

Uncertainty About Evidence

Extended Abstract

Adam Bjorndahl

Carnegie Mellon University
Pittsburgh, USA
abjorn@andrew.cmu.edu

Aybüke Özgün

ILLC, University of Amsterdam
Amsterdam, the Netherlands
Arché, University of St. Andrews
St. Andrews, Scotland
a.ozgun@uva.nl

We develop a logical framework for reasoning about knowledge and evidence in which the agent may be uncertain about how to interpret their evidence. Rather than representing an evidential state as a fixed subset of the state space, our models allow the set of possible worlds that a piece of evidence corresponds to to vary from one possible world to another, and therefore itself be the subject of uncertainty. Such structures can be viewed as (epistemically motivated) generalizations of topological spaces. In this context, there arises a natural distinction between what is *actually* entailed by the evidence and what the agent *knows* is entailed by the evidence—with the latter, in general, being much weaker. We provide a sound and complete axiomatization of the corresponding bi-modal logic of knowledge and evidence entailment, and investigate some natural extensions of this core system, including the addition of a belief modality and its interaction with evidence interpretation and entailment, and the addition of a “knowability” modality interpreted via a (generalized) interior operator.

1 Introduction

In everyday speech, when we claim, say, that the grass being wet is evidence for its having rained recently, the intended meaning seems to be that seeing the wet grass provides some sort of partial, imperfect, defeasible reason to believe or consider it more likely that it rained recently. In sharp contrast to this, many formal models of information update interpret evidence as being essentially infallible or factive. Standard Bayesian updating, for example, tells us that to update a belief (i.e., a probability measure π) on the basis of an observation (i.e., a subset E of the background state space, the “evidence”), we should condition π on E , after which E is assigned probability 1 (see, e.g., [21, Chapter 3]). For another example, AGM-style belief revision updates an initial state of knowledge/belief (captured by a set of formulas) on the basis of some new information (i.e., a particular formula φ , the “evidence”) to produce a new set of formulas that always contains the input formula φ [1, 20]. And a variety of logical models for evidence and belief update assume that each piece of evidence corresponds to a set of possible worlds U and entails exactly those propositions φ such that $U \subseteq \llbracket \varphi \rrbracket$ [24, 14, 26, 9, 7, 8]. Of special note are those models in which the collection of evidence is assumed to take the structure of a *topology* [6, 12, 11, 2, 25]; the framework we propose can be viewed as a generalization of this paradigm.

In this paper we develop a logic for reasoning about evidence that is founded on distinguishing what a given piece of evidence *actually* entails from what it is *believed* (perhaps erroneously) to entail. Thus, in our models, evidence is factive but its *interpretation* can be uncertain. This is accomplished by allowing the set of possible worlds that a given piece of evidence corresponds to to vary across possible worlds, and therefore itself be the subject of uncertainty. Viewed as a generalization of topology, roughly

speaking this corresponds to replacing each individual open set with a parametrized family of sets (with the parameter taken from the underlying space itself).

This paves the way for a natural representation of several closely related phenomena including calibration error (where an agent receives a signal from a measurement device but is uncertain or mistaken about how that signal is related to the measured quantity), evidence “introspection” failure (where the agent in fact has evidence for φ but lacks evidence *that* their evidence entails φ), and uncertain margins of error (where the agent has taken a measurement with a certain margin of error, but is unsure what exactly that margin is).

This paper is organized as follows. In the next section we briefly review some of the syntax and notation we will rely on in the rest of the paper. Section 3 introduces *evidence models*, motivating the definitions with intuitions and examples (Examples 1 and 2) and situating the framework in the context of existing paradigms (Observations 1 and 2). We also provide a sound and complete axiomatization of the corresponding logic of evidence and knowledge. In Section 4 we extend evidence models to incorporate belief, discuss some possible relationships between belief and evidence, and axiomatize the resulting logics. Section 5 introduces a dynamic component to the models in the form of an evidence combination operation, and explores how this additional structure can be used to define a notion of *knowability*; once again, we provide an axiomatization of the relevant logics. Section 6 concludes with a discussion of ongoing work. Omitted proofs appear in the full version.

2 Syntactic Preliminaries

We first specify a class of logical languages appropriate for the kinds of reasoning that concern us in this paper. Given unary modalities \star_1, \dots, \star_k , let $\mathcal{L}_{\star_1, \dots, \star_k}$ denote the propositional language recursively generated by

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \star_i \varphi,$$

where $p \in \text{PROP}$, the (countable) set of *primitive propositions*, and $1 \leq i \leq k$. Our focus in this paper is on the trimodal languages $\mathcal{L}_{K,E,B}$ and $\mathcal{L}_{K,E,\square}$, and various bi-modal and uni-modal fragments thereof, where we read $K\varphi$ as “the agent knows φ ”, $E\varphi$ as “the evidence entails φ ”, $B\varphi$ as “the agent believes φ ”, and finally $\square\varphi$ as “ φ is knowable” or “the agent could come to know φ ”. The Boolean connectives \vee , \rightarrow , and \leftrightarrow are defined as usual, and \perp is defined as an abbreviation for $p \wedge \neg p$. We also employ \hat{K} as an abbreviation for $\neg K \neg$, \hat{E} for $\neg E \neg$, \hat{B} for $\neg B \neg$, and \diamond for $\neg \square \neg$.

(K \star)	$\vdash \star(\varphi \rightarrow \psi) \rightarrow (\star\varphi \rightarrow \star\psi)$	Distribution
(D \star)	$\vdash \star\varphi \rightarrow \neg \star \neg \varphi$	Consistency
(T \star)	$\vdash \star\varphi \rightarrow \varphi$	Factivity
(4 \star)	$\vdash \star\varphi \rightarrow \star\star\varphi$	Positive introspection
(5 \star)	$\vdash \neg \star \varphi \rightarrow \star \neg \star \varphi$	Negative introspection
(Nec \star)	from $\vdash \varphi$ infer $\vdash \star\varphi$	Necessitation

Table 1: Some axiom schemes and a rule of inference for \star

Let CPL denote an axiomatization of classical propositional logic. Then, following standard naming

conventions, we define the following logical systems:

$$\begin{aligned}
K_* &= \text{CPL} + (K_*) + (\text{Nec}_*) \\
KT_* &= K_* + (T_*) \\
S4_* &= KT_* + (4_*) \\
S5_* &= S4_* + (5_*) \\
KD45_* &= K_* + (D_*) + (4_*) + (5_*).
\end{aligned}$$

3 Evidence Models for Evidence Entailment and Knowledge

An **evidence space** is a tuple (X, \mathcal{E}, I) where X is a nonempty set of *worlds*, \mathcal{E} is a nonempty set of *evidence states*, and $I = \{I_e\}_{e \in \mathcal{E}}$ is a parametrized family of functions $I_e : X \rightarrow 2^X$. An **evidence model** $\mathcal{M} = (X, \mathcal{E}, I, \nu)$ (over PROP) is an evidence space (X, \mathcal{E}, I) equipped with a *valuation function* $\nu : \text{PROP} \rightarrow 2^X$.

