Hyperintensional Logics for Evidence, Knowledge and Belief

MSc Thesis *(Afstudeerscriptie)*

written by

Maximilian Siemers
(born August 1st, 1994 in Traunstein, Germany)

under the supervision of **Aybüke Özgün**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
August 23, 2021

Members of the Thesis Committee:
dr. Maria Aloni (Chair)
dr. Aybüke Özgün (Supervisor)
dr. Malvin Gattinger
dr. Anthia Solaki

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

Traditional epistemic logics generally define knowledge as truth in all epistemic alternatives. This approach has two shortcomings: first, the link between knowledge and justification is not represented. Second, epistemic agents are highly idealised and suffer from the defect of logical omniscience. In this thesis, we propose a framework that aims at bringing together knowledge, belief and evidence possession of less idealised epistemic agents.

Our contribution consists of a combination and enrichment of existing approaches that address each of these issues individually. On the one hand, there are possible-worlds based semantic representations of evidence and corresponding notions of knowledge and belief based on this conception of evidence. On the other hand, it has been argued that epistemic attitudes are topic-sensitive, which makes them hyperintensional.

Our framework combines the semantic representation of evidence with strategies to formalise hyperintensional knowledge and belief. We suggest that evidence itself should be understood as a hyperintensional concept. We demonstrate how some of the closure principles for knowledge and belief related to logical omniscience are consequences of the purely possible worlds-based approach of interpreting the underlying evidence. In particular, we suggest that the missing component that cannot be captured by these approaches is the relation that holds between a piece of evidence and a proposition whenever the former is relevant for the latter.

Based on existing frameworks modelling subject matters, we develop a topic-sensitive notion of evidence and show that knowledge and belief can be defined based on this novel kind of evidence. As a consequence of grounding knowledge and belief entirely in hyperintensional evidence, our target notions of knowledge and belief are themselves hyperintensional. In particular, our approach circumvents many of the defects pertaining to logical omniscience.

Our main technical contribution consists of a sound and complete axiomatisation of a logic for hyperintensional evidence and knowledge, which is expressive enough to define all evidence, knowledge and belief modalities developed throughout this thesis. We also provide a separate sound and complete axiomatisation of a belief fragment.

Acknowledgements

First of all, I would like to thank my supervisor, Aybüke Özgün. Aybüke, without you this thesis could have never become what it is. Your support, encouragement, and inspiration continuously exceeded my hopes and expectations. Thanks to you, this thesis was as much fun as it was demanding. I cannot imagine how one could supervise better than you did.

A very special thanks goes to Thomas Randriamahazaka. Thomas, this thesis was so much more fun because you were there for me with your unbelievable brain in the moments of greatest logical despair. Your love for logic has infected me, and provided a guiding star throughout the Master of Logic. It is amazing to be friends with you.

I would like to thank the members of the defense committee: Maria Aloni, Malvin Gattinger, and Anthia Solaki, for their time to read this thesis and the pleasant and stimulating defense.

Being a student in the Master of Logic with a background in psychology was as inspiring as it was challenging. I could only write this thesis because the programme gave me the chance to acquire the necessary philosophical and mathematical background. I am extremely thankful for the opportunities that the programme provides to students of all backgrounds to dive into new areas and to intellectually re-invent themselves. I would like to thank the student mentors, my academic mentor, Ex Falso, and everyone who contributed to the ILLC community for their countless efforts to create this wonderful supportive and inspiring atmosphere, which makes the ILLC such a unique and amazing place to study and conduct research.

Personally, I would like to thank my friends: Johannes, it feels like we are taking this journey through both life and academia together. Thank you for always being there to talk. Heiko, thank you for your friendship and wisdom. I cannot imagine being where I am without you in my life. Rachel, thank you for all the beautiful moments that we've had together and that are yet to come. Anna, Fergus, Grace, Iason, Ion, Martin, Oana, Pepijn, Sam, Terence and Tex: thank you for your friendship. You made Amsterdam feel like home to me.

Last but not least, I would like to thank my parents. Without your never-ending support and belief in me, none of this would ever have been possible.

Contents

Abstract **i**

Acknowledgements **ii**

Contents **iii**

1 Introduction **1**
1.1 Evidence à la van Benthem & Pacuit 2
1.2 Hyperintensional evidence 4
1.3 Fragmentation 5
1.4 Knowledge and belief 7
1.5 Brief comparison with existing approaches 8
1.6 In this thesis 10

2 A hyperintensional logic of evidence, knowledge and belief **11**
2.1 Topic-sensitive evidence 12
2.2 Fragmented evidence 15
2.3 Semantics 16
2.4 Hyperintensionality revisited 19
2.5 Fragmentation revisited 26
2.6 Logical omniscience 27
2.7 Overview 30

3 Technical results: soundness and completeness **32**
3.1 The factive evidence-knowledge fragment $\mathcal{L}_{[\forall]\Box_k K_k}$ 32
3.2 The belief fragment $\mathcal{L}_{[\forall]B_k}$ 51

4 Comparison with other approaches **58**
4.1 Topological evidence models 58
4.2 Question-based approaches 60
4.3 Awareness logics 62
4.4 Justification logics 63

5 Conclusions and further research **65**
5.1 Conjunctive parts 66
5.2 Computationally bounded agents 68

References **68**

Chapter 1

Introduction

This thesis is concerned with the formalisation of certain notions of evidence, knowledge and belief. Since Hintikka (1962), it has become the standard approach in epistemic logic to interpret knowledge as truth in all epistemic alternatives. This formalisation of knowledge suffers from two limitations: first, it is not connected with any kind of *evidence* or *justification*. Second, it generally idealises an agent's capacity to derive logical consequences of what they already know; this phenomenon is known as the problem of *logical omniscience*.

Each of these issues have been separately addressed in a number of works. However, with a few exceptions mentioned in Chapter 4, there are currently no frameworks addressing both problems combined. The goal of this thesis is to provide a unifying framework that grounds knowledge and belief in evidence in a way that gives rise to less idealised epistemic agents.

The evidence we are interested in is the kind of evidence that realistic agents receive in normal, every day life. In particular, evidence does not necessarily specify the complete state of the world, but may leave open several ways of how things could be. Furthermore, we take evidence to be possibly false or misleading. To give a semantic representation of this kind of evidence, we capitalise on existing approaches and interpret evidence as sets of possible worlds, to which an additional component is added: a *subject matter* or *topic*, which represents what the evidence is *about*, or the context in which the evidence is received and processed by the agent.

We also take inspiration from the existing literature in our treatment of the problem of logical omniscience. In particular, we treat knowledge and belief as *hyperintensional* concepts: they sometimes treat intensionally equivalent contents in different ways. Thus, possible worlds are not sufficient to give accurate and realistic formalisations of knowledge and belief. Instead, we take them to be attitudes towards propositions whose content consists of both an *intensional* (i.e., possible worlds-based) component and a *subject matter* or *topic* component in the tradition of Yablo (2014) (see also Hawke, 2018).

The conceptual novelty in our approach is that we locate the source of the hyperintensionality of both knowledge and belief at the level of evidence. Instead of representing evidence purely as sets of possible worlds, we equip evidence with topics. Based on this semantic representation, we offer a definition of what it means for a piece of evidence to be *relevant* for a proposition. The relation of evidential relevance is hyperintensional, and we argue that knowledge and belief inherit their hyperintensionality from the underlying evidence.

The resulting notions of knowledge and belief already block some of the problematic closure principles pertaining to logical omniscience: for example, not every tautology is known *because* the agent does not necessarily possess relevant evidence for every tautology. However, we shall see that this approach does not suffice to circumvent all problematic closure principles. For example, hyperintensional knowledge is still closed under taking conjunctions. To improve on this situation, we take the concept of *fragmentation* on board, according to which the total body of evidence possessed by an agent may be scattered across different fragments. These can be thought of as different contexts in which evidence is processed, as different sources from which the evidence is received, or as distinct questions whose answers the agent tries to obtain using her evidence.

In the remainder of the introduction, we briefly introduce the core notions upon which our framework is based: the *semantic representation of evidence*, the *hyperintensional relation of evidential relevance*, and *fragmentation*. We also give a brief sketch of the problem of *logical omniscience*, indicate the most important existing frameworks we took as inspiration, and finally give an outline of the structure of this thesis.

1.1 Evidence à la van Benthem & Pacuit

In this section, we present the *evidence models* originally presented in van Benthem and Pacuit (2011) and developed further in van Benthem, Fernández-Duque, and Pacuit (2012) and van Benthem, Fernández-Duque, and Pacuit (2014). These models laid the cornerstone of a whole research agenda centring around the semantic representation of evidence as sets of possible worlds, and they also provide the starting point of our journey. Taking evidence as unstructured propositions, that is, as sets of worlds rather than formulas in a fixed language allows us to remain flexible with respect to the different forms that evidence can take. In particular, it enables us to represent non-linguistic evidence, such as some kinds of sensory percepts or memories.

Definition 1.1 (Evidence Models (van Benthem and Pacuit 2011)). An *evidence model* is a tuple $\mathfrak{M} = (W, \mathcal{E}_0, V)$, where

i) W is a non-empty set of *possible worlds*;

ii) $\mathcal{E}_0 \subseteq \mathcal{P}(W)$ is a set of sets of worlds called *basic pieces of evidence*, satisfying $W \in \mathcal{E}_0$ and $\varnothing \notin \mathcal{E}_0$; and

iii) $V : \text{PROP} \to \mathcal{P}(W)$ is a valuation map.

The evidence models originally defined in van Benthem and Pacuit (2011) are more general: they allow the collection of evidence $\mathcal{E}_0$ to vary between different worlds. Here, however, we restrict attention to the "uniform" evidence models that assign the same evidence sets to all worlds.

Given an evidence model $\mathfrak{M} = (W, \mathcal{E}_0, V)$, a basic piece of evidence $e \in \mathcal{E}_0$ and a formula[1] φ, we say that e *entails* or *supports* φ if φ is true throughout all worlds contained in e.

If we pick out a world $w \in W$ as the *actual world*, we are in the position to distinguish between *truthful* or *factive* evidence and *false* evidence: if a basic piece of evidence $e \in \mathcal{E}_0$ contains w, then e is truthful; otherwise, it is false.

Given two pieces of basic evidence $e, f \in \mathcal{E}_0$, we say that e and f are *consistent* with each other if they have a non-empty intersection. In this case, there exists a world which is compatible with both e and f.

Evidence combination

Two basic pieces of evidence e and f that are consistent with each other may be *combined* by taking their intersection. The resulting piece of evidence $e \cap f$ is a stronger piece of evidence in the sense that it supports at least as many propositions as e and f individually.

From now on, we will by default assume that our agent automatically combines all mutually consistent available evidence. Formally speaking, this amounts to assuming that the set of evidence is closed under non-empty finite intersections. To be clear about this difference to the original models from van Benthem and Pacuit (2011), we will denote the collection of evidence closed under non-empty finite intersections by $\mathcal{E}$ instead of $\mathcal{E}_0$. Clearly, the assumption that an agent always forms all possible finite combinations of available evidence is an idealisation. Later, we shall introduce the concept of *fragmentation* to weaken this idealisation.

Based on the van Benthem-Pacuit evidence models, we can define two modalities E and $\Box$ for evidence possession. These have the following intended readings:

[1] Of some propositional language; we will introduce our own full language of evidence, knowledge and belief in Chapter 2.

$E\varphi$	The agent has (combined) evidence for φ.
$\Box\varphi$	The agent has truthful (combined) evidence for φ.

Evidence operators

The following semantics for evidence operators were developed in van Benthem and Pacuit (2011) and Baltag et al. (2016). Given an evidence model $\mathfrak{M} = (W, \mathcal{E}, V)$ (note that we assume $\mathcal{E}$ to be closed under non-empty finite evidence combination) and a world $w \in W$, the semantics for operators E and $\Box$ for (factive) evidence possession is defined as follows:

$$\begin{array}{lll} \mathfrak{M}, w \Vdash E\varphi & \text{iff} & \exists e \in \mathcal{E}\ (e \subseteq [\![\varphi]\!]); \\ \mathfrak{M}, w \Vdash \Box\varphi & \text{iff} & \exists e \in \mathcal{E}\ (w \in e \subseteq [\![\varphi]\!]). \end{array}$$

In words, an agent *has evidence for* φ iff there exists a (combined) piece of evidence entailing φ; and an agent *has factive evidence for* φ iff there exists a (combined) *truthful* piece of evidence entailing φ.

These definitions formalise the relation e *is evidence for* φ in terms of purely intensional entailment. We now turn to an example of a counter-intuitive consequence not only of the formalisation of van Benthem and Pacuit (2011), but of *any* purely intensional representation of evidence.

1.2 Hyperintensional evidence

Consider the following two examples from Özgün and Berto (2020):

Example 1.2.

(1) All bachelors are unmarried.

(2) No three positive integers x, y and z satisfy $x^n + y^n = z^n$ for any integer value of n greater than 2.

Both of these sentences express necessary truths. As such, they are true at every world in every evidence model and are therefore entailed by every piece of evidence. In general, the following holds on intensional evidence models:

Every piece of evidence constitutes evidence for every necessary truth.

This seems counter-intuitive: a radio announcement of Joe Biden having won the state of Pennsylvania might provide evidence for the outcome of the elections, but has *nothing to do* with the marital status of bachelors or Fermat's Last Theorem. Intuitively, then, what a piece of evidence e is *about* should be related to what P is about if e constitutes evidence for P. In a slogan:

(Evidential relevance)	*The content of any evidence for a proposition P should be* relevant *for P.*

As the previous example suggests, the relevant notion of *content* cannot be captured solely by the intensional content of evidence and propositions. That is, our target notion of evidence content is *hyperintensional*: we need something else besides possible worlds to capture evidential relevance.

Our proposal for a solution is presented in the next chapter, where we also see some more examples of intuitively problematic principles that are validated by intensional accounts of evidence, but should be rejected on grounds of *evidential relevance*.

1.3 Fragmentation

For the van Benthem-Pacuit evidence models, we already introduced the notion of *evidence combination*: given two consistent pieces of evidence, an agent may take their intersection to obtain a new piece of evidence, which may support more propositions. Arguably, it is an idealisation to assume that an agent always combines *all* available evidence. Consider the following modified example, originally from Lewis (1982, p. 436), as cited in Yalcin (2020, p. 11):

Example 1.3. "I used to think that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south; and that the two were roughly parallel."

Let's add the following to Lewis' story: his belief concerning the orientation of Nassau Street was based on his frequent shopping tours through the city centre after the end of his lectures, when the evening sun shone at the street from the west. His belief concerning the orientation of the railroad stemmed from a lake that he could see when he took the train to New York and of which he knew that it was north of the city. Finally, his belief that the street and the railroad were parallel was based on his memory of passing a train in his car while driving out of the city on Nassau Street. We return to Lewis' story:

"So each sentence in an inconsistent triple was true according to my beliefs, but not everything was true according to my beliefs. Now, what about the blatantly inconsistent conjunction of the three sentences? I say that it was not true according to my beliefs. My system of beliefs was broken into (overlapping) fragments. Different fragments came into action in different situations, and the whole system of beliefs never manifested itself all at once. [...] The inconsistent conjunction of all three did not belong to, was in no way implied by, and was not true according to, any one fragment. That is why it was not true according to my system of beliefs taken as a whole. Once the fragmentation was healed, straightway my beliefs changed: now I think that Nassau Street and the railroad both run roughly northeast-southwest."

We offer the following addition to Lewis' explanation: the different fragments into which his belief system was broken correspond to fragments across which his evidence was distributed. One fragment might have included his evidence from his shopping tours that justified his belief concerning the street, another fragment included the evidence from his train trips that justified his belief concerning the railroad, and a third fragment included the evidence supporting the belief that the street and the railroad ran in parallel.

Why would an agent's body of evidence be broken into fragments? An intuitive answer is based on the *cognitive limitations* of agents that operate in finite time and space: we constantly receive a large amount of information from different sources, and only some of that information is relevant to us. Realistically, we cannot always combine all of these pieces of information and derive all their logical consequences. Indeed, Yalcin (2020) argued that it may be *rational* for epistemic agents to have fragmented beliefs.

A second explanation focuses on *conceptual understanding*: an agent may possess two pieces of evidence and think that they have nothing to do with each other; she associates them with different concepts. Especially given the resource constraints mentioned in the previous paragraph, it makes sense to assume that an agent only combines pieces of evidence that from her perspective belong to similar concepts.

A third way of looking at fragmentation is centred around *questions*. On this view, evidence is used to answer questions raised by the agent. The different questions that are relevant to an agent at the same time can be grouped into different *agendas* or *inquiries*: overarching big questions with several associated smaller questions. When two pieces of evidence seem to provide answers to entirely unrelated questions, or questions belonging to different inquiries, an agent may not see the need to combine them with each other.

We have seen three different motivations for fragmentation: First, an agent's mind can be seen as fragmented due to her *failure to combine information from different contexts with each other*; on this interpretation, the different fragments are like "logical echo chambers" between which no information permeates. Second, an agent may *understand different concepts without being able to combine them to larger concepts*. Third, an agent may *pursue inquiries related to different questions*, or aiming at different goals. On the last interpretation, it is natural to assume that for each of her goals, she consults different information sources, and again may fail to combine information from different sources.

In this thesis, we will take inspiration from all of these perspectives and take the idea of fragmentation on-board. Given our evidence-first perspective, we will split our agent's total body of evidence into different *frames*, where each frame is associated with an overarching topic, which intuitively represents the connection between all evidence in that frame.

This gives as a way of relaxing the idealised assumption that an agent always combines all available evidence:

(Evidential fragmentation) *An agent's body of evidence is fragmented. Only evidence belonging to the same fragment can be combined.*

1.4 Knowledge and belief

Knowledge and belief grounded in evidence

The seminal work of Hintikka (1962) laid the groundwork for modern epistemic logic: knowledge is interpreted as truth in all epistemic alternatives as a *necessity operator* of modal logic. Formally, this means that a formula φ is known at a world w if φ is true at all worlds that are considered epistemically possible at w.

This formalisation of knowledge trails behind the developments in epistemology. Already in Plato's *Theaetetus* (Ichikawa and Steup 2018) we can find an argument that true belief is not sufficient for knowledge: the presence of some kind of *justification* should be added as a necessary condition for knowledge. In the famous counterexamples by Gettier (1963), Plato's conception of knowledge as *justified true belief* was rebutted. In the subsequent epistemological literature, a wide variety of proposals was developed. One example is the so-called *Defeasibility Theory of Knowledge* of Lehrer and Paxson (1969) and Lehrer (2018), according to which knowledge is true evidence-based belief which is stable under acquisition of certain kinds of new evidence.

