



Evidence and plausibility in neighborhood structures



Johan van Benthem^{a,b}, David Fernández-Duque^c, Eric Pacuit^{d,*}

^a University of Amsterdam, The Netherlands

^b Stanford University, United States

^c University of Seville, Spain

^d University of Maryland, United States

ARTICLE INFO

Article history:

Available online 27 August 2013

MSC:

03C35

03B44

03B45

03B62

Keywords:

Modal logic

Neighborhood semantics

Plausibility ordering

Evidence-based reasoning

Doxastic logic

ABSTRACT

The intuitive notion of evidence has both semantic and syntactic features. In this paper, we develop an *evidence logic* for epistemic agents faced with possibly contradictory evidence from different sources. The logic is based on a neighborhood semantics, where a neighborhood N indicates that the agent has reason to believe that the true state of the world lies in N . Further notions of relative plausibility between worlds and beliefs based on the latter ordering are then defined in terms of this evidence structure, yielding our intended models for evidence-based beliefs. In addition, we also consider a second more general flavor, where belief and plausibility are modeled using additional primitive relations, and we prove a representation theorem showing that each such general model is a p -morphic image of an intended one. This semantics invites a number of natural special cases, depending on how uniform we make the evidence sets, and how coherent their total structure. We give a structural study of the resulting ‘uniform’ and ‘flat’ models. Our main result are sound and complete axiomatizations for the logics of all four major model classes with respect to the modal language of evidence, belief and safe belief. We conclude with an outlook toward logics for the dynamics of changing evidence, and the resulting language extensions and connections with logics of plausibility change.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

It has become standard practice in Artificial Intelligence and Game Theory to use possible-worlds models to describe the knowledge and beliefs of a group of agents. In such models, the agents’ knowledge is based on what is true throughout the set of epistemically accessible worlds, the current information range. Following a similar pattern, the agents’ beliefs are based on what is true in the set of most “plausible” worlds. Now, it is often implicitly assumed that the agents arrived at these structures through some process of investigation, but these details are no longer present in the models.

In a number of areas, ranging from epistemology to computer science and decision theory, the need has been recognized for models that keep track of the “reasons”, or the *evidence* for beliefs and other

* Corresponding author at: University of Maryland, United States.

E-mail address: epacuit@umd.edu (E. Pacuit).

informational attitudes (cf. [23,20,3,34]). Encoding evidence as the current range of worlds the agent considers possible ignores how the agent arrived at this epistemic state. This also ignores the fine-structure of evidence that allows us to consider or modify just parts of it. One extreme for recording this additional structure are models with complete syntactic details of what the agent has learned so far (including the precise formulation and sources for each piece of evidence) (cf. [39]). In this paper, we will explore a middle ground in between ranges of possible worlds and syntactic fine-structure, viz. neighborhood models, where the available evidence is recorded as a family of sets of worlds. Our models are not unlike some earlier proposals in the study of conditionals and belief revision (cf. [21,40,13,28]), but we quickly strike out in other directions, and provide a more in-depth logical treatment.

This paper is a continuation of our earlier work in [8,38] on a new evidence interpretation of neighborhood models.¹ In a neighborhood model, each state is assigned a collection of subsets of the set of states. We view these different collections of subsets as the evidence that the agent has acquired – allowing the agent to have different evidence at different states. Given such an explicit description of the evidence that the agent has acquired, one can explore different notions of beliefs and related cognitive attitudes over neighborhood models. The logical systems that arise on natural model classes of this sort with modalities for evidence and belief have been axiomatized completely in [38]. In this paper, we go one step further, and add some further crucial structure to the neighborhood structures.

In general, there are two ways we can enrich neighborhood structures with descriptions of the agent’s beliefs. The first approach is to add new accessibility relations corresponding to each epistemic or doxastic attitude in the neighborhood structure. We then impose constraints on these new relations to ensure that they are “appropriately grounded” on the available evidence. The second approach is to define the agent’s beliefs and related cognitive attitudes directly using no more than the given evidence structure in the neighborhood structures. The latter “intended models” may be considered as a special case of the former “general models”. This paper offers a careful study of these two approaches.

Our second new contribution is to elaborate on the relationship between our neighborhood models and another general framework for belief change in the modal tradition. Originally used as a semantics for conditionals (cf. [22]), *plausibility models* are wide-spread in modal logics of belief [35,36,6,15]. The main idea is to endow epistemic ranges with an ordering $w \preceq v$ of relative *plausibility* on worlds (usually uniform across epistemic equivalence classes): “(according to the agent) world v is at least as plausible as w ”.² Plausibility orders are typically assumed to be reflexive and transitive, and often even *connected*, making every two worlds comparable. Connections between evidence structure and plausibility order will occur throughout this paper, and their reflection in logical axioms will be determined.

In all, we shall consider four variants of the logic of evidence-based belief, which depend on some fundamental assumptions one may make about evidence models, in terms of uniformity of evidence across worlds, and total coherence of maximally consistent sets of evidence. For each of the resulting logics, we prove two main results. The first is a characterization theorem for general evidence models as being *p-morphic images* of intended models. The second result determines 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.

Anyone familiar with Sergei Artemov’s work will have seen a similarity by now. It is very natural to attach to every believed proposition a “justification” for that proposition. This idea was first studied in Artemov’s seminal paper [2] and applied to epistemic logic in [4] (see [3] for an overview and pointers to the relevant literature). In particular, $t : \varphi$ is intended to mean that the agent believes φ and that t is the justification for this belief. Here t is a proof term, and sophisticated logical systems have been developed that

¹ Neighborhood models have been used to provide a semantics for both normal and non-normal modal logics. See [29] for an early discussion of neighborhood models and their logics, and [18,25,19] for modern motivations and mathematical details.

² In conditional semantics, such plausibility or ‘similarity’ orders are typically world-dependent.

extend traditional “provability logic” to reason explicitly about the logical structure of these justifications.³ Moreover, an appealing semantics exists for these systems (cf. [12]) that combines ideas from possible worlds semantics and the syntax of proof terms. The approach in the current paper is more coarse-grained, since we do not give ourselves proof or evidence terms that can be manipulated in our deductive systems. Nevertheless, we hope to show that even at our chosen level of modeling evidence, an amazing amount of fine-structure exists. A deeper comparison between our approach and justification logic seems a well-worth effort, but it is one step beyond the horizon of this paper.⁴

This paper is organized as follows. The next section (Section 2) introduces the logical systems that we will study. Each of these logical systems is defined using sublanguages of a single modal language, which includes both a non-normal modality (the “evidence-for” modality) and normal modalities (two belief modalities and a universal modality). One important theme in this paper is that there are two different types of models that can be used as a semantics for this language. The first type of model is a neighborhood structure where the relations used to interpret the normal modalities are *derived* from the neighborhood function. The second type of model extends neighborhood models with relations, one for each belief operator. Our first technical contribution is to clarify the relationship between these two types of models (Sections 2.2 and 4). More generally, we compare these models with the more standard *plausibility models* often used to represent a rational agent’s belief in Section 3. Sections 5 and 6 contain our second main contribution of this paper: proofs that our axiom systems are complete with respect to their intended class of models. Section 7 returns to the original motivation for studying evidence logic from [8]: developing dynamic logics of evidence management. We offer some concluding remarks in Section 8.

2. Logics of evidence and belief

We start by presenting the logical system that we will study in this paper. 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 following modal language can be used to describe what the agent believes given her available evidence (cf. [8]).

Definition 1. Let At be a fixed set of atomic propositions. Let $\mathcal{L} = \mathcal{L}_{ABE\preceq}$ be the smallest set of formulas generated by the following grammar

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid [B]\varphi \mid [E]\varphi \mid [A]\varphi \mid [\preceq]\varphi$$

where $p \in \text{At}$. The propositional connectives ($\wedge, \rightarrow, \leftrightarrow$) are defined as usual and the duals⁵ of $[E], [B], [\preceq]$ and $[A]$ are $\langle E \rangle, \langle B \rangle, \langle \preceq \rangle$ and $\langle A \rangle$, respectively.

A *signature* is a subset of $\{A, B, E, \preceq\}$, written usually as a string rather than a set. Sublanguages $\mathcal{L}' \subseteq \mathcal{L}$ will be denoted by \mathcal{L}_σ , where σ is a signature, and \mathcal{L}' contains only formulas with modalities in the signature.

The intended interpretation of $[E]\varphi$ is “the agent has evidence for φ ” and $[B]\varphi$ says that “the agents believes that φ is true.” The modality $[\preceq]\varphi$ refers to the plausibility order to be introduced below, and it has become standard to read it as “the agent *safely believes* that φ is true”, though we will really explore

³ By now, these systems have been applied to a variety of issues in epistemology and game theory, and other fields: see [3] for the relevant references.

⁴ For the moment, we can only refer to [5] for a system that merges features of justification logic with plausibility models and syntax dynamics in the style of “dynamic-epistemic logic” [36], an enterprise close to the spirit of the present paper.

⁵ In other words, $\langle B \rangle = \neg[B]\neg$, and similarly for other operators.

this interpretation only in Section 3. We also include the universal modality ($[A]\varphi$: “ φ is true in all states”) for technical convenience.⁶

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 ($[E]\varphi$) is not a normal modal operator. That is, the agent may have evidence for φ and evidence for ψ ($[E]\varphi \wedge [E]\psi$) without having evidence for their conjunction ($\neg[E](\varphi \wedge \psi)$). Of course, the two belief and universal operators are normal modal operators. So, the logical systems we study in this paper *combine* a non-normal modal logic with a normal one.

2.1. Models for \mathcal{L} and its fragments

Our background assumption is that the agent uses the evidence she has gathered at each state to form her beliefs. There are two ways to make this precise. The first is to ground the agent’s beliefs (defined as relations on the set of states) on the evidence available at each state via a number of technical assumptions. The second way is to define the agent’s beliefs directly in terms of the evidence available at each state. These different options lead us to consider several different signatures. This motivates the following general definition of a *model*:

Definition 2. Let σ be a signature containing E and not containing⁷ A . A σ -**structure** is a tuple $\mathcal{M} = \langle W, \langle I(m) \rangle_{m \in \sigma}, V \rangle$, where W is a non-empty set of worlds and

- if $E \in \sigma$ then $I(E) \subseteq W \times \wp(W)$ is an *evidence relation*,
- if $\preceq \in \sigma$ then $I(\preceq)$ is a plausibility order on W : $I(\preceq)$ is reflexive and transitive.⁸
- if $B \in \sigma$ then $I(B)$ is an (arbitrary) binary relation on W , and⁹
- $V : \text{At} \rightarrow \wp(W)$ is a valuation function.

We will not distinguish between $m \in \sigma$ and its denotation $I(m)$ in what follows, and we will write $E(w)$ for the set $\{X \mid wEX\}$.

A σ -structure is a σ -**model** if it further satisfies the following constraints:

1. for each $w \in W$, $\emptyset \notin E(w)$ and $W \in E(w)$ and
2. if $\preceq \in \sigma$ then for all $w, v, u \in W$, if $w \preceq v$ and $w \in X \in E(u)$ it follows that $v \in X$ and
3. whenever $w \preceq v$ and uBw , it follows that uBv .

An **evidence model** is a σ -model where $\sigma = E$. We may also write simply **model** instead of σ -model when the value of σ is clear from context.

The following basic assumptions are implicit in the above definition:

- 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.

⁶ A natural interpretation of $[A]\varphi$ in the single-agent context of this paper is as “the agent knows that φ ”.

⁷ The universal modality will *always* be interpreted according to its intended meaning and need not be represented explicitly in the models. In addition, we consider evidence the most fundamental notion in our setting and assume it always in our signature.

⁸ In the literature \preceq is usually assumed to be converse well-founded, but we shall work in a more general setting. Note that the plausibility relation need not be *connected*.

