

Evidence Logic: A New Look at Neighborhood Structures

Johan van Benthem¹

*Institute for Logic, Language and Computation
University of Amsterdam
Department of Philosophy
Stanford University*

David Fernández-Duque²

*Department of Computer Science and Artificial Intelligence
University of Seville*

Eric Pacuit³

*Department of Philosophy
University of Maryland, College Park*

Abstract

Two of the authors (van Benthem and Pacuit) recently introduced *evidence logic* as a way to model epistemic agents faced with possibly contradictory evidence from different sources. For this the authors used neighborhood semantics, where a neighborhood N indicates that the agent has reason to believe that the true state of the world lies in N . A normal belief modality is defined in terms of the neighborhood structure. In this paper we consider four variants of evidence logic which hold for different classes of evidence models. For each of these logics we give a representation theorem using *extended evidence models*, where the belief operator is replaced by a standard relational modality. With this, we axiomatize all four logics, and determine whether each has the finite model property.

Keywords: Neighborhood Models, Logics of Belief, Combining Logics

1 Introduction

Neighborhood models are a generalization of the usual relational semantics for modal logic. In a neighborhood model, each state is assigned a collection of subsets of the set of states. Such structures provide a semantics for both normal

¹ johan.vanbenthem@uva.nl

² dfduque@us.es

³ epacuit@umd.ed

and non-normal modal logics. See [17] for an early discussion of neighborhood semantics for modal logic, and [6,12,7] for modern motivations and mathematical details. Concrete uses of neighborhood models include logics for knowledge [22], players' powers in games [13,14], concurrent PDL [15], and beliefs in the epistemic foundations of game theory [24,10].

In this paper, we study yet another concrete interpretation of neighborhood models. The idea is to interpret the neighborhood functions as describing the evidence that an agent has accepted (in general, we assume the agent accepts different evidence at different states). The agent then uses this evidence to form her beliefs. A dynamic extension of this logic of evidence-based belief was introduced in [21]. The main technical contribution of that paper is a series of *relative* completeness results in the style familiar to much of the literature on dynamic epistemic logic: validity in a language with dynamic modalities is reduced to a modal language without the dynamic modalities via the validity of so-called *recursion axioms* (see [19] for a general discussion of this technique). In this paper, we continue the project started in [21] focusing on the underlying static logics of belief and evidence.

To do this, we shall consider four variants of evidence logic, which depend on the fundamental assumptions one may make about evidence models. For each of the resulting logics, we prove two main results.

The first is a characterization theorem in terms of *extended evidence models*. These differ from evidence models only in that belief is interpreted via an explicit accessibility relation rather than in terms of the neighborhood structure. We show that, up to p -morphism, the class of extended evidence models and the class of evidence models are equivalent for each logic considered.

The second is to give a complete deductive calculus for each logic. Here our representation using extended evidence models is crucial, since it permits us to employ familiar techniques from modal logic.

2 A Logic of Evidence and Belief

We start by presenting our formal framework for evidence logics, leaving a more detailed discussion of its motivation to the end of this section. Given a set W of possible worlds or states, one of which represents the "actual" situation, an agent gathers evidence about this situation from a variety of sources. To simplify things, we assume these sources provide *binary* evidence, i.e., subsets of W which (may) contain the actual world. The agent uses this evidence (i.e., collection of subsets of W) to form her beliefs.

The following modal language can be used to describe what the agent believes given her available evidence (cf. [21]).

Definition 2.1 *Let At be a fixed set of atomic propositions. Let \mathcal{L} be the smallest set of formulas generated by the following grammar*

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid B\varphi \mid \Box\varphi \mid A\varphi$$

where $p \in \text{At}$. Additional propositional connectives ($\wedge, \rightarrow, \leftrightarrow$) are defined as

usual and the duals⁴ of \Box, B and A are \Diamond, \widehat{B} and \widehat{A} , respectively.

The intended interpretation of $\Box\varphi$ is “the agent has evidence for φ ” and $B\varphi$ says that “the agents believes that φ is true.” We also include the universal modality ($A\varphi$: “ φ is true in all states”) for technical convenience.⁵

Since we do not assume that the sources of the evidence are jointly consistent (or even that a single source is guaranteed to be consistent and provide *all* the available evidence), the “evidence for” operator ($\Box\varphi$) is not a normal modal operator. That is, the agent may have evidence for φ and evidence for ψ ($\Box\varphi \wedge \Box\psi$) without having evidence for their conjunction ($\neg\Box(\varphi \wedge \psi)$). Of course, both the belief and universal operators are normal modal operators. So, the logical system we study in this paper *combines* a non-normal modal logic with a normal one.

2.1 Neighborhood Models for \mathcal{L}

In the intended interpretation of evidence logic, there are many possible states of the world, and the agent possesses evidence for these states in the form of neighborhoods; all other epistemic operators are derived from this neighborhood structure.

Thus we define *evidence models* as follows:

Definition 2.2 *An evidence model is a tuple $\mathcal{M} = \langle W, E, V \rangle$, where W is a non-empty set of worlds, $E \subseteq W \times \wp(W)$ is an evidence relation, $V : \text{At} \rightarrow \wp(W)$ is a valuation function. We write $E(w)$ for the set $\{X \mid wEX\}$. Two constraints are imposed on the evidence sets: For each $w \in W$, $\emptyset \notin E(w)$ and $W \in E(w)$. A **uniform evidence model** is an evidence model where $E(\cdot)$ is a constant function (each state has the same set of evidence).*

We do not assume that the collection of evidence sets $E(w)$ is closed under supersets. Also, even though evidence pieces are non-empty, their combination through the obvious operations of taking *intersections* need not yield consistent evidence: we allow for disjoint evidence sets, whose combination may lead (and should lead) to trouble. But importantly, even though an agent may not be able to consistently combine *all* of her evidence, there will be maximal collections of admissible evidence that she can safely put together to form *scenarios*:

Definition 2.3 *A w -scenario is a maximal collection $\mathcal{X} \subseteq E(w)$ that has the fip (i.e., the finite intersection property: for each finite subfamily $\{X_1, \dots, X_n\} \subseteq \mathcal{X}$, $\bigcap_{1 \leq i \leq n} X_i \neq \emptyset$). A collection is called a **scenario** if it is a w -scenario for some state w .*

Truth of formulas in \mathcal{L} is defined as follows:

Definition 2.4 *Let $\mathcal{M} = \langle W, E, V \rangle$ be an evidence model. Truth of a formula $\varphi \in \mathcal{L}$ is defined inductively as follows:*

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ ($p \in \text{At}$)

⁴ In other words, $\widehat{B} = \neg B \neg$, and similarly for other operators.

⁵ A natural interpretation of $A\varphi$ in the context of this paper is “the agent knows that φ ”.

- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models \Box\varphi$ iff there exists X such that wEX and for all $v \in X$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models B\varphi$ for each w -scenario \mathcal{X} and for all $v \in \bigcap \mathcal{X}$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models A\varphi$ iff for all $v \in W$, $\mathcal{M}, v \models \varphi$

The truth set of φ is the set $\llbracket \varphi \rrbracket_{\mathcal{M}} = \{w \mid \mathcal{M}, w \models \varphi\}$. The standard logical notions of **satisfiability** and **validity** are defined as usual.

Our notion of having evidence for φ need not imply that the agent *believes* φ . In order to believe a proposition φ , the agent must consider *all* her evidence for or against φ . The idea is that each w -scenario represents a maximally consistent theory based on (some of) the evidence collected at w .⁶ Note that the definition of truth of the “evidence for” operator builds in monotonicity. That is, the agent has evidence for φ at w provided there is some evidence available at w that implies φ .

The class of evidence models we have described gives the most general setting such an agent may face. However, there are natural additional assumptions one may consider:

Definition 2.5 *An evidence model \mathcal{M} is **flat** if every scenario on \mathcal{M} has non-empty intersection. In the interest of brevity, we may write **b-evidence model** instead of flat evidence model.*

*An evidence model $\mathcal{M} = \langle W, E, V \rangle$ is **uniform** if E is a constant. In this case, we shall treat E as a set (of neighborhoods) rather than a function.*

Flatness and uniformity are natural assumptions which, as we will see, may be captured by adding different axioms to our logic.