What all these accounts have in common is that they take justifications as required for knowledge. The first of our two pivotal objectives in developing formalisations of knowledge and belief is to incorporate this epistemological tradition by grounding knowledge and belief in evidence. On the technical side, our approach takes inspiration from Baltag et al. (2016), who interpret knowledge and belief on topological models based on the evidence models by van Benthem and Pacuit (2011). We give our proposal for a definition of knowledge and belief in the next chapter.

Besides grounding knowledge and belief in evidence, our second objective is to ensure that our definitions of knowledge and belief do not give rise to overly idealises epistemic agents. To be clear about what kinds of idealisation we mean here, we now briefly specify the problem of *logical omniscience*.

Logical omniscience

Both the Hintikkan and the intensional evidence-based approaches to formalising knowledge and belief have in common that the epistemic attitudes of their agents are closed under certain types of logical inference. The following core principles are usually associated with logical omniscience (our list is based on Fagin et al. (1995) and Solaki (2017)):

- *Knowledge of validities:* If φ is valid, then φ is known.
- *Closure under logical entailment:* If φ logically entails ψ and φ is known, then ψ is known.
- *Closure under logical equivalence:* If φ and ψ are logically equivalent and φ is known, then ψ is known.
- *Closure under known implication:* If both φ and $\varphi \to \psi$ are known, then ψ is known.
- *Closure under conjunction:* If both φ and ψ are known, then $\varphi \land \psi$ is known.
- *Closure under disjunction:* If φ is known, then $\varphi \lor \psi$ is known.

Phrased in terms of evidence, the principles *knowledge of validities* and *closure under conjunction* are already familiar to us: In Example 1.2, we saw that every piece of evidence supports every necessary truth; and in Example 1.3, we discussed the idealised assumption that an agent always finitely combines all available evidence.

1.5 Brief comparison with existing approaches

Our goal is to develop a hyperintensional, fragmented notion of evidence and define knowledge and belief in such a way that they are (a) grounded in evidence and (b) give rise to less idealised epistemic agents. Our framework, which we will present in the next chapter, capitalises on the concepts developed in three lines of research:

The first line is concerned with the representation of evidence in the tradition of van Benthem and Pacuit (2011). We take particular inspiration from the works of Özgün (2017) and Baltag et al. (2016), which exploit the topological structure implicit in the van Benthem-Pacuit evidence models to develop topological, evidence-based notions of knowledge and belief.

Solutions to the problem of logical omniscience have been developed in, among others, Hawke, Özgün, and Berto (2020), Özgün and Berto (2020), and Berto and Hawke (2018), who argue that knowledge and belief should be interpreted as hyperintensional and specifically topic-sensitive notions. However, these proposals come without any notion of evidence.

The concept of fragmentation was already developed by Fagin and Halpern (1987), who expressed the idea that an epistemic agent does not feature one single mind which combines information in a coherent, centralised manner, but rather a *society of minds*, whose members form their own beliefs. Yalcin (2018) puts the idea this way: Belief is not one map by which we steer, but rather an *atlas*, or a collection of maps. Depending on the situation we're in, we choose the appropriate map from our atlas to navigate through the world. Finally, Hawke, Özgün, and Berto (2020) point to fragmentation as one of the central strategies to tackle the problem of logical omniscience. All of these approaches, however, locate fragmentation at the level of knowledge and belief, and do not include evidence.

Our contribution is a combination of these three lines of research. In particular, we argue that the notion of evidence should itself be understood as hyperintensional and fragmented, and extend the approach of van Benthem and Pacuit (2011) by equipping evidence with topics as they were developed for knowledge and belief in Hawke, Özgün, and Berto (2020) and Özgün and Berto (2020). We also adopt the notion of fragmentation for bodies of evidence. Building on this novel concept of topic-sensitive and fragmented evidence, we define notions of knowledge and belief that are entirely grounded in evidence using an approach resembling the topology-based formalisation in Özgün (2017) and Baltag et al. (2016). Thus, we present a unification of these authors' works which unravels the hyperintensional conception of knowledge and belief by including a layer of hyperintensional and fragmented evidence, which gives rise to hyperintensional and fragmented knowledge and belief.

Both hyperintensionality and fragmentation are needed Hawke, Özgün, and Berto (2020) argue that logical omniscience arises from at least two characteristics of classical epistemic logics: First, the definition of epistemic attitudes is entirely based on intensional content. That is, these approaches fail to capture the *hyperintensional* character of knowledge and belief. Second, these approaches assume that an agent always combines all available information. Therefore, *both* hyperintensionality and fragmentation are required to fully address the problem of logical omniscience. Our target notions of knowledge and belief will thus be both hyperintensional and fragmented.

1.6 In this thesis

In Chapter 2, we motivate and present our proposal for a logic of hyperintensional and fragmented evidence, knowledge and belief. We first introduce our language and provide the intended readings of our modalities. Following that, we develop our semantic representations for topic-sensitivity and fragmentation. After stating the full definitions of our models and evidence, knowledge and belief modalities, we demonstrate in examples how our formalisation of evidence circumvents several intuitively problematic principles that are validated by intensional conceptions of evidence. We extend this discussion to knowledge and belief by showing that they do not exhibit many of the principles associated with logical omniscience. The chapter is concluded with an overview of epistemologically relevant (in-)validities for all epistemic modalities introduced in this thesis.

Chapter 3 contains the main technical contributions of this thesis. We provide sound and complete axiomatisations of two fragments of our full logic of hyperintensional and fragmented evidence, knowledge and belief. The first fragment is the *factive evidence-knowledge fragment* and is expressive enough to define all epistemic modalities introduced in the thesis. The second fragment is the *belief fragment*, for which we give a sound and complete axiomatisation in order to provide a separate logical specification of our notion of belief.

In Chapter 4, we compare our framework with similar existing approaches. First, we state the difference between our notion of belief and its topological counterpart of Baltag et al. (2016). We also explain why our notion of belief does not exploit the full topological structure of evidence. After, we give a brief presentation of a question-based conception of subject matters, which provides an alternative to the more abstract notion of topics used by us. We sketch what a question-sensitive notion of evidence could look like and point to limitations in such an approach. Finally, we compare our framework to awareness logics and justification logics, which both share some of the objectives pursued here.

Finally, in Chapter 5, we summarise what we hope to have contributed with this thesis, and conclude with two possible directions for further research.

Chapter 2

A hyperintensional logic of evidence, knowledge and belief

In this chapter, we motivate and present the semantic machinery to capture hyperintensional and fragmented evidence, knowledge and belief. We begin by introducing the language and modalities we will work with. In Sections 2.1 and 2.2, we define topic-sensitivity and fragmentation for evidence. In Section 2.3, we give the full semantics of our logic, including the definition of our models and the semantic clauses for all evidence, knowledge and belief modalities. In Sections 2.4 and 2.5, we demonstrate how several problematic properties of intensional evidence models are not exhibited by our notion of evidence. In Section 2.6, we show how our formalisation of knowledge and belief inherit the hyperintensional and fragmented properties of the underlying evidence and in particular exhibit reduced logical omniscience. Finally, Section 2.7 provides a concise overview of relevant validities and non-validities for all our modalities.

Language of evidence, knowledge and belief

We first introduce the language we will work with. Let PROP be a countable set of *propositional variables*, and let $\mathcal{F} = \{1, ..., n\}$ be a non-empty, finite set of *frame symbols.* These frame symbols correspond to the fragments of an agent's evidence. The full language $\mathcal{L}$ of hyperintensional fragmented evidence, knowledge and belief is defined recursively by the grammar

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\forall]\varphi \mid E\varphi \mid \Box\varphi \mid B\varphi \mid K\varphi \mid \\ E_k\varphi \mid \Box_k\varphi \mid B_k\varphi \mid K_k\varphi \quad \text{for } k \in \mathcal{F}.$$

The remaining Boolean connectives $\vee$, $\rightarrow$, and $\leftrightarrow$ are defined as abbreviations in the usual way: $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$; $\varphi \rightarrow \psi := \neg\varphi \vee \psi$; and $\varphi \leftrightarrow \psi := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. For any formula $\varphi \in \mathcal{L}$, let $Var(\varphi)$ denote the set of propositional variables occurring in φ.

The intended meaning of the modalities for evidence, belief and knowledge is given in Table 2.1.

$E_k\varphi$	The agent has evidence for φ in the kth evidential frame.
$E\varphi$	The agent has evidence for φ (in some evidential frame).
$\Box_k\varphi$	The agent has *factive* evidence for φ in the kth evidential frame.
$\Box\varphi$	The agent has *factive* evidence for φ (in some evidential frame).
$B_k\varphi$	The agent believes φ in the kth evidential frame.
$B\varphi$	The agent believes φ (in some evidential frame).
$K_k\varphi$	The agent knows φ in the kth evidential frame.
$K\varphi$	The agent knows φ (in some evidential frame).

Table 2.1: Intended readings of epistemic modalities

$[\forall]$ denotes the *global modality*: $[\forall]\varphi$ says that φ is true in *all possible worlds*. This operator will serve mostly as a technical aid. Epistemically, we may interpret it as an *a priori* modality.

2.1 Topic-sensitive evidence

Our target notion of evidence should satisfy the following desideratum of *evidential relevance* from Section 1.2, which we repeat here:

(Evidential relevance) *The content of any evidence for a proposition P should be* relevant *for P.*

This requirement is what we identified as one of the missing components we target in the purely intensional account of evidence. In this section, we first introduce the necessary technical machinery and then provide a formal definition of the notion of *evidential relevance*. In Example 1.2, we saw that a piece of evidence concerning the outcome of the elections *should not* constitute evidence for Fermat's Last Theorem because it is not relevant for it; however, if the intensional content of a piece of evidence e and a proposition P is the only component in determining whether e is evidence for P, then this relevance condition cannot be captured.

Thus, we need to enrich the content of both propositions and of evidence to be able to capture evidential relevance. Intuitively, a radio announcement concerning the outcome of an election should be dismissed as irrelevant for the truth of Fermat's Last Theorem because it is about an entirely different subject matter. Before giving a formal definition of the notion of a topic, we turn to an example to get a better grip of the nature of subject matters.

Example 2.1.

(i) P: “Pia is a professor.”

(ii) Q: “Quita is a queen.”

What are the topics of these propositions, i.e. what are these propositions *about*? The following topic ascriptions have intuitive force:

(a) P is about both *Pia* and *Pia's profession*.

(b) Q is about both *Quita* and *Quita's profession*.

Aboutness relations exhibit some intuitive structure. For example, anything that is about *Pia's profession* is also about *Pia*, but not vice versa. These interrelations can be captured by a mereological structure of subject matters: for example, the topic *Pia's profession* is part of the larger topic *Pia*. The topic of $P \wedge Q$ is *Pia's and Quita's professions*, which is the fusion of the respective topics of P and Q. The mereological structure exhibited by topics is captured by the following definition:

Definition 2.2 (Mereology of Topics (Berto 2019))**.** A *mereology of topics* is a tuple $\mathfrak{T} = (\mathcal{T}, \oplus)$, where

i) $\mathcal{T}$ is a non-empty set of *topics*; and

ii) $\oplus : \mathcal{T} \times \mathcal{T} \to \mathcal{T}$ is a binary operation on topics called *topic fusion* assumed to be *idempotent*, *commutative* and *associative*.

The fusion $a \oplus b$ of two topics a and b is the smallest topic that both a and b are part of. Based on topic fusion $\oplus$, we define a binary relation $\sqsubseteq$ called *topic parthood* on $\mathcal{T}$ as follows:

$$\text{for all } a, b \in \mathcal{T}: \quad a \sqsubseteq b \text{ iff } a \oplus b = b.$$

Then $(\mathcal{T}, \oplus)$ is a *join semilattice* and $(\mathcal{T}, \sqsubseteq)$ a *poset*.

We are now equipped with a precise notion of topics, but we still need a way of assigning topics to sentences. We map propositional variables to topics and make sure that the topic assigned to complex sentences is systematically built up from the topics of its propositional components. We assume *topic-transparency of logical connectives*: logical connectives do not modify the topic of a sentence. For example, the topic of *Pia is a professor* is the same as the topic of *Pia is* not *a professor*. The topic of *Pia is a professor and Quita is a queen* is the fusion of the topics of *Pia is a professor* and of *Quita is a queen*.

Definition 2.3 (Topic assignment function)**.** Given a mereology of topics $\mathfrak{T} = (\mathcal{T}, \oplus)$, a *topic assignment function* $t : \textsc{prop} \to \mathcal{T}$ is a function that assigns topics from $\mathcal{T}$ to propositional variables in $\textsc{prop}$.

t is extended to the whole language $\mathcal{L}$ by putting $t(\varphi) = t(p_1) \oplus \cdots \oplus t(p_n)$, where $\{p_1, ..., p_n\} = Var(\varphi)$ are the propositional variables occurring in φ. We say that φ *is about* $t(\varphi)$.

Returning to the previous example, the underlying mereology of topics is the following[1]:

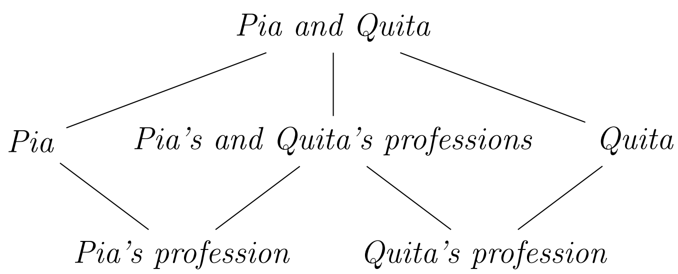

Figure 2.1: Topic structure of Example 2.1. Edges represent topic parthood $\sqsubseteq$: larger topics are further to the top

We are now in a position to formally capture aboutness relations of sentences. For example, $t(P) =$ *Pia's profession*, so P is about *Pia's profession.* Since $t(P) \sqsubseteq$ *Pia*, we have that P is also about *Pia.* And since

$$\begin{aligned} t(P \wedge Q) &= t(P) \oplus t(Q) \\ &= \textit{Pia's profession} \oplus \textit{Quita's profession} \\ &= \textit{Pia's and Quita's professions} \\ &\sqsubseteq \textit{Pia and Quita}, \end{aligned}$$

we have that $P \wedge Q$ is about both *Pia's and Quita's professions* and *Pia and Quita.*

Assigning topics to evidence We now have a topic assignment for sentences, but we still need to equip pieces of evidence with topics. The basic idea is simple: instead of representing a piece of evidence *only* as a set of worlds, we take it as a pair (e, a) consisting of a set of worlds e and a topic a. However, we also need to accommodate evidence combination: given two topic-equipped pieces of evidence (e, a) and (f, b) such that e and f are consistent (i.e., $e \cap f \neq \varnothing$), we need to make sure that the combined evidence $e \cap f$ receives an appropriate topic. Intuitively, if we combine a piece

[1]Obviously, there are more topics that could be depicted here, for example *Pia's hobbies*, *Quita's family*, or generally anything else related to *Pia* or *Quita.* The topic structure in Figure 2.1 is to be understood as a fragment of the structure of *all* topics.

of evidence concerning *Pia's profession* with another concerning *Quita's profession*, then the resulting piece of evidence concerns *Pia's and Quita's profession.* Thus, the combination of (e, a) and (f, b) is $(e \cap f, a \oplus b)$, i.e. a combined piece of evidence with the intersection of e and f as its intensional component, and the fusion of a and b as its topic component.

The next definition summarises the assignment of topics to evidence and completes our toolbox for representing hyperintensional evidence.

Definition 2.4 (Topic-indexed piece of evidence)**.** Given an evidence model $\mathfrak{M} = (W, \mathcal{E}, V)$ and a mereology of topics $\mathfrak{T} = (\mathcal{T}, \oplus)$, a *topic-indexed piece of evidence* is an element (e, a) of the set $\mathcal{E} \times \mathcal{T}$.

Two topic-indexed pieces of basic evidence (e, a) and (f, b) are called *consistent* if $e \cap f \neq \varnothing$. For any two consistent topic-indexed pieces of evidence (e, a) and (f, b), their *evidence combination* is defined as $(e \cap f, a \oplus b)$. If (e, a) is a topic-indexed (combined) piece of evidence, we say that (e, a) is *about* a and often simply write e_a instead of (e, a).

Evidential relevance Provided these definitions of the topic component of both sentences and of evidence, we are now in a position to formally define what it means for a piece of evidence to be relevant for a sentence:

Definition 2.5 (Evidential relevance)**.** If $\mathfrak{M} = (W, \mathcal{E}, V)$ is an evidence model, $\mathfrak{T} = (\mathcal{T}, \oplus)$ a mereology of topics, $e_a \in \mathcal{E} \times \mathcal{T}$ a topic-indexed piece of evidence and $\varphi \in \mathcal{L}$ a formula, we say that e_a is *relevant* for φ iff $t(\varphi) \sqsubseteq a$, i.e. iff the topic of φ is included in the topic a of e_a.

2.2 Fragmented evidence

We now turn to the second conceptual component of our framework: *fragmentation.* We recall our desideratum from Section 1.3:

(Evidential fragmentation) *An agent's body of evidence is fragmented. Only evidence belonging to the same fragment can be combined.*

We interpret this idea in a straightforward manner. We assume a finite number of *evidence fragments*, each of which is represented by a number from the set $\mathcal{F} = \{1, ..., n\}$ of *frame symbols.* Interpreted on a model, each frame symbol will be associated with its own body of evidence, i.e. a set of topic-indexed pieces of evidence $\mathcal{E}_k$ called *evidential frame.* Each evidential frame $\mathcal{E}_k$ is assumed to be closed under non-empty, finite evidence combination, i.e. if $e_a, f_b \in \mathcal{E}_k$ and $e \cap f \neq \varnothing$, then $(e \cap f, a \oplus b) \in \mathcal{E}_k$.

Each frame symbol $k \in \mathcal{F}$ is assigned a topic by extending the definition of the topic assignment function t to frame symbols:

$$t : \text{PROP} \cup \mathcal{F} \to \mathcal{T}$$

t is extended to the whole language $\mathcal{L}$ as before (see Definition 2.3).

We place the following condition on each evidential frame $\mathcal{E}_k$:

$$\text{for each } e_a \in \mathcal{E}_k,\ a \sqsubseteq t(k).$$

This assumption corresponds to the idea that all evidence in the same evidential frame is connected with the topic of that frame. If we interpret different evidential frames as different sources from which the agent receives her evidence or as different contexts in which she processes it, then the topic of that frame corresponds to this source or context. Alternatively, the topic associated with an evidential frame can be interpreted as the concept connecting all evidence in that frame, or as the overarching inquiry or question which the agent tries to answer using the evidence in that frame.