⁹ For readers familiar with the more standard literature on modal logics of belief, it is natural to wonder why we include a separate relation for our belief modalities (rather than using the plausibility ordering to define belief in the usual way). The reason for working with this more general definition is discussed in Sections 2.2 and 3.

- 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.
- Despite the fact that sources may not be reliable or jointly inconsistent, they are all the agent has for forming beliefs.¹⁰

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 mild conditions, that were stated in Constraint 1 in the above definition of our models:

- 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 (see Definition 3).

Next, constraint 2 in the definition of our models ensures that the agent's evidence "coheres" with her opinions about the relative plausibility of the states. The intended interpretation of $w \preceq v$ is that, although the agent does not know which of w or v is the "actual situation", she considers v at least as plausible as w . Thus, if X is evidence that the actual state may be w (i.e., $w \in X$), then X is also evidence for all states that the agent considers at least as plausible as w . The underlying idea is that the agent uses her plausibility ordering to "fill-out" the sets of states identified as evidence. We will also consider the even more constrained situation where the agent defines her plausibility ordering directly from the evidence (this is discussed in Section 2.2).

Finally, constraint 3 on our evidence models generalizes the usual assumption that the agent believes what holds throughout the set of most plausible worlds: Any state the agent believes is possible must be among the most plausible states overall, though we will not require that the converse always holds.

Let us define *truth* for formulas of \mathcal{L} in a given model:

Definition 3. Let $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ be an evidence model. Truth of a formula $\varphi \in \mathcal{L}$ is defined inductively as follows:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ ($p \in \text{At}$)
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models [E]\varphi$ iff there exists X such that wEX and for all $v \in X$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models [B]\varphi$ if for all v , if wBv , then $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models [\preceq]\varphi$ iff for all v , if $v \succeq w$, then $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models [A]\varphi$ 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.

¹⁰ Modeling sources and agents' *trust* in these is quite feasible in our setting – but we will not pursue this topic here.

2.2. From general models to intended models

In this section, we consider an alternative way of interpreting our modal language \mathcal{L} , specializing the above “general models” to a special class where belief is entirely and explicitly evidence-driven. Rather than interpreting each modality as a new relation on the set of states, we will now *derive* the relevant relations from the evidence sets. As a result, we can interpret the entire language \mathcal{L} on an evidence model $\mathcal{M} = \langle W, E, V \rangle$.

Recall that 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 4. 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 $\{X_1, \dots, X_n\} \subseteq \mathcal{X}$ we have that $\bigcap_{1 \leq i \leq n} X_i \neq \emptyset$). A collection is called a **scenario** if it is a *w-scenario* for some state *w*.


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*.¹¹ 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 . This motivates the following definition:

Definition 5. Given an evidence function $E : W \rightarrow \wp(W)$, we define $B_E \subseteq W \times W$ by $w B_E v$ if $v \in \bigcap \mathcal{X}$ for some *w-scenario* \mathcal{X} .

In order to derive a plausibility ordering, we borrow an ubiquitous idea, occurring in point-set topology, but also in recent theories of relation merge (cf. [1,24]): the so-called *specialization (pre)-order*. Under this ordering, *v* is more plausible than *w* if every set that is evidence for *w* is also evidence for *v*. Formally, we define the following:

Definition 6. Given an evidence function $E : W \rightarrow \wp(W)$, we define $\preceq_E \subseteq W \times W$ by $w \preceq_E v$ if whenever u, X are such that $w \in X \in E(u)$, then $v \in X$.

To make this definition a bit more concrete, here is a simple illustration.



¹¹ Analogous ideas occur in semantics of conditionals [21,40] and belief revision [13,28].

Note that the derived relation \preceq_E is uniform throughout the model, even though evidence itself is not. It is indeed possible to define a point-wise variant of the specialization order with u as parameter, but we shall not explore this option in this paper.

Given an evidence model $\mathcal{M} = \langle W, E, V \rangle$, we now define the extended structure

$$\mathcal{M}^\Delta = \langle W, E, B_E, \preceq_E, V \rangle.$$

We extend the truth-definitions on \mathcal{M} to all of \mathcal{L} by setting $\llbracket \cdot \rrbracket_{\mathcal{M}} = \llbracket \cdot \rrbracket_{\mathcal{M}^\Delta}$. If $\mathcal{M} = \langle W, E, V \rangle^\Delta$, we say \mathcal{M} is an *intended model*.

We will be interested in the precise relationship between intended models \mathcal{M}^Δ and our earlier general models $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$. This issue will be discussed in detail in the remainder of the paper, but we conclude this section with a brief remark. Notice that, in general, $\mathcal{M} = \langle W, E, B, \preceq_E, V \rangle$ is not necessarily a model according to [Definition 2](#). The problem is that the constraint stating that if $wBv \preceq_E u$ then wBu is not necessarily satisfied. A particularly simple example is a uniform model where $W = \{w, v\}$ and $E(w) = E(v) = \{W\}$. If we take, for example, $B = \{(w, w)\}$, then it should be clear that $wBw \preceq_E v$ yet $w \not\preceq v$.

However, this can never happen over an intended model:

Lemma 1. *Suppose that $\mathcal{M} = \langle W, E, V \rangle$ is an evidence model, then \mathcal{M}^Δ is a model according to [Definition 2](#).*

Proof. Most conditions are obvious; the only one that we need to check is that if $uB_E w$ and $w \preceq_E v$ then $uB_E v$. Since $uB_E w$, there is a scenario \mathcal{X} such that $w \in \bigcap \mathcal{X}$. But since $w \preceq_E v$, it follows that also $v \in \bigcap \mathcal{X}$, and since \mathcal{X} was a u -scenario, we have that $uB_E v$. \square

2.3. p -morphisms

It is important to identify maps between structures which preserve truth-sets. In our setting, the following definition will be very useful:

Definition 7. Given evidence models $\mathcal{M}_1 = \langle W_1, B_1, E_1, \preceq_1, V_1 \rangle$ and $\mathcal{M}_2 = \langle W_2, B_2, E_2, \preceq_2, V_2 \rangle$, 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 _{B} if wB_1v then $\pi(w)B_2\pi(v)$
- back _{B} if $\pi(w)B_2u$ then there is $v \in \pi^{-1}(u)$ such that wB_1v
- forth _{E} if wE_1X then there is $Y \subseteq W_2$ such that $\pi(w)E_2Y$ and $Y \subseteq \pi[X]$
- back _{E} if $\pi(w)E_2Y$ then there is X such that $\pi[X] \subseteq Y$ and wE_1X
- forth _{\preceq} if $w \preceq_1 v$ then $\pi(w) \preceq_2 \pi(v)$
- back _{\preceq} if $\pi(w) \preceq_2 u$ then there is $v \in \pi^{-1}(u)$ such that $w \preceq_1 v$

We define p -morphism between structures of smaller signature by restricting the above conditions accordingly.

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

Theorem 1. *If π is a p -morphism between σ -models \mathcal{M}_1 and \mathcal{M}_2 and φ is any formula over the signature σ , then*

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

If a surjective p -morphism exists from \mathcal{M}_1 to \mathcal{M}_2 , we will write $\mathcal{M}_1 \gg \mathcal{M}_2$. In Section 4 we will show that, given a general evidence model \mathcal{M} , there is an evidence model $\mathcal{M}^\nabla = \langle W^\nabla, E^\nabla, V^\nabla \rangle$ such that $(\mathcal{M}^\nabla)^\Delta \gg \mathcal{M}$. Thus a general evidence model may always be represented as a p -morphic image of an intended evidence model. This intended model \mathcal{M}^∇ , however, will often be much larger.

2.4. Special classes of evidence models

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 8. Suppose that \mathcal{M} is a $(EB \preceq)$ -model. We say \mathcal{M} is **flat** if whenever wEX there is $v \in X$ such that wBv .

We say \mathcal{M} is **uniform** if E and B are constant (that is, given u, v, w , then wBv if and only if uBv , and given w, v , $E(w) = E(v)$), and whenever wBv it follows that v is \preceq -maximal. In this case, we shall treat E as a set \mathcal{E} (of neighborhoods) rather than a function and B as a set of points.

Finally, we say \mathcal{M} is **concise** if it is flat, uniform and, whenever w is \preceq -maximal then $w \in B$.

The reason for this terminology will show in the arguments to follow. The definition applies, *mutatis mutandis*, to models with different signatures (e.g., evidence models). Lemma 4 shows that if $\mathcal{M} = \langle W, \mathcal{E}, V \rangle$ is a uniform intended evidence model then every point in B_E is maximal, whereas if \mathcal{M}^Δ is flat as well, then B_E coincides exactly with the set of maximal points. Note, however, that in general the property of being concise is stronger than that of merely being flat and uniform.

Flatness may seem like an odd assumption but it is actually very natural because every finite intended model is flat, though infinite models need not be. Uniformity is also natural because one often wishes to model a situation where an agent either *has* or does *not* have evidence right at the beginning, rather than obtaining different evidence at each state.

2.5. 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. *If \mathcal{M} is a flat model, then $\mathcal{M} \models [E]\varphi \rightarrow \langle B \rangle \varphi$.*

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 a non-empty intersection. \square

Meanwhile, the formula $[E]\varphi \rightarrow \langle B \rangle \varphi$ is not valid in general. To see this, consider a uniform evidence model $\mathcal{M}_\infty = \langle W, E, V \rangle$ with domain $W = \mathbb{N}$ and evidence sets $E(w) = \mathcal{E} = \{[N, \infty) \mid N \in \mathbb{N}\} \cup \{W\}$ 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 \mathcal{E} , but $\bigcap \mathcal{E} = \emptyset$. Hence $\mathcal{M}_\infty \models [B]\perp$, i.e., $\mathcal{M}_\infty \not\models \langle B \rangle \top$; yet $\mathcal{M}_\infty \models [E]\top$ (this formula is universally valid), and we conclude

$$\mathcal{M}_\infty \not\models [E]\top \rightarrow \langle B \rangle \top.$$

From this we get the following corollary:

Corollary 1. *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 it validates $[E]\top \rightarrow \langle B \rangle \top$; but as we have just shown, this formula of our language is not valid over all uniform evidence models. \square

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

Name	Formula	Logic(s)
tautology	all propositional tautologies	all
S5_A	S5 axioms for A	all
K_B	K axioms for B	all
S4_≲	S4 axioms for $[\preceq]$	all
no empty evidence	$\langle E \rangle \top$	all
pullout	$[E]\varphi \wedge [A]\psi \leftrightarrow [E](\varphi \wedge [A]\psi)$	all
universality	$[A]\varphi \rightarrow [X]\varphi$ for $X = E, B, \preceq$	all
E-monotonicity	$\frac{\varphi \rightarrow \psi}{[E]\varphi \rightarrow [E]\psi}$	all
plausible evidence	$[E]\varphi \rightarrow [E](\varphi \wedge [\preceq]\varphi)$	all
B-monotonicity	$[B]\varphi \rightarrow [B][\preceq]\varphi$	all
flatness	$[E]\varphi \rightarrow \langle B \rangle \varphi$	b, bu
○-uniformity	$\bigcirc \varphi \rightarrow [A] \bigcirc \varphi$ for $\bigcirc = [B], \langle B \rangle, [E], \langle E \rangle$	u, bu
maximality	$\langle B \rangle (\langle \preceq \rangle \alpha \wedge \beta) \rightarrow \langle B \rangle (\alpha \wedge \langle \preceq \rangle \beta)$	u, bu
conciseness	$[B]\varphi \rightarrow \langle \preceq \rangle [\preceq]\varphi$	bu
MP	Modus Ponens	all
N_A	$\frac{\varphi}{[A]\varphi}$	all

The above table gives a family of axioms which in turn define a family of logics; given a signature σ , we let Ev^σ denote the logic which uses only those axioms and rules that fall within \mathcal{L}^σ and are marked “all”. The subscripts b, u denote the addition of the respective axioms; note that the logic Ev_{bu}^σ includes the additional conciseness axiom which is not present in any other logic. We will denote derivability in the logic λ by \vdash_λ , where λ is any one of the four combinations that we may form from a given σ , that is, $\lambda \in \{\text{Ev}^\sigma, \text{Ev}_b^\sigma, \text{Ev}_u^\sigma, \text{Ev}_{bu}^\sigma\}$. If Θ is a (possibly infinite) set of formulas, $\Theta \vdash_\lambda \varphi$ means that there is a finite $\Theta' \subseteq \Theta$ such that $\vdash_\lambda \bigwedge \Theta' \rightarrow \varphi$.