Intuitively, each $e \in \mathcal{E}$ represents a “state of evidence” the agent may be in—perhaps arising from having made some observation, performed some experiment, found some clue, etc. Crucially, the agent is not conceptualized as being uncertain about *which* state of evidence they find themselves in, but rather about the *interpretation* of any such $e \in \mathcal{E}$. In particular, the evidence e at world x rules out exactly those worlds outside of $I_e(x)$, so $I_e(x)$ tells us what the evidence e *actually entails at* x . Call $I_e(x)$ an *interpretation of* e . These models therefore differ from many standard representations of evidence: rather than representing evidence directly as subsets of the state space, states of evidence are treated as abstract objects, each of which gets associated, via I , to various possible subsets of the state space—that is, various possible interpretations.

Suppose x is the actual world and e is the evidence state the agent is in; then, intuitively, we should have $x \in I_e(x)$, since otherwise the evidence would rule out the actual world, which seems absurd. Say that x and e are **coherent** when $x \in I_e(x)$ and define

$$U_e = \{x \in X : x \in I_e(x)\},$$

the collection of worlds that cohere with e . Since these are precisely the worlds at which the true interpretation of e is compatible with the world, intuitively U_e consists of exactly those worlds at which e is a possible state of evidence. Thus, when x and e are coherent (i.e., when $x \in U_e$), we call the corresponding pair (x, e) an **evidence scenario**—we think of such pairs as being wholistic, self-consistent descriptions of the world and the agent’s state of evidence, analogous to the “epistemic scenarios” of subset space logic [24, 14, 26]. Similarly to subset space semantics, formulas will be interpreted in evidence models not at worlds x but at evidence scenarios (x, e) .

We are now in a position to formalize our notion of (actual) evidence entailment. Given an evidence model $\mathcal{M} = (X, \mathcal{E}, I, \nu)$ and an evidence scenario (x, e) in \mathcal{M} , we interpret \mathcal{L}_E in \mathcal{M} as follows:

$$\begin{aligned}
(\mathcal{M}, x, e) \models p & \quad \text{iff} \quad x \in \nu(p) \\
(\mathcal{M}, x, e) \models \neg \varphi & \quad \text{iff} \quad (\mathcal{M}, x, e) \not\models \varphi \\
(\mathcal{M}, x, e) \models \varphi \wedge \psi & \quad \text{iff} \quad (\mathcal{M}, x, e) \models \varphi \text{ and } (\mathcal{M}, x, e) \models \psi \\
(\mathcal{M}, x, e) \models E\varphi & \quad \text{iff} \quad I_e(x) \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}^e
\end{aligned}$$

where $\llbracket \varphi \rrbracket_{\mathcal{M}}^e = \{x \in U_e : (\mathcal{M}, x, e) \models \varphi\}$, the **truth set of φ with respect to** e . We omit mention of \mathcal{M} when the model is clear from the context. A formula φ is said to be **satisfiable** in an evidence model \mathcal{M} if there is some evidence scenario (x, e) of \mathcal{M} such that $(\mathcal{M}, x, e) \models \varphi$ and **valid** in \mathcal{M} if for all (x, e) of \mathcal{M} ,

we have $(\mathcal{M}, x, e) \models \varphi$. Note that, by definition, if (x, e) is an evidence scenario we have $x \in I_e(x)$, from which it follows that these semantics validate $E\varphi \rightarrow \varphi$; that is, actual evidence entailment is factive.

Knowledge is often identified in epistemic models with what follows from the agent's information/evidence. In the present framework, however, this is arguably far too strong, since what *actually* follows from the evidence is a fact about the world that the agent may not have any access to. Somewhat more precisely: at the world x the evidence e entails $I_e(x)$ (as a matter of fact), but of course the agent might be uncertain about which world is the true world, and therefore uncertain about what e actually entails. Nonetheless, even without knowing what the evidence e actually entails, they can at a minimum be certain that *whatever* it entails, the world is compatible with that: in other words, the world is somewhere in $\bigcup_{y \in X} I_e(y)$.¹ This motivates the following semantics for knowledge:

$$(\mathcal{M}, x, e) \models K\varphi \text{ iff } \bigcup_{y \in X} I_e(y) \subseteq \llbracket \varphi \rrbracket^e.$$

In other words, the agent knows φ just in case φ follows from *every* interpretation of the evidence. Note that under these semantics, the scheme $K\varphi \rightarrow E\varphi$ is valid but its converse is generally not. It is also easy to see that K is an S5 modality, since the set $\bigcup_{y \in X} I_e(y)$ does not depend on the state.

Consider now the following natural condition:

$$(E1) \quad y \in I_e(x) \Rightarrow y \in I_e(y).$$

This simply states that at every world x , the evidence entails that it coheres with the world. (E1) implies (in fact is equivalent to) the following:

$$\bigcup_{y \in X} I_e(y) = U_e.$$

Therefore, under (E1), given evidence e , the agent is in a position to know that the world coheres with that evidence. This also implies that in evidence models satisfying (E1), the above semantic clause for knowledge can be equivalently restated as

$$(\mathcal{M}, x, e) \models K\varphi \text{ iff } U_e \subseteq \llbracket \varphi \rrbracket^e.$$

Note that since $\llbracket \varphi \rrbracket^e \subseteq U_e$ by definition, this semantic clause for knowledge is in turn equivalent to $U_e = \llbracket \varphi \rrbracket^e$.

Next we observe that evidence models subsume standard relational (Kripke-style) semantics, as well as subset space semantics.