2.3 Semantics

We are now ready to give the full definition of our models for the language $\mathcal{L}$ of hyperintensional fragmented evidence, knowledge and belief. These models are effectively a fusion of the evidence models from van Benthem and Pacuit (2011), topic mereologies, and fragmentation.

2.3.1 Models

Definition 2.6 (Topic-sensitive evidence model with fragmentation)**.** A *topic-sensitive evidence model with fragmentation*, or *tsef-model* for short, is a tuple $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$, where

i) W is a non-empty set of *possible worlds*;

ii) $(\mathcal{T}, \oplus)$ is a *mereology of topics* (see Definition 2.2);
$\sqsubseteq$ denotes the associated *topic parthood relation* as defined there;

iii) $t : \text{PROP} \cup \mathcal{F} \to \mathcal{T}$ is a *topic assignment function* assigning a topic to each atomic proposition and to each frame symbol;

iv) for each $k \in \mathcal{F}$, $\mathcal{E}_k \subseteq \mathcal{P}(W) \times \mathcal{T}$ is a set of topic-indexed pieces of evidence subject to the following two conditions:

(1) for each $e_a \in \mathcal{E}_k$, $a \sqsubseteq t(k)$; and

(2) $\mathcal{E}_k$ is closed under non-empty, finite evidence combination, i.e. if $e_a, f_b \in \mathcal{E}_k$ and $e \cap f \neq \varnothing$, then $(e \cap f)_{a\oplus b} \in \mathcal{E}_k$; and

v) $V : \text{PROP} \to \mathcal{P}(W)$ is a *valuation function.*

Topic assignment t is extended to $\mathcal{L}$ by taking as the topic of a sentence φ the fusion of the topics of its propositional variables: $t(\varphi) := \oplus\{t(p) \mid p \in Var(\varphi)\} = t(p_1) \oplus \cdots \oplus t(p_n)$.

2.3.2 Evidence modalities

We now define the semantics for Boolean connectives, the global modality, and our evidence modalities.

Definition 2.7 (Semantics of evidence modalities)**.** Given a tsef-model $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ and a world $w \in W$, the satisfaction relation $\Vdash$ for the Boolean connectives, the global modality and the evidence operators is defined recursively as follows (we write $[\![\varphi]\!]^{\mathfrak{M}}$ for the truth set $\{w \in W \mid \mathfrak{M}, w \Vdash \varphi\}$ of φ in $\mathfrak{M}$ and shorten $[\![\varphi]\!]^{\mathfrak{M}}$ to $[\![\varphi]\!]$ whenever $\mathfrak{M}$ is clear from the context):

$\mathfrak{M}, w \Vdash p$	iff	$w \in V(p)$
$\mathfrak{M}, w \Vdash \neg\varphi$	iff	not $\mathfrak{M}, w \Vdash \varphi$
$\mathfrak{M}, w \Vdash \varphi \wedge \psi$	iff	$\mathfrak{M}, w \Vdash \varphi$ and $\mathfrak{M}, w \Vdash \psi$
$\mathfrak{M}, w \Vdash [\forall]\varphi$	iff	$W \subseteq [\![\varphi]\!]$
$\mathfrak{M}, w \Vdash E_k\varphi$	iff	$(\exists e_a \in \mathcal{E}_k)$ $(e \subseteq [\![\varphi]\!]$ and $t(\varphi) \sqsubseteq a)$
$\mathfrak{M}, w \Vdash E\varphi$	iff	$(\exists k \in \mathcal{F})$ $(\exists e_a \in \mathcal{E}_k)$ $(e \subseteq [\![\varphi]\!]$ and $t(\varphi) \sqsubseteq a)$
$\mathfrak{M}, w \Vdash \Box_k\varphi$	iff	$(\exists e_a \in \mathcal{E}_k)$ $(w \in e \subseteq [\![\varphi]\!]$ and $t(\varphi) \sqsubseteq a)$
$\mathfrak{M}, w \Vdash \Box\varphi$	iff	$(\exists k \in \mathcal{F})$ $(\exists e_a \in \mathcal{E}_k)$ $(w \in e \subseteq [\![\varphi]\!]$ and $t(\varphi) \sqsubseteq a)$

Given a tsef-model $\mathfrak{M}$, *validity of a formula φ on $\mathfrak{M}$*, written $\mathfrak{M} \Vdash \varphi$, is defined as $\mathfrak{M}, w \Vdash \varphi$ for all $w \in W$.

The fragmented evidence operators E and $\Box$ can equivalently be defined as the following abbreviations:

$$E\varphi := \bigvee_{k\in\mathcal{F}} E_k\varphi$$

$$\Box\varphi := \bigvee_{k\in\mathcal{F}} \Box_k\varphi$$

2.3.3 Knowledge and belief

In defining knowledge and belief based on our topic-sensitive, fragmented notion of evidence, we take what Kelly (2016) calls an *evidentialist* stance towards knowledge and belief. This means that we regard both epistemic notions as supervening on evidence: knowledge and belief are entirely determined by an agent's evidential state. Two agents' knowledge (and belief) can differ only if their evidential states differ.

We will take strong inspiration from the topological evidence models presented in Özgün (2017) and Baltag et al. (2016) and in particular follow their *coherentist* perspective towards evidence with respect to knowledge and belief. This means that evaluation of an agent's belief and knowledge takes into consideration *all* evidence available to that agent. However, we depart from the topological approach in a twofold manner: First, we ground knowledge and belief in our *topic-sensitive*, *fragmented* notion of evidence. Second, we exploit only part of the topological structure of evidence: we assume that evidence is closed under *finite intersections*, but not under *arbitrary unions*. We shall return to this difference in Section 4.2.

Semantics for knowledge and belief We say that φ is *believed* in evidential frame k iff there exists a piece of evidence e_a which supports φ and which is consistent with all evidence in $\mathcal{E}_k$. Knowledge in a frame k is defined similarly, but with a *factive* piece of evidence as a witness. In parallel with our notions of fragmented evidence, the fragmented versions of knowledge and belief are defined as knowledge (belief) in *some* evidential frame.

Definition 2.8 (Semantics of knowledge and belief modalities)**.** Given a tsef-model $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ and a world $w \in W$, the satisfaction relation $\Vdash$ for the knowledge and belief modalities is defined as follows:

$$\begin{array}{lll}
\mathfrak{M}, w \Vdash B_k\varphi & \text{iff} & (\exists e_a \in \mathcal{E}_k)\ (e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
& & \qquad \text{and } \forall e'_{a'} \in \mathcal{E}_k :\ e \cap e' \neq \varnothing) \\
\mathfrak{M}, w \Vdash B\varphi & \text{iff} & (\exists k \in \mathcal{F})\ (\exists e_a \in \mathcal{E}_k)\ (e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
& & \qquad \text{and } \forall e'_{a'} \in \mathcal{E}_k :\ e \cap e' \neq \varnothing) \\
\mathfrak{M}, w \Vdash K_k\varphi & \text{iff} & (\exists e_a \in \mathcal{E}_k)\ (w \in e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
& & \qquad \text{and } \forall e'_{a'} \in \mathcal{E}_k :\ e \cap e' \neq \varnothing) \\
\mathfrak{M}, w \Vdash K\varphi & \text{iff} & (\exists k \in \mathcal{F})\ (\exists e_a \in \mathcal{E}_k)\ (w \in e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
& & \qquad \text{and } \forall e'_{a'} \in \mathcal{E}_k :\ e \cap e' \neq \varnothing)
\end{array}$$

In parallel with the evidence modalities, the fragmented modalities B for belief and K for knowledge can then be defined as follows:

$$B\varphi := \bigvee_{k\in\mathcal{F}} B_k\varphi$$
$$K\varphi := \bigvee_{k\in\mathcal{F}} K_k\varphi$$

The modalities B and K are intended to capture our target notions of hyperintensional, fragmented belief and knowledge.

2.4 Hyperintensionality revisited

We now consider several examples demonstrating that our notion of topic-sensitive evidence blocks a number of intuitively problematic principles valid on intensional evidence models.

The pivotal point of all examples and principles discussed in this section is the *topic-sensitivity* of our notion of evidence. The concept of fragmentation plays no role here: all problematic principles are blocked entirely in virtue of the topic-sensitivity of evidence. Similarly, it is irrelevant whether we phrase the principles and examples in terms of *factive* or *non-factive* evidence. For the sake of brevity, we formulate all principles and examples in terms of *fragmented, factive* evidence (i.e., using the operator $\Box$), but the reader should note that all arguments can be given in a similar way when substituting $\Box_k$, E, or E_k for $\Box$. A full overview of the (non-)validities considered in in this section is given at the end of the chapter.

We begin by revisiting a principle already familiar from the introduction: the problem of evidence and necessary truths.

2.4.1 Necessary truths

The following generally holds on intensional evidence models:

Every piece of evidence constitutes evidence for every necessary truth.

Here is a counterexample demonstrating that this is not the case in our framework. Let p denote the proposition that *Biden won the elections*, and let e_a represent a piece of evidence supporting p, for example a radio announcement proclaiming his victory. Let q denote the necessary truth from Example 1.2 stating that *Bachelors are unmarried*.

Counterexample 2.9. Let $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ be a tsef-model s.t. $W = \{w, v\}$, $V(p) = \{w\}$, $V(q) = W$, $\mathcal{T} = \{a, b, c\}$ with a representing the topic *outcome of the elections*, b the topic *marital status of bachelors*, and c the topic *marital status of bachelors and outcome of the elections*; $t(p) = a$, $t(q) = b$, and $\oplus$ and the only evidential frame $\mathcal{E}_1$ as depicted:

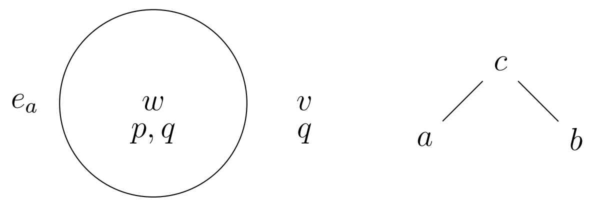

Figure 2.2: Counterexample $\mathfrak{M}$ for necessary truths: left hand side depicts evidence component consisting of one evidential frame $\mathcal{E}_1$, right hand side depicts topics component

In this example, there is no evidence for q, because the only available piece of evidence e_a is not relevant for q. Thus, validity of q together with the availability of some (factive) piece of evidence for p does not entail availability of evidence for q. Formally, we have

$$\mathfrak{M}, w \nVdash ([\forall]q \wedge \Box p) \rightarrow \Box q$$

Proof. Take the counterexample $\mathfrak{M}$ from Figure 2.2. Since $W \subseteq [\![q]\!]$, we have $\mathfrak{M}, w \Vdash [\forall]q$. As $w \in e \subseteq [\![p]\!]$ and $t(p) = a$, we have $\mathfrak{M}, w \Vdash \Box_1 p$ and therefore also $\mathfrak{M}, w \Vdash \Box p$. However, there is only one evidential frame $\mathcal{E}_1$ and only one piece of evidence $e_a \in \mathcal{E}_1$, and $t(q) = b \not\sqsubseteq a$, so we have $\mathfrak{M}, w \nVdash \Box_1 q$ and therefore also $\mathfrak{M}, w \nVdash \Box q$. □

2.4.2 Closure under logical entailment

Our next target is *closure under logical entailment*, which is one of the central principles pertaining to the problem of logical omniscience. Consider the following example (originally from Stalnaker (1984, p. 88), as cited in Hawke, Özgün, and Berto (2020, p. 734)), here phrased in terms of evidence:

Example 2.10. William III might have had evidence that England could avoid war with France, but he needn't have thereby had evidence that England could avoid nuclear war with France. He didn't have the concept nuclear, and so wasn't positioned to think about nuclear wars at all."

Since the implication *if war can be avoided, then nuclear war can be avoided* is necessarily true, any piece of evidence entailing that war can be avoided also entails that nuclear war can be avoided. Thus, evidence on van Benthem-Pacuit models satisfies:

Evidence entailment is closed under logical entailment.

Using the global modality $[\forall]$ and evidence modality $\Box$, this is expressed by the following principle, which is valid on intensional evidence models:

$$([\forall](\varphi \rightarrow \psi) \wedge \Box\varphi) \rightarrow \Box\psi$$

The next counterexample shows that this is not the case with our models. For the sake of brevity, we henceforth do not specify the full models, but only give pictorial presentations. The proofs are very similar in style to that of the previous example, so we omit them as well.

Counterexample 2.11. Let p denote the proposition that *England can avoid war with France*, and q the proposition that *England can avoid nuclear war with France*. The counterexample is provided in Figure 2.3.

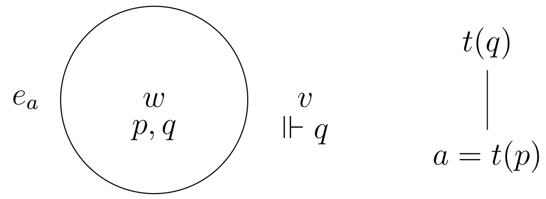

Figure 2.3: Counterexample $\mathfrak{M}$ for closure under logical entailment

We have:

$$\mathfrak{M}, w \nVdash ([\forall](p \to q) \wedge \Box p) \to \Box q$$

2.4.3 Closure under logical equivalence

The next principle is closely related to the previous one. Consider the following two examples:

Example 2.12.

(1) Amsterdam is beautiful.

(2) Fermat's Last Theorem is true.

Since (2) is a necessary truth, it holds in all possible worlds. Thus, *Amsterdam is beautiful* and the conjunction *Amsterdam is beautiful and Fermat's Last Theorem is true* express logically equivalent propositions, so they are true in the same worlds. In general, intensional evidence modals validate:

Evidence possession is closed under logical equivalence.

Formally, this is captured by validity of the following principle:

$$([\forall](\varphi \leftrightarrow \psi) \wedge \Box\varphi) \to \Box\psi$$

Counterexample 2.13. Let p denote the proposition *Amsterdam is beautiful*, and q the proposition *Fermat's Last Theorem is true*, then the model depicted in Figure 2.4 is a counterexample:

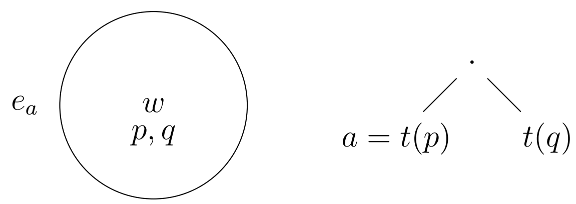

Figure 2.4: Counterexample $\mathfrak{M}$ for closure under logical equivalence

We have:

$$\mathfrak{M}, w \nVdash ([\forall](p \leftrightarrow (p \wedge q)) \wedge \Box p) \to \Box(p \wedge q)$$

2.4.4 Closure under disjunction

For the next problematic principle, we will revisit William III (example modified from Hawke, Özgün, and Berto (2020, p. 743)):

Example 2.14. "William III might have had evidence that France will go to war without having evidence that either they will go to war or develop a nuclear arsenal."

Any piece of evidence entailing φ also entails $\varphi \vee \psi$. Thus, van Benthem-Pacuit evidence models validate:

Evidence entailment is closed under disjunction.

Formally, this corresponds to validity of the following:

$$\Box\varphi \rightarrow \Box(\varphi \vee \psi)$$

Counterexample 2.15. Let p denote the proposition that France will go to war, and q the proposition that France will develop a nuclear arsenal. A counterexample is given in Figure 2.5:

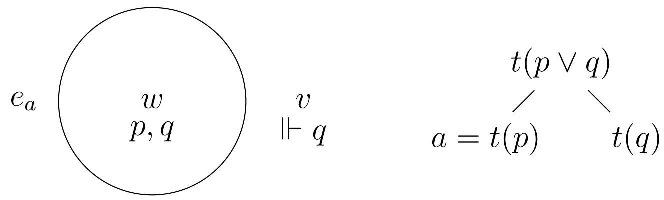

Figure 2.5: Counterexample $\mathfrak{M}$ for closure under disjunction

We have:

$$\mathfrak{M}, w \nVdash \Box p \rightarrow \Box(p \vee q)$$

2.4.5 Intensionally equivalent pieces of evidence

We now give an example of two intensionally equivalent pieces of evidence that provide evidence for different propositions. Unlike intensional evidence models, our framework can account for such situations. Consider the following two scenarios from Builes (2020, pp. 114 sqq.):

Example 2.16.

(1) **Finite Coins:** Suppose you are in a room with a countable infinity of people, and each of you flips a coin without looking at the result. You know that all the coin flips are fair and independent. I then inform you that something remarkable happened: *almost every coin landed tails*. More precisely, only finitely many coins landed heads. Now what should your credence be that your coin landed heads?

(2) **Finite Coins*:** Again, you are in a room with countably many people and each of you flips a coin without looking at the result. You know that all the coin flips are fair and independent. This time, you will only be told information about the *other* people in the room, excluding you. Let S be the set of these other people. I inform you of the following remarkable piece of information: *only finitely many people in S flipped heads*. What should your credence be that your coin landed heads?

In the first scenario, intuitively the credence of my coin being heads up should be 0, since I know that only a vanishingly small amount of all coins landed tails[2]. In the case of Finite Coins*, the intuitive answer is different: the credence of my coin being heads up should be 1/2, because the coin flips are fair and independent[3], and the information that only finitely many people in S flipped heads only concerns the coins of people in S, but not mine.

So if we call the piece of information we receive in Finite Coins evidence e, and the piece of information we receive in Finite Coins* evidence e^*, then we have the following situation: e constitutes evidence supporting the belief that the likelihood of my coin being heads up is 0, whereas e^* does not. This situation turns paradoxical once we realise that e and e^* are logically equivalent: the number of people whose coins landed heads is finite if and only if the number of people excluding myself whose coins landed heads is finite.

Since e and e^* are logically equivalent, a purely intensional representation of evidence would render them as identical sets of worlds. Thus, if p denotes the proposition that the credence of my coin being heads up is 0, these evidence models cannot account for our intuition that e should constitute evidence for p, whereas e^* shouldn't.

In general, any purely intensional account of evidence validates the following principle:

Intensionally equivalent pieces of evidence always constitute evidence for the same propositions.

This is not the case for our topic-sensitive notion of evidence:

Counterexample 2.17. Let *my coin* be the topic of p (*the credence of my coin being heads up is 0*), *everyone's coin* the topic assigned to e, and

[2]Builes (2020, pp. 115 sqq.) provides formal arguments for why it is rational to believe that the probability of my coin being heads up should be 0, but since these are irrelevant for our present purposes, we simply appeal to intuition.