These axioms are sound for their respective classes of models:

Theorem 2. *The logics $\text{Ev}^\sigma, \text{Ev}_b^\sigma, \text{Ev}_u^\sigma, \text{Ev}_{bu}^\sigma$ are sound for the classes of all σ -models, all uniform σ -models, all flat σ -models and all concise σ -models, respectively.*

Proof. We assume familiarity with neighborhood semantics and well-known modal logics, so we will restrict ourselves to commenting on the more unusual axioms.

Let \mathcal{M} be any model; by passing to \mathcal{M}^Δ if necessary, we may assume without loss of generality that $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$.

Let us begin by checking **plausible evidence**. Suppose that $w \in \llbracket [E]\varphi \rrbracket_{\mathcal{M}}$, so that there is $X \in E(w)$ such that $X \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}$. Pick $v \in X$ and suppose $v \preceq u$. Then, $u \in X$, so that $u \in \llbracket \varphi \rrbracket_{\mathcal{M}}$; since u was arbitrary, $v \in \llbracket \varphi \wedge [\preceq]\varphi \rrbracket_{\mathcal{M}}$ and thus $w \in \llbracket [E](\varphi \wedge [\preceq]\varphi) \rrbracket_{\mathcal{M}}$, as claimed.

The **B-monotonicity** axiom follows easily from the condition that if $u \preceq v$ and wBu , then wBv .

The **flatness** axiom is [Lemma 2](#) while **conciseness** follows from [Theorem 3](#), which we prove later, and it gives the slightly stronger $[B]\varphi \rightarrow [A]\langle \preceq \rangle [\preceq]\varphi$.

Finally, the **maximality** axiom uses the fact that, over concise models, every element of B is maximal. For indeed, if $\langle B \rangle (\alpha \wedge \langle \preceq \rangle \beta)$ holds, there is some $w \in B \cap \llbracket \alpha \wedge \langle \preceq \rangle \beta \rrbracket_{\mathcal{M}}$ and thus some $v \succ w$ with $v \in \llbracket \beta \rrbracket_{\mathcal{M}}$. Since w is maximal we must have that $v \preceq w$ and that v is maximal as well, so that $v \in B$ and v satisfies $\langle \preceq \rangle \alpha$. But it also satisfies β , so we have that $\langle B \rangle (\langle \preceq \rangle \alpha \wedge \beta)$ holds. \square

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

3. Belief and plausibility

Now we step aside, and consider a different tradition in modeling beliefs.

A *plausibility* model is a tuple $\langle W, \preceq, V \rangle$ where W is a non-empty 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 [36] for details and pointers to the relevant literature). Our evidence models 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? The answer is subtle (see [8, Section 5] for some initial observations).

On plausibility models, it is commonplace to say that an agent believes a formula if it is true in all the \preceq -maximal worlds. This makes belief definable in the language containing only the universal modality $[A]$ and the plausibility modality $[\preceq]$, namely, as $[A]\langle \preceq \rangle [\preceq]\varphi$.¹² This may look like a simple technical ploy, but Baltag and Smets [6] interpret $[\preceq]\varphi$ independently as “a *safe belief* in φ ”. Following [32], they show that this amounts to the beliefs the agent retains under all new true information about the actual world.¹³ Our first observation is that our belief operator $[B]$ is not similarly definable over the class of all evidence models.

Lemma 3. *The modality $[B]$ is not definable in terms of $[A], [E], [\preceq]$, even over uniform intended evidence models.*

Proof. The model \mathcal{M}_∞ from Corollary 1 will do. Observe that $\mathcal{M}_\infty^\Delta$ is bisimilar to the trivial model \mathcal{A}^Δ where \mathcal{A} has one point w and one evidence set $\{w\}$ (simply because all propositional variables are false). Yet $\mathcal{M}_\infty^\Delta \models [B]\perp$ while the model \mathcal{A}^Δ is finite, and hence it validates $[B]\top$. \square

The situation changes over the class of flat models. Here the following notion will be useful later:

Definition 9. Given a uniform evidence model $\mathcal{M} = \langle W, \mathcal{E}, V \rangle$, write $\mathcal{E}[w] = \{X \in \mathcal{E} \mid w \in X\}$.

This construction suggests an extension to our language that we will briefly discuss before moving on to the main result of this section. While what follows can be skipped without loss of continuity, it does give a more concrete idea of the many natural kinds of evidence structure to be found in our models.

Digression: reliable and unreliable evidence Suppose that $\mathcal{M} = \langle W, \mathcal{E}, V \rangle$ is a uniform evidence model. Then $\mathcal{E}[w]$ is the set of evidence that is “correct” or reliable at state w . Unlike the agent’s full evidential state at w , the set of reliable evidence at w can always be consistently combined, suggesting a new modality $[C]\varphi$ meaning “the agent’s reliable evidence entails φ ” or, for lack of a better term, “ φ is reliably believed”. The formal definition is:

$$\mathcal{M}, w \models [C]\varphi \quad \text{iff} \quad \text{for all } v \in \bigcap \mathcal{E}[w], \mathcal{M}, v \models \varphi.$$

There is one technical issue we need to address here, which has interesting consequences. Note that the set of reliable evidence differs from state to state. That is, even if the evidence function is assumed to be constant, the *reliable* evidence function will not be constant. This means that, unlike plain belief, even in

¹² This is not quite correct if there is more than one information cell (i.e., equivalence class of the knowledge relation). The formula $[A]\langle \preceq \rangle [\preceq]\varphi$ means that φ is believed at *all states in the model*. We are implicitly assuming that there is only one “information cell” for the agent, namely the set of all states W . In this case, the universal modality $[A]$ is best interpreted as “knowledge”.

¹³ For the same notion in the computer science literature on agency, see [31]. Also, see [33] for related discussions.

uniform evidence models, reliable belief does not satisfy introspection properties. So, what type of doxastic attitude is \Box^C ? We do not have the space here for an extensive discussion, but here are a few comments.

First, notice that \Box^C validates the truth axiom ($\Box^C\varphi \rightarrow \varphi$), so believing something based on reliable evidence implies truth. Of course, it is only the modeler, from a third person perspective, who can actually determine what the agent believes based only on reliable evidence. Since the agent does not have access to the actual world, she herself cannot determine which evidence is reliable and which is not. Second, the restriction to “truthful” evidence suggests that reliable belief might be *safe belief* on evidence models. However, the derived plausibility ordering is typically not connected, and on such models, safe belief quantifies over all worlds not strictly less plausible than the current world. In particular, if φ is safely believed at a world w in a non-connected plausibility model, then φ must be true at all worlds that are incomparable with w . Now, for a derived plausibility ordering, there are two reasons why a state v may be incomparable to the actual state w : either there is reliable evidence not containing v , or there is evidence containing v but not w . This suggests yet another modality.

Let $\mathcal{M} = \langle W, E, V \rangle$ be an evidence model and define $E^U(w) = \{X \in E(w) \mid w \notin X\}$. This is the *unreliable* or incorrect evidence at state w . The corresponding modality is \Box^U : “ φ follows from the unreliable evidence at w ”. Of course, the agent cannot necessarily consistently combine this evidence, but in the formal definition we can take the *union* of these sets:


$$\mathcal{M}, w \models [U]\varphi \text{ iff for all } v \in \bigcup E^U(w), \mathcal{M}, v \models \varphi.$$

It is the conjunction of these two operators that corresponds to safe belief on evidence models.

These new operators are not definable in our basic language \mathcal{L} .

Fact 1. *The operators $[C]$ and $[U]$ are not definable in evidence belief language $\mathcal{L}_{ABE\lessdot}$.*

Proof. Let \mathcal{M} and \mathcal{M}' be uniform evidence models with the evidence sets:



The dashed line is a p -morphism, so $\mathcal{M} \gg \mathcal{M}'$. So, \mathcal{M}, w and \mathcal{M}', v satisfy the same formulas of $\mathcal{L}_{ABE\lessdot}$. However, we have¹⁴ $\mathcal{M}, w \models [C]p \wedge [U]\neg p$ while $\mathcal{M}', v \models \neg[C]p \wedge \neg[U]\neg p$. \square

This ends our digression on reliable and unreliable evidence as an example of richer languages supported by our evidence models. This theme of language extensions will return in Section 7, but for now, we return to the earlier source of variety, the existence of natural subclasses of evidence models.

We are now ready to clarify the relationship between uniform, flat evidence models and plausibility models. It will be helpful to start with the following observation:

Lemma 4. *If $\mathcal{M} = \langle W, \mathcal{E}, V \rangle$ is a uniform evidence model, then every $w \in B_{\mathcal{E}}$ is maximal and every non-empty scenario \mathcal{X} is of the form $\mathcal{E}[w]$ for some $w \in B_{\mathcal{E}}$, so that*

¹⁴ In general, it need not be the case that the agent reliably believes φ iff the agent unreliably believes $\neg\varphi$. We do not discuss the logic of these operators here, but the reader might note that the set of reliable evidence at a state is a *filter* while the set of unreliable evidence is an *ideal*.

$$\bigcap \mathcal{X} = \{v: w \preceq_{\mathcal{E}} v \text{ and } v \preceq_{\mathcal{E}} w\}.$$

If \mathcal{M} is flat as well as uniform, then $B_{\mathcal{E}}$ is exactly the set of maximal points and every collection of the form $\mathcal{E}[w]$ with $w \preceq_{\mathcal{E}}$ -maximal is a scenario.

Proof. First assume that \mathcal{M} is a uniform evidence model and $w \in \bigcap \mathcal{X}$ for some scenario \mathcal{X} . It follows that $\mathcal{X} \subseteq \mathcal{E}[w]$ (since w lies in every element of \mathcal{X}), and by maximality of \mathcal{X} , $\mathcal{X} = \mathcal{E}[w]$. But if we had $w \prec_{\mathcal{E}} v$ for some v , then clearly $\mathcal{E}[w] \subsetneq \mathcal{E}[v]$, which would contradict the maximality of \mathcal{X} and thus w is $\preceq_{\mathcal{E}}$ -maximal.

Furthermore, the elements of $\bigcap \mathcal{E}[w]$ are precisely those v such that $w \preceq_{\mathcal{E}} v$ (by definition of $\preceq_{\mathcal{E}}$ and $\mathcal{E}[w]$), but from the maximality of w we have that $v \preceq_{\mathcal{E}} w$ as well.

For the other direction, assume \mathcal{M} is flat and uniform and w is maximal. By Zorn’s lemma, $\mathcal{E}[w]$ can be extended to a scenario \mathcal{X} . By flatness, we have that $\bigcap \mathcal{X} \neq \emptyset$. If $w \notin \bigcap \mathcal{X}$, there is some $v \in \bigcap \mathcal{X} \subseteq \bigcap \mathcal{E}[w]$. But then v lies in every evidence set containing w , yet there is some $Y \in \mathcal{X}$ with $v \in Y$ and $w \notin Y$, so by definition $w \prec_{\mathcal{E}} v$, which contradicts the maximality of w . We conclude that $\mathcal{X} = \bigcap \mathcal{E}[w]$, and thus $w \in B_{\mathcal{E}}$. \square

With this, we may describe the plausibility orders of evidence models.

Definition 10. Let \preceq be a plausibility order over W . Say $D \subseteq W$ is *directed* if any two elements of D have an upper bound in D .

A plausibility order \preceq satisfies the *boundedness condition* if every directed set D has an upper bound (not necessarily in D).