Observation 1. Standard *relational models* (see, e.g., [13, 18]) of the form (X, R_E, ν) , where R_E is the accessibility relation for a unary modality E , arise as a special case of evidence models when the accessibility relation is reflexive: simply take $\mathcal{E} = \{e\}$ and define $I_e(x) = \{y \in X : xR_E y\}$. Then it is easy to see that (E1) is satisfied (since R_E is reflexive) and $x \models E\varphi$ (in the relational model, i.e., when E is interpreted by universal quantification over all R_E -accessible states) just in case $(x, e) \models E\varphi$ (in the evidence model). Moreover, since $U_e = X$ in this case, the knowledge modality in the evidence

¹“But what if the agent doesn't know that the actual evidence is e ?” one might object. But the intent is for *all* uncertainty about the evidence to be encoded in the state space—we take e to be an abstract description of the evidence that is broad enough to be compatible with *every possible interpretation thereof* (as represented by the sets $\{I_e(y) : y \in X\}$). For example, we might imagine an agent who has pointed their measuring device at a phenomenon they wish to measure, and as a result they now see the number 11.1 on a little display window. Perhaps they do not know the margin of error of this device, or to what degree it is calibrated; perhaps they don't even know what exactly it is measuring! But what they *do* know (plausibly, if we restrict our attention to non-skeptical scenarios), is that the window reads 11.1.

model coincides with the universal modality in the relational model, that is, $(x, e) \models K\phi$ (in the evidence model) iff for all $x \in X$, $x \models \phi$ (in the relational model). This correspondence will play a crucial role in our completeness proof for $\mathcal{L}_{K,E}$ with respect to evidence models. \square

Observation 2. *Subset space models* [24, 14, 26] can also naturally be viewed as special cases of evidence models. Given a subset space model $\mathcal{X} = (X, \mathcal{S}, v)$,² we can take $\mathcal{E} = \{e_U : U \in \mathcal{S}\}$, so \mathcal{E} consists of one evidence state for each $U \in \mathcal{S}$, and define $I_{e_U}(x) = U$ for all $x \in X$. Thus, each e_U is interpreted uniformly as corresponding to the subset U . Then $\mathcal{M} = (X, \mathcal{E}, I, v)$ is an evidence model which clearly satisfies (E1) (since each I_{e_U} is constant) and we have $U_{e_U} = U$. Therefore, (x, U) is an epistemic scenario of \mathcal{X} if and only if (x, e_U) is an evidence scenario of \mathcal{M} , and moreover:

$$(\mathcal{X}, x, U) \models K\phi \text{ iff } (\mathcal{M}, x, e_U) \models K\phi \text{ iff } (\mathcal{M}, x, e_U) \models E\phi. \quad \square$$

Example 1. Consider Williamson’s famous clock example [29]: you look at a clock and have a perceptual experience that seems to indicate to you that the minute hand is somewhere on the righthand side of the clock. Let’s index the possible positions of the minute hand with the interval $C = [0, 60)$ in the obvious way, so for example the state 15 corresponds to it being quarter-past, 30 to half-past, etc. The perceptual experience you have is supposed to constitute evidence of some sort, with presumably some margin of error involved. That is, if we call this perceptual experience e , we want to say that e doesn’t tell us the exact position of the minute hand, but rather guarantees that it must lie in some interval containing the true position in its interior. Call this the *margin of error principle*.

In our framework, we can and will incorporate a further type of uncertainty, namely, uncertainty *about the margin of error*. This seems a very natural type of ignorance to model—after all, we may be sure that our perceptions are not exact without being sure of exactly how inexact they are! One way of capturing this scenario using a simple evidence model $M = (C, \mathcal{E}, I, v)$ is to define

$$I_e(c) = \begin{cases} \left(\frac{c}{2}, \frac{c+30}{2}\right) & \text{if } c \in (0, 30) \\ \emptyset & \text{otherwise.} \end{cases}$$

Then it is easy to see that $c \in I_e(c)$ iff $c \in (0, 30)$, and moreover for all $c \in C$, $I_e(c) \subset (0, 30)$, from which (E1) follows. Note also that in every state c , $I_e(c)$ contains c in its interior, so this model satisfies the margin of error principle.

Suppose our primitive propositions include those in the set $\{pos_t : t \in \{0, 1, \dots, 59\}\}$, where pos_t is read “the minute hand is t minutes past twelve” and v is defined in the obvious way: $v(pos_t) = \{t\}$. Clearly, evidence is not “introspective” in this model, in the sense that the principle $E\phi \rightarrow EE\phi$ can fail; for instance, it is easy to see that $(15, e) \models E\neg pos_6$ since $6 \notin (7.5, 22.5) = I_e(15)$, but $(15, e) \not\models EE\neg pos_6$, since for example $6 \in (4, 19) = I_e(8)$ and $8 \in (7.5, 22.5) = I_e(15)$. That is, at state 15, the evidence *in fact* rules out that it’s 6 minutes past twelve, but doesn’t itself guarantee that it rules this out. On the other hand, it is easy to see that $U_e = (0, 30)$, hence for all $c \in (0, 30)$, $(c, e) \models K\neg pos_{45}$ —you are in a position to *know* that the minute hand is not pointing directly to the left, even though you don’t know exactly what your evidence entails.³ \square

²That is, X is a nonempty set of states, \mathcal{S} a collection of subsets of X called *epistemic ranges*, and v a valuation function; formulas are evaluated with respect to *epistemic scenarios* of the form (x, U) where $x \in U \in \mathcal{S}$, and $(x, U) \models K\phi \Leftrightarrow U \subseteq \llbracket \phi \rrbracket^U := \{y \in X : (y, U) \models \phi\}$.

³Incidentally, in this model you are also in a position to know that your evidence is compatible with the hand pointing directly to the right, that is, for all $c \in (0, 30)$, $(c, e) \models K\neg E\neg pos_{15}$. By contrast, for each $t \neq 15$, there is a $c \in (0, 30)$ such that $(c, e) \models E\neg pos_t$.

Of course, in building an epistemic model of this scenario, there is no reason to assume that C itself constitutes the *epistemic* state space. Indeed, doing so leads to the potentially problematic implication that *all* the uncertainty the agent may face is indexed by the position of the clock's minute hand.

Example 2. We consider again the Williamson clock case, except this time we expand the epistemic state space to include not only the possible positions of the minute hand C , but also an additional parameter that captures variation in the margin for error. This allows us to “de-couple” the margin for error from the actual position of the hand, representing a richer space of epistemic possibilities in which the position of the hand and the margin for error can to some extent vary independently of one another. More precisely, define an evidence model $M' = (C \times (0, 1), \mathcal{E}', I', v')$ where

$$I'_e(c, \mu) = \begin{cases} ((1 - \mu)c, (1 - \mu)c + 30\mu) \times (0, 1) & \text{if } c \in (0, 30) \\ \emptyset & \text{otherwise.} \end{cases}$$

Intuitively, μ captures the (actual) precision of the observation: lower values of μ correspond to higher precision, and higher values of μ correspond to lower precision (it is easy to see that the length of the interval $I'_e(c, \mu)$ is just 30μ). The intervals defined in the previous example arise as the special case where $\mu = \frac{1}{2}$, since $I'_e(c, \frac{1}{2}) = (\frac{c}{2}, \frac{c+30}{2}) \times (0, 1)$. Notice also that this model assumes that your observation of the clock provides *no* evidence at all pertaining to the precision of that observation as captured by the value of μ , since every possible value of μ is compatible with every interpretation of e . This, of course, is not a required constraint of the present framework, but merely one we find plausible and convenient for the current scenario.