2.2 The Logics

We now turn to logics for reasoning about distinct classes of evidence models. Our first observation is that the language \mathcal{L} is sensitive to flatness:

Lemma 2.6 *If \mathcal{M} is a b-evidence model, then $\mathcal{M} \models \Box\varphi \rightarrow \widehat{B}\varphi$.*

Proof. If $X \in E(w)$ is an evidence set witnessing φ (i.e., $X \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}$), then the singleton $\{X\}$ can be extended to a w -scenario using Zorn’s Lemma, which, in flat structures, has non-empty intersection. \square

Meanwhile, the formula $\Box\varphi \rightarrow \widehat{B}\varphi$ is not valid in general: Consider a uniform evidence model $\mathcal{M}_{\infty} = \langle W, E, V \rangle$ with domain $W = \mathbb{N}$ and evidence sets $E(w) = \{\{N, \infty\} \mid N \in \mathbb{N}\}$ for each $w \in W$. The valuation is unimportant, so we may let $V = \emptyset$. Clearly, the only scenario on \mathcal{M}_{∞} is all of E , but $\bigcap E = \emptyset$. Hence $\mathcal{M}_{\infty} \models B\perp$, i.e., $\mathcal{M}_{\infty} \not\models \widehat{B}\top$; yet $\mathcal{M}_{\infty} \models \Box\top$ (this formula is universally valid), and we conclude

$$\mathcal{M}_{\infty} \not\models \Box\top \rightarrow \widehat{B}\top.$$

⁶ Analogous ideas occur in semantics of conditionals [9,23] and belief revision [3,16].

From this we get the following corollary:

Corollary 2.7 *The logic of evidence models does not have the finite model property, nor does the logic of uniform evidence models.*

Proof. Every finite model is flat, and hence validates $\Box\top \rightarrow \widehat{B}\top$; but as we have just shown, this formula is not valid over all uniform evidence models. \square

With this in mind, we state a list of axioms and rules for evidence logics:

taut	all propositional tautologies
$S5_A$	S5 axioms for A
K_B	K axioms for B
\top-evidence	$\Box\top$
pullout	$\Box\varphi \wedge A\psi \leftrightarrow \Box(\varphi \wedge A\psi)$
universal belief	$A\varphi \rightarrow BA\varphi$
\Box-monotonicity	$\frac{\varphi \rightarrow \psi}{\Box\varphi \rightarrow \Box\psi}$
b	$\Box\varphi \rightarrow \widehat{B}\varphi$
\bigcirc-uniformity	$\bigcirc\varphi \rightarrow A\bigcirc\varphi$ for $\bigcirc = B, \widehat{B}, \Box, \Diamond$
MP	Modus Ponens
N_{\bigcirc}	Necessitation for $\bigcirc = A, B$

We let **Log** denote the logic which uses all axioms and rules *except* for **b** or the uniformity axioms. The subscripts **b**, **u** denote the addition of the respective axioms. We will denote derivability in the logic λ by \vdash_λ , where λ is any one of the four combinations that we may form, that is, $\lambda \in \{\mathbf{Log}, \mathbf{Log}_b, \mathbf{Log}_u, \mathbf{Log}_{bu}\}$.

The weakest logic **Log** will be called *general* evidence logic, while $\mathbf{Log}_b, \mathbf{Log}_{bu}$ will be called *flat logics* and $\mathbf{Log}_u, \mathbf{Log}_{bu}$ will be called *uniform logics*. We will write λ -*consistency* for consistency over the logic λ .

2.3 Extended Neighborhood Models for \mathcal{L}

The finite model property fails for evidence logic, a fact which may be rather inconvenient. Fortunately, there is a way to sidestep this problem.

In evidence models, the belief operator is interpreted using the neighborhoods by taking intersections of scenarios. There is another natural class of models for the language \mathcal{L} which avoids the use of scenarios. The key idea is

to extend a neighborhood model with a relation R that will be used to interpret the belief modality. These models will allow us, later, to employ standard techniques from modal logic to prove completeness.

Definition 2.8 *An extended evidence model is a structure $\mathfrak{M} = \langle W, R, E, V \rangle$ where such that W is a set, R an (arbitrary) binary relation on W , $E \subseteq W \times \wp(W)$ is a relation such that $w E W$ and $w \not E \emptyset$ for all $w \in W$, and $V : \text{At} \rightarrow \wp(W)$.*

Truth in an extended evidence model is defined in the standard way: Boolean connectives are as usual, $\mathfrak{M}, w \models A\varphi$ iff for all $v \in W$, $\mathcal{M}, v \models \varphi$; $\mathfrak{M}, w \models B\varphi$ iff for all $v \in W$, if wRv then $\mathcal{M}, v \models \varphi$; and $\mathfrak{M}, w \models \Box\varphi$ iff there exists a $X \subseteq W$ such that $w E X$ and for all $v \in X$, $\mathcal{M}, v \models \varphi$. We write $\llbracket \varphi \rrbracket_{\mathfrak{M}} = \{w \mid \mathfrak{M}, w \models \varphi\}$ for the truth set of φ in \mathfrak{M} .

As we shall see below, for every evidence model \mathcal{M} there is an extended evidence model which is naturally associated to it and satisfies the same \mathcal{L} -formulas. But first, let us also define a notion of “flatness” for extended evidence models:

Definition 2.9 *Given $w E X$, we will say X is **flat for w** if there is $v \in X$ such that $w R v$. The set X is **flat** if it is flat for all $w \in W$ such that $w E X$. The extended evidence model \mathfrak{M} is **flat** if every evidence set is flat.*

For uniform extended evidence models, we will also demand that the accessibility relation be independent of the current state:

Definition 2.10 *An extended evidence model $\mathfrak{M} = \langle W, R, E, V \rangle$ is uniform if both R and E are constant; that is, given $w, v, u \in W$ and $X \subseteq W$, $w R u$ if and only if $v R u$ and, likewise, $w E X$ if and only if $v E X$.*

2.4 Motivating the Logics

Having stated our formal framework, we digress from the main technical goal of this paper to discuss the intended interpretation of evidence models in some more detail. The reader who is interested in technical aspects only can safely skip ahead to the next section.

In a number of areas, the need has long been recognized for models that keep track of the reasons, or the *evidence* for beliefs and other informational attitudes (cf. [11,8]). At one extreme, the evidence is encoded as the current range of worlds the agent considers possible. However, this ignores how the agent arrived at this epistemic state. At the other extreme, models record the complete syntactic details of what the agent has learned so far (including the precise formulation and sources for each piece of evidence). In this paper, we explore an intermediate level, viz. neighborhood structures, where evidence is recorded as a family of sets of worlds. In particular, we want to mention three general issues:

We start by making explicit the underlying assumptions motivating the logical framework defined in Section 2.1. Let W be a set of states (or possible worlds) one of which represents the “actual” state. We are interested in a situation where an agent gathers evidence about this state from a variety of

sources. To simplify things, we assume these sources provide *binary* evidence, i.e., subsets of W which (may) contain the actual world. The following basic assumptions are implicit in the above definitions:

- (i) Sources may or may not be *reliable*: a subset recording a piece of evidence need not contain the actual world. Also, agents need not know which evidence is reliable.
- (ii) The evidence gathered from different sources (or even the same source) may be jointly inconsistent. And so, the intersection of all the gathered evidence may be empty.
- (iii) Despite the fact that sources may not be reliable or jointly inconsistent, they are all the agent has for forming beliefs.⁷

The *evidential state* of the agent is the set of all propositions (i.e., subsets of W) identified by the agent’s sources. In general, this could be any collection of subsets of W ; but we do impose some constraints:


- No evidence set is empty (evidence per se is never contradictory),
- The whole universe W is an evidence set (agents know their ‘space’).

In addition, one might expect a ‘monotonicity’ assumption:

If the agent has evidence X and $X \subseteq Y$ then the agent has evidence Y .

To us, however, this is a property of propositions supported by evidence, not of the evidence itself. Therefore, we model this feature differently through the definition of our “evidence for” modality (\square).

