarrow) Let $T \in \|\Box_0 \psi\|^{\mathcal{M}}$. It is easy to see that $[\exists]\Box_0 \psi \in T$ (since $\vdash \Box_0 \psi \rightarrow [\exists]\Box_0 \psi$), and so also $[\exists]\Box_0 \psi \in T_0$ (since $T \in X$, thus, $T \sim T_0$). This means that the set $e := \widehat{[\exists]\Box_0 \psi} = \|\Box_0 \psi\|^{\mathcal{M}} \in \mathcal{E}_0$ is an evidence set in the canonical model (see Definition 11), and since $\Box_0 \psi \in \Sigma$, we conclude that $e^f \in \mathcal{E}_0^f$. We obviously have $T \in e$, and so $|T| \in e^f$. Since $\vdash \Box_0 \psi \rightarrow \psi$, we have $e = \|\Box_0 \psi\|^{\mathcal{M}} \subseteq \|\psi\|^{\mathcal{M}}$, and hence $e^f \subseteq \{|S| \mid S \in \|\psi\|^{\mathcal{M}}\} = \|\psi\|^{\mathcal{M}^f}$ (by the induction hypothesis). Thus, we have found $e^f \in \mathcal{E}_0^f$ such that $|T| \in e^f \subseteq \|\psi\|^{\mathcal{M}^f}$, i.e., shown that $|T| \in \|\Box_0 \psi\|^{\mathcal{M}^f}$. \square

Theorem 6 $\text{Log}_{\forall\Box_0}$ has strong finite quasi-model property.

Proof Let φ_0 be a $\text{Log}_{\forall\Box_0}$ -consistent formula. Then, by Lindenbaum’s Lemma (Lemma 6), φ_0 can be extended to a maximally consistent set T_0 such that $\varphi_0 \in T_0$. We can then construct a canonical quasi-model $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ for T_0 as in Definition 11, and by Lemma 11 obtain that $T_0 \in \|\varphi_0\|^{\mathcal{M}}$. Then, by Lemma 13, we have $|T_0| \in \|\varphi_0\|^{\mathcal{M}^f}$, where \mathcal{M}^f is the filtration of \mathcal{M} through the finite set Σ that is obtained by closing $\{\varphi_0\}$ under the

closure conditions (1)-(5). By Lemma 12, we know that \mathcal{M}^f is a finite model whose size is bounded by $2^{|\Sigma|}$, therefore we conclude that $\text{Log}_{\forall \square \square_0}$ has the strong finite quasi-model property. \square

Step 3: equivalence of finite quasi-models and finite alexandroff quasi-models In this section, we prove that every finite quasi-model is modally equivalent to a finite Alexandroff quasi-model, and therefore, to a topo-e-model with respect to the language $\mathcal{L}_{[\forall] \square_0 \square}$.

Let $\mathcal{M} = (X, \mathcal{E}_0, \leq, V)$ be a finite quasi-model. We form a new structure $\tilde{\mathcal{M}} = (\tilde{X}, \tilde{\mathcal{E}}_0, \tilde{\leq}, \tilde{V})$, by putting:

- $\tilde{X} = X \times \{0, 1\}$;
- $\tilde{V}(p) = V(p) \times \{0, 1\}$;
- $(x, i) \tilde{\leq} (y, j)$ iff $x \leq y$ and $i = j$;
- $\tilde{\mathcal{E}}_0 = \{e_i \mid e \in \mathcal{E}_0, i \in \{0, 1\}\} \cup \{e_i^y \mid y \in e \in \mathcal{E}_0, i \in \{0, 1\}\} \cup \{\tilde{X}\}$, where we used notations
 - $e_i = e \times \{i\} = \{(x, i) \mid x \in e\}$, and
 - $e_i^y = \uparrow y \times \{i\} \cup e \times \{1 - i\} = \{(x, i) \mid y \leq x\} \cup e_{1-i}$.

Lemma 14 $\tilde{\mathcal{M}}$ is a finite quasi-model.