[3]One might object that this example is inconsistent: on a frequentist conception of probability, the possibility of an event in which a coin lands tails up infinitely often seems to collide with its property of being fair. We ask the reader who shares this intuition to bear with us for the sake of the example.

everyone else's coin the topic of e^*. The corresponding topic structure is as depicted in Figure 2.6:

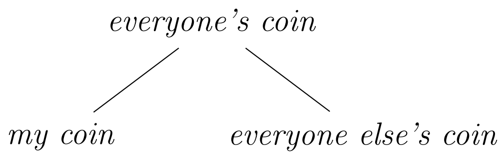

Figure 2.6: Counterexample for intensionally equivalent pieces of evidence

Since $t(p) = \textit{my coin} \sqsubseteq \textit{everyone's coin}$, evidence e with topic *everyone's coin* is relevant for p. In contrast, since *everyone else's coin* $\not\sqsubseteq t(p) =$ *my coin*, evidence e^* with topic *everyone else's coin* is *not* relevant for p.

2.4.6 Conjunction elimination

We saw various examples of principles that are intuitively problematic and successfully blocked in our framework. However, we should make sure that our hyperintensional evidence logic also accounts for some inferences that an agent intuitively *should* be able to make. One such example is *conjunction elimination*: if I have evidence that *it is sunny and warm*, then intuitively I should also have evidence that *it is sunny.* And indeed, the following is valid for $\Box$ and in fact all other evidence modalities as well:

$$\Box(\varphi \land \psi) \to (\Box\varphi \land \Box\psi)$$

Proof. We only prove the $\Box$ version; the proofs for $\Box_k$, E and E_k are similar. Let $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ be a tsef-model and $w \in W$ a world, and suppose that $\mathfrak{M}, w \Vdash \Box(\varphi \land \psi)$. Then by the semantics of $\Box$, there is an evidential frame $k \in \mathcal{F}$ and a witnessing piece of evidence $e_a \in \mathcal{E}_k$, i.e. $w \in e \subseteq [\![\varphi\land\psi]\!]$ and $t(\varphi\land\psi) \sqsubseteq a$. Now since $[\![\varphi\land\psi]\!] \subseteq [\![\varphi]\!]$, we have $w \in e \subseteq [\![\varphi]\!]$. Moreover, by definition of t and $\oplus$, we have $t(\varphi) \sqsubseteq t(\varphi) \oplus t(\psi) = t(\varphi \land \psi)$. Thus, $e_a \in \mathcal{E}_k$ is a witnessing piece of evidence for $\mathfrak{M}, w \Vdash \Box\varphi$, and in a similar way we can show that $\mathfrak{M}, w \Vdash \Box\psi$. Therefore, $\mathfrak{M}, w \Vdash \Box\varphi \land \Box\psi$.

(Note that the same proof works for B_k, B, K_k and K modalities: In all cases, the witnessing piece of evidence for (fragmented) knowledge/belief in $\varphi\land\psi$ is consistent with all evidence in the same evidential frame, so it is also a witness for (fragmented) knowledge/belief of φ and ψ, respectively.) $\Box$

There is an interesting symmetry in the previous proof: If e_a *entails* $\varphi \land \psi$, then it also entails φ; and if e_a is *relevant* for $\varphi \land \psi$, then it is also relevant for φ. Intuitively explained, this means that having evidence for φ *is part of* having evidence for $\varphi \land \psi$.

2.4.7 Closure under conjunction within a fragment

A second positive example of a principle which should follow from our conceptual motivation is the following: if an agent has a piece of evidence for φ and another piece of evidence for ψ in the *same evidential frame* and both pieces of evidence are truthful, then she can combine them to a piece of evidence for $\varphi \wedge \psi$. That is, the following is valid:

$$\Box_k \varphi \wedge \Box_k \psi \rightarrow \Box_k (\varphi \wedge \psi)$$

This corresponds to our assumption that each evidential frame is closed under finite, non-empty evidence combination. For the proof, we refer the reader to Proposition 3.6. Note that the non-factive E_k-version of this principle is *not* valid, because the relevant pieces of evidence might be inconsistent with each other.

2.4.8 Topic-sensitive closure under logical entailment

Hyperintensional evidence possession is not closed under logical entailment because a piece of evidence for φ might not be relevant for ψ, even if φ logically entails ψ. But what if φ logically entails ψ and our agent has both a piece of evidence for φ *and* a piece of evidence which is relevant for ψ in the same evidential frame? No matter whether the second piece of evidence actually entails ψ, if the two pieces of evidence are consistent with each other, then the agent should be able to combine them to a piece of evidence for ψ.

Indeed, there is a topic-sensitive version of *closure under logical entailment*, which is valid on our models. Before stating it, we need to introduce some auxiliary notation. Given a formula φ, let

$$\overline{\varphi} := \bigwedge_{p \in Var(\varphi)} (p \vee \neg p).$$

Note that $\overline{\varphi}$ is a tautology for any formula φ, so $\overline{\varphi}$ is always true at every world. Thus, $\Box_k \overline{\varphi}$ expresses that the agent has a piece of evidence which is *relevant* for φ, even though this evidence might not *entail* φ.

Using this notation, we can state the following valid topic-sensitive version of *closure under logical entailment*:

$$([\forall](\varphi \rightarrow \psi) \wedge \Box_k \varphi \wedge \Box_k \overline{\psi}) \rightarrow \Box_k \psi$$

For the proof, we refer to Proposition 3.6. Note that the *non-factive* version of this principle obtained by substituting E_k for $\Box_k$ is *not* valid, because the witnessing pieces of evidence for $E_k \varphi$ and $E_k \overline{\psi}$ might be inconsistent with each other. Likewise, the $\Box$ and E versions of this principle are not valid either, because the relevant pieces of evidence may belong to different evidential frames.

2.5 Fragmentation revisited

2.5.1 Fragmented closure under conjunction

In the introduction, we encountered Lewis and his beliefs about the local geography of Princeton. We observed that an agent's evidential state may be broken into several fragments, in between which no combination of evidence is performed. However, Lewis' evidence supporting his beliefs concerning the orientation of Nassau Street and the nearby railroad were unfortunately false, and we already noted in Section 2.4.7 that closure under conjunction for non-factive evidence operators E_k and E may fail simply because pieces of non-factive evidence may be inconsistent with each other.

To fully appreciate the fragmentation component of our framework, it will therefore be instructive to additionally consider an example in which an agent fails to combine two *factive* pieces of evidence.

The following is not valid on our models:

$$(\Box\varphi \wedge \Box\psi) \rightarrow \Box(\varphi \wedge \psi)$$

Here is an intuitive example:

Counterexample 2.18. Suppose you have a friend Pieter, who sends you a letter inviting you to visit him. Unfortunately, you don't know where he lives, so you try hard to consider all your evidence to find out where you need to travel. You know that Pieter sent you the invitation letter from his home, and since the letter has a stamp depicting the king of the Netherlands, you have (let's say) truthful evidence that Pieter currently lives in the Netherlands. A while ago, a colleague of yours told you that Pieter recently got married and moved in together with your common friend Renzo, of whom you know that he currently lives somewhere in the Caribbean. Thus, you have another (truthful) piece of evidence entailing that Pieter is in the Caribbean. Unfortunately, you associate your colleague and everything he tells you with your workplace, and you try to keep your work and private life apart. Thus, while pondering the whereabouts of Pieter, you fail to combine your two pieces of evidence to correctly infer that you should travel to Curaçao.

We formally represent this counterexample as follows:

Let p denote the proposition that *Pieter lives in the Netherlands*, and q the proposition that *Pieter lives in the Caribbean.* Let a represent the topic *holiday* and b the topic *work*, and let e_a denote the evidence in form of the invitation letter, and f_b the evidence your colleague has given to you. Our counterexample is the tsef-model $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_1, \mathcal{E}_2\}, V)$ with $W = \{w, v, u\}$, $V(p) = \{w, v\}$, $V(q) = \{v, u\}$, $\mathcal{T} = \{a, b, c\}$, and the two evidential frames and topic structure as depicted in Figure 2.7:

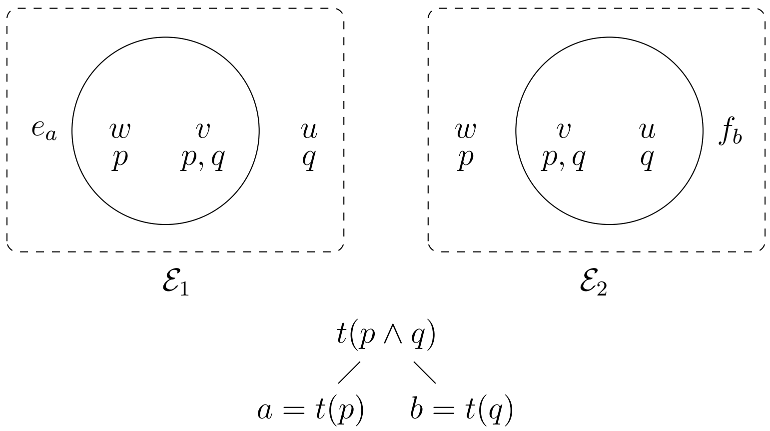

Figure 2.7: Counterexample $\mathfrak{M}$ for fragmented closure under conjunction

Then $\mathfrak{M}, w \nVdash (\Box p \wedge \Box q) \rightarrow \Box(p \wedge q)$.

Proof. Since $w \in e \subseteq [\![p]\!]$ and $t(p) \sqsubseteq a$, $e_a \in \mathcal{E}_1$ is a witness for $\mathfrak{M}, w \Vdash \Box p$. Similarly, $f_b \in \mathcal{E}_2$ is a witness for $\mathfrak{M}, w \Vdash \Box q$. Thus, we have $\mathfrak{M}, w \Vdash \Box p \wedge \Box q$. However, there is no evidential frame with a piece of evidence entailing $p \wedge q$. Therefore, $\mathfrak{M}, w \nVdash \Box(p \wedge q)$. □

2.6 Logical omniscience

In this section, we demonstrate that the notions of knowledge and belief which we defined as grounded in hyperintensional, fragmented evidence give rise to agents that are less idealised than those assumed in purely intensional, non-fragmented frameworks. In particular, our agents do not suffer from full logical omniscience. These properties of our knowledge and belief are direct consequences of the kind of evidence they are based on.

2.6.1 Topic-sensitivity

We shall see that both our fragmented and non-fragmented knowledge and belief operators all block the problematic principles discussed in Section 2.4 pertaining to hyperintensionality. This symmetry is no coincidence: hyperintensionality is inherited from the level of evidence to that of knowledge and belief.

Just as in the case of evidence, fragmentation plays no role in the discussion of these principles in the context of knowledge and belief. Unless otherwise stated, the following (in-)validities all hold for both the fragmented and non-fragmented versions of our knowledge and belief modalities. In fact, not even the distinction between knowledge and belief themselves matters for this discussion, so for the sake of brevity we will state only the B_k-versions of the relevant principles. A full overview of all (in-)validities can be found at the end of this chapter.

In the discussion of the evidence-related principles of the previous section, we chose all counterexamples in such a way that they also provide counterexamples for the cases of knowledge and belief. The proofs can be easily adapted by noting that in each counterexample, there is only one piece of evidence (per evidential frame), so this piece of evidence is vacuously consistent with all other evidence in the same evidential frame and is therefore a witness for knowledge and belief.

Necessary truths Our agents do not know/believe all necessary truths. That is, the following does not hold (see Counterexample 2.9):

$$\text{from } \varphi \text{ infer } B_k\varphi$$

Closure under logical entailment Knowledge and belief are not closed under logical entailment, i.e. the following does not hold (see Counterexample 2.11):

$$\text{from } \varphi \to \psi \text{ infer } B_k\varphi \to B_k\psi$$

Closure under logical equivalence Knowledge and belief are not closed under logical equivalence, i.e. the following does not hold (see Counterexample 2.13):

$$\text{from } \varphi \leftrightarrow \psi \text{ infer } B_k\varphi \to B_k\psi$$

Closure under disjunction Knowledge and belief are not closed under disjunction, i.e. the following is not valid (see Counterexample 2.15):

$$B_k\varphi \to B_k(\varphi \vee \psi)$$

This concludes our analysis of problematic principles that are blocked on the level of knowledge and belief due to the topic-sensitivity of the underlying evidence. Before we move on to fragmentation, we show that just as in the case of evidence, two principles are validated by knowledge and belief: *conjunction elimination* and *closure under conjunction within a fragment*.

Conjunction elimination Just as in the evidence case, our notions of fragmented and non-fragmented knowledge and belief are closed under conjunction elimination. This corresponds to the intuition that believing φ is part of believing $\varphi \wedge \psi$. Correspondingly, the following are valid:

$$B_k(\varphi \wedge \psi) \to (B_k\varphi \wedge B_k\psi)$$

$$B(\varphi \wedge \psi) \to (B\varphi \wedge B\psi)$$
$$K_k(\varphi \wedge \psi) \to (K_k\varphi \wedge K_k\psi)$$
$$K(\varphi \wedge \psi) \to (K\varphi \wedge K\psi)$$

The proof can be found in Section 2.4.6.

Closure under conjunction within a fragment The reason why we took fragmentation on board was to allow for the possibility that an agent sometimes fails to combine two pieces of evidence, even if they are factive. Thus, factive fragmented evidence (and knowledge and belief, we will come to that in the next section) are not closed under conjunction. However, it is interesting to note that the non-fragmented versions of knowledge and belief *are* closed under conjunction, i.e. the following are valid:

$$(B_k\varphi \wedge B_k\psi) \to B_k(\varphi \wedge \psi)$$
$$(K_k\varphi \wedge K_k\psi) \to K_k(\varphi \wedge \psi)$$

The interesting part of the proof is to show that the *combination* of a piece of evidence e_a witnessing $B_k\varphi$ and another piece of evidence witnessing $B_k\psi$ is non-empty and consistent with all evidence in evidential frame $\mathcal{E}_k$. For the proofs of the K_k- and B_k-versions of the principle, we refer to Propositions 3.6 and 3.29, respectively.

Topic-sensitive logical entailment Just as for the $\Box_k$ modality, the following topic-sensitive version of *closure under logical entailment* is valid:

$$([\forall](\varphi \to \psi) \wedge B_k\varphi \wedge B_k\overline{\psi}) \to B_k\psi$$

The fragmented knowledge K_k-version of this principle is valid as well. For the respective proofs, we refer to the soundness proof in Proposition 3.6. Note, however, that fragmentation blocks the B- and K-variants.

2.6.2 Fragmentation

Fragmented closure under conjunction As the reader might expect, fragmented knowledge and belief are not closed under conjunction. That is, the following are not valid (see Counterexample 2.18):

$$(B\varphi \wedge B\psi) \to B(\varphi \wedge \psi)$$
$$(K\varphi \wedge K\psi) \to K(\varphi \wedge \psi)$$

Fragmented closure under known/believed implication Closure under known implication, which corresponds to the (K) axiom $(\Box\varphi \wedge \Box(\varphi \to \psi)) \to \Box\psi$ of normal modal logics, is likewise invalidated by fragmentation. That is, the following are not valid (the proof is similar to the one given in Counterexample 2.18):

$$(B\varphi \wedge B(\varphi \to \psi)) \to B(\psi)$$
$$(K\varphi \wedge K(\varphi \to \psi)) \to K(\psi)$$

2.6.3 Positive and negative introspection

Two additional principles often mentioned in the epistemology literature are *positive* and *negative* introspection. In the context of knowledge, positive introspection says that if an agent *knows* φ, then she *knows that she knows* φ; whereas negative introspection expresses that if an agent *doesn't* know φ, then she *knows that she doesn't know* φ.

Positive introspection All of the modalities B_k, B, K_k, and K have the property of positive introspection. That is, the following (expressed in terms of B_k) is valid:

$$B_k\varphi \to B_k B_k\varphi$$

In the case of belief, this follows from (a) the fact that we assume all logical operators, including the epistemic ones, to be topic-transparent, i.e. $t(\varphi) = t(B_k\varphi) = t(B_k B_k\varphi)$; and (b) that belief is a *world-independent modality*, i.e. it either holds at all worlds of a given model or at none.

The full proof for positive introspection of K_k is given in Proposition 3.6; the proof for B_k in Proposition 3.29. It is easy to see how positive introspection for B and K follows from positive introspection for the respective non-fragmented versions.

Negative introspection None of our belief or knowledge modalities are negatively introspective. That is, the following (expressed in terms of B_k) is *not* valid:

$$\neg B_k\varphi \to B_k\neg B_k\varphi$$

Proof. Let $\mathfrak{M}$ be the tsef-model from Counterexample 2.9. The only piece of evidence in $\mathcal{E}_1$ is e_a and we have $t(q) \not\sqsubseteq a$ and $t(\neg B_k q) \not\sqsubseteq a$, therefore $\mathfrak{M}, w \not\Vdash \neg B_k\varphi \to B_k\neg B_k q$. □

2.7 Overview

Table 2.2 provides an overview of relevant (in-) validities for all epistemic modalities introduced in this chapter and used throughout this thesis.

		E_k	E	$\Box_k$	$\Box$	B_k	B	K_k	K
(KV)	$([\forall]\varphi \wedge \star\psi) \rightarrow \star\varphi$	✗	✗	✗	✗	✗	✗	✗	✗
(CLEn)	$[\forall](\varphi \rightarrow \psi) \rightarrow (\star\varphi \rightarrow \star\psi)$	✗	✗	✗	✗	✗	✗	✗	✗
(CLEq)	$[\forall](\varphi \leftrightarrow \psi) \rightarrow (\star\varphi \rightarrow \star\psi)$	✗	✗	✗	✗	✗	✗	✗	✗
(CE)	$\star(\varphi \wedge \psi) \rightarrow (\star\varphi \wedge \star\psi)$	✓	✓	✓	✓	✓	✓	✓	✓
(CC)	$(\star\varphi \wedge \star\psi) \rightarrow \star(\varphi \wedge \psi)$	✓	✗	✓	✗	✓	✗	✓	✗
(K)	$(\star\varphi \wedge \star(\varphi \rightarrow \psi)) \rightarrow \star\psi$	✓	✗	✓	✗	✓	✗	✓	✗
(CD)	$\star\varphi \rightarrow \star(\varphi \vee \psi)$	✗	✗	✗	✗	✗	✗	✗	✗
(PI)	$\star\varphi \rightarrow \star\star\varphi$	✓	✓	✓	✓	✓	✓	✓	✓
(NI)	$\neg\star\varphi \rightarrow \star\neg\star\varphi$	✗	✗	✗	✗	✗	✗	✗	✗

$\star$s represent placeholder for modalities

(KV)	Knowledge of validities
(CLEn)	Closure under logical entailment
(CLEq)	Closure under logical equivalence
(CE)	Conjunction elimination
(CC)	Closure under conjunction
(K)	Closure under known implication
(CD)	Closure under disjunction
(PI)	Positive introspection
(NI)	Negative introspection

Table 2.2: Overview of relevant (in-)validities for evidence, knowledge and belief modalities

Chapter 3

Technical results: soundness and completeness

This chapter contains the main technical contributions of this thesis. We provide sound and complete axiomatisations of two fragments of the full logic of hyperintensional and fragmented evidence, knowledge and belief. The first fragment is the *factive evidence-knowledge fragment* and is expressive enough to define all epistemic modalities introduced in the previous chapter (see Table 2.1 for a full list including the intended readings). Whereas the soundness proof is routine, the completeness proof involves a quasi-model construction and a non-trivial translation from quasi-models to our topic-sensitive evidence models with fragmentation. The second fragment is the *belief fragment*, for which we give a sound and complete axiomatisation in order to provide a separate logical specification of our notion of belief.