This brings us to a second point of discussion. The evidence models discussed in this paper do not *directly* represent the agent’s sources of evidence. However, the neighborhood models can be used to distinguish between a wide range of evidential situations. Consider the following three evidential states:



In each state, the agent has evidence that $p \vee q$ is true (but not both). However, the evidential situation underlying each state is very different. In the first situation, the agent has received the information from a single source (which the agent “trusts”) that the actual state satisfies either $p \wedge \neg q$ or $\neg p \wedge q$. In the second situation, the agent has received the same information from two

⁷ Modeling sources and agents’ *trust* in these is possible – but we will not pursue this here.

different sources (or perhaps the same source reinforced its evidence). The sources agree that either $p \wedge \neg q$ or $\neg p \wedge q$ is true, but disagree about the conditions under which these formulas are true. Finally, the most interesting case is the third one where the agent received information from two sources that agree on the conditions under which $p \wedge \neg q$ is true, but disagree about the conditions needed to make $\neg p \wedge q$ true. In this case, the agent believes p , whereas in the first two situations the agent only believes the weaker proposition that $p \vee q$ is true. Note that our language (see Definition 2.1) can distinguish between the third situation and the first two, but cannot distinguish between the first two evidential situations.

Our final point of discussion concerns the relationship between evidence models and existing modal logics of knowledge and belief based on so-called “plausibility models”.⁸ The models proposed here are not intended to *replace* plausibility models, but rather to complement them. So, what exactly is the relationship between these two frameworks for modeling beliefs? This question is explored in detail in [21, Section 5], but we only mention the highlights here. Every plausibility model can be transformed into a *uniform* evidence model where the set of evidence are all the downward closed subsets (according to the plausibility ordering). This, in part, motivates our interest in the logic of uniform evidence models. There is also a translation from evidence models to plausibility models (using the well-known definition of a *specialization (pre)-order*). However, it is easy to see that not every evidence model, not even a uniform one, comes from a plausibility model. So, evidence models generalize the standard plausibility models which have been successfully used to represent an agent’s knowledge and different flavors of belief. Once again, this shows the additional level of generality provided by neighborhood models.⁹

3 Representation Theorems

It is generally more convenient to work with extended evidence models than with standard evidence models, since it is easier to control accessibility relations than scenarios. Fortunately, as we shall show in this section, the two classes of models are equivalent with respect to our logics.

As it turns out, any evidence model may be represented as an extended evidence model with the same truth sets. The converse does not hold; yet, every extended evidence model is a p -morphic image of an evidence model.

The first claim is straightforward to check given the following construction.

Definition 3.1 *Given an evidence model $\mathcal{M} = \langle W, E, V \rangle$, define an extended evidence model $\mathfrak{M}^* = \langle W, R_E, E, V \rangle$ where $w R_E v$ if and only if v lies in the intersection of some w -scenario.*

⁸ A *plausibility* model is a tuple $\langle W, \preceq, V \rangle$ where W is a nonempty set, V is a valuation function and \preceq is a reflexive, transitive and well-founded order on W . We assume the reader is familiar with these well-studied models and the modal languages used to reason about them (see [19] for details and pointers to the relevant literature).

⁹ Both transformations extend to ternary world-dependent plausibility relations.

The following result is obvious by the definition of R_E and we present it without proof:

Theorem 3.2 *Given any evidence model \mathcal{M} and any formula φ ,*

$$\llbracket \varphi \rrbracket_{\mathcal{M}^*} = \llbracket \varphi \rrbracket_{\mathcal{M}}.$$

Further, it is immediate that if \mathcal{M} is uniform then so is \mathcal{M}^* . Flatness is also preserved by this operation; if $w E X$, then use Zorn's lemma to extend $\{X\}$ to a w -scenario \mathcal{X} . Then, if \mathcal{M} is flat, $\bigcap \mathcal{X} \neq \emptyset$, and hence for any $v \in \bigcap \mathcal{X}$ we have that $w R_E v \in X$.

Thus every (flat, uniform) evidence model can be represented as a (flat, uniform) extended evidence model. The opposite is true as well, but a bit more subtle. For this, we first recall the definition of a p-morphism:

Definition 3.3 *Given extended evidence models $\mathfrak{M}_1 = \langle W_1, R_1, E_1, V_1 \rangle$ and $\mathfrak{M}_2 = \langle W_2, R_2, E_2, V_2 \rangle$, we say a function $\pi : W_1 \rightarrow W_2$ is a p-morphism if the following conditions hold:*

atoms $V_1 = \pi^{-1}V_2$

forth $_R$ if $w R_1 v$ then $\pi(w) R_2 \pi(v)$

back $_R$ if $\pi(w) R_2 u$ then there is $v \in \pi^{-1}(u)$ such that $w R_1 v$

forth $_E$ if $w E_1 X$ then there is $Y \subseteq W_2$ such that $\pi(w) E_2 Y$ and $\pi[X] = Y$

back $_E$ if $\pi(w) E_2 Y$ then there is X such that $\pi[X] = Y$ and $w E_1 X$

Then we obtain the following familiar result which we present without proof:

Theorem 3.4 *If π is a p-morphism between extended evidence models \mathfrak{M}_1 and \mathfrak{M}_2 and φ is any formula, then*

$$\llbracket \varphi \rrbracket_{\mathfrak{M}_1} = \pi^{-1}[\llbracket \varphi \rrbracket_{\mathfrak{M}_2}].$$

If a surjective p-morphism exists from \mathfrak{M}_1 to \mathfrak{M}_2 , we will write $\mathfrak{M}_1 \gg \mathfrak{M}_2$. Our goal is to show now that, given an extended evidence model \mathfrak{M} , there is an evidence model \mathfrak{M}^+ such that $\mathfrak{M}^+ \gg \mathfrak{M}$; or, more precisely, $(\mathfrak{M}^+)^* \gg \mathfrak{M}$. Thus an extended evidence model may also be represented as a p-morphic image of an evidence model. The latter evidence model, however, is often much larger.

Definition 3.5 *Given an extended evidence model $\mathfrak{M} = \langle W, R, E, V \rangle$, we define an evidence model $\mathfrak{M}^+ = \langle W^+, E^+, V^+ \rangle$ and a map $\pi : W^+ \rightarrow W$ as follows:*

- (i) W^+ is the set of all triples $\langle w, X, n \rangle$ such that $X \subseteq W$ and either
 - (a) X is flat and $n \in \{0, 1\}$ or
 - (b) X is not flat and $n \in \mathbb{N}$.
- (ii) $\pi(\langle w, X, n \rangle) = w$
- (iii) $V^+(p) = \pi^{-1}[V(p)]$
- (iv) $E^+(\langle w, X, n \rangle) = \bigcup_{w E Y} \mathcal{B}^Y(w) \cup \{W^+\}$, where

(a) if Y is flat for w , then $\mathcal{B}^Y(w) = \{U_0, U_1\}$, with

$$U_i = \{\langle v, Y, 0 \rangle \mid \text{there is a } v \text{ with } w R v \text{ and } v \in Y\} \cup \{\langle v, Y, i \rangle \mid v \in Y\};$$

(b) if Y is not flat for w , then $\mathcal{B}^Y(w) = \{V_n\}_{n \in \mathbb{N}}$, where

$$V_N = \{\langle v, Y, n \rangle \mid v \in Y, n \geq N\}.$$

As a simple example, consider two extended evidence models, $\mathfrak{M}_1 = \langle W, R, E, V \rangle$ and $\mathfrak{M}_2 = \langle W, R', E, V \rangle$, so that the two are very similar, differing only on the accessibility relation. Suppose that $W = \{w\}$, $w R w$ but $R' = \emptyset$, and $E(w) = \{W\}$. The valuations are not too important for this example, so we may assume $V \equiv \emptyset$.

Then, \mathfrak{M}_1 is flat, but \mathfrak{M}_2 is not. Thus we have that $\mathfrak{M}_1^+ = \langle W_1^+, E_1^+, V_1^+ \rangle$, where W_1^+ consists of two copies of w ($\langle w, W, 0 \rangle$ and $\langle w, W, 1 \rangle$), but let us call them w_0, w_1 for simplicity). Both points have the same evidence sets, namely $\{w_0\}$ and $\{w_0, w_1\}$; thus the only scenario is $\{\{w_0\}, \{w_0, w_1\}\}$ which has intersection $\{w_0\}$.

Meanwhile, \mathfrak{M}_2^+ consists of countably many copies of w (of the form $\langle w, W, n \rangle$, but let us call them w_n) and is isomorphic to the model \mathcal{M}_∞ used in the proof of Corollary 2.7. As we saw then, the only scenario on \mathfrak{M}_2^+ has empty intersection.