Proof It is easy to see that $\tilde{\mathcal{M}}$ is finite, in fact, it is of size $2 \cdot |X|$. It is guaranteed by definition that $\tilde{X} \in \tilde{\mathcal{E}}_0$ and $\emptyset \notin \tilde{\mathcal{E}}_0$. To show that every element of $\tilde{\mathcal{E}}_0$ is an upset of $(\tilde{X}, \tilde{\leq})$, let $\tilde{e} \in \tilde{\mathcal{E}}_0$ and $(x, i), (y, j) \in \tilde{X}$ such that $(x, i) \in \tilde{e}$ and $(x, i) \tilde{\leq} (y, j)$. Then, by the definition of $\tilde{\leq}$, we know that $x \leq y$ and $i = j$. We have two cases: if $\tilde{e} = e \times \{i\}$ for some $e \in \mathcal{E}_0$, then $y \in e$ (since $e \in \text{Up}_{\leq}(X)$, $x \in e$, and $x \leq y$), therefore, $(y, i) \in e \times \{i\} = \tilde{e}$. If $\tilde{e} = e_k^z$ for some $z \in X$ and $k \in \{0, 1\}$, we again have two cases. If $k = 1 - i$, then the result follows as in the first case. If $k = i$, then $\uparrow z \times i \subseteq \tilde{e}$. Since $(x, i) \in \tilde{e}$, we obtain that $z \leq x$, and thus, $z \leq y$ (since \leq is transitive). We therefore conclude that $(y, i) \in \uparrow z \times i \subseteq \tilde{e}$. \square

Notation: For any set $\tilde{Y} \subseteq \tilde{X}$, put $\tilde{Y}_X := \{y \in X \mid (y, i) \in \tilde{Y} \text{ for some } i \in \{0, 1\}\}$ for the set consisting of first components of all members of \tilde{Y} . It is easy to see that we have $(\tilde{Y} \cup \tilde{Z})_X = \tilde{Y}_X \cup \tilde{Z}_X$, and $\tilde{X}_X = X$.

Lemma 15 If $y \in e \in \mathcal{E}_0, i \in \{0, 1\}$, and $\tilde{e} \in \{e_i, e_i^y\}$, then we have:

1. $\tilde{e}_X = e$;
2. $e_i^y \cap e_i = \uparrow(y, i)$, where $\uparrow(y, i) = \{\tilde{x} \in \tilde{X} \mid (y, i) \tilde{\leq} \tilde{x}\} = \{(x, i) \mid y \leq x\}$.

Proof (1): If $\tilde{e} = e_i$, then $\tilde{e}_X = (e \times \{i\})_X = e$. If $\tilde{e} = e_i^y$, then $\tilde{e}_X = (\uparrow y \times \{i\})_X \cup (e \times \{1 - i\})_X = \uparrow y \cup e = e$ (since $e \in \text{Up}_{\leq}(X)$ and $y \in e$, so $\uparrow y \subseteq e$).

(2): $e_i^y \cap e_i = (\uparrow y \times \{i\} \cup e \times \{1 - i\}) \cap (e \times \{i\}) = (\uparrow y \cap e) \times \{i\} = \uparrow y \times \{i\} = \uparrow(y, i)$ (since $\uparrow y \subseteq e$). \square

Lemma 16 $\tilde{\mathcal{M}}$ is an Alexandroff quasi-model (and thus also a topo-e-model).

Proof By Proposition 10, it is enough to show that, for every $(y, i) \in \tilde{X}$, the upset $\uparrow(y, i)$ is open in the topology $\tau_{\tilde{\mathcal{E}}_0}$ generated by $\tilde{\mathcal{E}}_0$: this follows directly from Lemma 15.2.. \square

Lemma 17 (Modal Equivalence Lemma) For all $\varphi \in \mathcal{L}_{[\forall] \square_0 \square}$, $\|\varphi\|^{\tilde{\mathcal{M}}} = \|\varphi\|^{\mathcal{M}} \times \{0, 1\}$.

Proof The proof follows by subformula induction on φ ; cases for the propositional variables, the Boolean connectives, and the modalities $\forall\psi$ and $\Box\psi$ are straightforward. We only prove here the inductive case for $\varphi := \Box_0\psi$.