Lemma 5. *If an evidence model is flat, then its derived plausibility relation satisfies the boundedness condition.*

Proof. Assume that \mathcal{M} is flat and let D be any directed set. Consider the family $\mathcal{Y} = \{Y: Y \cap D \neq \emptyset\}$. We claim that \mathcal{Y} has the *fp*. Indeed, let $Y_1, \dots, Y_n \in \mathcal{Y}$ and $y_i \in Y_i \cap D$. Let y be an upper bound for all y_i (it is an easy exercise to show that directed sets have upper bounds for finite subsets). Then, by definition, $y \in \bigcap_{i \leq n} Y_i$.

Thus \mathcal{Y} can be extended to a scenario \mathcal{X} . Since \mathcal{W} is flat, then $\bigcap \mathcal{X}$ is non-empty. But every $w \in \bigcap \mathcal{X}$ is an upper bound for D . \square

Corollary 2. *If \mathcal{M} is flat and \preceq is its derived plausibility relation then for every w there is $v \succ w$ such that v is maximal.*

Proof. If \preceq is directed it satisfies the conditions of Zorn’s lemma, and we may use it to find maximal elements above any given w . \square

Now we are ready to show that belief is definable over flat models:

Theorem 3. *Over the class of uniform evidence models with derived plausibility relation, $[A]\langle \preceq \rangle [\preceq] \varphi$ implies $[B]\varphi$.*

Over the class of models that are moreover flat, they are equivalent.

Proof. First assume that $[A]\langle \preceq \rangle [\preceq] \varphi$ and let \mathcal{X} be any scenario. Let $w \in \bigcap \mathcal{X}$. By Lemma 4, w is maximal; but since it satisfies $\langle \preceq \rangle [\preceq] \varphi$, it follows that w satisfies φ .

Now assume that the model is flat and $[B]\varphi$ holds and let $w \in W$. We have to show that there is $v \succ w$ satisfying $[\preceq] \varphi$. Use Corollary 2 to find a maximal $v \succ w$. By Lemma 4 we have that v lies in $\bigcap \mathcal{X}$ for some

scenario \mathcal{X} . Meanwhile, if $u \succ v$ then we also have $u \in \bigcap \mathcal{X}$, and by the assumption that $[B]\varphi$ holds, u also satisfies φ , i.e. v satisfies $[\preceq]\varphi$, as desired. \square

4. The representation theorem

It is generally convenient to work with models that are not necessarily intended, 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. For indeed, while it is not the case that every model is intended, we do have that every model is the p -morphic image of an intended model.

Our goal is then to define an intended model \mathcal{M}^∇ given an arbitrary model \mathcal{M} . First, let us define our domain, which depends on the underlying logic. Below, recall that

$$B[W] = \{w \in W : \exists v \in W \text{ such that } vBw\}.$$

Definition 11. Let $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ be a model. Define the following sets:

- $W^\nabla = W_u^\nabla$ is the set of all tuples (w, x, f) where $w \in W$, $x \in B[W] \cup \mathbb{N}$ and $f : W \rightarrow \{0, 1\}$.
- W_b^∇ is the set of all $(w, x, f) \in W^\nabla$ with $x \in B[W]$.
- W_{bu}^∇ is the set of all $(w, 0, f)$ such that $w \in W$ and $f : W \rightarrow \{0, 1\}$.

Define $\pi(w, x, f) = w$.

These will be the domains of our intended models. Note that W_λ^∇ is defined even if \mathcal{M} is not a λ -model, but in this case it will generally not be very useful. We remark that the parameter 0 in the concise case is simply a dummy parameter which we use for uniformity with the other cases.

Let us begin with a simple observation about the size of W^∇ :

Lemma 6. *The set W_λ^∇ is finite if and only if \mathcal{M} is finite and λ is flat.*

Before we define evidence sets, which are somewhat more elaborate, let us define the accessibility relations.

Definition 12. Given a model $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ and an evidence logic λ , define a relation \preceq_λ^∇ on W_λ^∇ by $(w, x, f) \preceq_\lambda^\nabla (v, y, g)$ if and only if $w \preceq v$, either $x = y \in W$ or $x \leq y \in \mathbb{N}$, and $f \uparrow v = g \uparrow v$, where $\uparrow v = \{v' \mid v \preceq v'\}$ is the *upset* of v .

The function π as we have defined it is a p -morphism on the restricted structures:

Lemma 7. *The function $\pi : W_\lambda^\nabla \rightarrow W$ is a surjective p -morphism with respect to the signature \preceq .*

Proof. The *forth* and *atoms* conditions are straightforward from the definitions so we focus on *back* $_{\preceq}$. But if $w \preceq v$ and $w = \pi(w, x, f)$, then $(w, x, f) \preceq_\lambda^\nabla (v, x, f)$ and $\pi(v, x, f) = v$. \square

We can now define the neighborhoods we shall use.

Definition 13. We define the following sets:

1. Given $v \in W$, $f : W \rightarrow \{0, 1\}$ and $X \subseteq W$, define $X^f(v)$ to be the set of all $(w, x, g) \in W_b^\nabla$ such that $x = v$ and $g \uparrow w = f \uparrow w$.

2. Given $n \in \mathbb{N}$, $f : W \rightarrow \{0, 1\}$ and $X \subseteq W$, define $X_\lambda^f(n)$ to be the set of all $(w, m, g) \in W_\lambda^\nabla$ such that $m \geq n$ and $g \upharpoonright w = f \upharpoonright w$.

Our evidence relation will be divided into several parts:

1. Given $(w, x, f) \in W^\nabla$ and $v \in W$, let $E^v(w, x, f)$ be the collection of all sets of the form $X^g(v)$ such that wEX and $g \upharpoonright v = f \upharpoonright v$.
2. Given $(w, x, f) \in W^\nabla$, $n \in \mathbb{N}$, let $E_\lambda^n(w, x, f)$ be the collection of all sets of the form $X_\lambda^g(n)$ such that wEX and X is not flat for w . Note that $E_b^n(w, x, f) = \emptyset$.

If $\mathcal{M} = \langle W, \mathcal{E}, B, \preceq \rangle$ is uniform, define

1. Given $v \in W$, \mathcal{E}^v to be the collection of all sets of the form $X^g(v)$ such that $v \in X \in \mathcal{E}$ and $g(v') = 1$ for all $v' \succ v$.
2. Given $n \in \mathbb{N}$, \mathcal{E}_u^n to be the collection of all sets of the form $X_\lambda^g(n)$ such that $X \in \mathcal{E}$ and X is not flat.
3. Define \mathcal{E}_{bu}^0 to be the collection of all sets of the form $X_\lambda^g(0)$ such that $X \in \mathcal{E}$.

For non-uniform λ let

$$E_\lambda^\nabla(w, x, f) = \{W_\lambda^\nabla\} \cup \bigcup_{wBv} E^v(w, x, f) \cup \bigcup_{n \in \mathbb{N}} E_\lambda^n(w, x, f).$$

Finally, let

$$\mathcal{E}_u^\nabla = \{W_u^\nabla\} \cup \bigcup_{v \in B} \mathcal{E}^v \cup \bigcup_{n \in \mathbb{N}} \mathcal{E}_u^n \quad \text{and} \quad \mathcal{E}_{bu}^\nabla = \{W_{bu}^\nabla\} \cup \mathcal{E}_{bu}^0.$$

We remark that $E_\lambda^\nabla(w, x, f)$ does not depend on x but we write it this way to conform to the definition of an evidence relation. It will take some work to check that the belief relation is intended, but less so to check that plausibility is.

Lemma 8. *If $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ is any model then $\preceq_\lambda^\nabla = \preceq_{E^\nabla}$.*

If moreover \mathcal{M} is uniform then $\preceq_\lambda^\nabla = \preceq_{\mathcal{E}^\nabla}$ for uniform λ .

Proof. First let us check that $\preceq_\lambda^\nabla \subseteq \preceq_{E^\nabla}$. Suppose that $(w, x, f) \preceq_\lambda^\nabla (v, y, g)$ and $(w, x, f) \in X \in E^\nabla(u, z, h)$. If $X = W^\nabla$ there is nothing to do, but otherwise consider three cases:

1. If $X = Y^{h'}(u')$ with uBu' , then we must have $x = y = u'$, f agrees with h' on $\uparrow w$ and hence g does on $\uparrow v$. It follows that $(v, y, g) \in Y^{h'}(u')$.
2. If $X = Y^h(k)$, as before the condition on g is satisfied and since $x \leq y$ we also have $(v, y, g) \in Y^h(k)$.

For the other direction, if $(w, x, f) \not\preceq_\lambda^\nabla (v, y, g)$, consider now three cases.

1. If $w \not\leq v$, then define h to be identical to f except that $h(v) \neq g(v)$ and define X as follows: if there is u such that wBu , then $X = W^h(u)$; otherwise, W is not flat for w and we choose $X = W_\lambda^h(0)$. In either case $X \in E(w, x, f)$, f agrees with h on $\uparrow w$ so $(w, x, f) \in X \in E(w, x, h)$ yet $(v, y, g) \notin X$.
2. If $x \neq y$ and at least one of them belongs to W (for example, x), find u such that uBx . Then, $(w, x, f) \in W^f(x) \in E(u, x, f)$ but $(u, y, g) \notin W^f(x)$. The cases where either $y \in W$ but $x \in \mathbb{N}$ or $x > y \in \mathbb{N}$ are very similar.

3. If $f \uparrow v \neq g \uparrow v$ and $x \in W$, let x' be such that $x'Bx$ if $x \in W$. Setting $X = W^f(x')$ we see that $(w, x, f) \in X \in E(w, x, f)$ but $(v, y, g) \notin W^f(x')$. If instead $x \in \mathbb{N}$, we can set $X = W^f(0)$.

We conclude that $\preceq_\lambda^\nabla = \preceq_{E^\nabla}$.

The uniform cases are very similar (and simpler). \square

We can now define our intended models:

Definition 14. Let $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ be a λ -model.

Define $V_\lambda^\nabla = W_\lambda^\nabla \cap \pi^{-1}V$ and set

$$\mathcal{M}_\lambda^\nabla = \langle W_\lambda^\nabla, E_\lambda^\nabla, V_\lambda^\nabla \rangle.$$

These models are convenient because scenarios are easy to identify. Below, let $E_\lambda^z(w, x, f) = E^z(w, x, f)$ if λ is not uniform, $E_u^z(w, x, f) = \mathcal{E}^z$.

Lemma 9. *Given a non-concise λ , a λ -model $\mathcal{M} = \langle W, E, B, \preceq, V \rangle$ and $\mathcal{M}_\lambda^\nabla = \langle W_\lambda^\nabla, E_\lambda^\nabla, V_\lambda^\nabla \rangle$:*

1. *for every set \mathcal{X} of the form*

$$\mathcal{X} = \{W_\lambda^\nabla\} \cup E_\lambda^v(w, x, f), \tag{1}$$

where wBv , \mathcal{X} is a (w, x, f) -scenario, and $\pi(\bigcap \mathcal{X}) = \uparrow v$;

2. *if \mathcal{X} is a (w, x, f) -scenario that is not of the above form, then*

$$\mathcal{X} \subseteq \{W_\lambda^\nabla\} \cup \bigcup_{n \in \mathbb{N}} E_\lambda^n(w, x, f) \tag{2}$$

and $\bigcap \mathcal{X} = \emptyset$.

Proof. We shall only consider the non-uniform case; the uniform one is analogous.

If $(v, y, g) \in X \in E^z(w, x, f)$, then either $y = z \in W$ or $y, z \in \mathbb{N}$. It follows that if $X, Y \in \mathcal{X}$ are such that $X \in E^z(w, x, f)$ and $Y \in E^{z'}(w, x, f)$, either $z = z' \in W$, $z, z' \in \mathbb{N}$, or $X \cap Y = \emptyset$, where the latter is impossible if \mathcal{X} is a scenario.