More generally, it is always the case that flat, finite extended evidence models give rise to finite models, and uniform extended evidence models give rise to uniform evidence models:

Lemma 3.6 *If \mathfrak{M} is flat and finite, then \mathfrak{M}^+ is finite. Further, if \mathfrak{M} is uniform, then so is \mathfrak{M}^+*

Proof. If all evidence sets are flat, we have $W^+ = W \times \wp(W) \times \{0, 1\}$, which is a finite set (provided W is finite).

The second claim is easy to check using the definition of E^+ . \square

The following key lemma shows a close relationship between the accessibility relations on \mathfrak{M} and $(\mathfrak{M}^+)^*$.

Lemma 3.7 *Let $\mathfrak{M} = \langle W, R, E, V \rangle$ be an extended evidence model with associated evidence model \mathfrak{M}^+ . Let $\alpha \in W^+$ and $v \in W$ be arbitrary.*

Then, $\pi(\alpha) R v$ if and only if there is $\beta \in \pi^{-1}(v)$ such that $\alpha R_{E^+} \beta$.

The proof can be found in Appendix A. With this one can check that π is a p-morphism and $(\mathfrak{M}^+)^* \gg \mathfrak{M}$:

Theorem 3.8 *If \mathfrak{M} is an extended evidence model and π is the associated map for \mathfrak{M}^+ , then π is a surjective p-morphism between $(\mathfrak{M}^+)^*$ and \mathfrak{M} .*

Proof. The atoms clause holds by the definition of V^+ and the clauses for R hold by Lemma 3.7.

It remains to check that forth_E and back_E hold as well.

To check that forth_E holds, note that neighborhoods of $\alpha = \langle w, X, n \rangle$ are all of the form

$$N = \{\langle v, Y, m \rangle : v \in Y \text{ and } w E Y\};$$

but then, $\pi[N] = Y$ and hence $\pi(\alpha) = w E Y$.

Meanwhile, back_E holds because if $w E Y$ and $\pi(\alpha) = w$, we note that $\alpha E N$, where

$$N = Y \times \{Y\} \times \{0\}$$

and $\pi[N] = Y$. □

Corollary 3.9 *Given a extended evidence model \mathfrak{M} with map $\pi : W^+ \rightarrow W$ and a formula $\varphi \in \mathcal{L}_0$, $\llbracket \varphi \rrbracket_{\mathfrak{M}^+} = \pi^{-1}[\llbracket \varphi \rrbracket_{\mathfrak{M}}]$.*

Proof. Immediate from Theorems 3.8 and 3.4. □

From this we obtain the following, very useful result:

Theorem 3.10 *A set of formulas Φ is satisfiable on a (flat, uniform) extended evidence model if and only if it is satisfiable on a (flat, uniform) evidence model.*

4 Completeness

In view of the previous section, it suffices to build extended evidence models for consistent sets of formulas in order to prove completeness. Extended evidence models are much closer to standard semantics of modal logic than evidence models and hence we can apply familiar techniques.

We assume that all formulas are in ‘negation-normal form’ in which negations are only applied to propositional variables. Of course, in order to do this we must allow for dual operators ($\diamond, \widehat{B}, \widehat{A}$) to appear. Henceforth, we assume all formulas are in this form unless we explicitly indicate otherwise. Note that this convention is only for the sake of exposition, as we allow negation in our calculus and dual operators are really abbreviations.

We denote the normal-form negation of φ by $\sim\varphi$. The **closure** of φ , denoted $cl(\varphi)$ contains all subformulas of φ , is closed under normal-form negation (if $\psi \in cl(\varphi)$ then $\sim\psi \in cl(\varphi)$), and we stipulate that $\top, \square\top, \diamond\top$ and $B\top$ are all in $cl(\varphi)$. If Ω is a set of formulas, then $cl(\Omega) = \bigcup\{cl(\varphi) \mid \varphi \in \Omega\}$. As usual, the states in our canonical extended evidence model are maximally consistent sets of formulas (which we call types).

In the remainder of this paper, by an *evidence logic* we will mean exclusively an element of $\{\text{Log}, \text{Log}_s, \text{Log}_u, \text{Log}_{bu}\}$.

Definition 4.1 *Let Ω be a set of formulas and λ an evidence logic. An (Ω, λ) -type is a maximal λ -consistent subset of $cl(\Omega)$.*

A set of formulas Γ is a λ -type if it is a (Γ, λ) -type, i.e., Γ is λ -consistent and for each $\psi \in cl(\Gamma)$, either $\psi \in \Gamma$ or $\sim\psi \in \Gamma$.

Note that (Ω, λ) -types may be finite or infinite, depending on whether Ω is. We denote the set of (Ω, λ) -types by $\text{type}_\lambda(\Omega)$. Given an (Ω, λ) -type Φ , define Φ^A as $\{\psi \mid A\psi \in \Phi\}$, and similarly for the other modalities.

Of course, in a given model all points must satisfy the same universal formulas, so it is convenient to consider the collection of all such types.

Given a set of formulas Γ , if $\lambda \in \{\mathbf{Log}, \mathbf{Log}_b\}$, define $\Gamma^\lambda = A\Gamma^A \cup \widehat{A}\Gamma^{\widehat{A}}$; if $\lambda \in \{\mathbf{Log}_u, \mathbf{Log}_{bu}\}$, define

$$\Gamma^\lambda = \bigcup \{ \bigcirc \Gamma^\bigcirc : \bigcirc = A, \widehat{A}, B, \widehat{B}, \square, \diamond \}.$$

Definition 4.2 *Given a λ -type Φ , we define*

- (i) $\text{type}_\lambda^A(\Phi)$ to be the set of all (Φ, λ) -types Ψ such that $\Psi^\lambda = \Phi^\lambda$;
- (ii) $\text{type}_\lambda^B(\Phi) = \{ \Psi \in \text{type}_\lambda^A(\Phi) \mid \Phi^B \subseteq \Psi \}$.

Now we need to define evidence sets on our extended evidence models.

Definition 4.3 *Given a λ -type Φ with $\square\alpha \in \Phi$, we define the α -neighborhood of Φ as*

$$\mathcal{N}_\lambda^\alpha(\Phi) = \{ \Psi \in \text{type}_\lambda^A(\Phi) \mid \alpha \in \Psi \}.$$

We are now ready to define a canonical extended evidence model :

Definition 4.4 *Let Φ be a λ -type. The λ -canonical extended evidence model for Φ is the extended evidence model $\mathfrak{M}_\lambda(\Phi) = \langle W_\lambda^\Phi, E_\lambda^\Phi, R_\lambda^\Phi, V_\lambda^\Phi \rangle$ where*

- (i) $W_\lambda^\Phi = \text{type}_\lambda^A(\Phi)$
- (ii) For each $p \in \text{At} \cap \Gamma$, $\Gamma \in V_\lambda^\Phi(p)$ iff $p \in \Gamma$,
- (iii) $E_\lambda^\Phi(\Gamma) = \{ \mathcal{N}_\lambda^\alpha(\Gamma) \mid \square\alpha \in \Gamma \}$.
- (iv) $\Gamma R_\lambda^\Phi \Delta$ iff $\Delta \in \text{type}_\lambda^B(\Gamma)$ (i.e., $\Gamma^B \subseteq \Delta$).

Note that our ‘canonical’ extended evidence models are not unique as they depend on the type Φ we wish to satisfy and the logic λ we wish to work in; in this sense they could be considered ‘semicanonical’. Considering different extended evidence models is unavoidable, given that our language includes a universal modality and it is impossible to satisfy all types on a single model. However, we are also taking advantage of this in order to obtain the finite (extended) model property directly, since canonical extended evidence models for finite types are themselves finite.

Let us observe that canonical models for uniform types are uniform:

Lemma 4.5 *If $\lambda \in \{\mathbf{Log}_u, \mathbf{Log}_{bu}\}$ and Φ is λ -consistent, then the model $\mathfrak{M}_\lambda(\Phi)$ is uniform.*

Proof. Given $\Psi \in W_\lambda^\Phi$, we have that both the accessible states and the neighborhoods of Ψ depend only on Ψ^λ , which is constant amongst all of W_λ^Φ . \square

Our goal will now be to prove a version of the Truth Lemma which will imply that the canonical extended evidence models we have defined satisfy the right formulas. For this, we need some preliminaries. We start by gathering the main results about the axiom system that we use to prove the Truth Lemma. We defer proofs to the Appendix.