Despite the extra richness in the epistemic state space, this evidence model shares several key properties with the one described in Example 1. As before, one can easily check that $(c, \mu) \in I'_e(c, \mu)$ iff $c \in (0, 30)$, and for all $c \in C$, $I'_e(c, \mu) \subset (0, 30) \times (0, 1)$, so (E1) holds. Moreover, in every state (c, μ) , the first component of $I'_e(c, \mu)$ contains c in its interior, so this model also satisfies the margin of error principle. Furthermore, if we interpret pos_t in M' in the obvious way (i.e., by setting $v'(pos_t) = \{t\} \times (0, 1)$), then the scheme $E\phi \rightarrow EE\phi$ does not hold here either: for example, $((15, .3), e) \models E\neg pos_{10}$ since $10 \notin (10.5, 19.5)$ and $I'_e(15, .3) = (10.5, 19.5) \times (0, 1)$; on the other hand, $((15, .3), e) \not\models EE\neg pos_{10}$, since $(15, 0.4) \in I'_e(15, .3)$, and $(10, .3) \in (9, 21) \times (0, 1) = I'_e(15, 0.4)$. And finally, analogously to the previous example, we have $U_e = (0, 30) \times (0, 1)$, so for all $c \in (0, 30)$ and all $\mu \in (0, 1)$, $((c, \mu), e) \models K\neg pos_{45}$. \square

3.1 Soundness and Completeness for $\mathcal{L}_{K,E}$

When we interpret $\mathcal{L}_{E,K}$ in the class of evidence models satisfying (E1), the logic of evidence entailment and knowledge we obtain is a compound of two familiar logics together with one simple interaction axiom:

$$\text{EK} = \text{S5}_K + \text{KT}_E + (\text{KE}),$$

where (KE) denotes the axiom scheme $K\phi \rightarrow E\phi$.⁴

Theorem 1. *EK is a sound axiomatization of $\mathcal{L}_{E,K}$ with respect to the class of evidence models satisfying (E1).*

Our completeness proof relies on a standard Kripke-style interpretation of $\mathcal{L}_{E,K}$ in relational models and the completeness results pertaining thereto. We therefore begin with a brief review of these notions.

⁴This axiom system, as an extension of the normal modal logic KT with the universal modality, has previously been studied in [19] within the standard relational framework. The completeness results obtained therein will help us prove completeness with respect to evidence models.

A **relational evidence frame** is a pair (X, R_E) where X is a non-empty set and R_E is a reflexive, binary relation on X . A **relational evidence model** is a relational evidence frame equipped with a valuation function $v : \text{PROP} \rightarrow 2^X$. The language $\mathcal{L}_{E,K}$ is interpreted in a relational evidence model $M = (X, R_E, v)$ by extending the valuation function via the standard recursive clauses for the Boolean connectives together with the following:

$$\begin{aligned} (M, x) \models E\varphi & \text{ iff } R_E(x) \subseteq \|\varphi\|_M \\ (M, x) \models K\varphi & \text{ iff } X = \|\varphi\|_M, \end{aligned}$$

where $R_E(x) = \{y \in X : xR_E y\}$ and $\|\varphi\|_M = \{x \in X : (M, x) \models \varphi\}$. Thus E is interpreted by universal quantification over the R_E -accessible states (as usual for a box-type modality in standard relational semantics), while K is interpreted as a universal modality, as might be expected from Observation 1. We omit mention of M when the model is clear from context.

Theorem 2 ([19]). *EK is a sound and complete axiomatization of $\mathcal{L}_{E,K}$ with respect to the class of relational evidence models.*

Given a relational evidence model $M = (X, R_E, v)$, consider the tuple $\mathcal{M}_M = (X, \{e\}, I, v)$ where $I = \{I_e\}$ and for all $x \in X$, $I_e(x) = R_E(x)$. As shown in Observation 1, \mathcal{M}_M is an evidence model satisfying (E1), and $U_e = X$. In particular, every pair (x, e) is an evidence scenario in \mathcal{M}_M .

Lemma 3. *Let $M = (X, R_E, v)$ be a relational evidence model. Then for all $\varphi \in \mathcal{L}_{E,K}$ and $x \in X$, we have*

$$M, x \models \varphi \text{ iff } \mathcal{M}_M, (x, e) \models \varphi,$$

where $\mathcal{M}_M = (X, \{e\}, I, v)$ as described above.

Corollary 4. *EK is a complete axiomatization of $\mathcal{L}_{E,K}$ with respect to the class of evidence models satisfying (E1).*

Proof. This follows from Theorem 2 and Lemma 3: if $\varphi \in \mathcal{L}_{E,K}$ is such that $\not\models_{EK} \varphi$, then by Theorem 2 there is a relational evidence model M that refutes φ at some state x . Then, by Lemma 3, φ is also refuted in \mathcal{M}_M at the epistemic scenario (x, e) , which completes the proof. \square

4 Evidence Models for Belief

It is natural to wish to extend the framework we have developed to include a representation not only for knowledge but also *belief*. Defining this extension is relatively straightforward, as it parallels a similar construction from previous work [12]. The interest here arises not in the definition itself, but from the subsequent investigation into the interplay between belief and uncertainty about the interpretation of evidence.