(\Rightarrow) Suppose that $(x, i) \in \|\Box_0\psi\|^{\mathcal{M}}$. Then there exists some $\tilde{e} \in \tilde{\mathcal{E}}_0$ such that $(x, i) \in \tilde{e} \subseteq \|\psi\|^{\mathcal{M}} = \|\psi\|^{\mathcal{M}} \times \{0, 1\}$ (we use the induction hypothesis for ψ in the last step). From this, we obtain that $x \in \tilde{e}_X \subseteq (\|\psi\|^{\mathcal{M}} \times \{0, 1\})_X = \|\psi\|^{\mathcal{M}}$. But by the construction of $\tilde{\mathcal{E}}_0$, $\tilde{e} \in \tilde{\mathcal{E}}_0$ means that either $\tilde{e} = \tilde{X}$ or there exist $e \in \mathcal{E}_0$, $y \in e$ and $j \in \{0, 1\}$ such that $\tilde{e} \in \{e_j, e_j^y\}$. If the former is the case, we have $x \in \tilde{e}_X = X \subseteq \|\psi\|^{\mathcal{M}}$. Since $X \in \mathcal{E}_0$, by the semantics of \Box_0 , we obtain $x \in \|\Box_0\psi\|_{\mathcal{M}}$. If the latter is the case, by Lemma 15.1., we have $\tilde{e}_X = e$, so we conclude that $x \in \tilde{e}_X = e \subseteq \|\psi\|^{\mathcal{M}}$. Therefore, again by the semantics of \Box_0 , we have $x \in \|\Box_0\psi\|^{\mathcal{M}}$.

(\Leftarrow) Suppose that $x \in \|\Box_0\psi\|^{\mathcal{M}}$. Then, there exists some $e \in \mathcal{E}_0$ such that $x \in e \subseteq \|\psi\|^{\mathcal{M}}$. Take now the set $e_i = e \times \{i\} \in \tilde{\mathcal{E}}_0$. Clearly, we have $(x, i) \in e_i \subseteq \|\psi\|^{\mathcal{M}} \times \{i\} \subseteq \|\psi\|^{\mathcal{M}} \times \{0, 1\} = \|\psi\|^{\mathcal{M}}$ (we use the induction hypothesis for ψ in the last step), i.e., we have $(x, i) \in \|\Box_0\psi\|^{\mathcal{M}}$. □

Theorem 7 *Every finite quasi-model is modally equivalent to a finite Alexandroff quasi-model, therefore, to a topo-e-model with respect to the language $\mathcal{L}_{\forall\Box_0}$.*

Proof The proof immediately follows from Lemma 17 and Proposition 11: the same formulas are satisfied at x in \mathcal{M} as at (x, i) in $\tilde{\mathcal{M}}$. □

Proof of Theorem 5 Theorem 5 (completeness and finite model property for topo-e-models) is thus obtained as an immediate corollary of Proposition 12, Theorems 6 and 7. □

7 Conclusions and further directions

We have studied a topological semantics for various notions of *evidence*, *evidence-based justification*, *argument*, (*conditional*) *belief*, and *knowledge*. We have done so by using topological structures based on the (uniform) evidence models of van Benthem and Pacuit (2011). Several soundness, completeness, finite model property, and decidability results concerning the logics of belief, knowledge, and evidence on all topological (evidence) models have been provided.

This project has been of both technical and conceptual interest. Philosophically speaking, we have shown that the topological perspective on epistemic logic enables refined representations of the aforementioned notions, and, in turn, can account for subtle distinctions pertaining to, e.g., (possibly true, but) *misleading* and *non-misleading evidence*. The rich formal framework afforded by topologically interpreted modal logics allows one to clarify and address some key issues regarding debates on the relationship between knowledge and belief, defeasibility theories of knowledge, and Stalnaker’s view on belief as subjective certainty. Mathematically, our framework takes a significant step toward developing formal systems in which we can talk about evidential grounds of knowledge and belief, both at the syntactic and semantic level (via evidence modalities and topological structures, respectively). This, in turn, offers novel modal languages that can potentially express some properties of a subspace of a topological space, enriching modal logics of space that are designed to talk about only topologies or topological bases (more on this below).

Table 7 Matching epistemic and topological notions

Epistemology	Topology
Basic evidence	Subbasis of a topology (\mathcal{E}_0)
(Combined) evidence	Basis of a topology (\mathcal{E})
Arguments	Open sets ($\tau_{\mathcal{E}}$)
Justifications	Dense open sets
Belief	Dense interior (nowhere dense complement)
Knowledge (of P at x)	$x \in \text{Int}(P)$ and $\text{Int}(P)$ is dense