Throughout the chapter, let PROP be a countable set of propositional variables and $\mathcal{F} = \{1, ..., n\}$ a non-empty, finite set of n *frame symbols*. We use k as a metavariable for evidential frames. For any formula φ, we use $Var(\varphi)$ to denote the propositional variables occurring in φ, and let $\overline{\varphi}$ denote the tautology $\bigwedge_{p \in Var(\varphi)}(p \vee \neg p)$.

3.1 The factive evidence-knowledge fragment $\mathcal{L}_{[\forall]\Box_k K_k}$

3.1.1 Syntax

Definition 3.1 (Syntax of $\mathcal{L}_{[\forall]\Box_k K_k}$).

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\forall]\varphi \mid \Box_k\varphi \mid K_k\varphi \quad \text{for } k \in \mathcal{F}$$

We define $[\exists]\varphi := \neg[\forall]\neg\varphi$ as an abbreviation for the dual of $[\forall]\varphi$.

3.1.2 Semantics

Definition 3.2 (Semantics for $\mathcal{L}_{[\forall]\Box_k K_k}$ on tsef-models)**.** Given a tsef-model $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ and a world $w \in W$, the $\Vdash$-semantics for $\mathcal{L}_{[\forall]\Box_k K_k}$ is defined recursively as follows (we write $[\![\varphi]\!]^{\mathfrak{M}}$ for the truth set $\{w \in W \mid \mathfrak{M}, w \Vdash \varphi\}$ of φ in $\mathfrak{M}$ and shorten $[\![\varphi]\!]^{\mathfrak{M}}$ to $[\![\varphi]\!]$ when $\mathfrak{M}$ is clear from the context):

$$
\begin{array}{lll}
\mathfrak{M}, w \Vdash p & \text{iff} & p \in V(p) \\
\mathfrak{M}, w \Vdash \neg\varphi & \text{iff} & \text{not } \mathfrak{M}, w \Vdash \varphi \\
\mathfrak{M}, w \Vdash \varphi \wedge \psi & \text{iff} & \mathfrak{M}, w \Vdash \varphi \text{ and } \mathfrak{M}, w \Vdash \psi \\
\mathfrak{M}, w \Vdash [\forall]\varphi & \text{iff} & W \subseteq [\![\varphi]\!] \\
\mathfrak{M}, w \Vdash \Box_k\varphi & \text{iff} & \exists e_a \in \mathcal{E}_k \ (w \in e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a) \\
\mathfrak{M}, w \Vdash K_k\varphi & \text{iff} & \exists e_a \in \mathcal{E}_k \ (w \in e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
 & & \qquad\qquad \text{and } \forall e'_{a'} \in \mathcal{E}_k : \ e \cap e' \neq \varnothing)
\end{array}
$$

Validity of a formula φ on a tsef-model $\mathfrak{M}$ is defined in Definition 2.7. If $\Gamma \cup \{\varphi\}$ is a set of formulas over $\mathcal{L}_{[\forall]\Box_k K_k}$, we say that Γ is *satisfied* at w in $\mathfrak{M}$ and write $\mathfrak{M}, w \Vdash \Gamma$ iff $\mathfrak{M}, w \Vdash \varphi$ for all $\varphi \in \Gamma$; moreover, we say that φ is a *local semantic consequence of* Γ and write $\Gamma \Vdash \varphi$ if for all tsef-models $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ and all worlds $w \in W$, if $\mathfrak{M}, w \Vdash \Gamma$, then $\mathfrak{M}, w \Vdash \varphi$.

3.1.3 Expressivity

The virtue of the factive evidence-knowledge fragment $\mathcal{L}_{[\forall]\Box_k K_k}$ is its expressivity: all our (fragmented) evidence, belief and knowledge modalities can be defined in it. The definitions of the modalities not covered by the semantics above are as follows:

$$
\begin{aligned}
E_k\varphi &:= [\exists]\Box_k\varphi \\
B_k\varphi &:= [\exists]K_k\varphi \\
E\varphi &:= \bigvee_{k\in\mathcal{F}} E_k\varphi \\
B\varphi &:= \bigvee_{k\in\mathcal{F}} B_k\varphi \\
K\varphi &:= \bigvee_{k\in\mathcal{F}} K_k\varphi
\end{aligned}
$$

The validity of these definitions under the full semantics given in Definitions 2.7 and 2.8 follow directly from the respective semantic clauses.

3.1.4 Axiomatisation

Table 3.1 provides a sound and complete axiomatisation of the logic HEK of hyperintensional evidence and knowledge over $\mathcal{L}_{[\forall]\Box_k K_k}$. The proofs for soundness and completeness can be found in Proposition 3.6 and Theorem 3.7, respectively.

(I)	(CPL)	all classical propositional tautologies and modus ponens
(II)	(S5) **axioms and rules for** $[\forall]$**:**	
	$(\mathsf{K}_{[\forall]})$	$([\forall]\varphi \wedge [\forall](\varphi \to \psi)) \to [\forall]\psi$
	$(\mathsf{T}_{[\forall]})$	$[\forall]\varphi \to \varphi$
	$(\mathsf{4}_{[\forall]})$	$[\forall]\varphi \to [\forall][\forall]\varphi$
	$(\mathsf{5}_{[\forall]})$	$\neg[\forall]\varphi \to [\forall]\neg[\forall]\varphi$
	$(\mathsf{Nec}_{[\forall]})$	from φ, infer $[\forall]\varphi$
(III)	**Axioms for** $\heartsuit_k$ **with** $\heartsuit \in \{\Box, K\}$ **and** $k \in \mathcal{F}$**:**	
	$(\mathsf{T}_{\heartsuit_k})$	$\heartsuit_k\varphi \to \varphi$
	$(\mathsf{4}_{\heartsuit_k})$	$\heartsuit_k\varphi \to \heartsuit_k\heartsuit_k\varphi$
	$(\mathsf{C}_{\heartsuit_k})$	$\heartsuit_k(\varphi \wedge \psi) \leftrightarrow (\heartsuit_k\varphi \wedge \heartsuit_k\psi)$
	$(\mathsf{Ax1}_{\heartsuit_k})$	$\heartsuit_k\varphi \to \heartsuit_k\overline{\varphi}$
(IV)	**Axioms connecting** $[\forall]$ **and** $\heartsuit_k$ **with** $\heartsuit \in \{\Box, K\}$ **and** $k \in \mathcal{F}$**:**	
	$(\mathsf{Ax2}_{\heartsuit_k})$	$([\forall](\varphi \to \psi) \wedge \heartsuit_k\varphi \wedge \heartsuit_k\overline{\psi}) \to \heartsuit_k\psi$
	$(\mathsf{Ax3}_k)$	$([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to [\exists]\Box_k(\varphi \wedge \psi)$
	(F_k)	$K_k\varphi \to \Box_k\varphi$

Table 3.1: Sound and complete axiomatisation of the logic HEK of hyperintensional evidence and knowledge over $\mathcal{L}_{[\forall]\Box_k K_k}$

Proposition 3.3. *The following are derivable in* HEK*:*

(i) $\heartsuit_k\overline{\varphi} \leftrightarrow \bigwedge_{p\in Var(\varphi)} \heartsuit_k\overline{p}$ *for* $\heartsuit \in \{\Box, K\}$

(ii) $\heartsuit_k\overline{\varphi} \to \heartsuit_k\overline{\psi}$*, if* $Var(\psi) \subseteq Var(\varphi)$ *for* $\heartsuit \in \{\Box, K\}$

(iii) $([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to [\exists](K_k\varphi \wedge \Box_k\psi)$

Proof.

(i) $\heartsuit_k\overline{\varphi} \leftrightarrow \bigwedge_{p\in Var(\varphi)} \heartsuit_k\overline{p}$:

By definition of $\overline{\varphi}$, $\heartsuit_k\overline{\varphi} = \heartsuit_k \bigwedge_{p\in Var(\varphi)} \overline{p}$. Now, by $(\mathsf{C}_{\heartsuit_k})$, we have

$$\vdash_{\mathsf{HEK}} \heartsuit_k \bigwedge_{p\in Var(\varphi)} \overline{p} \leftrightarrow \bigwedge_{p\in Var(\varphi)} \heartsuit_k\overline{p},$$

i.e. $\vdash_{\mathsf{HEK}} \heartsuit_k\overline{\varphi} \leftrightarrow \bigwedge_{p\in Var(\varphi)} \heartsuit_k\overline{p}$.

(ii) $\heartsuit_k\overline{\varphi} \to \heartsuit_k\overline{\psi}$, if $Var(\psi) \subseteq Var(\varphi)$ for $\heartsuit \in \{\Box, k\}$:

Suppose $Var(\psi) \subseteq Var(\varphi)$.

$$\vdash_{\mathsf{HEK}} \heartsuit_k\overline{\varphi} \to \bigwedge_{p \in Var(\varphi)} \heartsuit_k\overline{p} \qquad \text{(by (i))}$$

$$\vdash_{\mathsf{HEK}} \bigwedge_{p \in Var(\varphi)} \heartsuit_k\overline{p} \to \bigwedge_{p \in Var(\psi)} \heartsuit_k\overline{p}$$

$$\text{(by } Var(\psi) \subseteq Var(\varphi) \text{ and } \vdash_{\text{CPL}} (\varphi \wedge \psi) \to \varphi)$$

$$\vdash_{\mathsf{HEK}} \bigwedge_{p \in Var(\psi)} \heartsuit_k\overline{p} \leftrightarrow \heartsuit_k\overline{\psi} \qquad \text{(by (i))}$$

$$\vdash_{\mathsf{HEK}} \heartsuit_k\overline{\varphi} \to \heartsuit_k\overline{\psi} \qquad \text{(by CPL)}$$

Therefore, $\vdash_{\mathsf{HEK}} \heartsuit_k\overline{\varphi} \to \heartsuit_k\overline{\psi}$.

(iii) $([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to [\exists](K_k\varphi \wedge \Box_k\psi)$:

$\vdash_{\mathsf{HEK}} [\exists]K_k\varphi \to [\exists]K_kK_k\varphi$ (by $(\mathsf{4}_{K_k})$ and $\mathsf{S5}_{[\forall]}$)
$\vdash_{\mathsf{HEK}} [\exists]\Box_k\psi \to [\exists]\Box_k\Box_k\psi$ (by $(\mathsf{4}_{\Box_k})$ and $\mathsf{S5}_{[\forall]}$)
$\vdash_{\mathsf{HEK}} ([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to ([\exists]K_kK_k\varphi \wedge [\exists]\Box_k\Box_k\psi)$
$\vdash_{\mathsf{HEK}} ([\exists]K_kK_k\varphi \wedge [\exists]\Box_k\Box_k\psi) \to [\exists]\Box_k(K_k\varphi \wedge \Box_k\psi)$ (by $(\mathsf{Ax3}_k)$)
$\vdash_{\mathsf{HEK}} ([\exists]K_kK_k\varphi \wedge [\exists]\Box_k\Box_k\psi) \to [\exists](K_k\varphi \wedge \Box_k\psi)$ (by $(\mathsf{T}_{\Box_k})$)
$\vdash_{\mathsf{HEK}} ([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to [\exists](K_k\varphi \wedge \Box_k\psi)$

□

3.1.5 Soundness

Definition 3.4 (Derivability and Consistency)**.** We say that a formula $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ *is a theorem of* HEK and write $\vdash_{\mathsf{HEK}} \varphi$ if $\varphi \in \mathsf{HEK}$. If Γ is a set of formulas over $\mathcal{L}_{[\forall]\Box_k K_k}$, we say that φ *is* HEK*-derivable from* Γ, and write $\Gamma \vdash_{\mathsf{HEK}} \varphi$, if $\vdash_{\mathsf{HEK}} \varphi$ or there are formulas $\psi_1, ..., \psi_n \in \Gamma$ s.t. $\vdash_{\mathsf{HEK}} (\psi_1 \wedge \cdots \wedge \psi_n) \to \varphi$. We say that a set of formulas Γ is HEK-*inconsistent* if there is some formula $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $\Gamma \vdash_{\mathsf{HEK}} \varphi \wedge \neg\varphi$, and HEK-*consistent* otherwise. Finally, a formula $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ is HEK-inconsistent if $\{\varphi\}$ is HEK-inconsistent; otherwise φ is HEK-consistent.

Lemma 3.5. *Let $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k \in \mathcal{F}}, V)$ be a tsef-model and $e_a, f_b \in \mathcal{E}_k$ s.t. $e \cap f \neq \varnothing$, $e \cap g \neq \varnothing$ and $f \cap g \neq \varnothing$ for all $g_c \in \mathcal{E}_k$. Then $(e \cap f) \cap g \neq \varnothing$ for all $g_c \in \mathcal{E}_k$.*

Proof. Suppose towards a contradiction that e_a, f_b are as described and there is $g_c \in \mathcal{E}_k$ s.t. $(e \cap f) \cap g = \varnothing$. By assumption, $f \cap g \neq \varnothing$. Since $\mathcal{E}_k$ is closed under non-empty finite evidence combination, we have $(f \cap g)_{b \oplus c} \in \mathcal{E}_k$. This implies $e \cap (f \cap g) \neq \varnothing$, which is a contradiction. □

Proposition 3.6 (Soundness of HEK)**.** *The logic* HEK *of hyperintensional evidence and knowledge is sound with respect to the class of tsef-models. That is, for all formulas* $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ *and all tsef-models* $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k \in \mathcal{F}}, V)$, $\vdash_{\mathsf{HEK}} \varphi$ *implies* $\mathfrak{M}, w \Vdash \varphi$ *for all* $w \in W$.

Proof. Soundness is proved as usual by showing the validity of the axioms and the preservation of soundness under the inference rules. We omit the axioms and inference rules of classical propositional logic. The fact that the global modality $[\forall]$ is an (S5)-operator is standard. The validity of $(\mathsf{T}_{\heartsuit_k})$ and (F_k) follows immediately from the definition of the $\Vdash$-semantics. It remains to show the validity of the axioms $(\mathsf{4}_{\heartsuit_k})$, $(\mathsf{C}_{\heartsuit_k})$, $(\mathsf{Ax1}_{\heartsuit_k})$, $(\mathsf{Ax2}_{\heartsuit_k})$, and $(\mathsf{Ax3}_{\heartsuit_k})$. For all of these axioms, we only show validity of the $\heartsuit_k = K_k$ case; the proofs for the $\Box_k$ cases are similar (and easier). For each of the following cases, let $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k \in \mathcal{F}}, V)$ be an arbitrary tsef-model and $w \in W$.

– Validity of $(\mathsf{4}_{K_k})$ $K_k\varphi \to K_k K_k \varphi$:

  Suppose $\mathfrak{M}, w \Vdash K_k\varphi$, then there is $e_a \in \mathcal{E}_k$ s.t. $w \in e \subseteq [\![\varphi]\!]$, $t(\varphi) \sqsubseteq a$, and $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$. It is easy to see that $e \subseteq [\![K_k\varphi]\!]$. Moreover, we have $t(K_k\varphi) = t(\varphi) \sqsubseteq a$, so e_a is a witness for $\mathfrak{M}, w \Vdash K_k K_k \varphi$.

– Validity of (C_{K_k}) $K_k(\varphi \wedge \psi) \leftrightarrow (K_k\varphi \wedge K_k\psi)$:

  (⇒:) Suppose $\mathfrak{M}, w \Vdash K_k(\varphi \wedge \psi)$. Then there is $e_a \in \mathcal{E}_k$ s.t. $w \in e \subseteq [\![\varphi \wedge \psi]\!]$, $t(\varphi \wedge \psi) \sqsubseteq a$, and $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$. Now we simply observe that $e \subseteq [\![\varphi]\!]$ and $t(\varphi) \sqsubseteq t(\varphi) \oplus t(\psi) = t(\varphi \wedge \psi) \sqsubseteq a$, so e_a is also a witness for $\mathfrak{M}, w \Vdash K_k\varphi$; similarly for $\mathfrak{M}, w \Vdash K_k\psi$.

  (⇐:) Suppose $\mathfrak{M}, w \Vdash K_k\varphi \wedge K_k\psi$, then there are

  * $e_a \in \mathcal{E}_k$ witnessing $\mathfrak{M}, w \Vdash K_k\varphi$, i.e. s.t.
    · $w \in e \subseteq [\![\varphi]\!]$;
    · $t(\varphi) \sqsubseteq a$;
    · $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$;
  * $f_b \in \mathcal{E}_k$ witnessing $\mathfrak{M}, w \Vdash K_k\psi$, i.e. s.t.
    · $w \in f \subseteq [\![\psi]\!]$;
    · $t(\psi) \sqsubseteq b$;
    · $f \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$.

We claim that $(e \cap f)_{a \oplus b}$ is a witness for $\mathfrak{M}, w \Vdash K_k(\varphi \wedge \psi)$. First of all, since $w \in e \cap f$, we have $e \cap f \neq \varnothing$, so $(e \cap f)_{a \oplus b} \in \mathcal{E}_k$ because $\mathcal{E}_k$ is closed under non-empty finite evidence combination. Next, we have $t(\varphi \wedge \psi) \sqsubseteq a \oplus b$: As $t(\varphi) \sqsubseteq a$ and $t(\psi) \sqsubseteq b$, we have $t(\varphi), t(\psi) \sqsubseteq a \oplus b$. Now $t(\varphi \wedge \psi) = t(\varphi) \oplus t(\psi)$ is by definition the *least* upper bound of $\{t(\varphi), t(\psi)\}$, so we have $t(\varphi \wedge \psi) \sqsubseteq a \oplus b$. It remains to show that $(e \cap f) \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$, but this follows from Lemma 3.5. So $\mathfrak{M}, w \Vdash K_k(\varphi \wedge \psi)$.