A **doxastic evidence model** is simply an evidence model $\mathcal{M} = (X, \mathcal{E}, I, v)$ in which truth is evaluated with respect to **doxastic evidence scenarios**, which are tuples of the form (x, e, V) where (x, e) is an evidence scenario and $\emptyset \neq V \subseteq U_e$. The subset V is meant to capture the beliefs of the agent: that is, each $y \in V$ is a world the agent (subjectively) considers possible. Given an evidence model $\mathcal{M} = (X, \mathcal{E}, I, v)$ and a doxastic evidence scenario (x, e, V) , the semantic clauses for the primitive propositions and Boolean connectives are as before, while for the modalities we have:

$$\begin{aligned} (\mathcal{M}, x, e, V) \models E\varphi & \text{ iff } I_e(x) \subseteq \llbracket \varphi \rrbracket^{e, V} \\ (\mathcal{M}, x, e, V) \models K\varphi & \text{ iff } U_e \subseteq \llbracket \varphi \rrbracket^{e, V} \\ (\mathcal{M}, x, e, V) \models B\varphi & \text{ iff } V \subseteq \llbracket \varphi \rrbracket^{e, V} \end{aligned}$$

where $[[\varphi]]_{\mathcal{M}}^{e,V} = \{x \in U_e : (\mathcal{M}, x, e, V) \models \varphi\}$. Thus, K and E are interpreted essentially as before, while the belief modality B quantifies universally over V . Intuitively, the set V might be interpreted as the agent’s “conjecture” about how the world is, given the evidence e , and the requirements that $V \neq \emptyset$ and $V \subseteq U_e$ guarantee that the agent does not believe inconsistencies and they believe that their evidence coheres with the world, respectively.⁵ This corresponds to the very standard “knowledge-implies-belief” principle: that is, it makes valid the scheme $K\varphi \rightarrow B\varphi$. Note also that, just like U_e , the doxastic range V is state-independent, which guarantees the validity of the strong introspection principles given in Table 2 and Lemma 5.⁶

Another constraint one might impose on doxastic evidence scenarios (x, e, V) is the following:

$$(E2) \quad y \in V \Rightarrow I_e(y) \subseteq V.$$

Condition (E2) essentially stipulates that the agent takes evidence entailment seriously: if they consider it possible that the state of evidence leaves open those worlds in $I_e(y)$, then they consider each such world possible too. This validates the scheme $B\varphi \rightarrow BE\varphi$: if the agent believes φ then they believe that the evidence entails φ . So it’s a kind of “have responsible beliefs” constraint: you should only believe that which you believe is entailed by the evidence.

This condition bears a close resemblance to a principle suggested by Stalnaker [27], which he called “strong belief”, namely: $B\varphi \rightarrow BK\varphi$, if you believe φ then you believe that you know it. This essentially makes belief subjectively indistinguishable from knowledge. In the special context of subset space models (Observation 2), or more generally whenever the K and E modalities collapse, our (E2) principle just *is* Stalnaker’s “strong belief” principle. But in general (E2) is weaker: you may believe many things without believing that you *know* them—that is, that they are entailed by *every* interpretation of the evidence—instead, what (E2) says is that anything you believe is entailed by *those interpretations of the evidence that you consider possible*.

Even this weaker form of Stalnaker’s principle may seem too restrictive, however. Interestingly, it is possible to drop it as a *constraint* on doxastic evidence scenarios without abandoning the intuition entirely. Suppose (x, e, V) is a doxastic evidence scenario; let $V^1 = V$, and define, for $k > 1$,

$$V^k = \bigcup_{y \in V^{k-1}} I_e(y)$$

and

$$V^\infty = \bigcup_{k=1}^{\infty} V^k.$$

Then it is easy to see that $V^1 \subseteq V^2 \subseteq \dots$ is a nested increasing sequence of sets, and V^∞ actually does satisfy (E2)—in fact it’s the smallest set containing V with this property. We might then interpret V^∞ as representing the agent’s most “conservative” beliefs (so the fact that they satisfy the “responsibility” constraint, (E2), makes some sense), whereas $V = V^1$ represents the agent’s least conservative “conjecture”, with the sequence $V^2 \subseteq V^3 \subseteq \dots$ bridging the gap between these extremes in a series of discrete jumps or “levels” of belief. This is related to the idea of using *plausibility rankings* on possible worlds in order to produce a sequence of beliefs, starting with the “strongest” beliefs and gradually weakening them by including less plausible (though still possible) worlds [17, 22]. For example, if we apply this idea to

⁵One can also study agents with possibly inconsistent beliefs in a similar way by simply eliminating the requirement $V \neq \emptyset$.

⁶More general semantics for $\mathcal{L}_{E,K,B}$ that do not validate these introspection principles can be obtained by interpreting B with respect to a family of parametrized relations $R = \{R_e\}_{e \in \mathcal{E}}$, where each $R_e \subseteq X^2$, rather than a fixed V given in a doxastic evidence scenario. Due to the page limit, we leave the details of such a generalization for the extended version of this paper.

Example 1 starting with the initial conjecture $V^1 = (10, 20)$ (corresponding to the belief that the hand of the clock is pointing between the 2 and the 4), it is easy to see that $V^2 = (5, 25)$, $V^3 = (2.5, 27.5)$, \dots , and $V^\infty = (0, 30)$. A systematic development of this “ranked belief” framework in the context of evidence models is left to the full paper.

4.1 Soundness and Completeness for $\mathcal{L}_{E,K,B}$

In order to distinguish the semantics of $\mathcal{L}_{E,K,B}$ given with respect to doxastic evidence scenarios from those proposed for $\mathcal{L}_{E,K}$ in Section 3, we call the former *doxastic-evidence semantics*. Satisfiability and validity of a formula in doxastic-evidence semantics is defined the same way as given in Section 3.

The weakest logic of evidence, knowledge, and belief we consider in this paper, denoted EKB, is obtained by strengthening EK with the additional axiom schemes given in Table 2.

(K_B)	$B(\varphi \rightarrow \psi) \rightarrow (B\varphi \rightarrow B\psi)$	Distribution of belief
(D_B)	$B\varphi \rightarrow \neg B\neg\varphi$	Consistency of belief
(sPI)	$B\varphi \rightarrow KB\varphi$	Strong positive introspection
(KB)	$K\varphi \rightarrow B\varphi$	Knowledge implies belief

Table 2: Additional axiom schemes for EKB

Lemma 5. *Nec_B and $\neg B\varphi \rightarrow K\neg B\varphi$ (strong negative introspection) are derivable in EKB.*

Theorem 6. *EKB is a sound axiomatization of $\mathcal{L}_{E,K,B}$ with respect to the class of evidence models satisfying (E1) under doxastic-evidence semantics.*

The completeness proof again relies on a standard Kripke-style interpretation of $\mathcal{L}_{E,K,B}$ in relational models and the corresponding relational completeness result.