Moreover, our topological approach contributes to the evidence setting of van Benthem and Pacuit (2011); van Benthem et al. (2012, 2014) in many ways. First of all, this topological approach gives mathematically more natural meanings to the epistemic/doxastic modalities we considered by providing a precise match between epistemic and topological notions. The list of the epistemic notions studied together with their topological counterparts is given in Table 7 below. Besides, our proposal yields a notion of belief that coincides with the one of van Benthem and Pacuit (2011) in “good” cases, and that behaves better in general. More precisely, our justified belief is always consistent, in fact, it satisfies the axioms and rules of the standard belief system KD_{45}_B on all topological spaces (Sect. 6.3). It moreover admits a natural topological reading in terms of dense-open sets (or equivalently, in terms of nowhere dense sets) as “truth in most states of the model”, where “most” refers to “everywhere but a nowhere dense part”. We have also shown that the logic of evidence models under our proposed semantics has the finite model property, whereas this was not the case in (van Benthem & Pacuit, 2011; van Benthem et al., 2012, 2014).

The formalism developed in this paper improves also on our own work (Baltag et al., 2013, 2019b) where another topological semantics for Stalnaker’s epistemic-doxastic system was proposed. While in Baltag et al. (2013, 2019b) we could talk about evidential grounds of knowledge and belief only on a semantic level, the current setting provides syntactic representations of evidence, therefore, makes the notion of evidence part of the logic. Moreover, we showed that knowledge and belief can be interpreted on arbitrary topological spaces (rather than on extremally disconnected or hereditarily extremally disconnected spaces), without changing their logic. To this end, the semantics of knowledge and belief proposed in this paper generalizes the setting of Baltag et al. (2013, 2019b).

In the rest of this section, we name a few directions for future research.

Connection to topological formal learning theory Various ideas motivating the use of topological spaces to model epistemic notions and information change guide the research program in *topological* formal learning theory, as initiated by Kevin Kelly and others (Baltag et al., 2011, 2015; Gierasimczuk et al., 2014; Kelly, 1996, 2014; Kelly et al., 1995; Schulte & Juhl, 1996). The topologically interpreted logics developed in this paper provide a framework naturally suited to the representation of reasoning about inductive learning from successful observations and therefore constitute a bridge between modal epistemic/doxastic logics and formal learning theory. Investigations focusing on this connection, aiming at bringing learning and logic into closer proximity, have already been initiated by some of the co-authors of this paper and their colleagues (Baltag et al., 2019c, 2020; Vargas Sandoval, 2020).

Another line of inquiry towards this direction involves adding to the semantic structure a larger set $\mathcal{E}_0^\diamond \supseteq \mathcal{E}_0$ of *potential evidence*, meant to encompass all the evidence that might be learnt in the future. A formal setting that involves both actual evidence \mathcal{E}_0 and potential evidence $\mathcal{E}_0^\diamond \supseteq \mathcal{E}_0$ would combine coherentist justification with predictive learning. A logical syntax appropriate for this setting could be obtained by extending our language with operators borrowed from topo-logic (Moss & Parikh, 1992), such as an operator $\diamond\varphi$, expressing the fact that φ can become true after more evidence is learnt. Inductive *learnability* of φ is then captured by the formula $\diamond K\varphi$, where K is our defeasible knowledge (rather than the absolutely certain knowledge operator of topo-logic).