– Validity of $(\mathsf{Ax1}_{K_k})$ $K_k\varphi \to K_k\overline{\varphi}$:

  Suppose $\mathfrak{M}, w \Vdash K_k\varphi$, then there is $e_a \in \mathcal{E}_k$ s.t.

  · $w \in e \subseteq [\![\varphi]\!]$;
  · $t(\varphi) \sqsubseteq a$;
  · $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$.

  We claim that e_a is also a witness for $\mathfrak{M}, w \Vdash K_k\overline{\varphi}$:

  First, note that $\overline{\varphi} = \bigwedge_{p \in Var(\varphi)}(p \vee \neg p)$ is a tautology, so $w \in e \subseteq [\![\overline{\varphi}]\!] = W$. It remains to show that $t(\overline{\varphi}) \sqsubseteq a$, but this follows from $Var(\overline{\varphi}) = Var(\varphi)$, which implies $t(\overline{\varphi}) = t(\varphi)$.

– Validity of $(\mathsf{Ax2}_{K_k})$ $([\forall](\varphi \to \psi) \wedge K_k\varphi \wedge K_k\overline{\psi}) \to K_k\psi$:

  Suppose $\mathfrak{M}, w \Vdash [\forall](\varphi \to \psi) \wedge K_k\varphi \wedge K_k\overline{\psi}$, then by the semantics of $[\forall]$ we have $[\![\varphi]\!] \subseteq [\![\psi]\!]$; and moreover, there are

  – $e_a \in \mathcal{E}_k$ witnessing $\mathfrak{M}, w \Vdash K_k\varphi$, i.e.
    * $w \in e \subseteq [\![\varphi]\!]$;
    * $t(\varphi) \sqsubseteq a$;
    * $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$;
  – $f_b \in \mathcal{E}_k$ witnessing $\mathfrak{M}, w \Vdash K_k\overline{\psi}$, i.e.
    * $w \in f \subseteq [\![\overline{\psi}]\!]$;
    * $t(\overline{\psi}) \sqsubseteq b$;
    * $f \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$.

  We claim that $(e \cap f)_{a \oplus b}$ is a witness for $\mathfrak{M}, w \Vdash K_k\psi$: First, note that $w \in e \cap f$, so $e \cap f \neq \varnothing$. Since $\mathcal{E}_k$ is closed under non-empty finite evidence combination, we have $(e \cap f)_{a \oplus b} \in \mathcal{E}_k$. Second, since $[\![\varphi]\!] \subseteq [\![\psi]\!]$, we have $e \cap f \subseteq e \subseteq [\![\varphi]\!] \subseteq [\![\psi]\!]$. Third, we have $t(\psi) = t(\overline{\psi}) \sqsubseteq b \sqsubseteq a \oplus b$. Finally, it remains to show that $(e \cap f) \cap g \neq \varnothing$ for all $g_c \in \mathcal{E}_k$, but this was shown in Lemma 3.5. Therefore, $(e \cap f)_{a \oplus b}$ is indeed a witness for $\mathfrak{M}, w \Vdash K_k\psi$.

- Validity of ($\mathsf{Ax3}_k$) $([\exists]K_k\varphi \wedge [\exists]\Box_k\psi) \to [\exists]\Box_k(\varphi \wedge \psi)$:

  Suppose $\mathfrak{M}, w \Vdash [\exists]K_k\varphi \wedge [\exists]\Box_k\psi$. Then, since $\mathfrak{M}, w \Vdash [\exists]K_k\varphi$, there are $u \in W$ and $e_a \in \mathcal{E}_k$ s.t.

  - $u \in e \subseteq [\![\varphi]\!]$;
  - $t(\varphi) \sqsubseteq a$; and
  - $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$.

  Moreover, since $\mathfrak{M}, w \Vdash [\exists]\Box_k\psi$, there are $v \in W$ and $f_b \in \mathcal{E}_k$ s.t.

  - $v \in f \subseteq [\![\psi]\!]$ and
  - $t(\psi) \sqsubseteq b$.

  From $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k$ it follows that $e \cap f \neq \varnothing$, so $(e \cap f)_{a \oplus b} \in \mathcal{E}_k$. Let $x \in e \cap f$. Note that we have $x \in (e \cap f) \subseteq [\![\varphi \wedge \psi]\!]$. Moreover, since $t(\varphi \wedge \psi) = t(\varphi) \oplus t(\psi)$ is the least upper $\sqsubseteq$-bound of $\{t(\varphi), t(\psi)\}$, it follows from $t(\varphi) \sqsubseteq a$ and $t(\psi) \sqsubseteq b$ that $t(\varphi \wedge \psi) \sqsubseteq a \oplus b$. Thus, $(e \cap f)_{a \oplus b}$ is a witness for $\mathfrak{M}, w \Vdash \Box_k(\varphi \wedge \psi)$. Therefore, $\mathfrak{M}, w \Vdash [\exists]\Box_k(\varphi \wedge \psi)$.

This concludes the soundness proof of HEK. □

3.1.6 Completeness

Theorem 3.7 (Completeness of HEK)**.** *The logic* HEK *of hyperintensional evidence and knowledge is strongly complete with respect to the class of tsef-models. That is, for any set of formulas* $\Gamma \cup \{\varphi\}$ *of the language* $\mathcal{L}_{[\forall]\Box_k K_k}$*, if* $\Gamma \Vdash \varphi$*, then* φ *is derivable in* HEK *from* Γ*.*

Before delving into the completeness proof, we need a few preliminary definitions and results.

Definition 3.8 (Maximal HEK-consistent set)**.** Let Γ be a set of formulas over HEK. Γ is called *maximally* HEK*-consistent* (or, for short, a HEK-MCS) if Γ is HEK-consistent, and for any set of formulas Γ', if $\Gamma \subsetneq \Gamma'$, then Γ' is HEK-inconsistent. We denote by $\mathcal{MCS}_{\mathsf{HEK}}$ the set of all maximally HEK-consistent sets.

Lemma 3.9 (Properties of HEK-MCS)**.** *The following hold for all* HEK-MCS Γ*:*

(i) Γ *is closed under modus ponens: if* $\varphi \in \Gamma$ *and* $\varphi \to \psi \in \Gamma$*, then* $\psi \in \Gamma$*;*

(ii) if $\vdash_{\mathsf{HEK}} \varphi$*, then* $\varphi \in \Gamma$*;*

(iii) *for all formulas* φ, $\varphi \in \Gamma$ *or* $\neg\varphi \in \Gamma$*;*

(iv) *for all formulas* φ, ψ*:* $\varphi \wedge \psi \in \Gamma$ *iff* $\varphi \in \Gamma$ *and* $\psi \in \Gamma$*.*

Proof. Standard (e.g., Proposition 4.16 in Blackburn, Rijke, and Venema (2001)). □

Lemma 3.10 (Lindenbaum's Lemma)**.** *If* Φ *is a* HEK*-consistent set of formulas, then there is a* HEK-MCS Φ^+ *s.t.* $\Phi \subseteq \Phi^+$*.*

Proof. Standard (e.g., Lemma 4.17 in Blackburn, Rijke, and Venema (2001)). □

Proposition 3.11. HEK *is strongly complete with respect to the class of tsef-models iff every* HEK*-consistent set of formulas is satisfiable on some world in some tsef-model* $\mathfrak{M}$*.*

Proof. Standard (e.g., Proposition 4.12 in Blackburn, Rijke, and Venema (2001)). □

To prove strong completeness with respect to tsef-models, we prove the equivalent statement that every HEK-consistent set of formulas is satisfiable on some tsef-model. The equivalence is a standard fact and stated in Proposition 3.11. However, we need to make a detour and prove this in two steps: First, we show that the logic of HEK is strongly complete with respect to a canonical *quasi-model*, i.e. that every HEK-consistent set of formulas is satisfiable on the canonical quasi-model. We then proceed to show that for every quasi-model, there is a modally equivalent tsef-model, which yields the desired result that every HEK-consistent set of formulas is satisfiable on a tsef-model.

Step 1: Strong completeness for quasi-models

Definition 3.12 (Quasi-model for $\mathcal{L}_{[\forall]\Box_k K_k}$)**.** A quasi-model is a tuple $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$, where W, V, $\mathcal{T}$, $\oplus$ and t are as in Definition 2.6 of tsef-models; and

i) for each $k \in \mathcal{F}$, $\mathcal{E}_{\Box_k} \subseteq \mathcal{P}(W) \times \mathcal{T}$ is a set of topic-indexed pieces of evidence s.t. for all $e_a \in \mathcal{E}_{\Box_k}$, we have $a \sqsubseteq t(k)$; additionally, we require that $\mathcal{E}_{\Box_k}$ is closed under non-empty evidence combination, i.e. for all $e_a, f_b \in \mathcal{E}_{\Box_k}$, if $e \cap f \neq \varnothing$, then $(e \cap f, a \oplus b) \in \mathcal{E}_{\Box_k}$; and

ii) for each $k \in \mathcal{F}$, $\mathcal{E}_{K_k} \subseteq \mathcal{E}_{\Box_k}$ is a set of topic-indexed pieces of evidence closed under non-empty evidence combination, i.e. for all $e_a, f_b \in \mathcal{E}_{K_k}$, if $e \cap f \neq \varnothing$, then $(e \cap f, a \oplus b) \in \mathcal{E}_{K_k}$; and moreover, we require that for all $e_a \in \mathcal{E}_{K_k}$ and $f_b \in \mathcal{E}_{\Box_k}$, $e \cap f \neq \varnothing$.

Definition 3.13 (Semantics for $\mathcal{L}_{[\forall]\Box_k K_k}$ on quasi-models)**.** Given a quasi-model $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ and a world $w \in W$, the $\Vdash^*$-semantics for $\mathcal{L}_{[\forall]\Box_k K_k}$ is defined recursively as follows (we write $[\varphi]^{\mathfrak{Q}}$ for the truth set $\{w \in W \mid \mathfrak{Q}, w \Vdash \varphi\}$ of φ in $\mathfrak{Q}$ and shorten $[\varphi]^{\mathfrak{Q}}$ to $[\varphi]$ when $\mathfrak{Q}$ is clear from the context):

$$
\begin{array}{lll}
\mathfrak{Q}, w \Vdash^* p & \text{iff} & p \in V(p) \\
\mathfrak{Q}, w \Vdash^* \neg\varphi & \text{iff} & \text{not } \mathfrak{Q}, w \Vdash^* \varphi \\
\mathfrak{Q}, w \Vdash^* \varphi \wedge \psi & \text{iff} & \mathfrak{Q}, w \Vdash^* \varphi \text{ and } \mathfrak{Q}, w \Vdash^* \psi \\
\mathfrak{Q}, w \Vdash^* [\forall]\varphi & \text{iff} & W \subseteq [\varphi] \\
\mathfrak{Q}, w \Vdash^* \Box_k\varphi & \text{iff} & \exists e_a \in \mathcal{E}_{\Box_k} \ (x \in e \subseteq [\varphi] \text{ and } t(\varphi) \sqsubseteq a) \\
\mathfrak{Q}, w \Vdash^* K_k\varphi & \text{iff} & \exists e_a \in \mathcal{E}_{K_k} \ (x \in e \subseteq [\varphi] \text{ and } t(\varphi) \sqsubseteq a)
\end{array}
$$

Definition 3.14. For any HEK-MCS Γ, we define

$$\Gamma[\forall] := \{\varphi \in \mathcal{L}_{[\forall]\Box_k K_k} \mid [\forall]\varphi \in \Gamma\}.$$

We define a binary relation $\sim_{[\forall]}$ on $\mathcal{MCS}_{\mathsf{HEK}}$ as follows: for any $\Gamma, \Delta \in \mathcal{MCS}_{\mathsf{HEK}}$,

$$\Gamma \sim_{[\forall]} \Delta \quad \text{iff} \quad \Gamma[\forall] \subseteq \Delta.$$

Lemma 3.15. $\sim_{[\forall]}$ *is an equivalence relation.*

Proof. We need to show that $\sim_{[\forall]}$ is reflexive, symmetric, and transitive.

- *Reflexivity:* Take any $\Gamma \in \mathcal{MCS}_{\mathsf{HEK}}$. By $(\mathsf{T}_{[\forall]})$ and Lemma 3.9, $\Gamma[\forall] \subseteq \Gamma$, so $\Gamma \sim_{[\forall]} \Gamma$.
- *Symmetry:* Take any $\Gamma, \Delta \in \mathcal{MCS}_{\mathsf{HEK}}$ and suppose $\Gamma \sim_{[\forall]} \Delta$. i.e. $\Gamma[\forall] \subseteq \Delta$. We show that $\Delta[\forall] \subseteq \Gamma$. Take any $\varphi \in \Delta[\forall]$, then $[\forall]\varphi \in \Delta$. Towards a contradiction, suppose $\varphi \notin \Gamma$. By Lemma 3.9, $\neg\varphi \in \Gamma$. By $(\mathsf{T}_{[\forall]})$ and Lemma 3.9, $[\exists]\neg\varphi \in \Delta$, which contradicts $[\forall]\varphi \in \Delta$. So $\Delta[\forall] \subseteq \Gamma$, which means that $\Delta \sim_{[\forall]} \Gamma$.
- *Transitivity:* Take any $\Gamma, \Delta, \mathrm{E} \in \mathcal{MCS}_{\mathsf{HEK}}$ and suppose that $\Gamma \sim_{[\forall]} \Delta$ and $\Delta \sim_{[\forall]} \mathrm{E}$. Take any $\varphi \in \Gamma[\forall]$, then $[\forall]\varphi \in \Gamma$. By $(\mathsf{4}_{[\forall]})$ and Lemma 3.9, we have $[\forall][\forall]\varphi \in \Gamma$. Since $\Gamma[\forall] \subseteq \Delta$, we have $[\forall]\varphi \in \Delta$, and since $\Delta[\forall] \subseteq \mathrm{E}$, we have $\varphi \in \mathrm{E}$, as required.

□

Definition 3.16 (Canonical quasi-model for Γ_0)**.** If Γ_0 is a HEK-MCS, then the canonical quasi-model for Γ_0 is a tuple $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k\in\mathcal{F}}, V^c)$, where

i) $W^c := \{\Gamma \in \mathcal{MCS}_{\mathsf{HEK}} \mid \Gamma_0 \sim_{[\forall]} \Gamma\}$;

ii) $\mathcal{T}^c := \mathcal{P}(\text{PROP})$;

iii) $\oplus^c := \cup$ (it follows that $\sqsubseteq^c = \subseteq$);

iv) $t^c : \text{PROP} \cup \mathcal{F} \to \mathcal{T}$ s.t.

- $t^c(p) = \{p\}$ for any $p \in \text{PROP}$;
- $t^c(k) = \{p \in \text{PROP} \mid [\exists]\Box_k \overline{p} \in \Gamma_0\}$ for any $k \in \mathcal{F}$; and
- t^c is extended to $\mathcal{L}_{[\forall]\Box_k K_k}$ by putting $t^c(\varphi) := Var(\varphi)$;

v) for each $k \in \mathcal{F}$, $\mathcal{E}^c_{\Box_k} := \{(\widehat{\Box_k \varphi}, Var(\varphi)) \mid [\exists]\Box_k \varphi \in \Gamma_0\}$;

vi) for each $k \in \mathcal{F}$, $\mathcal{E}^c_{K_k} := \{(\widehat{K_k \varphi}, Var(\varphi)) \mid [\exists]K_k \varphi \in \Gamma_0\}$; and

vii) $V^c : \text{PROP} \to \mathcal{P}(W^c)$ s.t. $V^c(p) := \widehat{p}$,
where $\widehat{\varphi} := \{\Gamma \in W^c \mid \varphi \in \Gamma\}$ for all $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$.

Lemma 3.17. *For any $n \in \mathbb{N}$, if $\bigwedge_{i \le n} [\forall]\varphi_i \to \psi$ is a theorem of $\mathsf{S5}_{[\forall]}$, then so is $\bigwedge_{i \le n} [\forall]\varphi_i \to [\forall]\psi$.*

Proof.

$$\vdash_{\mathsf{S5}_{[\forall]}} \bigwedge_{i \le n} [\forall]\varphi_i \to \psi \qquad \text{(Assumption)}$$

$$\Rightarrow \vdash_{\mathsf{S5}_{[\forall]}} [\forall](\bigwedge_{i \le n} [\forall]\varphi_i \to \psi) \qquad (\mathsf{Nec}_{[\forall]})$$

$$\Rightarrow \vdash_{\mathsf{S5}_{[\forall]}} [\forall]\bigwedge_{i \le n} [\forall]\varphi_i \to [\forall]\psi \qquad (\mathsf{K}_{[\forall]})$$

$$\Rightarrow \vdash_{\mathsf{S5}_{[\forall]}} [\forall][\forall]\bigwedge_{i \le n} \varphi_i \to [\forall]\psi \qquad (\mathsf{C}_{[\forall]})$$

$$\Rightarrow \vdash_{\mathsf{S5}_{[\forall]}} [\forall]\bigwedge_{i \le n} \varphi_i \to [\forall]\psi \qquad (\mathsf{4}_{[\forall]})$$

$$\Rightarrow \vdash_{\mathsf{S5}_{[\forall]}} \bigwedge_{i \le n} [\forall]\varphi_i \to [\forall]\psi \qquad (\mathsf{C}_{[\forall]})$$

where $(\mathsf{C}_{[\forall]})$ denotes $[\forall](\varphi \wedge \psi) \leftrightarrow ([\forall]\varphi \wedge [\forall]\psi)$, which is derivable from $\mathsf{S5}_{[\forall]}$. $\Box$

Lemma 3.18 (Existence lemma for $[\forall]$)**.** *Let Γ_0 be a HEK-MCS and $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k \in \mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k \in \mathcal{F}}, V^c)$ the canonical quasi-model for Γ_0. Then, for all $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$, we have*

$$\widehat{\varphi} \neq \varnothing \textit{ iff } \widehat{[\exists]\varphi} \neq \varnothing.$$

Proof.

(⇒:) Suppose $\widehat{\varphi} \neq \varnothing$. Then there is $\Delta \in \widehat{\varphi}$, i.e. $\Delta \in W^c$ and $\varphi \in \Delta$. Since $\varphi \to [\exists]\varphi$ is a theorem of $\mathsf{S5}_{[\forall]}$, we have by Lemma 3.9 and modus ponens that $[\exists]\varphi \in \Delta$. Therefore, $\widehat{[\exists]\varphi} \neq \varnothing$.