Lemma 4.6 *Let Γ, Δ be sets of formulas and λ an evidence logic.*

- (i) *Suppose that $\varphi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \psi$. Then, $\Box\varphi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \Box\psi$.*
- (ii) *Suppose that $\Phi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \psi$. Then, $B\Phi, A\Gamma, \widehat{A}\Delta \vdash_\lambda B\psi$.*

It will be very useful to draw some dual conclusions from the above result.

Lemma 4.7 *Let Γ be a set of formulas and λ an evidence logic.*

- (i) *If Γ is λ -consistent, $\alpha \in \Gamma^\square$ and $\delta \in \Gamma^\diamond$, then $\{\alpha, \delta\} \cup \Gamma^\lambda$ is λ -consistent.*
- (ii) *If Γ is λ -consistent and $\psi \in \Gamma^{\widehat{B}}$, then $\{\psi\} \cup \Gamma^B \cup \Gamma^\lambda$ is λ -consistent.*
- (iii) *If $\lambda \in \{\text{Log}_b, \text{Log}_{bu}\}$, Γ is λ -consistent and $\psi \in \Gamma^\square$, then $\{\psi\} \cup \Gamma^B \cup \Gamma^\lambda$ is λ -consistent.*

The next lemmas show the key steps in the proof of the Truth Lemma.

Lemma 4.8 *Let Γ be a (Φ, λ) -type.*

- (i) *If $\Box\alpha, \diamond\beta \in \Gamma$, there is $\Delta \in \mathcal{N}_\lambda^\alpha(\Gamma)$ with $\beta \in \Delta$.*
- (ii) *If $\lambda \in \{\text{Log}_b, \text{Log}_{bu}\}$ and $\alpha \in \Gamma^\square$, then $\mathcal{N}_\lambda^\alpha(\Gamma) \cap \text{type}_\lambda^B(\Gamma)$ is non-empty.*
- (iii) *Suppose that $\widehat{B}\alpha \in \Gamma$. Then, there is $\Delta \in \text{type}_\lambda^B(\Gamma)$ with $\beta \in \Delta$.*

The proofs of the above three lemmas can be found in Appendix B. For what follows, it will be convenient to flesh out some of the model-theoretic consequences of Lemma 4.8.

Lemma 4.9 *Given a λ -type Φ with a λ -canonical extended evidence model $\mathfrak{M}_\lambda(\Phi) = \langle W_\lambda^\Phi, R_\lambda^\Phi, E_\lambda^\Phi, V_\lambda^\Phi \rangle$,*

- (i) *for each $\Gamma \in W_\lambda^\Phi$, $\Gamma E_\lambda^\Phi W_\lambda^\Phi$;*
- (ii) *if $\Box\alpha \in \Phi$, $\mathcal{N}_\lambda^\alpha(\Phi)$ is non-empty; and*
- (iii) *if Φ is flat and $\Gamma E_\lambda^\Phi X$, there is $\Delta \in X$ such that $\Gamma R_\lambda^\Phi \Delta$.*

Proof. Suppose Γ is a (Φ, λ) -type and let $\mathfrak{M}_\lambda(\Phi)$ be as above.

- (i) Recall that we stipulated that $\Box\top, \top \in \text{cl}(\Phi)$, and these formulas are valid so it follows that $\top \in \Delta$ for all $\Delta \in W_\lambda^\Phi$; thus $\mathcal{N}_\lambda^\top(\Phi) = W_\lambda^\Phi$. Similarly, $\Box\top \in \Gamma$, so that $\Gamma E_\lambda^\Phi W_\lambda^\Phi$, as claimed.
- (ii) Put $\beta = \top$ in Lemma 4.8(i).
- (iii) Immediate from the definition of the accessibility relation and neighborhoods on canonical extended evidence models and Lemma 4.8(ii). □

Putting everything together, we have:

Corollary 4.10 *If λ is an evidence logic and Φ is λ -consistent, then $\mathfrak{M}_\lambda(\Phi)$ is a extended λ -evidence model.*

Proof. From Lemma 4.9(i) we see that, given $\Gamma \in W_\lambda^\Phi$, $\Gamma E W_\lambda^\Phi$, while from Lemma 4.9(ii) we obtain $\Gamma E \emptyset$. It follows that $\mathfrak{M}_\lambda(\Phi)$ is a extended evidence model.

If further we have that $\lambda \in \{\text{Log}_b, \text{Log}_{bu}\}$, by Lemma 4.9(iii), every neighborhood on $\mathfrak{M}_\lambda(\Phi)$ is flat, i.e., $\mathfrak{M}_\lambda(\Phi)$ is flat.

Finally, if $\lambda \in \{\text{Log}_u, \text{Log}_{bu}\}$, we use Lemma 4.5 to see that the model $\mathfrak{M}_\lambda(\Phi)$ is uniform. \square

We are now ready to prove our own version of the ‘Truth Lemma’:

Proposition 4.11 (Truth Lemma) *For every formula $\psi \in \text{cl}(\Phi)$ and every set $\Gamma \in \text{type}_\lambda^A(\Phi)$,*

$$\psi \in \Gamma \Rightarrow \Gamma \in \llbracket \psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}.$$

Proof. The proof is by induction on the structure of ψ . The only interesting cases are the modalities \square and B (and their duals).

Suppose first that $\square\psi \in \Gamma$, and let $X = \mathcal{N}_\lambda^\psi(\Gamma)$. By definition, we have $\psi \in \Theta$ for every $\Theta \in X$, and by the induction hypothesis this implies that for each $\Theta \in X$, $\Theta \in \llbracket \psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$. Thus, X is a neighborhood of Γ ($X \in E^\Phi(\Gamma)$) with $X \subseteq \llbracket \psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$, and so, $\Gamma \in \llbracket \square\psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$, as desired.

Now, assume that $\diamond\psi \in \Gamma$. Let X be any neighborhood of Γ . Then $X = \mathcal{N}^\alpha(\Gamma)$ for some α . Hence, $\square\alpha \in \Gamma$ and $\diamond\psi \in \Gamma$. By Lemma 4.8.1, there is $\Theta \in \mathcal{N}^\alpha(\Gamma)$ with $\psi \in \Theta$. By the induction hypothesis, $\Theta \in \llbracket \psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$, and, since X was arbitrary, we conclude that $\Gamma \in \llbracket \diamond\psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$.

Assume that $B\psi \in \Gamma$. Now, suppose that $\Gamma R_\lambda^\Phi \Delta$. Then, by definition $\psi \in \Delta$, which by the induction hypothesis implies $\Delta \in \llbracket \psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$. Thus, $\Gamma \in \llbracket B\psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$.

Assume that $\widehat{B}\psi \in \Gamma$. By Lemma 4.8.3, $\psi \in \Theta$ for some $\Theta \in \text{type}_\lambda^B(\Gamma)$. But then $\Gamma R_\lambda^\Phi \Theta$, and thus $\Gamma \in \llbracket \widehat{B}\psi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$. \square

4.1 Proof of the Main Theorem

With the work developed in previous sections, we are ready to state and prove our main results.

Theorem 4.12 *Each evidence logic λ is sound and strongly complete both for the class of λ -evidence models and the class of extended λ -evidence models.*

Proof. Let Φ be λ -consistent and let $\mathfrak{M}_\lambda(\Phi)$ be the canonical extended evidence model for Φ . By Corollary 4.10, $\mathfrak{M}_\lambda(\Phi)$ is a λ -extended evidence model, and Φ can be extended to a (Φ, λ) -type Γ . By Proposition 4.11, given $\varphi \in \Gamma$, $\Gamma \in \llbracket \varphi \rrbracket_{\mathfrak{M}_\lambda(\Phi)}$, so that Φ is satisfiable in $\mathfrak{M}_\lambda(\Phi)$.

By Lemma 3.9, $(\mathfrak{M}(\Phi))^+$ is a λ -model also satisfying Φ . \square

Theorem 4.13 *The flat evidence logics $\text{Log}_b, \text{Log}_{bu}$ have the finite model property. In fact, if φ has length ℓ , then it has a λ -model of size at most¹⁰ $2^\ell \cdot 2^{2^\ell} \cdot 2$.*

Non-flat evidence logics do not have the finite model property, but they do have the finite extended evidence model property.