Multi-agent extensions Another line of research involves extending our framework to a *multi-agent* setting. It is straightforward to generalize our semantics to multiple agents, though obtaining a completeness result might not be that easy. However, the real interesting challenge comes when we look at notions of *group knowledge*, for some group G of agents. For *common knowledge*, there are at least two different natural options: (1) the standard Lewis-Aumann concept of the infinite conjunctions of “everybody knows that everybody knows etc.” (Aumann, 1976; Lewis, 1969), and (2) a stronger concept, based on *shared evidence* (the intersection $\bigcap_{a \in G} \mathcal{E}_0^a$ of the evidence families \mathcal{E}_0^a of all agents $a \in G$). The two concepts differ in general. This is related to Barwise’s older observation on the distinction of concepts of common knowledge in a topological framework (Barwise, 1988), in contrast to Kripke models, where all the different versions collapse to the same notion (see also van Benthem and Sarenac, 2004 and Bezhanishvili and van der Hoek, 2014, Sect. 12.4.2.5 for a discussion on the different formalizations of common knowledge on topological spaces). Similarly, in this evidence-based setting, the standard notion of *distributed knowledge* does not seem appropriate to capture a group’s *epistemic potential*. Standardly, a group of agents G is said to have distributed (implicit) knowledge of φ if φ is implied by the knowledge of all individuals in G pooled together (see, e.g., Fagin et al., 1995, Chapter 2 for a standard treatment of distributed knowledge based on relation models). In our setting though, a natural way to think about a group’s epistemic potential is to let the agents share all their evidence, and compute their knowledge based on the evidence family obtained by taking the union $\mathcal{E}_0^G = \bigcup_{a \in G} \mathcal{E}_0^a$ of all the evidence families \mathcal{E}_0^a of all agents a in G . This corresponds to moving to the smallest topology that includes all agents’ evidential topologies τ^a , which also gives us a natural way to define a consistent notion of (potential) *group belief*. However, this setting has some apparent ‘defects’, that is, some facts known by one individual in the group might be defeated by another member’s false or misleading evidence, therefore, the individual knowledge of these facts will be lost after the group members share all their evidence. This is in contrast with the standard notion of distributed knowledge that is group monotonic: the distributed knowledge of a larger group always includes the distributed knowledge of any of its subgroups, and so, in particular, it includes everything known by any member of the group. One option is to simply give up the dogma that groups are always wiser than their members and retain the evidence-based model of group knowledge as providing a better representation of the epistemic potential of a group. Learning from others might not always be epistemically beneficial: it all depends on the quality of the others’ evidence. There are also ways to avoid this conclusion, pursued by Ramirez (2015), via natural modifications of our models and by defining knowledge to be undefeated by any potential evidence that the agent may learn. This way Ramirez (2015) re-establishes group monotonicity, but showing completeness for the resulting logic possesses technical challenges (see Ramirez, 2015, for details).

Reliability sensitive accounts of evidence We often have different levels of trust in our information sources based on, for example, personal experience (thus, with respect to a subjective measure), statistical results reflecting the truthfulness of our sources, or the given error range of the measure devices used in an experiment (thus, with respect to a more objective measure). Receiving information from a variety of sources directs us toward weighing the evidence gathered from these sources with respect to their level of reliability in the process of forming belief and knowledge based on this evidence. The approach proposed in this paper takes every evidence piece the agent has on a par, thus, is not sufficiently fine-grained to account for *the degree of reliability and credence of the evidence sources*. A fruitful formal account of evidence-based belief and information dynamics capturing such a situation can be constructed by extending our qualitative topological models of evidence with a probabilistic component or a preference ordering representing the degree of reliability of/trust in evidence sources.

Modal logics of space So far our work on the relationship between topology and modal logic has been motivated by the search for formal models that help advance our understanding of the epistemic notions and phenomena in question. The work of relating topology and modal logic can however be approached from another direction, as has traditionally been the way: the primary interest lies in spatial structures and building modal logics as tools to reason about them. Inspired by the celebrated topological completeness results of McKinsey and Tarski (1944) for the language of basic modal logic, this approach paved the way for a whole new area of spatial logics, establishing a long standing connection between modal logic and topology (see, e.g., Aiello et al., 2007; van Benthem & Bezhanishvili, 2007 for a survey on this topic). To illustrate, one can (as is often done) define a topological space from its subbasis and, in the presence of our basic evidence modalities, study the logics of the subbasis of a particular topological space (see, e.g., Baltag et al., 2019a; Fernández González, 2018).

Acknowledgements Versions of this paper were presented at the following workshops, conferences, and seminars: Lunchtime Talk, Center for Formal Epistemology, Department of Philosophy, Carnegie Mellon University (2016); Logic and Interactive Rationality Seminar (LIRa), ILLC, University of Amsterdam (2016); Dynamics in Logic IV, Delft University of Technology (2016); LogiCICWorkshop: The Logical Structure of Correlated Information Change, ILLC, University of Amsterdam (2016); The 23rd Workshop on Logic, Language and Computation (WoLLIC 2016), Department of Computer Science, Benemérita Universidad Autónoma. We are grateful to the organizers and participants of these events who have provided us with valuable feedback on this work in the last years, and two anonymous referees of this journal for helpful comments and remarks.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aiello, M., Pratt-Hartmann, I., & van Benthem, J. (2007). *Handbook of spatial logics*. Springer.
- Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6), 1236–1239.