Hence every scenario is either contained in a collection of the form (1) or is of the form (2). Let us first check that collections of the form (1) are indeed scenarios and have the appropriate intersection.

So assume that \mathcal{X} is of form (1) for some v, f with wBv . First note that if $v \preceq u$ then $(u, v, f) \in \bigcap \mathcal{X}$. For clearly, $(u, v, f) \in W_\lambda^\nabla$ and every set in $E^v(w, x, f)$ is of the form X_g^v for some g which agrees with f on $\uparrow v$, thus on $\uparrow u$, which means that $(u, v, f) \in X^f(v)$. It follows that $\uparrow v \subseteq \pi(\bigcap \mathcal{X})$.

Moreover, if $v \not\preceq u$, then any element of $\pi^{-1}(u)$ is of the form (u, u', g) . Define a function h which is identical to f except that $f(u) \neq g(u)$. Then, by definition $(u, u', g) \notin W^h(v)$, yet wEW and h agrees with g on $\uparrow v$ so that $W^h(v) \in E^v(w, x, f)$ and $(u, u', g) \notin \bigcap \mathcal{X}$. Since u', g were arbitrary, we conclude that $u \notin \pi(\bigcap \mathcal{X})$. Thus, $\uparrow v \supseteq \pi(\bigcap \mathcal{X})$, and our claim follows.

Now let us check that if \mathcal{X} is of form (2), then it has an empty intersection. We begin with three simple observations:

1. \mathcal{X} contains at least one element of the form $X^f(n)$ with $n \in \mathbb{N}$. For indeed, as discussed above, if \mathcal{X} is not of form (1) then it does not intersect any $E^v(w, x, f)$ with $v \in W$, and the collection $\{W_\lambda^\nabla\}$ by itself cannot be a (w, x, f) -scenario, for either wBv for some v , in which case $\{W_\lambda^\nabla, W^f(v)\}$ would still

have a non-empty intersection (it includes (v, v, f)), or there is no such v so that W is not flat for x , and $\{W_\lambda^\nabla, W_\lambda^f(0)\}$ would be a proper extension with a non-finite intersection (it includes $(w, 0, f)$).

2. Any element of $\bigcap \mathcal{X}$ must be of the form (v, n, g) with $n \in \mathbb{N}$ (for $X^f(n)$ contains only elements of this form).
3. If $n < k$ and $(v, n, g) \in X \in \mathcal{X}$ it easily follows that $(v, k, g) \in X$ by case-by-case inspection.

With this, towards a contradiction we assume that $(v, n, g) \in \bigcap \mathcal{X}$ and pick $X^f(m) \in \mathcal{X}$. Thus, $(v, n + 1, g) \in X^f(n + 1) \cap \bigcap \mathcal{X}$, from which it follows that $(v, n, g) \notin X^f(n + 1) \cap \bigcap \mathcal{X} \neq \emptyset$ and thus $X^f(n + 1) \notin \mathcal{X}$; but this contradicts the maximality of \mathcal{X} , and we conclude that there may be no such point in the intersection and hence it must be empty. \square

An immediate corollary of the above is the following:

Corollary 3. *If \mathcal{M} is flat, then so is \mathcal{M}_b^∇ .*

Proof. By definition, \mathcal{M}_b^∇ cannot contain any scenarios of form (2), hence every scenario must be of form (1) and thus have a non-empty intersection. \square

We may now collect our results into the following:

Theorem 4. *Given a λ -model \mathcal{M} , there exists an evidence model $\mathcal{M}_\lambda^\nabla = \langle W_\lambda^\nabla, E_\lambda^\nabla, V_\lambda^\nabla \rangle$ such that \mathcal{M} is a p-morphic image of $(\mathcal{M}_\lambda^\nabla)^\Delta$.*

Further, if \mathcal{M} is (i) flat and finite, (ii) flat, (iii) uniform or (iv) finite and concise, respectively, then so is \mathcal{M}^∇ .

Proof. We have explicitly stated and proven most required properties. One exception is the B -clauses of a p-morphism, but this follows immediately from Lemma 9, and in the concise case from Lemma 4, using the fact that finite models are flat, so that $B_{\mathcal{E}_{bu}^\nabla}$ is precisely the set of maximal worlds. \square

5. Syntactic properties of evidence logics

We now turn our attention to establishing our main completeness results. In order to do this, we shall collect in this section the preliminary syntactic results we shall need. Recall that if Θ is a (possibly infinite) set of formulas, $\Theta \vdash_\lambda \varphi$ means that there is finite $\Theta' \subseteq \Theta$ such that $\vdash_\lambda \bigwedge \Theta' \rightarrow \varphi$. In this section and the next we will work exclusively over the full signature $\sigma = EBA \preceq$.

Lemma 10. *Let Γ be a set of formulas and λ an evidence logic over $\sigma = EBA \preceq$.*

1. *If $\varphi \wedge [\preceq]\varphi, [A]\Gamma \vdash_\lambda \psi$, then*

$$[E]\varphi, [A]\Gamma \vdash_\lambda [E]\psi.$$

2. *If $\Phi, [A]\Gamma \vdash_\lambda \psi$, then*

$$[B]\Phi, [A]\Gamma \vdash_\lambda [B]\psi.$$

3. *If $\Phi, [\preceq]\Delta, [A]\Gamma \vdash_\lambda \psi$, then*

$$[\preceq]\Phi, [\preceq]\Delta, [A]\Gamma \vdash_\lambda [\preceq]\psi.$$

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

$$\bigwedge [A]\Gamma \leftrightarrow [A] \bigwedge \Gamma$$

is derivable in S5, so we can replace $[A]\Gamma$ by a single formula $[A]\gamma$.

1 If $\varphi \wedge [\preceq]\varphi, [A]\gamma \vdash_\lambda \psi$, then

$$\vdash_\lambda (\varphi \wedge [\preceq]\varphi) \wedge [A]\gamma \rightarrow \psi.$$

By the monotonicity rule

$$\vdash_\lambda [E]((\varphi \wedge [\preceq]\varphi) \wedge [A]\gamma) \rightarrow [E]\psi.$$

Applying the **pullout** axiom we get

$$\vdash_\lambda [E](\varphi \wedge [\preceq]\varphi) \wedge [A]\gamma \rightarrow [E]\psi,$$

that is,

$$[E](\varphi \wedge [\preceq]\varphi), [A]\gamma \vdash_\lambda [E]\psi.$$

Finally, using **plausible evidence** we see that

$$[E]\varphi, [A]\gamma \vdash_\lambda [E]\psi.$$

2 If $\Phi, [A]\gamma \vdash_\lambda \psi$, since $[B]$ is a normal operator, we also have

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

But since $[A]\alpha \rightarrow [A][A]\alpha$ and $[A][A]\alpha \rightarrow [B][A]\alpha$ are axioms, we also get

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

3 As before we may assume that Δ is finite. Also, $[\preceq]$ is a normal operator and we obtain

$$[\preceq]\Phi, [\preceq][\preceq]\Delta, [\preceq][A]\gamma \vdash_\lambda [X]\psi$$

and, from this,

$$[\preceq]\Phi, [\preceq][\preceq]\Delta, [A]\gamma \vdash_\lambda [X]\psi.$$

Finally, we use the transitivity of $[\preceq]$ to obtain

$$[\preceq]\Phi, [\preceq]\Delta, [A]\gamma \vdash_\lambda [X]\psi,$$

as desired. \square

It will be very useful to draw some dual conclusions from the above result. Given a set of formulas Φ , define $\Phi^{[A]}$ as $\{\psi \mid [A]\psi \in \Phi\}$, and similarly for the other modalities. For a set of formulas Γ and a signature σ containing A , if $\lambda \in \{\text{Ev}^\sigma, \text{Ev}_b^\sigma\}$, define $\Gamma^\lambda = [A]\Gamma^{[A]} \cup \langle A \rangle \Gamma^{\langle A \rangle}$; if $\lambda \in \{\text{Ev}_u^\sigma, \text{Ev}_{bu}^\sigma\}$, define

$$\Gamma^\lambda = \bigcup \{ \bigcirc \Gamma^\bigcirc : \bigcirc = [X], \langle X \rangle \text{ with } X \in \{E, B, A\} \}.$$

Lemma 11. *Let Γ be a set of formulas and λ an evidence logic over $\sigma = EBA \preceq$.*

1. *If Γ is λ -consistent, $\alpha \in \Gamma^{[E]}$ and $\delta \in \Gamma^{\langle E \rangle}$, then $\{\alpha \wedge [\preceq]\alpha, \delta\} \cup \Gamma^\lambda$ is λ -consistent.*
2. *If Γ is λ -consistent and $\psi \in \Gamma^{\langle B \rangle}$, then $\{\psi\} \cup \Gamma^{[B]} \cup \Gamma^\lambda$ is λ -consistent.*
3. *If Γ is λ -consistent and $\psi \in \Gamma^{\langle \preceq \rangle}$, then $\{\psi\} \cup [\preceq]\Gamma^{[\preceq]} \cup \Gamma^\lambda$ is λ -consistent.*
4. *If $\lambda \in \{\text{Ev}_b^\sigma, \text{Ev}_{bu}^\sigma\}$, Γ is λ -consistent and $\psi \in \Gamma^{[E]}$, then*

$$\{\psi \wedge [\preceq]\psi\} \cup \Gamma^{[B]} \cup \Gamma^\lambda$$

is λ -consistent.

Proof. Assume Γ is λ -consistent. As before, we may assume without loss of generality that Γ is finite.

Note that we have $\vdash_\lambda \bigwedge \Gamma^\lambda \leftrightarrow [A] \bigwedge \Gamma^\lambda$ independently of λ ; over the non-uniform case, this follows from the axiom $\langle A \rangle \varphi \rightarrow [A] \langle A \rangle \varphi$, while over the uniform case (i.e., if $\lambda \in \{\text{Ev}_u, \text{Ev}_{bu}\}$), this is because $\bigcirc \gamma$ is equivalent to $[A] \bigcirc \gamma$ over λ for any modality $\bigcirc = [X], \langle X \rangle$ with $X \in \{E, B, A\}$. Thus we may replace Γ^λ by a single formula $[A]\gamma$ equivalent to $\bigwedge \Gamma^\lambda$ over λ .

1 Suppose otherwise. Then, we would have

$$\alpha \wedge [\preceq]\alpha, [A]\gamma \vdash_\lambda \neg \delta.$$

Thus by [Lemma 10.1](#),

$$[E]\alpha, [A]\gamma \vdash_\lambda [E]\neg \delta,$$

so that

$$\{\neg[E]\neg \delta, [E]\alpha, [A]\gamma\}$$

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

2 Suppose $\psi \in \Gamma^{\langle B \rangle}$. If the claim were false, then we would have

$$\Gamma^{[B]}, [A]\gamma \vdash_\lambda \neg \psi,$$

and hence, by [Lemma 10.2](#),

$$[B]\Gamma^{[B]}, [A]\gamma \vdash_\lambda [B]\neg \psi.$$

Thus

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

would be λ -inconsistent, but this contradicts the consistency of Γ .

3 This is analogous to the previous item but uses [Lemma 10.3](#).

4 If $\psi \in \Gamma^{[E]}$, the claim follows from item 2 above using the derivability of $[E]\psi \rightarrow \langle B \rangle (\psi \wedge [\preceq]\psi)$ by **plausible evidence** and **flatness**. \square

6. Completeness

Throughout this section we will work with logics in the full signature. Recall that by an *evidence logic* we will mean exclusively a logic of one of the forms $\text{Ev}^\sigma, \text{Ev}_v^\sigma, \text{Ev}_u^\sigma, \text{Ev}_{vu}^\sigma$. A λ -type is a maximal λ -consistent set of formulas. We state the well-known Lindenbaum lemma without proof:

Lemma 12. *If Γ is λ -consistent then there is a λ -type Δ such that $\Gamma \subseteq \Delta$.*

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.