Proof. If φ is λ -consistent, we can extend it to a $(\{\varphi\}, \lambda)$ -type Φ and let \mathfrak{M} be the canonical extended λ -evidence model for Φ . Clearly, \mathfrak{M} has at most 2^ℓ states, where ℓ is the length of φ , and thus it is finite.

¹⁰This bound could be improved to $2^\ell \cdot \ell \cdot 2$.

If φ is further \flat -consistent, then the model \mathfrak{M}^+ is finite as well, and it has at most $2^\ell \cdot 2^{2^\ell} \cdot 2$ points. \square

Corollary 4.14 *Both flat and general evidence logic are decidable.*

Our analysis does not yield the exact computational complexity of deciding validity or satisfiability. Note however that the general evidence logic is a conservative extension of K_u (i.e., K with a universal modality) and hence the validity problem is EXPTIME-hard; the same should be true of Log_\flat . Meanwhile, we expect the uniform logics to be much simpler, as they extend S5 rather than K , but do not have a specific conjecture on their complexity.

5 Language Extensions

The logical systems studied in the previous sections are interesting in their own right as they combine features of both normal and non-normal modal logics. An additional appealing feature of the logical systems studied here and in [21] is the fresh new interpretation of neighborhood structures. In particular, interpreting neighborhoods as bodies of evidence suggests a number of new and interesting modalities beyond the usual repertoire studied in modal neighborhood logics. In this section, we briefly explore this rich landscape of extensions of our basic language, and point out some resulting problems of axiomatization.

Two operators that immediately suggest themselves in anticipation of the dynamic extensions found in [21] are conditional versions of our evidence and belief operators. The conditional belief operator ($B^\varphi\psi$: “the agent believes ψ conditional on φ ”) is well-known, but is given a new twist in our setting. Some of the agent’s current evidence may be inconsistent with φ (i.e., disjoint with $\llbracket\varphi\rrbracket_{\mathcal{M}}$). If one is restricting attention to situations where φ is true, then such inconsistent evidence must be “ignored”. Here is how we do this:


Definition 5.1 (Relativized maximal overlapping evidence.) *Suppose that $X \subseteq W$. Given a collection \mathcal{X} of subsets of W (i.e., $\mathcal{X} \subseteq \wp(W)$), the relativization of \mathcal{X} to X is the set $\mathcal{X}^X = \{Y \cap X \mid Y \in \mathcal{X}\}$. We say that a collection \mathcal{X} of subsets of W has the **finite intersection property relative to X** (X -f.i.p.) if, for each $\{X_1, \dots, X_n\} \subseteq \mathcal{X}^X$, $\bigcap_{1 \leq i \leq n} X_i \neq \emptyset$. We say that \mathcal{X} has the **maximal X -f.i.p.** if \mathcal{X} has X -f.i.p. and no proper extension \mathcal{X}' of \mathcal{X} has the X -f.i.p.*

To simplify notation, when X is the truth set of formula φ , we write “maximal φ -f.i.p.” for “maximal $\llbracket\varphi\rrbracket_{\mathcal{M}}$ -f.i.p.” and “ \mathcal{X}^φ ” for “ $\mathcal{X}^{\llbracket\varphi\rrbracket_{\mathcal{M}}}$ ”. Now we define a natural notion of conditional belief:

- $\mathcal{M}, w \models B^\varphi\psi$ iff for each maximal φ -f.i.p. $\mathcal{X} \subseteq E(w)$, for each $v \in \bigcap \mathcal{X}^\varphi$, $\mathcal{M}, v \models \psi$

This notion suggests a logical investigation beyond what we have provided so far. First of all, strikingly, $B\varphi \rightarrow B^\psi\varphi$ is not valid. One can compare this to the failure of monotonicity for antecedents in conditional logic. In our more general setting which allows inconsistencies among accepted evidence, we also see that even the following variant is not valid: $B\varphi \rightarrow (B^\psi\varphi \vee B^{\neg\psi}\varphi)$. To see

this, consider an evidence model with $E(w) = \{X_1, Y_1, X_2, Y_2\}$ where the sets are defined as follows:



Then, $\mathcal{M}, w \models Bq$; however, $\mathcal{M}, w \not\models B^p q \vee B^{\neg p} q$. This is interesting as it is valid on *connected* plausibility models for conditional belief (cf. [2] for a complete modal logic of conditional belief on such models). Extending our completeness proof from the previous section requires new ideas as there is no obvious way to interpret conditional belief operators on our auxiliary “extended evidence models”.

The conditional evidence operator $\Box^\varphi\psi$ (“the agent has evidence that ψ is true conditional on φ being true”) is a new modality on neighborhoods. We say that $X \subseteq W$ is **consistent (compatible) with φ** if $X \cap \llbracket\varphi\rrbracket_{\mathcal{M}} \neq \emptyset$. Truth of conditional evidence can then be defined as follows:

- $\mathcal{M}, w \models \Box^\varphi\psi$ iff there exists an evidence set $X \in E(w)$ which is consistent with φ such that for all worlds $v \in X \cap \llbracket\varphi\rrbracket_{\mathcal{M}}$, $\mathcal{M}, v \models \psi$.

In particular, if there is no evidence consistent with φ , then $\Box^\varphi\psi$ is false. This, in turn means that $\Box^\varphi\psi$ is not equivalent to $\Box(\varphi \rightarrow \psi)$.¹¹ Indeed, a simple bisimulation argument shows that no definition exists for conditional evidence in the language with absolute evidence and belief. The conditional evidence operators satisfy the monotonicity inference rule (from $\varphi \rightarrow \psi$ infer $\Box^\alpha\varphi \rightarrow \Box^\alpha\psi$) and, for example, the axiom scheme $\Box^\alpha\varphi \rightarrow \Box^\alpha\alpha$; however, a complete logic will be left for future work.

Both of these operators are special cases of more general modalities that were discovered in [21] as static counterparts to natural dynamic modalities of adding or removing evidence.

For the conditional evidence operator $\Box^\varphi\psi$, we can require the witnessing evidence set to be “compatible” with a sequence of formulas. Let $\bar{\varphi} = (\varphi_1, \dots, \varphi_n)$ be a finite sequence of formulas. We say that a set of states X is **compatible with $\bar{\varphi}$** provided that, for each formula φ_i , $X \cap \llbracket\varphi_i\rrbracket_{\mathcal{M}} \neq \emptyset$. Then, we define the general conditional evidence operator $\Box_{\bar{\varphi}}^\alpha\psi$ as follows:

- $\mathcal{M}, w \models \Box_{\bar{\varphi}}^\alpha\psi$ iff there is $X \in E(w)$ compatible with $\bar{\varphi}$, α such that $X \cap \llbracket\alpha\rrbracket_{\mathcal{M}} \subseteq \llbracket\psi\rrbracket_{\mathcal{M}}$.

¹¹To see this, consider a model where φ is false at all worlds. Then $\Box^\varphi\psi$ is also false at all worlds, but $\Box(\varphi \rightarrow \psi)$ will be true at all worlds, since $\varphi \rightarrow \psi$ is true everywhere.

The conditional belief operator can be generalized in two ways. The first way is to incorporate the above notion of compatibility with a sequence of formulas. The intended interpretation of $B_{\vec{\gamma}}^{\varphi}\psi$ is “the agent believes χ conditional on φ assuming compatibility with each of the γ_i ”. The formal definition is a straightforward generalization of the earlier definition of $B^{\varphi}\psi$. A maximal f.i.p. set \mathcal{X} is **compatible with** a sequence of formulas $\vec{\varphi}$ provided for each $X \in \mathcal{X}$, X is compatible with $\vec{\varphi}$. Then,

- $\mathcal{M}, w \models B_{\vec{\varphi}}^{\alpha}\psi$ iff for each maximal α -f.i.p. \mathcal{X} compatible with $\vec{\varphi}$, we have that $\bigcap \mathcal{X}^{\alpha} \subseteq \llbracket \psi \rrbracket_{\mathcal{M}}$.