- Aumann, R. J. (1999). Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3), 263–300.
- Baltag, A., Bezhanishvili, N., & Fernández González, S. (2019a). The McKinsey-Tarski theorem for topological evidence logics. In R. Iemhoff, M. Moortgat, & R. de Queiroz (Eds.), *Logic, language, information, and computation* (pp. 177–194). Springer.
- Baltag, A., Bezhanishvili, N., Özgün, A., & Smets, S. (2013). The topology of belief, belief revision and defeasible knowledge. In *Proceedings of the 4th international workshop on logic, rationality and interaction (LORI 2013)* (pp. 27–40). Springer.
- Baltag, A., Bezhanishvili, N., Özgün, A., & Smets, S. (2016). Justified belief and the topology of evidence. In *Proceedings of the 23rd workshop on logic, language, information and computation (WOLLIC 2016)* (pp. 83–103).
- Baltag, A., Bezhanishvili, N., Özgün, A., & Smets, S. (2019). A topological approach to full belief. *Journal of Philosophical Logic*, 48, 205–244.
- Baltag, A., Gierasimczuk, N., Özgün, A., Vargas Sandoval, A. L., & Smets, S. (2019). A dynamic logic for learning theory. *Journal of Logical and Algebraic Methods in Programming*, 109, 100485.
- Baltag, A., Gierasimczuk, N., & Smets, S. (2011). Belief revision as a truth-tracking process. In *Proceedings of the 13th conference on theoretical aspects of rationality and knowledge (TARK 2011)* (pp. 187–190). ACM.
- Baltag, A., Gierasimczuk, N., & Smets, S. (2015). On the solvability of inductive problems: A study in epistemic topology. In *Proceedings of 15th conference on theoretical aspects of rationality and knowledge (TARK 2015)* (pp. 81–98). Electronic Proceedings in Theoretical Computer Science.
- Baltag, A., Özgün, A., & Vargas Sandoval, A. L. (2020). The logic of AGM learning from partial observations. In L. S. Barbosa, & A. Baltag (Eds.), *Dynamic logic. New trends and applications* (pp. 35–52). Springer
- Baltag, A., van Ditmarsch, H. P., & Moss, L. S. (2008). Epistemic logic and information update. In P. Adriaans & J. van Benthem (Eds.), *Handbook of the philosophy of information*. Elsevier
- Barwise, J. (1988). Three views of common knowledge. In *Proceedings of the 2nd conference on theoretical aspects of reasoning about knowledge (TARK 1988)* (pp. 365–379).
- Berto, F., & Hawke, P. (2021). Knowability relative to information. *Mind*, 130(517), 1–33.
- Bezhanishvili, N., & van der Hoek, W. (2014). Structures for epistemic logic. In A. Baltag & S. Smets (Eds.), *Johan van Benthem on logic and information dynamics* (pp. 339–380). Springer.
- Bjorndahl, A., & Özgün, A. (2020). Logic and topology for knowledge, knowability, and belief. *The Review of Symbolic Logic*, 13(4), 748–775.
- Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal logic. Cambridge tracts in the theoretical computer SCIE* (Vol. 53). Cambridge University Press
- Chagrov, A. V. & Zakharyashev, M. (1997). *Modal logic. Oxford logic guides* (Vol. 35). Oxford University Press.
- Chellas, B. F. (1980). *Modal logic*. Cambridge University Press.
- Clark, M. (1963). Knowledge and grounds: A comment on Mr. Gettier's paper. *Analysis*, 24(2), 46–48.
- DeRose, K. (2009). *The case for contextualism* (1st ed.). Oxford University Press.
- Dretske, F. (2014). Is knowledge closed under known entailment? The case against closure. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology*, 2nd ed. (pp. 27–40). Wiley.
- Dretske, F. I. (2016). Epistemic operators. In H. Arló-Costa, V. F. Hendricks, & J. van Benthem (Eds.), *Readings in formal epistemology: Sourcebook* (pp. 553–566). Springer
- Dugundji, J. (1965). *Topology. Allyn and Bacon series in advanced mathematics*. Prentice Hall.
- Engelking, R. (1989). *General topology*, 2nd ed. (Vol. 6). Heldermann
- Fagin, R., & Halpern, J. Y. (1987). Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1), 39–76.
- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about knowledge*. MIT Press.
- Fernández González, S. (2018). *Generic models for topological evidence logic*. Master's thesis, ILLC, University of Amsterdam.
- Fiutek, V. (2013). *Playing with knowledge and belief*. PhD thesis, University of Amsterdam.
- Foley, R. (2012). *When is true belief knowledge?* Princeton University Press.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Gierasimczuk, N., Hendricks, V. F., & de Jongh, D. (2014). Logic and learning. In A. Baltag & S. Smets (Eds.), *Johan van Benthem on logic and information dynamics* (pp. 267–288). Springer.
- Goldman, A. (1967). A causal theory of knowing. *Journal of Philosophy*, 64, 355–372.