Definition 15. Given λ with signature $\sigma = EBA \preceq$ and a λ -type Φ , we define

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

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

Definition 16. Given a λ -type Φ with $[E]\alpha \in \Phi$, we define the α -neighborhood of Φ , denoted $\mathcal{N}_\lambda^\Phi(\alpha)$, as

$$\{\Psi \in \text{type}_\lambda^A(\Phi) \mid \alpha, [\preceq]\alpha \in \Psi\}.$$

Observe that $\mathcal{N}_\lambda^\Phi(\alpha)$ depends only on Φ^λ and not on the whole of Φ . We are now ready to define a canonical extended evidence model:

Definition 17. Let λ be an evidence logic over $\sigma = EAB \preceq$ and Φ be a λ -type. The **λ -canonical extended evidence model for Φ** is the extended evidence model $\mathcal{M}_\lambda(\Phi) = \langle W_\lambda^\Phi, E_\lambda^\Phi, B_\lambda^\Phi, \preceq_\lambda^\Phi, V_\lambda^\Phi \rangle$ where

1. $W_\lambda^\Phi = \text{type}_\lambda^A(\Phi)$,
2. for each $p \in \text{At} \cap \Gamma$, $\Gamma \in V_\lambda^\Phi(p)$ iff $p \in \Gamma$,
3. $E_\lambda^\Phi(\Gamma) = \{\mathcal{N}_\lambda^\Phi(\alpha) \mid [E]\alpha \in \Gamma\}$,
4. $\Gamma B_\lambda^\Phi \Delta$ if $\Gamma^{[B]} \subseteq \Delta$,
5. $\Gamma \preceq_\lambda^\Phi \Delta$ if $\Gamma^{[\preceq]} \subseteq \Delta^{[\preceq]}$.

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 in this setting is unavoidable, given that our language includes a universal modality and it is impossible to satisfy all types on a single model.

Let us begin by showing that canonical models are models in the sense of [Definition 2](#):

Lemma 13. *Given a consistent λ -type Φ with a λ -canonical extended evidence model $\mathcal{M}_\lambda(\Phi)$,*

1. For each $\Gamma \in W_\lambda^\Phi$, $\Gamma E_\lambda^\Phi W_\lambda^\Phi$;
2. If $[E]\alpha \in \Phi$, $\mathcal{N}_\lambda^\Phi(\alpha)$ is non-empty;

3. The relation \preceq_λ^Φ is a preorder, and whenever $\Gamma \preceq_\lambda^\Phi \Delta$ and $\Gamma \in X \in E(\Theta)$ it follows that $\Delta \in X$.
4. If $\Gamma \preceq_\lambda^\Phi \Delta$ and $\Theta B_\lambda^\Phi \Gamma$ then $\Theta B_\lambda^\Phi \Delta$.

Proof. Suppose that Φ is a λ -type.

1 The formulas $[E]\top, \top$ are derivable, so it follows that $\top \in \Delta$ for all $\Delta \in W_\lambda^\Phi$; thus $\mathcal{N}_\lambda^\Phi(\top) = W_\lambda^\Phi$. Similarly, $[E]\top \in \Gamma$, so that $\Gamma E_\lambda^\Phi W_\lambda^\Phi$, as claimed.

2 Put $\delta = \top$ in [Lemma 11.1](#) and observe that any λ -type Δ extending $\{\alpha \wedge [\preceq]\alpha, \top\} \cup \Gamma^\lambda$ belongs to $\mathcal{N}_\lambda^\Phi(\alpha)$.

3 The transitivity of \preceq_λ^Φ follows by definition. For the second condition, let $X = \mathcal{N}_\lambda^\Phi(\alpha)$ be any neighborhood of Θ . Since $\Gamma \in X$ it follows that $\alpha, [\preceq]\alpha \in \Gamma$ so that $[\preceq]\alpha \in \Delta$. But then, $\alpha \in \Delta$ as well by the truth axiom, so $\Delta \in X$.

4 If $\Gamma \preceq_\lambda^\Phi \Delta$ and $\Theta B_\lambda^\Phi \Gamma$, we need to check that, given $[B]\varphi \in \Theta$, we have $\varphi \in \Delta$. But by **B -monotonicity** we have $[B][\preceq]\varphi \in \Theta$, so that $[\preceq]\varphi \in \Gamma$ and thus $[\preceq]\varphi \in \Delta$, which by the truth axiom implies $\varphi \in \Delta$. \square

The next lemmas are the key ingredients in the proof of the Truth Lemma.

Lemma 14. *Let λ be an evidence logic over the signature $EBA \preceq$, and Φ be a λ -type.*

1. If $[E]\alpha, \langle E \rangle \beta \in \Gamma \in W_\lambda^\Phi$, there is $\Delta \in \mathcal{N}_\lambda^\Phi(\alpha)$ with $\beta \in \Delta$.
2. If $\langle X \rangle \alpha \in \Gamma$ with $X \in \{B, \preceq\}$ there is Δ such that $\Gamma X_\lambda^\Phi \Delta$ and $\alpha \in \Delta$.
3. If $\lambda \in \{Ev_b, Ev_{bu}\}$ and $\alpha \in \Gamma^{[E]}$, then $\mathcal{N}_\lambda^\Phi(\alpha) \cap \text{type}_\lambda^B(\Gamma)$ is non-empty.

Proof. Let Φ be a λ -type.

1 Since Γ is λ -consistent and $[E]\alpha, \langle E \rangle \beta \in \Gamma$, we have, by [Lemma 11.1](#), that

$$\{\alpha \wedge [\preceq]\alpha, \beta\} \cup \Gamma^\lambda$$

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

2 For $X = B$, by [Lemma 11.2](#), since Φ 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 λ -type Δ . Evidently, $\Gamma B_\lambda^\Phi \Delta$ and $\beta \in \Delta$.

The case $X = \preceq$ is analogous but uses [Lemma 11.3](#).

3 By [Lemma 11.4](#), 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,

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

Now let us check that canonical models fall within the appropriate class of models for each logic.

Lemma 15. *Let λ be an evidence logic over $EBA \preceq$ and Φ be λ -consistent. Then,*

1. If λ is uniform, then so is $\mathcal{M}_\lambda(\Phi)$.
2. If λ is flat, then so is $\mathcal{M}_\lambda(\Phi)$.
3. If $\lambda = Ev_{bu}^\sigma$, then $\mathcal{M}_\lambda(\Phi)$ is concise.

Proof. 1 In the uniform case, 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 . It remains to check that if $\Gamma \in B_\lambda^\Phi$ then Γ is \preceq_λ^Φ -maximal.

So suppose that $\Gamma \preceq_\lambda^\Phi \Delta$ and $[\preceq]\alpha \in \Delta$. It follows that $\langle \preceq \rangle [\preceq]\alpha \in \Gamma$. If moreover $\langle \preceq \rangle \neg\alpha \in \Gamma$, then $\langle \preceq \rangle [\preceq]\alpha \wedge \langle \preceq \rangle \neg\alpha \in \Gamma$ and hence $\langle B \rangle (\langle \preceq \rangle [\preceq]\alpha \wedge \langle \preceq \rangle \neg\alpha) \in \Phi$. But then, by the **maximality** axiom and transitivity of \preceq , $\langle B \rangle ([\preceq]\alpha \wedge \langle \preceq \rangle \neg\alpha) \in \Phi$, which cannot be, since Φ is consistent.

We then conclude that $[\preceq]\alpha \in \Gamma$ and, since α was arbitrary, $\Delta \preceq_\lambda^\Phi \Gamma$, as required.

2 This amounts to showing that if $\Gamma E_\lambda^\Phi X$, there is $\Delta \in X$ such that $\Gamma B_\lambda^\Phi \Delta$. But this is immediate from the definition of the accessibility relation and neighborhoods on canonical extended evidence models and [Lemma 14.2](#).

3 As we have already seen that $\mathcal{M}_\lambda(\Phi)$ is both flat and uniform, it remains to check that if Γ is \preceq_λ^Φ -maximal then $\Gamma \in B_\lambda^\Phi$.

Whenever $[B]\varphi \in \Gamma$ it follows that $\langle \preceq \rangle [\preceq]\varphi \in \Gamma$ by **conciseness**. If Γ is \preceq_λ^Φ -maximal then it also follows that $\varphi \in \Gamma$; otherwise $[\preceq]\varphi \notin \Gamma$, while by [Lemma 14.2](#) there would be Δ with $\Gamma \preceq_\lambda^\Phi \Delta$ and $[\preceq]\varphi \in \Delta$, so that $\Gamma \prec_\lambda^\Phi \Delta$. But this would contradict the maximality of Γ . Thus $\Gamma B_\lambda^\Phi \Gamma$, and since B_λ^Φ is constant this means that $\Gamma \in B_\lambda^\Phi$. \square

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

Proposition 1 (Truth Lemma). *Given a λ -type Φ we have that, for every formula ψ and every $\Gamma \in \text{type}_\lambda^A(\Phi)$,*

$$\psi \in \Gamma \quad \Leftrightarrow \quad \Gamma \in \llbracket \psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}.$$

Proof. The proof is by induction on the structure of ψ . The only interesting cases are the modalities $[E]$, $[B]$ and $[\preceq]$ (and their duals).

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

Now, assume that $\langle E \rangle \psi \in \Gamma$. Let X be any neighborhood of Γ . Then $X = \mathcal{N}_\lambda^\Phi(\alpha)$ for some α . Hence, $[E]\alpha \in \Gamma$ and $\langle E \rangle \psi \in \Gamma$. By [Lemma 14.1](#), there is $\Theta \in \mathcal{N}_\lambda^\Phi(\alpha)$ with $\psi \in \Theta$. By the induction hypothesis, $\Theta \in \llbracket \psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}$, and, since X was arbitrary, we conclude that $\Gamma \in \llbracket \langle E \rangle \psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}$.

Let us now consider $[B]\psi \in \Gamma$. Suppose that $\Gamma B_\lambda^\Phi \Delta$. Then, by definition $\psi \in \Delta$, which by the induction hypothesis implies $\Delta \in \llbracket \psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}$. Thus, $\Gamma \in \llbracket [B]\psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}$.

Assume next that $\langle B \rangle \psi \in \Gamma$. By [Lemma 14.3](#), $\psi \in \Theta$ for some $\Theta \in \text{type}_\lambda^B(\Gamma)$. But then $\Gamma B_\lambda^\Phi \Theta$, and thus $\Gamma \in \llbracket \langle B \rangle \psi \rrbracket_{\mathcal{M}_\lambda(\Phi)}$.

The cases for $[\preceq]\psi$, $\langle \preceq \rangle \psi$ are similar and we skip them. \square

We recall that for \mathcal{M}_{bu}^∇ to be concise, we require that \mathcal{M} be concise *and* finite. Thus we need to produce finite models; we shall do so by a filtration argument, which requires only a slight modification from the standard construction.

Definition 18. Given a λ -type Φ and a formula φ , we define $\mathcal{M}_\lambda(\Phi)/\varphi = \langle W_\lambda^{\Phi/\varphi}, E_\lambda^{\Phi/\varphi}, B_\lambda^{\Phi/\varphi}, \preceq_\lambda^{\Phi/\varphi}, V_\lambda^{\Phi/\varphi} \rangle$ as follows:

1. Write $\Gamma \equiv_\varphi \Delta$ if
 - 1.1. for every subformula ψ of φ , $\psi \in \Gamma$ if and only if $\psi \in \Delta$, and
 - 1.2. Γ is \preceq_λ^Φ -maximal if and only if Δ is.

Then, let $W_\lambda^{\Phi/\varphi}$ be the set of equivalence classes of $\text{type}^A(\Phi)$ modulo \equiv_φ . We denote the equivalence class of Γ by $[\Gamma]_\varphi$.