A **relational doxastic evidence model** $M = (X, R_E, R_B, \nu)$ is a relational evidence model (X, R_E, ν) equipped with an additional binary relation R_B on X such that for all $x, y \in X$, $R_B(x) \neq \emptyset$ and $R_B(x) = R_B(y)$. The language $\mathcal{L}_{E,K,B}$ is interpreted in a relational doxastic evidence model $M = (X, R_E, R_B, \nu)$ as before for E and K (see Section 3.1); for B we have:

$$(M, x) \models B\varphi \quad \text{iff} \quad R_B(x) \subseteq \|\varphi\|_M.$$

As usual we omit mention of the model when it is clear from context.

Theorem 7. *EKB is a sound and complete axiomatization of $\mathcal{L}_{E,K,B}$ with respect to the class of relational doxastic evidence models.*

Proof. While soundness is a matter of routine validity check, completeness follows from a fairly straightforward canonical model construction where (sPI) guarantees that for all $x, y \in X$, $R_B(x) = R_B(y)$ in the canonical model (see, e.g., [13, Chapters 4 & 7]). For a similar construction for topological subset space semantics, see [12, pp. 20-21, full paper]. \square

Given a doxastic relational evidence model $M = (X, R_E, R_B, \nu)$, we construct the evidence model $\mathcal{M}_M = (X, \{e\}, I, \nu)$ satisfying (E1) exactly the same way as in Section 3.1. Let $V = R_B(x)$ for any $x \in X$ and recall that $U_e = X$. Therefore, as $\emptyset \neq V \subseteq U_e$, every tuple of the form (x, e, V) is a doxastic evidence scenario in \mathcal{M}_M .

Lemma 8. *Let $\mathcal{M} = (X, R_E, R_B, \nu)$ be a relational doxastic evidence model. Then, for all $\varphi \in \mathcal{L}_{E,K,B}$ and $x \in X$, we have*

$$M, x \models \varphi \text{ iff } \mathcal{M}_M, (x, e, V) \models \varphi.$$

Corollary 9. *EKB is a complete axiomatization of $\mathcal{L}_{E,K,B}$ with respect to the class of evidence models satisfying (E1) under doxastic-evidence semantics.*

Proof. Similar to the proof of Corollary 4, by Theorem 7 and Lemma 8. \square

We also provide an axiomatization of $\mathcal{L}_{E,K,B}$ for evidence models that satisfy (E2) in addition to (E1).

Theorem 10. *EKB + $(B\varphi \rightarrow BE\varphi)$ is a sound and complete axiomatization of $\mathcal{L}_{E,K,B}$ with respect to the class of evidence models satisfying (E1) and (E2) under doxastic-evidence semantics.*

5 Evidence Models for Knowability

The logics we have considered so far have been *static* in the sense that they include no mechanism for an agent to update their information in any way. As a first step toward introducing a dynamic component to our setting, we consider a simple mechanic for *changing* the state of evidence. Perhaps the simplest intuition comes from the case of an agent who takes multiple successive measurements—assuming they remember the results of previous measurements, it seems reasonable to represent the final state of evidence as a combination $e_1 \oplus \dots \oplus e_k$, where e_i is the evidence state corresponding to the i th observation.

This is captured formally in the definition of an **evidence interaction model**, which is a tuple $\mathcal{M} = (X, \mathcal{E}, \oplus, I, \nu)$ where (\mathcal{E}, \oplus) is a meet-semilattice, (X, \mathcal{E}, I, ν) is an evidence model satisfying (E1), and for all $x \in X$ and finite $\mathcal{E}' \subseteq \mathcal{E}$, $I_{\oplus \mathcal{E}'}(x) = \bigcap_{e \in \mathcal{E}'} I_e(x)$. A notion of *evidence parthood*, denoted by \leq , is given by

$$\forall e', e \in \mathcal{E} (e' \leq e \text{ iff } e' \oplus e = e'). \quad (\text{EP})$$

Moreover, it is not difficult to see that for all finite $\mathcal{E}' \subseteq \mathcal{E}$, $U_{\oplus \mathcal{E}'} = \bigcap_{e \in \mathcal{E}'} U_e$.

Note the analogy with topological spaces. A topological space has the form (X, \mathcal{T}) , where \mathcal{T} is a collection of subsets of X called *opens*, often conceived of as the results of possible measurements. Evidence spaces effectively replace the topology \mathcal{T} with the structure (\mathcal{E}, I) , so that in place of open subsets U of X we have *families* of subsets $\{I_e(x) : x \in X\}$ of X , one for each $e \in \mathcal{E}$. Loosely speaking, topological spaces might be viewed as special cases of evidence spaces where each I_e is a constant function (cf. Observation 2).

This analogy is taken a step further with evidence interaction models, since the closure of \mathcal{E} under the meet operation \oplus parallels the closure of \mathcal{T} under intersection. Thus, it may be easier to think of (\mathcal{E}, \oplus, I) as the analog of a *basis* for X , rather than a full topology. We can also define a kind of *generalized interior operator* in evidence interactions models, and use it to articulate a notion of *measurability* corresponding to what the agent *could come to know* after taking a sufficiently good measurement or otherwise obtaining a sufficiently strong piece of evidence (see [11, 12, 10]). Given an evidence interaction model $(X, \mathcal{E}, \oplus, I, \nu)$ and an evidence scenario (x, e) , we interpret the propositional variables, Boolean connectives, K , and E as before, and for \square we define

$$(\mathcal{M}, x, e) \models \square \varphi \text{ iff } \exists e' \in \mathcal{E} (x \in U_{e \oplus e'} \subseteq \llbracket \varphi \rrbracket^e).$$

Thus, $\square \varphi$ holds just in case there is some piece of evidence e' that, when combined with the agent's current evidence e , would result in knowledge of $\llbracket \varphi \rrbracket^e$.

5.1 Soundness and Completeness for $\mathcal{L}_{E,K,\square}$

The logic of evidence, knowledge, and knowability is obtained by strengthening EK as follows

$$\text{EKK} = \text{EK} + \text{S4}_{\square} + (\text{K}\square),$$

where $(\text{K}\square)$ denotes the axiom scheme $K\phi \rightarrow \square\phi$.