The second generalization focuses on the conditioning operation. Note that $B^{\varphi}\psi$ may be true at a state w without having any w -scenarios that imply φ (i.e., there is no w -scenario \mathcal{X} such that $\bigcap \mathcal{X} \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}$). A more general form of conditioning is $B^{\varphi, \alpha}\psi$ where “the agent believe ψ , after having settled on α and conditional on φ ”. Formally,

- $\mathcal{M}, w \models B^{\varphi, \psi}\chi$ iff for all maximally φ -compatible sets $\mathcal{X} \subseteq E(w)$, if $\bigcap \mathcal{X} \cap \llbracket \varphi \rrbracket_{\mathcal{M}} \subseteq \llbracket \psi \rrbracket_{\mathcal{M}}$, then $\bigcap \mathcal{X} \cap \llbracket \varphi \rrbracket_{\mathcal{M}} \subseteq \llbracket \chi \rrbracket_{\mathcal{M}}$.

Note that $B^{+\varphi}$ can be defined as $B^{\varphi, \top}$.

This splitting of notions shows that neighborhood structures are a good vehicle for exploring finer epistemic and doxastic distinctions than those found in standard relational models for modal logic. Moreover, they support interesting matching forms of reasoning beyond standard axioms. Validities include interesting connections between varieties of conditional belief such as: $B^{\varphi}\psi \rightarrow B(\varphi \rightarrow \psi)$ and $B(\varphi \rightarrow \psi) \rightarrow B^{\top, \varphi}\psi$.

At the same time, this richness means that new techniques may be needed in the logical analysis of this richer form of neighborhood semantics. Model-theoretically, we need stronger notions of bisimulation and -morphism matching these more expressive languages, while proof-theoretically, we need to lift our earlier completeness technique to this setting.

6 Conclusion and Future Work

There are two main contributions in this paper. First, our completeness theorems solve an important problem left open in [21] contributing to the general study of basic evidence logic and its dynamics. Second, in doing so, we develop a new perspective on neighborhood models that suggests extensions of the usual systems. For the modal logician, the pleasant surprise of our language extensions is that there is a lot of well-motivated new modal structure to be explored on these simple models.

Clearly, many open problems remain. These start at the base level of the motivations for our framework, touched upon lightly earlier on. For instance, our framework still needs to be related to other modal logics of evidence [18,5], justification [1] and argumentation [4]. Here are a few more specific technical avenues for future research:

- *Further interpretations:* By imposing additional constraints on the evidence relations (i.e., the neighborhood functions), we get evidence models that are topological spaces. Can we give a spatial interpretation to our belief operator and the new modalities discussed in Section 5? This suggests new, richer spatial logics for reasoning about topological spaces (cf. [20]).
- *Computational complexity:* Neighborhood logics are often NP-complete, while basic modal logics on relational models are often Pspace-complete. What about mixtures of the two? By Corollary 4.14 we know that the validity problem for flat and general evidence logic is decidable. What is the precise complexity?
- *Extended model theory:* Combining existing notions of bisimulation for relational models and neighborhood models takes us only so far. What new notions of bisimulation directly on evidence models will match our extended modal languages? And in this setting, can the key representation method of this paper (cf. Theorem 3.8) be extended to deal with such richer languages, perhaps starting from richer extended evidence models carrying plausibility relations?
- *Extended proof theory:* How can we axiomatize the richer modal logic of neighborhood models that arises in Section 5? Will modifying existing techniques, including our approach in this paper, still work, or do we need a new style of analysis of canonical models?

Appendix

A Proof from Section 3

Lemma A.1 (Lemma 3.7) *Let $\mathfrak{M} = \langle W, R, E, V \rangle$ be an extended evidence model with associated evidence model \mathfrak{M}^+ . Let $\alpha \in W^+$ and $v \in W$ be arbitrary.*

Then, $\pi(\alpha) R v$ if and only if there is $\beta \in \pi^{-1}(v)$ such that $\alpha R_{E^+} \beta$.

Proof. Assume that $\alpha \in W^+$ and $\pi(\alpha) = w$. We first claim that every α -scenario on \mathfrak{M}^+ is of the form $\mathcal{B}^Y(\alpha) \cup \{W^+\}$.

For this, suppose that $\alpha E^+ X$ and $\alpha E^+ Y$ with $X \cap Y \neq \emptyset$ and $X, Y \neq W^+$. This means that there is a $\langle v, Z, n \rangle \in X \cap Y$. By inspecting the definition of E^+ , this implies $X, Y \in \mathcal{B}^Z(\alpha)$. Hence, any collection $\mathcal{X} \subseteq E^+(\alpha)$ with the finite intersection property must be contained in $\mathcal{B}^Z(\alpha) \cup \{W^+\}$ for some Z .

Meanwhile, it is easy to see that for every $Z \subseteq W$, $\mathcal{B}^Z(\alpha) \cup \{W^+\}$ already has the fip (by case-by-case inspection), hence every scenario is of this form.

Now, if Y is not flat for w , one can check that $\bigcap \mathcal{B}^Y(w) = \emptyset$. If Y is flat for w , then all elements of $\bigcap \mathcal{B}^Y(w)$ are of the form $\langle v, Y, 0 \rangle$ with $w R v$ and $v \in Y$. In either case, $\pi[\bigcap \mathcal{B}^Y(w)] \subseteq R(w)$. It follows that, whenever $\alpha R_{E^+} \beta$, then $w R \pi(\beta)$.

It is also straightforward to see that

$$\bigcap \mathcal{B}^W(w) = R(w) \times \{W\} \times \{0\},$$

so that $\pi[\bigcap \mathcal{B}^W(w)] = R(w)$. Thus, if $w R v$ then there is an α -scenario \mathcal{U} (namely $\mathcal{U} = \mathcal{B}^W(w) \cup \{W^+\}$) with $\langle v, W, 0 \rangle \in \bigcap \mathcal{U}$, so that $\alpha R_{E^+} \beta = \langle v, W, 0 \rangle$. \square

B Proofs from Section 4

Lemma B.1 (Lemma 4.6) *Let Γ, Δ be sets of formulas.*

- (i) *Suppose that $\varphi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \psi$. Then, $\Box\varphi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \Box\psi$.*
- (ii) *Suppose that $\Phi, A\Gamma, \widehat{A}\Delta \vdash_\lambda \psi$. Then, $B\Phi, A\Gamma, \widehat{A}\Delta \vdash_\lambda B\psi$.*

Without loss of generality we can assume Γ, Δ, Φ to be finite, since in general we have that $\Theta \vdash_\lambda \alpha$ if and only if for some finite $\Theta' \subseteq \Theta$, $\Theta' \vdash_\lambda \alpha$.

Proof. Note that

$$\bigwedge A\Gamma \wedge \bigwedge \widehat{A}\Delta \leftrightarrow A \left(\bigwedge A\Gamma \wedge \bigwedge \widehat{A}\Delta \right)$$

is derivable in **S5**, so we can replace $A\Gamma, \widehat{A}\Delta$ by a single formula $A\gamma$.

- (i) If $\varphi, A\gamma \vdash_\lambda \psi$, then

$$\vdash_\lambda \varphi \wedge A\gamma \rightarrow \psi.$$

By the monotonicity rule

$$\vdash_\lambda \Box(\varphi \wedge A\gamma) \rightarrow \Box\psi.$$

Applying the pullout axiom we get

$$\vdash_\lambda \Box\varphi \wedge A\gamma \rightarrow \Box\psi,$$

that is,

$$\Box\varphi, A\gamma \vdash_\lambda \Box\psi.$$

- (ii) If $\Phi, A\gamma \vdash_\lambda \psi$, since B is a normal operator, we also have

$$B\Phi, BA\gamma \vdash_\lambda B\psi.$$

But since $A\alpha \rightarrow BA\alpha$ is an axiom, we get

$$B\Phi, A\gamma \vdash_\lambda B\psi.$$

\square

Below, recall that, if $\lambda \in \{\mathbf{Log}, \mathbf{Log}_b\}$, we defined $\Gamma^\lambda = A\Gamma^A \cup \widehat{A}\Gamma^{\widehat{A}}$, while if $\lambda \in \{\mathbf{Log}_u, \mathbf{Log}_{bu}\}$,

$$\Gamma^\lambda = \bigcup \{ \bigcirc \Gamma^\bigcirc : \bigcirc = A, \widehat{A}, B, \widehat{B}, \Box, \diamond \}.$$

Lemma B.2 (Lemma 4.7) *Let Γ be a set of formulas.*

- (i) *If Γ is λ -consistent, $\alpha \in \Gamma^\Box$ and $\delta \in \Gamma^\diamond$, then $\{\alpha, \delta\} \cup \Gamma^\lambda$ is λ -consistent.*
- (ii) *If Γ is λ -consistent and $\psi \in \Gamma^{\widehat{B}}$, then $\{\psi\} \cup \Gamma^B \cup \Gamma^\lambda$ is λ -consistent.*

(iii) If λ is flat, Γ is λ -consistent and $\psi \in \Gamma^\square$, then $\{\psi\} \cup \Gamma^B \cup \Gamma^\lambda$ is λ -consistent.