2. For $X = B, \preceq$, $[Γ]_{\varphi} X_{\lambda}^{\Phi/\varphi} [\Delta]_{\varphi}$ if and only if there are $\Gamma' \equiv_{\varphi} \Gamma$ and $\Delta' \equiv_{\varphi} \Delta$ with $\Gamma' X_{\lambda}^{\Phi} \Delta'$.
3. $[Γ]_{\varphi} E_{\lambda}^{\Phi/\varphi} X$ if there is $\Gamma' \equiv_{\varphi} \Gamma$ and $Y \in E_{\lambda}^{\Phi}(\Gamma)$ such that $X = [Y]_{\varphi} = \{[\Delta]_{\varphi} : \Delta \in Y\}$.
4. If p is a subformula of φ then $V_{\lambda}^{\Phi/\varphi}(p) = \{[\Gamma]_{\varphi} : p \in \Gamma\}$; otherwise, $V_{\lambda}^{\Phi/\varphi}(p) = \emptyset$.

The following is standard and we present it without proof:

Lemma 16. *If ψ is a subformula of φ , then $\Gamma \in \llbracket \psi \rrbracket_{\mathcal{M}_{\lambda}(\Phi)}$ if and only if $[\Gamma]_{\varphi} \in \llbracket \psi \rrbracket_{\mathcal{M}_{\lambda}(\Phi)/\varphi}$.*

To check that filtered models preserve conciseness, we must make sure that they preserve maximal elements:

Lemma 17. *Given a λ -type Φ , a formula φ and $\Gamma \in W_{\lambda}^{\Phi}$, $[\Gamma]_{\varphi}$ is $\preceq_{\lambda}^{\Phi/\varphi}$ -maximal if and only if Γ is \preceq_{λ}^{Φ} -maximal.*

Proof. If Γ is maximal and $[\Gamma]_{\varphi} \preceq_{\lambda}^{\Phi/\varphi} [\Delta]_{\lambda}$, there are $\Gamma' \equiv_{\varphi} \Gamma$ and $\Delta' \equiv_{\varphi} \Delta$ such that $\Gamma' \preceq_{\lambda}^{\Phi} \Delta'$; but then, Γ' is maximal by definition so that $\Delta' \preceq \Gamma'$ and thus $[\Delta]_{\varphi} \preceq_{\lambda}^{\Phi/\varphi} [\Gamma]_{\varphi}$, establishing the maximality of $[\Gamma]_{\varphi}$.

For the other direction, it suffices to show that, given $\Gamma \in W_{\lambda}^{\Phi}$, there is $\Delta \succ \Gamma$ such that Δ is maximal; for then, if Γ is not maximal, then no $\Gamma' \equiv_{\varphi} \Gamma$ is maximal, whereas every $\Delta' \equiv_{\varphi} \Delta$ is maximal, so $\Gamma' \not\preceq \Delta'$ and $[\Delta]_{\varphi} \succ_{\lambda}^{\Phi/\varphi} [\Gamma]_{\varphi}$.

But if $X \subseteq W_{\lambda}^{\Phi/\varphi}$ is totally ordered, it is not hard to check that X has an upper bound (pick any λ -type extending $\preceq \cup X$). Thus by Zorn's lemma, the set of upper bounds of Γ has a maximal element, as needed. \square

With this we may prove the following:

Lemma 18. *Given a bu -type Φ and a formula φ , $\mathcal{M}_{bu}(\Phi)/\varphi$ is finite and concise.*

Proof. The finiteness is obvious because there can only be finitely many \equiv_{φ} -equivalence classes. It is not too hard to see that $\mathcal{M}_{bu}(\Phi)/\varphi$ is flat, for if $X = [Y]_{\varphi}$ is any neighborhood of $[\Gamma]_{\varphi}$, by flatness we have that $\Gamma B_{bu}^{\Phi} \Delta$ for some $\Delta \in Y$ and thus $[\Gamma]_{\varphi} B_{bu}^{\Phi/\varphi} [\Delta]_{\varphi} \in X$. It is also easy to see that $B_{bu}^{\Phi/\varphi}, E_{bu}^{\Phi/\varphi}$ are constants given that $B_{bu}^{\Phi}, E_{bu}^{\Phi}$ are.

It remains to check the maximality clause. But this is immediate from Lemma 17, for indeed $[\Gamma]_{\varphi}$ is maximal if and only if Γ is, which in turn is equivalent to $\Gamma \in B_{bu}^{\Phi}$ and thus to $[\Gamma]_{\varphi} \in B_{bu}^{\Phi/\varphi}$. \square

6.1. Completeness theorem

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

Theorem 5. *Each evidence logic λ over the full signature $EBA \preceq$ is sound and strongly complete for the class of λ -models. Moreover,*

1. Ev, Ev_b, Ev_{bu} are sound and strongly complete for their class of evidence models,
2. Ev_{bu} is sound and weakly complete for its class of finite evidence models.

Proof. Let Φ be λ -consistent and let Θ be a λ -type extending Φ . Then, let $\mathcal{M}_{\lambda}(\Theta)$ be its canonical λ -structure. By Lemma 15, $\mathcal{M}_{\lambda}(\Phi)$ is a λ -model, and by Proposition 1, given $\varphi \in \Theta$, $\Theta \in \llbracket \varphi \rrbracket_{\mathcal{M}_{\lambda}(\Theta)}$, so that Φ is satisfied by $\mathcal{M}_{\lambda}(\Theta)$. By Theorem 4, $\mathcal{M}_{\lambda}(\Theta)_{\lambda}^{\nabla}$ is a λ -evidence model also satisfying Θ if $\lambda \neq Ev_{bu}$.

The completeness of Ev_{bu} for its evidence models requires a filtration. Let φ be bu -consistent and Φ be a bu -type extending φ . Then, consider $\mathcal{M}_{bu}(\Phi)/\varphi$. By Lemma 18, $\mathcal{M}_{bu}(\Phi)/\varphi$ is finite, concise and by Lemma 16, satisfies φ . It then follows from Theorem 4 that $(\mathcal{M}_{bu}(\Phi)/\varphi)_{bu}^{\nabla}$ is a finite, concise evidence model which satisfies φ . \square

We know that Ev and Ev_u do not have the finite evidence model property, but the question of whether they have finite extended models remains open. Other important questions left open are: does Ev_b have the finite evidence model property? Is Ev_{bu} strongly complete for its class of evidence models? We conjecture affirmative answers to these questions, but leave them for future inquiry.

7. Dynamic 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 in a non-trivial way. An additional appealing feature of the systems studied here, following their original introduction in [8], is our fresh interpretation of neighborhood structures. Viewing neighborhoods as bodies of evidence suggested a number of new and interesting modalities beyond the usual repertoire studied in modal neighborhood logics. We will now show briefly how this process of language elicitation can be taken further when we take a further step, toward a dynamic perspective.

Evidence is not a static substance that we have once and for all. It is continually affected by new incoming information, and also by processes of internal re-evaluation. Taking a cue from recent dynamic logics of knowledge update [11,36] and belief revision [35,6], van Benthem and Pacuit [8] study a number of dynamic extensions of evidence models. They develop a rich repertoire of dynamic logics of “evidence management”. We will follow one such thread here, showing the new logical issues that arise.

The simplest type of information change is receiving information from an infallible source. This operation is called *public announcement* [26,14] and it transforms a model by removing all states where the announced formula is false. From the perspective of our evidence models, such a basic action in possible worlds models is actually a compound of various transformations. A public announcement of φ can be naturally “deconstructed” into a complex combination of three distinct operations:

1. **Evidence addition:** the agent accepts that φ is an “admissible” piece of evidence (perhaps on par with the other available evidence).
2. **Evidence removal:** the agent removes any evidence for $\neg\varphi$.
3. **Evidence modification:** the agent incorporates φ into each piece of evidence gathered so far, making φ the most important piece of evidence.

Our richer evidence models allows us to study each of these operations individually (cf. [8]). In this section, we focus on just one of these: accepting an input from a trusted source, leaving other dynamic operations for further study. Let \mathcal{M} be an evidence model and φ a new piece of evidence which the agent decides to *accept*. Here “acceptance” does not necessarily mean that the agent believes that φ is true, but rather that she agrees that φ should be considered when “weighing” her evidence. The formal definition of this action is straightforward:

Definition 19 (*Evidence addition*). Let $\mathcal{M} = \langle W, E, V \rangle$ be an evidence model, and φ a formula in \mathcal{L} . The model $\mathcal{M}^{+\varphi} = \langle W^{+\varphi}, E^{+\varphi}, V^{+\varphi} \rangle$ has $W^{+\varphi} = W$, $V^{+\varphi} = V$ and for all $w \in W$,

$$E^{+\varphi}(w) = E(w) \cup \{[\varphi]_{\mathcal{M}}\}.$$

This operation can be described explicitly in our modal language with a dynamic modality $[+\varphi]\psi$ which is intended to mean “ ψ is true after φ is accepted as an admissible piece of evidence”. The truth condition for this formula is straightforward:

$$(EA) \quad \mathcal{M}, w \models [+\varphi]\psi \text{ iff } \mathcal{M}, w \models E\varphi \text{ implies } \mathcal{M}^{+\varphi}, w \models \psi.$$

Here, since evidence sets cannot be empty, the precondition for evidence addition is that φ is true at some state. Compare this with the better-known precondition for public announcement which requires the accepted formula to be true. The logical analysis of even this simple operation is surprisingly subtle, suggesting a number of interesting extensions to the modal languages for evidence studied in this paper.

Let us start with laws of reasoning for the dynamic modalities, extending the logics that we had so far. For languages only containing the evidence and universal modality, we have the following new kind of validities, that may be viewed as “recursion laws” for evidence and knowledge modalities after an act of evidence addition:

$$\begin{aligned} [+\varphi]p &\leftrightarrow ((A)\varphi \rightarrow p) && (p \in \text{At}) \\ [+\varphi](\psi \wedge \chi) &\leftrightarrow ([+\varphi]\psi \wedge [+\varphi]\chi) \\ [+\varphi]\neg\psi &\leftrightarrow ((A)\varphi \rightarrow \neg[+\varphi]\psi) \\ [+\varphi][E]\psi &\leftrightarrow (E\varphi \rightarrow ([E][+\varphi]\psi \vee [A](\varphi \rightarrow [+\varphi]))) \\ [+\varphi]A\psi &\leftrightarrow ((A)\varphi \rightarrow [A][+\varphi]\psi). \end{aligned}$$


But things can quickly get more complicated. Finding a similar recursion law for the belief operator of this paper requires an extension to our basic language. This is the operator $B^\varphi\psi$ of *conditional belief*: “the agent believes ψ conditional on φ ”, which is well-known, but now with 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 can deal with this in our models:

Definition 20 (*Relativized scenarios*). 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} is a w -scenario relative to X when $\mathcal{X} \subseteq E(w)$ and \mathcal{X} has the **maximal X -f.i.p.**

To simplify notation, when X is the truth set of formula φ , we write “ \mathcal{X}^φ ” for “ $\mathcal{X}^{\llbracket\varphi\rrbracket_{\mathcal{M}}}$ ” and “ \mathcal{X} is a φ -scenario (for w)” for “ \mathcal{X} is a w -scenario relative to $\llbracket\varphi\rrbracket_{\mathcal{M}}$ ”. Now we define a natural notion of conditional belief:

- $\mathcal{M}, w \models [B]^\varphi\psi$ iff for each φ -scenario $\mathcal{X} \subseteq E(w)$, for each $v \in \bigcap \mathcal{X}^\varphi$, $\mathcal{M}, v \models \psi$.