Theorem 11. *EKK is a sound axiomatization of $\mathcal{L}_{E,K,\square}$ with respect to the class of evidence interaction models.*

Similarly to the previous completeness proofs, we prove the completeness of EKK via a detour to the standard relational interpretation of $\mathcal{L}_{E,K,\square}$ and its corresponding relational completeness. More precisely, we rely on the completeness of EKK—under the standard Kripke semantics—with respect to the class of finite models of the form $M = (X, R_E, R_{\square}, \nu)$ where R_E is reflexive and R_{\square} is reflexive and transitive. We call such structures **relational evidence and knowability models**. While K and E are interpreted in a relational evidence and knowability model as before, \square is interpreted, in the standard way, via the accessibility relation R_{\square} :

$$(M, x) \models \square\phi \quad \text{iff} \quad R_{\square}(x) \subseteq \|\phi\|_M.$$

Theorem 12. *EKK is a sound and complete axiomatization of $\mathcal{L}_{E,K,\square}$ with respect to the class of finite relational evidence and knowability models.*

We now construct an evidence interaction model $\mathcal{M}_M = (X, \mathcal{E}, \oplus, I, \nu)$ from a finite relational evidence and knowability model $M = (X, R_E, R_{\square}, \nu)$ in such a way that M and \mathcal{M}_M are point-wise modally equivalent with respect to $\mathcal{L}_{E,K,\square}$. While finiteness of the model is not essential, it will simplify our construction.

Let $M = (X, R_E, R_{\square}, \nu)$ be a finite relational evidence and knowability model and denote the set of all upsets of the preordered set (X, R_{\square}) by $Up(X)$. Since X is finite, $Up(X)$ is finite. Therefore, we can enumerate the elements of $Up(X)$ and write $Up(X) = \{E_1, \dots, E_n\}$. Note that $X \in Up(X)$, so, wlog, we let $X = E_n$. For each element in $Up(X)$, we put a corresponding element in $\mathcal{E} = \{e_1, \dots, e_n\}$ (so e_n is the evidence state corresponding to X —this will become clearer below). We then define an evidence parthood relation \leq on \mathcal{E} as

$$\forall e_i, e_j \in \mathcal{E} (e_i \leq e_j \text{ iff } E_i \subseteq E_j).$$

It is easy to see that (\mathcal{E}, \leq) is a poset with the top element e_n , that is, $e_j \leq e_n$ for all $e_j \in \mathcal{E}$. We can define the corresponding meet, $e_i \oplus e_j$, in a standard way as the greatest lower bound of $\{e_i, e_j\}$ with respect to \leq . More generally, for any finite $\mathcal{E}' \subseteq \mathcal{E}$, the element $\oplus \mathcal{E}'$ is the greatest lower bound of \mathcal{E}' with respect to \leq (see, e.g., [15] for a general introduction to lattice theory). Finally, for all $x \in X$ and $e_i \in \mathcal{E}$, set $I_{e_i}(x) = E_i \cap R_E(x)$. Notice that, as $E_n = X$, we have $I_{e_n}(x) = R_E(x)$.

Lemma 13. *Given a finite relational evidence and knowability model $M = (X, R_E, R_{\square}, \nu)$, the structure $\mathcal{M}_M = (X, \mathcal{E}, \oplus, I, \nu)$ constructed in the above described way is an evidence interaction model.*

Lemma 14. *Let $M = (X, R_E, R_{\square}, \nu)$ be a finite relational evidence and knowability model. Then, for all $\phi \in \mathcal{L}_{E,K,\square}$ and $x \in X$, we have*

$$M, x \models \phi \text{ iff } (\mathcal{M}_M, x, e_n) \models \phi.$$

Corollary 15. *EKK is a complete axiomatization of $\mathcal{L}_{E,K,\square}$ with respect to the class of evidence interaction models.*

Proof. Similar to the proof of Corollary 4, by Theorem 12 and Lemma 14. □

Theorem 16. *S4_{\square} is a sound and complete axiomatization of \mathcal{L}_{\square} with respect to the class of evidence interaction models.*

6 Further Work

We introduced *evidence models* as a means of representing agents who may be uncertain about what their evidence actually entails. We also explored some extensions of this framework that include belief and knowability. There are many interesting avenues to continue this line of work. From a philosophical angle, we believe the framework we have developed here is well-suited to the analysis of a variety of conceptual puzzles that arise when less flexible models of evidence entailment are implicitly relied upon, while on the more mathematical side, it is clear that the logical systems we defined have a variety of natural extensions.

In Section 4, for example, we outlined a way of using evidence interpretations to extend an agent’s initial conjecture to a graded notion of belief/plausibility. And in Section 5, the account of knowability we provided only scratched the surface of the potential for developing fully dynamic logics atop this foundation. Consider a public announcement style update mechanic in which knowability plays the role of the precondition of the corresponding announcement, as in [11]. The effect of an announcement is then manifested as a transition from the initial evidence state to a more informative one, without requiring global changes in the given model, as in logics of information dynamics interpreted on subset space models [28, 12, 16, 6, 3]. The enriched structure owing to the evidence states and their variable interpretations raises the question of whether such a dynamic logic can be reduced to a weaker, static logic, as is often the case in similar settings.

When we view evidence models as a generalization of subset space models (recall Observation 2), another natural dynamic extension suggests itself: adding the so-called *effort modality*, the trademark of subset space logics. The effort modality, denoted here by $\diamond\varphi$, is intended to capture a notion of “epistemic effort”, such as taking further measurements, and might be read in the present context as “ φ becomes true after some further evidence intake”. It can then be naturally interpreted on evidence interaction models as

$$(x, e) \models \diamond\varphi \text{ iff } (\exists e' \in \mathcal{E})(x \in U_{e \oplus e'} \text{ and } (x, e \oplus e') \models \varphi).$$

Incorporating such an operator in the current setting would provide a formal framework in which we could study a truly dynamic notion of knowability via the scheme $\diamond K\varphi$, as opposed to its static counterpart $\square\varphi$ (see also [11] for a discussion of the same issue in topological subset space semantics). Moreover, the relationship between \diamond and our static modalities K, B, E , and \square could help further the research on dynamic logics for topological formal learning theory [23], initiated by [4, 5] and further developed within subset space style logics in [3]. Such investigations are the subject of ongoing research.

Acknowledgements

We thank the anonymous reviewers of TARK 2019 for their valuable comments. Aybuke Özgün’s research was funded by the European Research Council (ERC CoG), Consolidator grant no. 681404, ‘The Logic of Conceivability’.

References

- [1] Carlos E. Alchourrón, Peter Gärdenfors & David Makinson (1985): *On the Logic of Theory Change: Partial Meet Contraction and Revision Functions*. *Journal of Symbolic Logic* 50(2), pp. 510–530, doi:10.2307/2274239.