Proof. Assume Γ is λ -consistent.

Note that we have $\Gamma^\lambda \equiv A\Gamma^\lambda$ independently of λ ; over the non-uniform case, this follows from the axiom $\widehat{A}\varphi \rightarrow A\widehat{A}\varphi$, while over the uniform case (i.e., if $\lambda \in \{\text{Log}_u, \text{Log}_{bu}\}$), this is because $\bigcirc\Gamma^\square$ is equivalent to $A\bigcirc\Gamma^\square$ over λ for any modality \bigcirc .

As before, we may assume Γ is finite. Thus we may replace Γ^λ by a single formula $A\gamma$ equivalent to $\bigwedge \Gamma^\lambda$ over λ .

(i) Suppose otherwise. Then, we would have

$$\alpha, A\gamma \vdash_\lambda \neg\delta.$$

Thus by Lemma 4.6.1,

$$\square\alpha, A\gamma \vdash_\lambda \square\neg\delta,$$

so that

$$\{\neg\square\neg\delta, \square\alpha, A\gamma\}$$

is λ -inconsistent. But this is a subset of Γ , contradicting our assumption that Γ is λ -consistent.

(ii) Suppose $\psi \in \Gamma^{\widehat{B}}$. If the claim was false, then we would have

$$\Gamma^B, A\gamma \vdash_\lambda \neg\psi,$$

and hence, by Lemma 4.6.2,

$$B\Gamma^B, A\gamma \vdash_\lambda B\neg\psi.$$

Thus

$$\{\neg B\neg\psi, A\gamma\} \cup B\Gamma^B$$

would be λ -inconsistent. If $\psi \in \Gamma^{\widehat{B}}$, this contradicts the consistency of Γ .

(iii) If $\psi \in \Gamma^\square$, the claim follows from item 2 above using the axiom $\square\psi \rightarrow \widehat{B}\psi$. \square

Lemma B.3 (Lemma 4.8) Let Γ be a (Φ, λ) -type.

- (i) If $\square\alpha, \diamond\beta \in \Gamma$, there is $\Delta \in \mathcal{N}_\lambda^\alpha(\Gamma)$ with $\beta \in \Delta$.
- (ii) If $\lambda \in \{\text{Log}_b, \text{Log}_{bu}\}$ and $\alpha \in \Gamma^\square$, then $\mathcal{N}_\lambda^\alpha(\Gamma) \cap \text{type}_\lambda^B(\Gamma)$ is non-empty.
- (iii) Suppose that $\widehat{B}\alpha \in \Gamma$. Then, there is $\Delta \in \text{type}_\lambda^B(\Gamma)$ with $\beta \in \Delta$.

Proof. Let Γ be a (Φ, λ) -type (for some λ -consistent set of formulas Φ).

(i) Since Γ is λ -consistent and $\square\alpha, \diamond\beta \in \Gamma$, we have, by Lemma 4.7.1, that

$$\{\alpha, \beta\} \cup \Gamma^\lambda$$

is λ -consistent. This can be extended to a Φ -type Δ . Then, by definition we have $\Delta \in \mathcal{N}_\lambda^\alpha(\Gamma)$ and $\beta \in \Delta$, as desired.

- (ii) By Lemma 4.7.3, if Γ is λ -consistent and λ is flat, then $\{\alpha\} \cup \Gamma^B \cup \Gamma^\lambda$ is λ -consistent as well, and hence can be extended to a Φ -type Δ . Obviously,

$$\Delta \in \mathcal{N}_\lambda^\alpha(\Gamma) \cap \text{type}_\lambda^B(\Gamma).$$

- (iii) By Lemma 4.7.2, since Γ is λ -consistent, we have that

$$\{\beta\} \cup \Gamma^B \cup \Gamma^\lambda$$

is λ -consistent as well. Thus, it can be extended to a Φ -type Δ . Evidently, $\Delta \in \text{type}_\lambda^B(\Gamma)$ and $\beta \in \Delta$. □

References

- [1] Artemov, S. and E. Nogina, *Introducing justification into epistemic logic*, Journal of Logic and Computation **15** (2005), pp. 1059–1073.
- [2] Board, O., *Dynamic interactive epistemology*, Games and Economic Behavior **49** (2004), pp. 49–80.
- [3] Gärdenfors, P., “Knowledge in Flux: Modeling the Dynamics of Epistemic States,” Bradford Books, MIT Press, 1988.
- [4] Grossi, D., *On the logic of argumentation theory*, in: W. van der Hoek, G. Kaminka, Y. Lesperance, M. Luck and S. Sandip, editors, *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS’10)*, 2010, pp. 409–416.
- [5] Halpern, J. and R. Pucella, *A logic for reasoning about evidence*, Journal of AI Research **26** (2006), pp. 1–34.
- [6] Hansen, H. H., “Monotonic Modal Logic,” Master’s thesis, Universiteit van Amsterdam (ILLC technical report: PP-2003-24) (2003).
- [7] Hansen, H. H., C. Kupke and E. Pacuit, *Neighbourhood structures: Bisimilarity and basic model theory*, Logical Methods in Computer Science **5** (2009), pp. 1–38.
- [8] Horty, J. F., *Reasons as defaults*, Philosophers’ Imprint **7** (2007), pp. 1–28.
- [9] Kratzer, A., *What must and can must and can mean*, Linguistics and Philosophy **1** (1977), pp. 337–355.
- [10] Lismont, L. and P. Mongin, *A non-minimal but very weak axiomatization of common belief*, Artificial Intelligence **70** (1994), pp. 363–374.
- [11] List, C. and F. Dietrich, *Reasons for (prior) belief in bayesian epistemology* (2012), unpublished manuscript.
- [12] Pacuit, E., *Neighborhood semantics for modal logic: An introduction* (2007), ESSLLI 2007 course notes (ai.stanford.edu/~epacuit/classes/).
- [13] Parikh, R., *The logic of games and its applications*, in: M. Karpinski and J. v. Leeuwen, editors, *Topics in the Theory of Computation*, Annals of Discrete Mathematics 24, Elsevier, 1985 pp. 111 – 140.
- [14] Pauly, M., *A modal logic for coalitional power in games*, Journal of Logic and Computation **12** (2002), pp. 149–166.
- [15] Peleg, D., *Concurrent dynamic logic*, Journal of the ACM **34** (1987), pp. 450–479. URL <http://doi.acm.org/10.1145/23005.23008>
- [16] Rott, H., “Change, Choice and Inference: A Study in Belief Revision and Nonmonotonic Reasoning,” Oxford University Press, 2001.
- [17] Segerberg, K., “An Essay in Classical Modal Logic,” Filosofiska Studier, Uppsala Universitet, 1971.
- [18] Shafer, G., “A Mathematical Theory of Evidence,” Princeton University Press, 1976.

- [19] van Benthem, J., “Logical Dynamics of Information and Interaction,” Cambridge University Press, 2011.
- [20] van Benthem, J. and G. Bezhanishvili, *Modal logics of space*, in: M. Aiello, I. Pratt-Hartmann and J. van Benthem, editors, *Handbook of Spatial Reasoning*, 2007, pp. 217–298.
- [21] van Benthem, J. and E. Pacuit, *Dynamic logics of evidence-based beliefs*, *Studia Logica* **99** (2011), pp. 61–92.
- [22] Vardi, M., *On epistemic logic and logical omniscience*, in: J. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference* (1986), pp. 293–305.
- [23] Veltman, F., *Prejudices, presuppositions and the theory of conditionals*, in: J. Groenendijk and M. Stokhof, editors, *Amsterdam Papers in Formal Grammar* (1976), pp. 248–281.
- [24] Zvesper, J., “Playing with Information,” Ph.D. thesis, ILLC University of Amsterdam Dissertation Series DS-2010-02 (2010).