- Goranko, V., & Passy, S. (1992). Using the universal modality: Gains and questions. *Journal of Logic and Computation*, 2(1), 5–30.
- Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17(2), 157–170.
- Hawke, P., Özgün, A., & Berto, F. (2020). The fundamental problem of logical omniscience. *Journal of Philosophical Logic*, 49, 727–766.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Cornell University Press.
- Ichikawa, J. J. & Steup, M. (2013). The analysis of knowledge. In E. N. Zalta, (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University, fall 2013 edition.
- Kelly, K. T. (1996). *The logic of reliable inquiry*. Oxford University Press.
- Kelly, K. T. (2014). A computational learning semantics for inductive empirical knowledge. In A. Baltag & S. Smets (Eds.), *Johan van Benthem on logic and information dynamics* (pp. 289–337). Springer International Publishing.
- Kelly, K. T., Schulte, O., & Hendricks, V. (1995). *Reliable belief revision* (pp. 383–398). Kluwer Academic Pub.
- Klein, P. (1971). A proposed definition of propositional knowledge. *Journal of Philosophy*, 68, 471–482.
- Klein, P. D. (1980). Misleading evidence and the restoration of justification. *Philosophical Studies*, 37(1), 81–89.
- Klein, P. (1981). *Certainty, a refutation of scepticism*. University of Minneapolis Press.
- Lehrer, K. (1990). *Theory of knowledge*. Routledge.
- Lehrer, K., & Paxson, T. J. (1969). Knowledge: Undefeated justified true belief. *Journal of Philosophy*, 66, 225–237.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Lewis, D. K. (1973). *Counterfactuals*. Blackwell.
- McKinsey, J. C. C., & Tarski, A. (1944). The algebra of topology. *Annals of Mathematics*, 45(1), 141–191.
- Moss, L. S. & Parikh, R. (1992). Topological reasoning and the logic of knowledge. In *Proceedings of 4th conference on theoretical aspects of computer science (TARK 1992)* (pp. 95–105). Morgan Kaufmann.
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
- Özgün, A. (2013). Topological models for belief and belief revision. Master's thesis, ILLC, University of Amsterdam.
- Özgün, A. (2017). *Evidence in epistemic logic: A topological perspective*. PhD thesis, University of Amsterdam and University of Lorraine.
- Özgün, A., & Berto, F. (2021). Dynamic hyperintensional belief revision. *The Review of Symbolic Logic*, 14(3), 766–811.
- Ramirez, A. I. R. (2015). Topological models for group knowledge and belief. Master's thesis, ILLC, University of Amsterdam.
- Rott, H. (2004). Stability, strength and sensitivity: Converting belief into knowledge. *Erkenntnis*, 61(2–3), 469–493.
- Schulte, O., & Juhl, C. (1996). Topology as epistemology. *The Monist*, 79(1), 141–147.
- Siemers, M. (2021). Hyperintensional logics for evidence, knowledge and belief. Master's thesis, ILLC, University of Amsterdam.
- Sosa, E. (1999). How to defeat opposition to Moore. *Nous*, 33, 141–153.
- Stalnaker, R. (2006). On logics of knowledge and belief. *Philosophical Studies*, 128(1), 169–199.
- van Benthem, J. (2007). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2), 129–155.
- van Benthem, J. & Bezhanishvili, G. (2007). Modal logics of space. In *Handbook of spatial logics* (pp. 217–298). Springer
- van Benthem, J., Fernández-Duque, D., & Pacuit, E. (2012). Evidence logic: A new look at neighborhood structures. In *Advances in modal logic* (Vol. 9, pp. 97–118). King's College Press.
- van Benthem, J., Fernández-Duque, D., & Pacuit, E. (2014). Evidence and plausibility in neighborhood structures. *Annals of Pure and Applied Logic*, 165(1), 106–133.
- van Benthem, J., & Pacuit, E. (2011). Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1), 61–92.
- van Benthem, J., & Sarenac, D. (2004). The geometry of knowledge. In *Aspects of universal Logic*, 17, 1–31.
- Vargas Sandoval, A. L. (2020). *On the path to the truth: Logical & computational aspects of learning*. PhD thesis, University of Amsterdam.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.