(⇐:) Suppose $\widehat{[\exists]\varphi} \neq \varnothing$, then there is $\Delta \in \widehat{[\exists]\varphi}$, i.e. $\Delta \in W^c$ s.t. $[\exists]\varphi \in \Delta$. Define $F := \{[\forall]\chi \mid [\forall]\chi \in \Delta\} \cup \{\varphi\}$.

We claim that F is consistent: towards a contradiction, suppose that F is inconsistent. Then there are $[\forall]\chi_1, ..., [\forall]\chi_n \in \Delta$ s.t. $\vdash_{\mathsf{HEK}} \bigwedge_{i \leq n} [\forall]\chi_i \to \neg\varphi$. By Lemma 3.17, it follows that $\vdash_{\mathsf{HEK}} \bigwedge_{i \leq n} [\forall]\chi_i \to [\forall]\neg\varphi$. Since $[\forall]\chi_1, ..., [\forall]\chi_n \in \Delta$, by Lemma 3.9 we have $\bigwedge_{i \leq n} [\forall]\chi_i \in \Delta$, so by Lemma 3.9 and modus ponens it follows that $[\forall]\neg\varphi \in \Delta$. But this contradicts $[\exists]\varphi \in \Delta$.

So F is consistent. By Lemma 3.10, there is a HEK-MCS Φ s.t. $F \subseteq \Phi$. We claim that $\Phi \in W^c$: we need to show that for any $[\forall]\chi \in \Gamma_0$, we have $\chi \in \Phi$, so let $[\forall]\chi \in \Gamma_0$. Since $[\forall]\chi \to [\forall][\forall]\chi$ is a theorem of $\mathsf{S5}_{[\forall]}$, it follows by Lemma 3.9 and modus ponens that $[\forall][\forall]\chi \in \Gamma_0$. By definition of W^c, this implies that $[\forall]\chi \in \Delta$, and by definition of F, we have $[\forall]\chi \in F \subseteq \Phi$. Since $[\forall]\chi \to \chi$ is a theorem of $\mathsf{S5}_{[\forall]}$, it follows by Lemma 3.9 and modus ponens that $\chi \in \Phi$, as required.

Thus, $\Phi \in W^c$. Since $\varphi \in F \subseteq \Phi$, we have $\widehat{\varphi} \neq \varnothing$.

□

Lemma 3.19. *For any* HEK-MCS Γ_0*, the canonical model* $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k \in \mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k \in \mathcal{F}}, V^c)$ *for* Γ_0 *is a quasi-model.*

Proof. Let Γ_0 be a HEK-MCS and $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k \in \mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k \in \mathcal{F}}, V^c)$ the canonical model for Γ_0. Then

(i) $W^c \neq \varnothing$ because by $(\mathsf{T}_{[\forall]})$, $\Gamma_0 \in W^c$;

(ii) $\mathcal{T}^c = \mathcal{P}(\text{PROP}) \neq \varnothing$ because PROP $\neq \varnothing$;

(iii) $\oplus^c = \cup$ is idempotent, commutative and associative;

(iv) $\mathcal{T}^c$ is closed under binary $\oplus^c$ because $\mathcal{P}(\text{PROP})$ is closed under unions;

(v) t^c is clearly well-defined; and moreover, we have $t^c(\varphi) = Var(\varphi) = \bigcup\{\{p\} \mid p \in Var(\varphi)\} = \bigoplus^c\{t^c(p) \mid p \in Var(\varphi)\}$;

(vi) for each $k \in \mathcal{F}$, $\mathcal{E}^c_{K_k} \subseteq \mathcal{E}^c_{\Box_k}$:

Let $e_a \in \mathcal{E}^c_{K_k}$, then $(e, a) = (\widehat{K_k\varphi}, Var(\varphi))$ for some $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]K_k\varphi \in \Gamma_0$. We need to show that $e_a \in \mathcal{E}^c_{\Box_k}$, i.e. that $(e, a) = (\widehat{\Box_k\psi}, Var(\psi))$ for some $\psi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\Box_k\psi \in \Gamma_0$. We establish the following two claims:

(1) $\widehat{K_k\varphi} = \widehat{\Box_k K_k\varphi}$:

(⊆:) $\widehat{K_k\varphi} \subseteq \widehat{K_k K_k\varphi} \subseteq \widehat{\Box_k K_k\varphi}$, where the first inclusion follows from $(\mathsf{4}_{K_k})$, and the second inclusion from (F_k), applying Lemma 3.9 throughout.

(⊇:) $\widehat{\Box_k K_k\varphi} \subseteq \widehat{K_k\varphi}$ by Axiom $(\mathsf{T}_{\Box_k})$ and Lemma 3.9.

(2) $[\exists]\Box_k K_k\varphi \in \Gamma_0$: Towards a contradiction, suppose $[\exists]\Box_k K_k\varphi \notin \Gamma_0$. Then, by Lemma 3.9, $\neg[\exists]\Box_k K_k\varphi \in \Gamma_0$. So $[\forall]\neg\Box_k K_k\varphi \in \Gamma_0$. We have $[\exists]K_k\varphi \in \Gamma_0$. Thus, $\widehat{[\exists]K_k\varphi} \neq \varnothing$, so by Lemma 3.18, $\widehat{K_k\varphi} \neq \varnothing$. Let $\Delta \in \widehat{K_k\varphi}$. Applying Lemma 3.9 throughout, we obtain by $(\mathsf{4}_{K_k})$ that $K_k K_k\varphi \in \Delta$. By (F_k), $\Box_k K_k\varphi \in \Delta$. Now, since $\Delta \in W^c$ and $[\forall]\neg\Box_k K_k\varphi \in \Gamma_0$, we have by definition of W^c that $\neg\Box_k K_k\varphi \in \Delta$, but this contradicts $\Box_k K_k\varphi \in \Delta$. So $[\exists]\Box_k K_k\varphi \in \Gamma_0$.

Since we also have $Var(K_k\varphi) = Var(\varphi)$, the two claims show that $K_k\varphi$ is the formula ψ we are looking for. That is, $(e,a) = (\widehat{K_k\varphi}, Var(\varphi)) = (\widehat{\Box_k K_k\varphi}, Var(K_k\varphi)) \in \mathcal{E}^c_{\Box_k}$.

(vii) for all $e_a \in \mathcal{E}^c_{\Box_k}$, we have $a \sqsubseteq^c t^c(k)$:

Let $e_a \in \mathcal{E}^c_{\Box_k}$, then (e,a) is of the form $(\widehat{\Box_k\varphi}, Var(\varphi))$ for some $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\Box_k\varphi \in \Gamma_0$. We claim that this implies $\bigwedge_{p\in Var(\varphi)}[\exists]\Box_k\overline{p} \in \Gamma_0$:

$$\begin{aligned} &\vdash_{\mathsf{HEK}} \Box_k\varphi \to \Box_k\overline{\varphi} && (\mathsf{Ax1}_{\Box_k})\\ \Rightarrow &\vdash_{\mathsf{HEK}} [\exists]\Box_k\varphi \to [\exists]\Box_k\overline{\varphi} && (\mathsf{S5}_{[\forall]})\\ \Rightarrow &\vdash_{\mathsf{HEK}} [\exists]\Box_k\varphi \to [\exists]\bigwedge_{p\in Var(\varphi)}\Box_k\overline{p} && \text{(by Proposition 3.3(i))}\\ \Rightarrow &\vdash_{\mathsf{HEK}} [\exists]\Box_k\varphi \to \bigwedge_{p\in Var(\varphi)}[\exists]\Box_k\overline{p} && (\mathsf{S5}_{[\forall]}) \end{aligned}$$

So by Lemma 3.9 it indeed follows that $\bigwedge_{p\in Var(\varphi)}[\exists]\Box_k\overline{p} \in \Gamma_0$. Therefore, applying Lemma 3.9 again, we have $Var(\varphi) \subseteq \{p \mid [\exists]\Box_k\overline{p} \in \Gamma_0\}$. So $a \sqsubseteq^c t^c(k)$.

(viii) for all $\heartsuit \in \{\Box, K\}$ and $e_a, f_b \in \mathcal{E}^c_{\heartsuit_k}$, $e\cap f \neq \varnothing$ implies $(e\cap f)_{a\oplus^c b} \in \mathcal{E}^c_{\heartsuit_k}$:
Let $e_a, f_b \in \mathcal{E}^c_{\heartsuit_k}$ and suppose that $e \cap f \neq \varnothing$. This means that

$(e,a) = (\widehat{\heartsuit_k\varphi}, Var(\varphi))$ for some $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\heartsuit_k\varphi \in \Gamma_0$ and
$(f,b) = (\widehat{\heartsuit_k\psi}, Var(\psi))$ for some $\psi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\heartsuit_k\psi \in \Gamma_0$.

Since $e \cap f \neq \varnothing$, we have $\widehat{\heartsuit_k\varphi} \cap \widehat{\heartsuit_k\psi} \neq \varnothing$, so there is $\Delta \in W^c$ s.t. $\heartsuit_k\varphi, \heartsuit_k\psi \in \Delta$. By Lemma 3.9, this implies that $\heartsuit_k\varphi \wedge \heartsuit_k\psi \in \Delta$, and by ($\mathsf{C}_{\heartsuit_k}$), we obtain $\heartsuit_k(\varphi \wedge \psi) \in \Delta$.

We claim that $[\exists]\heartsuit_k(\varphi \wedge \psi) \in \Gamma_0$. Towards a contradiction, suppose $[\exists]\heartsuit_k(\varphi \wedge \psi) \notin \Gamma_0$; then by Lemma 3.9, $\neg[\exists]\heartsuit_k(\varphi \wedge \psi) \in \Gamma_0$. This means that $[\forall]\neg\heartsuit_k(\varphi\wedge\psi) \in \Gamma_0$. Since $\Delta \in W^c$, we have $\neg\heartsuit_k(\varphi\wedge\psi) \in \Delta$, which contradicts $\heartsuit_k(\varphi \wedge \psi) \in \Delta$.

So $[\exists]\heartsuit_k(\varphi \wedge \psi) \in \Gamma_0$. Thus, by definition of $\mathcal{E}^c_{\heartsuit_k}$, we have

$$\begin{aligned}
\mathcal{E}^c_{\heartsuit_k} &\ni (\widehat{\heartsuit_k(\varphi \wedge \psi)}, Var(\varphi \wedge \psi)) \\
&= (\widehat{\heartsuit_k\varphi \wedge \heartsuit_k\psi}, Var(\varphi) \cup Var(\psi)) && \text{(by } (\mathsf{C}_{\heartsuit_k})) \\
&= (\widehat{\heartsuit_k\varphi} \cap \widehat{\heartsuit_k\psi}, Var(\varphi) \cup Var(\psi)) && \text{(by Lemma 3.9)} \\
&= (e \cap f, a \oplus^c b).
\end{aligned}$$

(ix) for all $e_a \in \mathcal{E}^c_{K_k}$ and $f_b \in \mathcal{E}^c_{\Box_k}$, $e \cap f \neq \varnothing$:

Let $e_a \in \mathcal{E}_{K_k}$ and $f_b \in \mathcal{E}_{\Box_k}$. This means that

$(e, a) = (\widehat{K_k\varphi}, Var(\varphi))$ for some $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]K_k\varphi \in \Gamma_0$ and

$(f, b) = (\widehat{\Box_k\psi}, Var(\psi))$ for some $\psi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\Box_k\psi \in \Gamma_0$.

Since $[\exists]K_k\varphi \in \Gamma_0$ and $[\exists]\Box_k\psi \in \Gamma_0$, by Lemma 3.9 it follows that $[\exists]K_k\varphi \wedge [\exists]\Box_k\psi \in \Gamma_0$. By Proposition 3.3(iii), modus ponens and Lemma 3.9, it follows that $[\exists](K_k\varphi\wedge\Box_k\psi) \in \Gamma_0$. Thus, $\widehat{[\exists](K_k\varphi \wedge \Box_k\psi)} \neq \varnothing$. By Lemma 3.18, we obtain $\widehat{K_k\varphi \wedge \Box_k\psi} \neq \varnothing$. By Lemma 3.9, we have $\widehat{K_k\varphi \wedge \Box_k\psi} = \widehat{K_k\varphi} \cap \widehat{\Box_k\psi} = e \cap f$. Therefore, $e \cap f \neq \varnothing$.

Therefore, $\mathfrak{Q}^c$ is indeed a quasi-model. □

Lemma 3.20 (Truth Lemma). *Let Γ_0 be a* HEK-MCS *and $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k\in\mathcal{F}}, V^c)$ the canonical quasi-model for Γ_0. Then, for all $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ and all $\Gamma \in W^c$, we have*

$$\mathfrak{Q}^c, \Gamma \Vdash^* \varphi \textit{ iff } \varphi \in \Gamma.$$

Proof. By induction on the complexity of φ. The Boolean cases are elementary and omitted. The cases for $\varphi = \heartsuit_k\psi$ are identical for $\heartsuit \in \{\Box, K\}$ and $k \in \mathcal{F}$, so we prove them in one go.

– Case $\varphi = [\forall]\psi$:

(⇒): Suppose $\mathfrak{Q}^c, \Gamma \Vdash^* [\forall]\psi$. Then $W^c \subseteq [\psi]$. Towards a contradiction, suppose $[\forall]\psi \notin \Gamma$. Then, by Lemma 3.9, $\neg[\forall]\psi \in \Gamma$, which means that $[\exists]\neg\psi \in \Gamma$. Thus, $\widehat{[\exists]\neg\psi} \neq \varnothing$. By Lemma 3.18, it follows that $\widehat{\neg\psi} \neq \varnothing$, so there is $\Delta \in W^c$ s.t. $\neg\psi \in \Delta$. By Lemma 3.9, this means that $\psi \notin \Delta$. By the induction hypothesis, we have $\mathfrak{Q}^c, \Delta \nVdash^* \psi$, but this contradicts $W^c \subseteq [\psi]$. Therefore, $[\forall]\psi \in \Gamma$.

($\Leftarrow$): Suppose $[\forall]\psi \in \Gamma$. By definition of W^c and Lemma 3.15, we have $W^c \subseteq \widehat{\psi}$. By the induction hypothesis, it follows that $W^c \subseteq [\psi]$, so $\mathfrak{Q}^c, \Gamma \Vdash^* [\forall]\psi$.

– Case $\varphi = \heartsuit_k\psi$, where $\heartsuit \in \{\Box, K\}$ and $k \in \mathcal{F}$:

($\Rightarrow$): Suppose $\mathfrak{Q}^c, \Gamma \Vdash^* \heartsuit_k\psi$. Then there is $e_a \in \mathcal{E}^c_{\heartsuit_k}$ s.t. $\Gamma \in e \subseteq [\psi]$ and $t^c(\psi) \sqsubseteq^c a$. This means that there is a formula $\chi \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\exists]\heartsuit_k\chi \in \Gamma_0$, $\Gamma \in \widehat{\heartsuit_k\chi} \subseteq [\psi]$, and $Var(\psi) \subseteq Var(\chi)$. Since we have $[\exists]\heartsuit_k\chi \in \Gamma_0$ and $[\exists]\heartsuit_k\chi \to [\forall][\exists]\heartsuit_k\chi$ is a theorem of $\mathsf{S5}_{[\forall]}$, it follows by Lemma 3.9 that $[\forall][\exists]\heartsuit_k\chi \in \Gamma_0$. Thus, by definition of W^c, we have $[\exists]\heartsuit_k\chi \in \Gamma$.

We proceed by establishing the following claims:

(1) $[\forall](\heartsuit_k\chi \to \psi) \in \Gamma$.
Proof. Define $D := \{[\forall]\xi \mid [\forall]\xi \in \Gamma\} \cup \{\heartsuit_k\chi \wedge \neg\psi\}$.
We claim that D is inconsistent: towards a contradiction, suppose that D is consistent. Then, by Lemma 3.10, there is a HEK-MCS Δ s.t. $D \subseteq \Delta$. Since $[\forall]\xi \in \Delta$ for all ξ s.t. $[\forall]\xi \in \Gamma$, by Lemma 3.15 we have $\Delta \in W^c$. Now $\heartsuit_k\chi \in \Delta$, so by $\widehat{\heartsuit_k\chi} \subseteq [\psi]$, we have $\mathfrak{Q}^c, \Delta \Vdash^* \psi$. Thus, by the induction hypothesis, we have $\psi \in \Delta$, which contradicts $\neg\psi \in \Delta$.
So D is inconsistent. This means that $\{[\forall]\xi \mid [\forall]\xi \in \Gamma\} \vdash_{\mathsf{HEK}} \heartsuit_k\chi \to \psi$. Thus, there are $\xi_1, ..., \xi_n \in \mathcal{L}_{[\forall]\Box_k K_k}$ s.t. $[\forall]\xi_1, ..., [\forall]\xi_n \in \Gamma$ and $\vdash_{\mathsf{HEK}} \bigwedge_{i\leq n}[\forall]\xi_i \to (\heartsuit_k\chi \to \psi)$. By Lemma 3.17, we have $\vdash_{\mathsf{HEK}} \bigwedge_{i\leq n}[\forall]\xi_i \to [\forall](\heartsuit_k\chi \to \psi)$. Since $[\forall]\xi_1, ..., [\forall]\xi_n \in \Gamma$, by Lemma 3.9 we have $\bigwedge_{i\leq n}[\forall]\xi_i \in \Gamma$. Therefore, by modus ponens, it follows that $[\forall](\heartsuit_k\chi \to \psi) \in \Gamma$.

(2) $\heartsuit_k\heartsuit_k\chi \in \Gamma$.
Proof. Since $\Gamma \in \widehat{\heartsuit_k\chi}$, we have $\heartsuit_k\chi \in \Gamma$. By ($\mathsf{4}_{\heartsuit_k}$), modus ponens and Lemma 3.9, it follows that $\heartsuit_k\heartsuit_k\chi \in \Gamma$.

(3) $\heartsuit_k\overline{\psi} \in \Gamma$.
Proof. Since $\heartsuit_k\chi \in \Gamma$, it follows from ($\mathsf{Ax1}_{\heartsuit_k}$), Lemma 3.9 and modus ponens that $\heartsuit_k\overline{\chi} \in \Gamma$. As $Var(\psi) \subseteq Var(\chi)$, by Proposition 3.3(ii) we have $\vdash_{\mathsf{HEK}} \heartsuit_k\overline{\chi} \to \heartsuit_k\overline{\psi}$, so using Lemma 3.9 and modus ponens again, we obtain $\heartsuit_k\overline{\psi} \in \Gamma$.