Conditional belief is an entirely natural notion in probability and belief revision theory. Still, this notion suggests a logical investigation beyond what we have provided in this paper. 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, 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, $\mathcal{M}, w \models [B]q$; however, $\mathcal{M}, w \not\models [B]^p q \vee [B]^{-p} q$. This counter-example is interesting, as our putative implication is indeed valid on *connected* plausibility models for conditional belief (cf. [9] for a complete modal logic of conditional belief on such models). Now in principle, our earlier axiomatization captures conditional belief as well, since there is a definition for it in terms of safe belief that arises from relativizing the earlier definition of belief in terms of safe belief and the universal modality (cf. [7,10,36]). But things get more subtle when we consider the dynamics of evidence addition once more. Clearly, to keep the total system in harmony, we must provide a recursion law, not just for pure beliefs, but also for conditional beliefs. And then the earlier static base language requires extension once more.

The notion of conditional belief alone does not yet allow us to state a recursion axiom. To see this, 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 appropriate to this setting is $B^{\varphi, \alpha} \psi$ where “the agent believes that 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 $B^{\varphi, \top}$. We also write $[B]^{-\varphi}$ for $[B]^{\top, \neg \varphi}$. With this operator added, we can indeed state a recursion law for our operators of absolute and conditional belief (the proofs that these axioms are valid can be found in [8]):

$$\begin{aligned}
 [+ \varphi][B]\psi &\leftrightarrow (\langle A \rangle \varphi \rightarrow ([B]^{+\varphi} [+ \varphi]\psi \wedge [B]^{-\varphi} [+ \varphi]\psi)) \\
 [+ \varphi][B]^{\psi, \alpha} \chi &\leftrightarrow (\langle A \rangle \varphi \rightarrow ([B]^{\varphi \wedge [+ \varphi]\psi, [+ \varphi]\alpha} [+ \varphi]\chi \wedge [B]^{+\varphi, \psi, \neg \varphi \wedge [+ \varphi]\alpha} [+ \varphi]\chi)).
 \end{aligned}$$

Finally, recursion laws for evidence addition and our plausibility modalities also raise interesting new issues. Recall that a condition on our models (Definition 2) is that the evidence sets are upwards closed under the plausibility relation. This means that the precondition for adding a new piece of evidence to an arbitrary model is that it is upwards closed under the plausibility ordering. Alternatively, we can introduce a new evidence addition operation $E^{\oplus X}(w) = E(w) \cup \{X'\}$ where $X' = \{u \mid \text{there is a } v \in X \text{ such that } v \preceq u\}$. We leave a full analysis of this new operation for future work.

What we have shown so far is how a logical analysis of evidence dynamics that changes given neighborhood models for agents uncovers the need for new modalities in our base language of evidence models. Clearly, there are many new questions of axiomatization resulting from this, but other kinds of issues emerge as well. We conclude with just a simple illustration.

There is a natural issue of “harmony” between two perspectives on evidence dynamics. As we have seen, on evidence models $\mathcal{M} = \langle W, E, V \rangle$, we can *derive* the plausibility ordering directly from the evidence set (see Definition 6). Viewed in this way, an addition to the agent’s evidence also changes the agent’s plausibility ordering. But this means that we can also look at evidence change in terms of transformations of models of a more familiar relational kind, described in Section 2.2: namely, plausibility orderings of epistemic ranges. Belief change on models of the latter kind has been studied extensively: cf. [6,36,24], and this suggests a parallel line of inquiry finding dynamic operations on plausibility orderings matching changes in the agent’s evidence.

For evidence addition, the corresponding technical question is which operation makes the following diagram commute:

$$\begin{array}{ccc}
 \mathcal{E} & \xrightarrow{+X} & \mathcal{E}^+ \\
 \downarrow & & \downarrow \\
 \preceq_E & \xrightarrow{??} & \preceq^+
 \end{array}$$

And the answer is not hard to find. Suppose that $\preceq^+ = \preceq - \{(w, v) \mid w \in X \text{ and } v \notin X\}$. Then,

$$\preceq^+ = \preceq_{\mathcal{E}^+} .$$

This simple observation is in line with the study of “two-level models” for preference dynamics in [24], to which we refer for further relevant results. The extent of these harmony phenomena remains to be investigated, where we expect our earlier distinction between general and intended models, suitably generalized to the dynamic setting, to be important once more.

Our discussion of the dynamics of evidence change shows one further way in which neighborhood structures are a good vehicle for exploring finer epistemic and doxastic distinctions than those found in standard relational models for modal logic. 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 *p*-morphism matching the more expressive languages that emerged in the formulation of dynamic recursion laws, while proof-theoretically, we need to lift our earlier completeness technique to this richer setting.

8. Concluding remarks

Starting from earlier work on evidence logics in neighborhood models, this paper has offered two main new contributions. First, our completeness theorems contribute 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, even without any specific evidence interest in mind. For the modal logician, the pleasant surprise of our analysis is that there is a lot of well-motivated new modal structure to be explored on models that are often considered totally explored.

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 [30,17] and argumentation [16]. Closer to what we have presented, however, here are a few more specific 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 also give a spatial interpretation to our belief operator and the new matching modalities discussed in Section 7? This would suggest new, richer modal logics for reasoning about topological spaces (cf. [37]).
- *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?
- *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 (Theorem 4) be extended to deal with such richer languages?

- *Extended proof theory*: How can we axiomatize the richer modal logic of neighborhood models that arises in our dynamic analysis? Will the techniques of our completeness proofs carry over, or do we need a new style of analysis of canonical models?

Once again, we end with Sergei Artemov’s work on justification logics. As we have said in the Introduction, our evidence models provide a coarser level of analysis somewhere above the syntactic detail of proof terms. In conclusion, we briefly highlight two further comparisons. In an evidence model, the agent has evidence for φ provided there is a proposition (i.e., a set of states) that the agent has identified as evidence which logically implies φ . An important feature of justification logic is that a *proof term* not only records the fact that certain beliefs are grounded in evidence, but also the proof of why that evidence justifies a particular belief. A second point of comparison is that our notion of evidence is weaker than what is found in the justification logic literature. Having a justification for φ means, in particular, that the agent believes that φ is true. Contrast this with our setting in which an agent only believes what is implied by *all* of her evidence. Clearly, this only scratches the surface and there is much more to say about the connection with justification logic, but this will be left for future work.

References

- [1] H. Andreka, M. Ryan, P.Y. Schobbens, Operators and laws for combining preference relations, *J. Logic Comput.* 12 (2002) 13–53.
- [2] S. Artemov, Explicit provability and constructive semantics, *Bull. Symb. Log.* 7 (2001) 1–36.
- [3] S. Artemov, M. Fitting, Justification logic, in: E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, 2012, fall 2012 edition.
- [4] S. Artemov, E. Nogina, Introducing justification into epistemic logic, *J. Logic Comput.* 15 (2005) 1059–1073.
- [5] A. Baltag, B. Renne, S. Smets, The logic of justified belief change, soft evidence and defeasible knowledge, in: L. Ong, R. de Queiroz (Eds.), *Proceedings of the 19th Workshop on Logic, Language, Information, and Computation (WoLLIC 2012)*, in: *Lecture Notes in Comput. Sci.*, vol. 7456, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 168–190.
- [6] A. Baltag, S. Smets, Conditional doxastic models: A qualitative approach to dynamic belief revision, in: G. Mints, R. de Queiroz (Eds.), *Proceedings of WOLLIC 2006*, in: *Lecture Notes in Comput. Sci.*, vol. 165, 2006, pp. 5–21.
- [7] A. Baltag, S. Smets, Dynamic belief revision over multi-agent plausibility models, in: G. Bonanno, W. van der Hoek, M. Wooldridge (Eds.), *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision (LOFT 2006)*, pp. 11–24.
- [8] J. van Benthem, E. Pacuit, Dynamic logics of evidence-based beliefs, *Studia Logica* 99 (2011) 61–92.
- [9] O. Board, Dynamic interactive epistemology, *Games Econom. Behav.* 49 (2004) 49–80.
- [10] C. Boutilier, Conditional logics for default reasoning and belief revision, Ph.D. thesis, University of Toronto, 1992.
- [11] H. van Ditmarsch, W. van der Hoek, B. Kooi, *Dynamic Epistemic Logic*, Synthese Library, Springer, 2007.
- [12] M. Fitting, The logic of proofs, semantically, *Ann. Pure Appl. Logic* 132 (2005) 1–25.
- [13] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, 1988.
- [14] J. Gerbrandy, Bisimulations on Planet Kripke, Ph.D. thesis, Institute for Logic, Language and Computation (DS-1999-01), 1999.
- [15] P. Girard, Modal logic for belief and preference change, Ph.D. thesis, Dissertation Series DS-2008-04, ILLC University of Amsterdam, 2008.
- [16] D. Grossi, On the logic of argumentation theory, in: W. van der Hoek, G. Kaminka, Y. Lesperance, M. Luck, S. Sandip (Eds.), *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS’10)*, 2010, pp. 409–416.
- [17] J. Halpern, R. Pucella, A logic for reasoning about evidence, *J. Artificial Intelligence Res.* 26 (2006) 1–34.
- [18] H.H. Hansen, Monotonic modal logic, Master’s thesis, Universiteit van Amsterdam, 2003 (ILLC technical report: PP-2003-24).
- [19] H.H. Hansen, C. Kupke, E. Pacuit, Neighbourhood structures: Bisimilarity and basic model theory, *Log. Methods Comput. Sci.* 5 (2009) 1–38.
- [20] J.F. Horty, Reasons as defaults, *Philos. Impr.* 7 (2007) 1–28.
- [21] A. Kratzer, What must and can must and can mean, *Linguist. Philos.* 1 (1977) 337–355.
- [22] D. Lewis, *Counterfactuals*, Blackwell Publishers, Oxford, 1973.
- [23] C. List, F. Dietrich, Reasons for (prior) belief in bayesian epistemology, 2012, Unpublished manuscript.
- [24] F. Liu, *Reasoning about Preference Dynamics*, *Synth. Libr.*, vol. 354, Springer-Verlag, 2008.
- [25] E. Pacuit, Neighborhood semantics for modal logic: An introduction, 2007.
- [26] J. Plaza, Logics of public communications, in: M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, Z. Ras (Eds.), *Proceedings, 4th International Symposium on Methodologies for Intelligent Systems*, 1989, pp. 201–216 (republished as [27]).
- [27] J. Plaza, Logics of public communications, *Synthese* 158 (2007) 165–179.
- [28] H. Rott, *Change, Choice and Inference: A Study in Belief Revision and Nonmonotonic Reasoning*, Oxford University Press, 2001.

- [29] K. Segerberg, *An Essay in Classical Modal Logic*, Filosofiska Stuiet, Uppsala Universitet, 1971.
- [30] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [31] Y. Shoham, K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2009.
- [32] R. Stalnaker, Knowledge, belief and counterfactual reasoning in games, *Econ. Philos.* 12 (1996) 133–163.
- [33] R. Stalnaker, On logics of knowledge and belief, *Philos. Stud.* 128 (2006) 169–199.
- [34] K. Su, A. Sattar, G. Governatori, Q. Chen, A computationally grounded logic of knowledge, belief and certainty, in: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '05, 2005*, pp. 149–156.
- [35] J. van Benthem, Dynamic logic for belief revision, *J. Appl. Non-Classical Logics* 14 (2004) 129–155.
- [36] J. van Benthem, *Logical Dynamics of Information and Interaction*, Cambridge University Press, 2011.
- [37] J. van Benthem, G. Bezhanishvili, Modal logics of space, in: M. Aiello, I. Pratt-Hartmann, J. van Benthem (Eds.), *Handbook of Spatial Logics*, 2007, pp. 217–298.
- [38] J. van Benthem, D. Fernández-Duque, E. Pacuit, Evidence logic: A new look at neighborhood structures, in: *Advances in Modal Logic*, College Publications, 2012.
- [39] J. van Benthem, F.R. Velázquez-Quesada, The dynamics of awareness, *Synthese* 177 (2010) 5–27.
- [40] F. Veltman, Prejudices, presuppositions and the theory of conditionals, in: J. Groenendijk, M. Stokhof (Eds.), *Amsterdam Papers in Formal Grammar*, University of Amsterdam Press, 1976, pp. 248–281.