- [2] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün & Sonja Smets (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, doi:10.1007/978-3-662-52921-8_6.
- [3] Alexandru Baltag, Nina Gierasimczuk, Aybüke Özgün, Ana Lucia Vargas-Sandoval & Sonja Smets (2018): *A Dynamic Logic for Learning Theory*. In Alexandre Madeira & Mário Benevides, editors: *Dynamic Logic. New Trends and Applications*, Springer International Publishing, Cham, pp. 35–54, doi:10.1007/978-3-319-73579-5_3.
- [4] Alexandru Baltag, Nina Gierasimczuk & Sonja Smets (2011): *Belief Revision as a Truth-Tracking Process*. In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2011)*, ACM, pp. 187–190, doi:10.1145/2000378.2000400.
- [5] Alexandru Baltag, Nina Gierasimczuk & Sonja Smets (2015): *On the Solvability of Inductive Problems: A Study in Epistemic Topology*. In R. Ramanujam, editor: *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015)*, pp. 81–98, doi:10.4204/EPTCS.215.7.
- [6] Alexandru Baltag, Aybüke Özgün & Ana Lucia Vargas-Sandoval (2017): *Topo-Logic as a Dynamic-Epistemic Logic*. In: *Proceedings of the 6th LORI*, Springer Berlin Heidelberg, pp. 330–346, doi:10.1007/978-3-662-55665-8_23.
- [7] Johan van Benthem, David Fernández-Duque & Eric Pacuit (2012): *Evidence logic: A new look at neighborhood structures*. In: *Advances in Modal Logic 9*, King's College Press, pp. 97–118. Available at <http://www.aiml.net/volumes/volume9/Benthem-Fernandez-Duque-Pacuit.pdf>.
- [8] Johan van Benthem, David Fernández-Duque & Eric Pacuit (2014): *Evidence and plausibility in neighborhood structures*. *Annals of Pure and Applied Logic* 165(1), pp. 106–133, doi:10.1016/j.apal.2013.07.007.
- [9] Johan van Benthem & Eric Pacuit (2011): *Dynamic Logics of Evidence-Based Beliefs*. *Studia Logica* 99(1), pp. 61–92, doi:10.1007/s11225-011-9347-x.
- [10] Adam Bjorndahl (2018): *The Epistemology of Nondeterminism*. In: *Proceedings of the 25th Workshop on Logic, Language, Information and Computation (WoLLIC 2018)*, pp. 145–162, doi:10.1007/978-3-662-57669-4_8.
- [11] Adam Bjorndahl (2018): *Topological Subset Space Models for Public Announcements*. In Hans van Ditmarsch & Gabriel Sandu, editors: *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics*, Springer International Publishing, Cham, pp. 165–186, doi:10.1007/978-3-319-62864-6_6.
- [12] Adam Bjorndahl & Aybüke Özgün (2017): *Logic and Topology for Knowledge, Knowability, and Belief - Extended Abstract*. In: *Proceedings of 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017), Electronic Proceedings in Theoretical Computer Science 251*, Open Publishing Association, pp. 88–101, doi:10.4204/EPTCS.251.7.
- [13] Patrick Blackburn, Maarten de Rijke & Yde Venema (2001): *Modal Logic*. *Cambridge Tracts in Theoretical Computer Science* 53, Cambridge University Press, Cambridge, doi:10.1017/CBO9781107050884.
- [14] Andrew Dabrowski, Lawrence S. Moss & Rohit Parikh (1996): *Topological reasoning and the logic of knowledge*. *Annals of Pure and Applied Logic* 78(1), pp. 73–110, doi:10.1016/0168-0072(95)00016-X.
- [15] Brian A. Davey & Hilary A. Priestley (1990): *Introduction to lattices and order*. Cambridge University Press, Cambridge, doi:10.1017/CBO9780511809088.
- [16] Hans van Ditmarsch, Sophia Knight & Aybüke Özgün (2017): *Announcement as effort on topological spaces*, doi:10.1007/s11229-017-1592-8. Forthcoming in *Synthese*.
- [17] Hans P. Van Ditmarsch (2005): *Prolegomena to Dynamic Logic for Belief Revision*. *Synthese* 147(2), pp. 229–275, doi:10.1007/s11229-005-1349-7.
- [18] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (1995): *Reasoning About Knowledge*. MIT Press, doi:10.7551/mitpress/5803.001.0001.
- [19] Valentin Goranko & Solomon Passy (1992): *Using the Universal Modality: Gains and Questions*. *Journal of Logic and Computation* 2(1), pp. 5–30, doi:10.1093/logcom/2.1.5.

- [20] Adam Grove (1988): *Two modellings for theory change*. *Journal of Philosophical Logic* 17(2), pp. 157–170, doi:10.1007/BF00247909.
- [21] Joseph Y. Halpern (2003): *Reasoning About Uncertainty*. MIT Press, Cambridge, MA, USA, doi:10.7551/mitpress/10951.001.0001.
- [22] Wiebe van der Hoek & John-Jules Meyer (1992): *Graded modalities in epistemic logic*. In Anil Nerode & Mikhail Taitlin, editors: *Logical Foundations of Computer Science — Tver '92*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 503–514, doi:10.1007/BFb0023902.
- [23] Kevin T. Kelly (1996): *The Logic of Reliable Inquiry*. Oxford University Press.
- [24] Lawrence S. Moss & Rohit Parikh (1992): *Topological Reasoning and The Logic of Knowledge*. In: *Proceedings of the 4th Conference on Theoretical Aspects of Computer Science (TARK 1992)*, Morgan Kaufmann, pp. 95–105. Available at <http://dblp.org/rec/html/conf/tark/MossP92>.
- [25] Aybüke Özgün (2017): *Evidence in Epistemic Logic: A Topological Perspective*. Ph.D. thesis, University of Amsterdam & Université de Lorraine. Available at <https://www.illc.uva.nl/Research/Publications/Dissertations/DS-2017-07.text.pdf>.
- [26] Rohit Parikh, Lawrence S. Moss & Chris Steinsvold (2007): *Topology and Epistemic Logic*. In: *Handbook of Spatial Logics*, Springer Verlag, pp. 299–341, doi:10.1007/978-1-4020-5587-4_6.
- [27] Robert Stalnaker (2006): *On Logics of Knowledge and Belief*. *Philosophical Studies* 128(1), pp. 169–199, doi:10.1007/s11098-005-4062-y.
- [28] Yi N. Wáng & Thomas Ågotnes (2013): *Subset Space Public Announcement Logic*. In: *Proceedings of the 5th Indian Conference on Logic and Its Applications (ICLA 2013)*, Springer, pp. 245–257, doi:10.1007/978-3-642-36039-8_22.
- [29] Timothy Williamson (2000): *Knowledge and its Limits*. Oxford University Press, doi:10.1093/019925656X.001.0001.