Now, by ($\mathsf{Ax2}_{\heartsuit_k}$) and Lemma 3.9, $([\forall](\heartsuit_k\chi \to \psi) \wedge \heartsuit_k\heartsuit_k\chi \wedge$

$\heartsuit_k\overline{\psi}) \to \heartsuit_k\psi \in \Gamma$. From claims (1) through (3) above together with modus ponens, it follows that $\heartsuit_k\psi \in \Gamma$.

(⇐:) Suppose $\heartsuit_k\psi \in \Gamma$. Define $(e, a) := (\widehat{\heartsuit_k\psi}, Var(\psi))$. We proceed by establishing the following claims:

(1) $(e, a) \in \mathcal{E}^c_{\heartsuit_k}$.
Proof. We have $\heartsuit_k\psi \in \Gamma$. We claim that $[\exists]\heartsuit_k\psi \in \Gamma_0$: towards a contradiction, suppose otherwise. Then, by Lemma 3.9, $\neg[\exists]\heartsuit_k\psi \in \Gamma_0$, which means $[\forall]\neg\heartsuit_k\psi \in \Gamma_0$. Since $\Gamma_0 \sim_{[\forall]} \Gamma$, it follows that $\neg\heartsuit_k\psi \in \Gamma$, contradicting the consistency of Γ. So $[\exists]\heartsuit_k\psi \in \Gamma_0$. Therefore, by definition of $\mathcal{E}^c_{\heartsuit_k}$, $(e, a) = (\widehat{\heartsuit_k\psi}, Var(\psi)) \in \mathcal{E}^c_{\heartsuit_k}$.

(2) $\Gamma \in e$.
Proof. Follows immediately from $\heartsuit_k\psi \in \Gamma$.

(3) $e \subseteq [\psi]$.
Proof. Let $\Gamma \in e = \widehat{\heartsuit_k\psi}$, then $\heartsuit_k\psi \in \Gamma$. From (T_{K_k}), Lemma 3.9 and modus ponens it follows that $\psi \in \Gamma$. By the induction hypothesis, we have $\mathfrak{Q}^C, \Gamma \Vdash^* \psi$, so $\Gamma \in [\psi]$.

(4) $t^c(\psi) \sqsubseteq^c a$.
Proof. $t^c(\psi) = Var(\psi) = a$.

Combining claims (1) through (4), we obtain that $(e, a) \in \mathcal{E}^c_{\heartsuit_k}$ is a witness for $\mathfrak{Q}^c, \Gamma \Vdash^* \heartsuit_k\psi$.

□

Step 2: Equivalence of quasi-models and tsef-models

We have now established that every HEK-consistent set of formulas is satisfiable on the canonical quasi-model. We now show that for every quasi-model, there exists a modally equivalent tsef-model. This step requires a construction that transforms quasi-models into tsef-models, which is given in Definition 3.21.

In order to build a tsef-model from a quasi-model $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$, we need to "fuse" the two evidence sets $\mathcal{E}_{\Box_k}$ and $\mathcal{E}_{K_k}$ into one for each $k \in \mathcal{F}$ in a way that preserves satisfaction of formulas of the form $\Box_k\varphi$ and $K_k\varphi$. Each piece of evidence in $\mathcal{E}_{K_k}$ is consistent with all evidence in $\mathcal{E}_{\Box_k}$ by definition of quasi-models, so intuitively, each piece of evidence in $\mathcal{E}_{K_k}$ potentially witnesses knowledge of some formula at some world. However, there may also be pieces of evidence that are in $\mathcal{E}_{\Box_k}$ but not in $\mathcal{E}_{K_k}$ that are consistent with all evidence in $\mathcal{E}_{\Box_k}$. In order to guarantee that the initial quasi-model and the tsef-model constructed from it are indeed equivalent, we need to ensure that none of these evidence pieces give rise to knowledge in the translated tsef-model. We achieve this by creating two copies of the original quasi-model and additionally copies of the evidence in $\mathcal{E}_{\Box_k}$ and $\mathcal{E}_{K_k}$ in such a way that all and only the evidence from $\mathcal{E}_{K_k}$ are consistent with *all* evidence in the constructed tsef-model.

Definition 3.21 (Translation of quasi-models to tsef-models)**.** Given a quasi-model $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$, we construct a tsef-model $\mathfrak{M}^{\mathfrak{Q}} := (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k\in\mathcal{F}}, V^{\mathfrak{Q}})$ as follows:

$$\begin{aligned}
W^{\mathfrak{Q}} &:= \{0,1\} \times W;\\
\mathcal{T}^{\mathfrak{Q}} &:= \mathcal{T};\\
\oplus^{\mathfrak{Q}} &:= \oplus;\\
t^{\mathfrak{Q}} &:= t;\\
\mathcal{E}_k^{\mathfrak{Q}} &:= \quad \{(\{0\} \times e, a) \mid e_a \in \mathcal{E}_{\Box_k}\}\\
&\quad\ \cup \{(\{1\} \times e, a) \mid e_a \in \mathcal{E}_{\Box_k}\}\\
&\quad\ \cup \{(\{0,1\} \times e, a) \mid e_a \in \mathcal{E}_{K_k}\};\\
V^{\mathfrak{Q}} &: \quad \text{PROP} \to W^{\mathfrak{Q}} \text{ such that } V^{\mathfrak{Q}}(p) := \{0,1\} \times V(p).
\end{aligned}$$

Lemma 3.22. *For any quasi-model* $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$*, the corresponding structure* $\mathfrak{M}^{\mathfrak{Q}} := (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k\in\mathcal{F}}, V^{\mathfrak{Q}})$ *is a tsef-model.*

Proof. Let $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ be a quasi-model and $\mathfrak{M}^{\mathfrak{Q}} = (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k\in\mathcal{F}}, V^{\mathfrak{Q}})$ its corresponding tsef-model as given in Definition 3.21. Since $\mathcal{T}^{\mathfrak{Q}} = \mathcal{T}$, $\oplus^{\mathfrak{Q}} = \oplus$ and $t^{\mathfrak{Q}} = t$, it suffices to check the following conditions:

(i) $W^{\mathfrak{Q}} \neq \varnothing$: follows from $W \neq \varnothing$;

(ii) for all $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$, we have $a \sqsubseteq^{\mathfrak{Q}} t^{\mathfrak{Q}}(k)$:

follows from the definition of $\mathcal{E}_{\Box_k}$ and the fact that $\sqsubseteq^{\mathfrak{Q}} = \sqsubseteq$;

(iii) $\mathcal{E}_k^{\mathfrak{Q}}$ is closed under non-empty finite evidence combination, i.e. for all $e_{1,a_1}, ..., e_{n,a_n} \in \mathcal{E}_k^{\mathfrak{Q}}$, $e_1 \cap \cdots \cap e_n \neq \varnothing$ implies $e_{1,a_1} \sqcap \cdots \sqcap e_{n,a_n} \in \mathcal{E}_k^{\mathfrak{Q}}$.

Proof. We show that if $e_a, e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$ and $e \cap e' \neq \varnothing$, then $(e \cap e')_{a \oplus^{\mathfrak{Q}} a'} \in \mathcal{E}_k^{\mathfrak{Q}}$. After establishing this, a simple inductive argument suffices to prove the desired result.

Let $e_a, e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$ and suppose that $e \cap e' \neq \varnothing$. Since $e \cap e' \neq \varnothing$ and without loss of generality, exactly one of the three following cases holds for some $i \in \{0,1\}$:

$$\begin{aligned}
(1) \qquad (e, a) &= (\{i\} \times f, a) && \text{for some } f_a \in \mathcal{E}_{\Box_k} \text{ and}\\
(e', a') &= (\{i\} \times f', a') && \text{for some } f'_{a'} \in \mathcal{E}_{\Box_k};
\end{aligned}$$

(2) $(e,a) = (\{i\} \times f, a)$ for some $f_a \in \mathcal{E}_{\Box_k}$ and
$(e',a') = (\{0,1\} \times f', a')$ for some $f'_{a'} \in \mathcal{E}_{K_k}$;

(3) $(e,a) = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$;
$(e',a') = (\{0,1\} \times f', a')$ for some $f'_{a'} \in \mathcal{E}_{K_k}$.

In the first case, since $e \cap e' \neq \varnothing$, we have $f \cap f' \neq \varnothing$. Since $f_a, f'_{a'} \in \mathcal{E}_{\Box_k}$, it follows that $(f \cap f', a \oplus a') \in \mathcal{E}_{\Box_k}$. Therefore,

$$\begin{aligned}(e \cap e', a \oplus^{\mathfrak{Q}} a') &= ((\{i\} \times f) \cap (\{i\} \times f'), a \oplus^{\mathfrak{Q}} a') \\ &= (\{i\} \times (f \cap f'), a \oplus^{\mathfrak{Q}} a') \in \mathcal{E}_k^{\mathfrak{Q}}.\end{aligned}$$

In the second case, since $f'_{a'} \in \mathcal{E}_{K_k}$, we have $f \cap f' \neq \varnothing$. Since $f'_{a'} \in \mathcal{E}_{K_k} \subseteq \mathcal{E}_{\Box_k}$ and $\mathcal{E}_{\Box_k}$ is closed under non-empty evidence combination, we have $(f \cap f', a \oplus a') \in \mathcal{E}_{\Box_k}$. Then

$$\begin{aligned}(e \cap e', a \oplus^{\mathfrak{Q}} a') &= (\{i\} \times f \cap (\{0,1\} \times f'), a \oplus^{\mathfrak{Q}} a') \\ &= (\{i\} \times (f \cap f'), a \oplus^{\mathfrak{Q}} a') \in \mathcal{E}_k^{\mathfrak{Q}},\end{aligned}$$

where the last step follows from the definition of $\mathcal{E}_k^{\mathfrak{Q}}$.

In the third case, since $e \cap e' \neq \varnothing$, we have $f \cap f' \neq \varnothing$. Moreover, since $f_a, f'_{a'} \in \mathcal{E}_{K_k}$, we have $(f \cap f', a \oplus a') \in \mathcal{E}_{K_k}$. Now

$$\begin{aligned}(e \cap e', a \oplus^{\mathfrak{Q}} a') &= ((\{0,1\} \times f) \cap (\{0,1\} \times f'), a \oplus^{\mathfrak{Q}} a') \\ &= (\{0,1\} \times (f \cap f'), a \oplus^{\mathfrak{Q}} a') \in \mathcal{E}_k^{\mathfrak{Q}},\end{aligned}$$

where the last step follows from the definition of $\mathcal{E}_k^{\mathfrak{Q}}$.

So $\mathcal{E}_k^{\mathfrak{Q}}$ is closed under non-empty finite evidence combination.

Therefore, $\mathfrak{M}^{\mathfrak{Q}}$ is indeed a tsef-model. □

Lemma 3.23. *If $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k \in \mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k \in \mathcal{F}}, V)$ is a quasi-model and $\mathfrak{M}^{\mathfrak{Q}} = (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k \in \mathcal{F}}, V^{\mathfrak{Q}})$ its corresponding tsef-model as given in Definition 3.21, then the following holds for every piece of evidence $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$:*

$$\begin{aligned}&e \cap e' \neq \varnothing \textit{ for all } e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}} \\ &\textit{iff } (e,a) \textit{ is of the form } (\{0,1\} \times f, a) \textit{ for some } f_a \in \mathcal{E}_{K_k}.\end{aligned}$$

Proof. If $\mathcal{E}_{\Box_k} = \varnothing$, then $\mathcal{E}_k^{\mathfrak{Q}} = \varnothing$, and the lemma holds vacuously. Thus, assume that $\mathcal{E}_{\Box_k} \neq \varnothing$.

(⇒:) Let $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$, and suppose that $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$. By assumption, we have $\mathcal{E}_{\Box_k} \neq \varnothing$, so let $f'_{b'} \in \mathcal{E}_{\Box_k}$. By definition of $\mathcal{E}_k^{\mathfrak{Q}}$, we have $(\{0\} \times f', b) \in \mathcal{E}_k^{\mathfrak{Q}}$ and $(\{1\} \times f', b) \in \mathcal{E}_k^{\mathfrak{Q}}$. Now, since by assumption $e \cap (\{0\} \times f') \neq \varnothing$ and $e \cap (\{1\} \times f') \neq \varnothing$, it follows from the definition of $\mathcal{E}_{K_k}$ that (e,a) must be of the form $(\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$, as required.

(⇐:) Let $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$, and suppose that $(e, a) = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$.

Since $f_a \in \mathcal{E}_{K_k}$, we have by definition of $\mathcal{E}_{K_k}$ that $f \cap f' \neq \varnothing$ for all $f'_{a'} \in \mathcal{E}_{\Box_k}$. Thus, $\{0\} \times (f \cap f') \neq \varnothing$ and $\{1\} \times (f \cap f') \neq \varnothing$ for all $f'_{a'} \in \mathcal{E}_{\Box_k}$. It is easy to see that this implies $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$, as required.

□

Lemma 3.24. *For any quasi-model* $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ *and its corresponding tsef-model* $\mathfrak{M}^{\mathfrak{Q}} = (W^{\mathfrak{Q}}, V^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k\in\mathcal{F}})$ *as given in Definition 3.21, and for any formula* $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ *and any* $w \in W$*, we have*

$$\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash \varphi \text{ iff } \mathfrak{M}^{\mathfrak{Q}}, (1, w) \Vdash \varphi.$$

Proof. Let $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ be a quasi-model and $\mathfrak{M}^{\mathfrak{Q}} = (W^{\mathfrak{Q}}, V^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}_k^{\mathfrak{Q}}\}_{k\in\mathcal{F}})$ its corresponding tsef-model as given in Definition 3.21. The proof proceeds by induction on the complexity of φ. The cases for the Boolean connectives $\neg$, $\wedge$ and the global modality $[\forall]$ are elementary, so we only present the cases for $\varphi = \Box_k\psi$ and $\varphi = K_k\psi$.

- Case $\varphi = \Box_k\psi$:

  (⇒:) Suppose $\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash \Box_k\psi$. Then there is $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$ s.t. $(0, w) \in e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$ and $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$. Since $(0, w) \in e$ and by definition of $\mathcal{E}_k^{\mathfrak{Q}}$, either one of the following two cases holds:

  (1) $(e, a) = (\{0\} \times f, a)$ for some $f_a \in \mathcal{E}_{\Box_k}$; or
  (2) $(e, a) = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$.

  Case (1): Since $\{0\} \times f = e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$, by the induction hypothesis, we have $\{1\} \times f \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$. From $f_a \in \mathcal{E}_{\Box_k}$ and the definition of $\mathcal{E}_k^{\mathfrak{Q}}$, it follows that $(\{1\} \times f, a) \in \mathcal{E}_k^{\mathfrak{Q}}$. As $(0, w) \in e$, we have $w \in f$, so $(1, w) \in \{1\} \times f$. Moreover, we have $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$. Therefore, $(\{1\} \times f, a) \in \mathcal{E}_k^{\mathfrak{Q}}$ is a witness for $\mathfrak{M}^{\mathfrak{Q}}, (1, w) \Vdash \Box_k\psi$.

  Case (2): Since $(0, w) \in e$, we have $w \in f$, so $(1, w) \in \{1\} \times f \subseteq e$. As $e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$ and $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$, $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$ is a witness for $\mathfrak{M}^{\mathfrak{Q}}, (1, w) \Vdash \Box_k\psi$.

  (⇐:) Similar.

- Case $\varphi = K_k\psi$:

  (⇒:) Suppose $\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash K_k\psi$. Then there is $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$ s.t. $(0, w) \in e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$, $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$ and $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$.

By Lemma 3.23, we have that $(e, a) = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$. Since $(0, w) \in e$, we have $w \in f$, so $(1, w) \in e$. We also have $e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$, $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$ and $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}^{\mathfrak{Q}}_k$. Therefore, $e_a \in \mathcal{E}^{\mathfrak{Q}}_k$ is a witness for $\mathfrak{M}^{\mathfrak{Q}}, (1, w) \Vdash K_k\psi$.

($\Leftarrow$:) Similar.

$\square$

Lemma 3.25. *For any quasi-model* $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ *and its corresponding tsef-model* $\mathfrak{M}^{\mathfrak{Q}} := (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}^{\mathfrak{Q}}_k\}_{k\in\mathcal{F}}, V^{\mathfrak{Q}})$ *as given in Definition 3.21, and for any formula* $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ *and* $w \in W$*, we have*

$$\mathfrak{Q}, w \Vdash^* \varphi \text{ iff } \mathfrak{M}^{\mathfrak{Q}}, (i, w) \Vdash \varphi \text{ for any } i \in \{0, 1\}.$$

Proof. Let $\mathfrak{Q} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}_{K_k}\}_{k\in\mathcal{F}}, V)$ be a quasi-model and $\mathfrak{M}^{\mathfrak{Q}} := (W^{\mathfrak{Q}}, \mathcal{T}^{\mathfrak{Q}}, \oplus^{\mathfrak{Q}}, t^{\mathfrak{Q}}, \{\mathcal{E}^{\mathfrak{Q}}_k\}_{k\in\mathcal{F}}, V^{\mathfrak{Q}})$ its corresponding tsef-model as given in Definition 3.21. By Lemma 3.24, it suffices to show that for any formula $\varphi \in \mathcal{L}_{[\forall]\Box_k K_k}$ and any $w \in W^{\mathfrak{Q}}$, $\mathfrak{Q}, w \Vdash^* \varphi$ iff $\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash \varphi$. The proof proceeds by induction on the complexity of φ. The cases for the Boolean connectives $\neg$, $\wedge$ and the global modality $[\forall]$ are elementary, so we only present the cases for $\varphi = \Box_k\psi$ and $\varphi = K_k\psi$.

– Case $\varphi = \Box_k\psi$:

($\Rightarrow$:) Suppose $\mathfrak{Q}, w \Vdash^* \Box_k\psi$, then there is $e_a \in \mathcal{E}_{\Box_k}$ s.t. $w \in e \subseteq [\psi]^{\mathfrak{Q}}$ and $t^{\mathfrak{Q}}(\psi) \sqsubseteq^{\mathfrak{Q}} a$. By definition of $\mathcal{E}^{\mathfrak{Q}}_k$, $(\{0\}\times e, a) \in \mathcal{E}^{\mathfrak{Q}}_k$. As $w \in e$, we have $(0, w) \in \{0\} \times e$. From $e \subseteq [\psi]^{\mathfrak{Q}}$ and the induction hypothesis, it follows that $\{0\} \times e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$. Moreover, since $t^{\mathfrak{Q}} = t^{\mathfrak{M}^{\mathfrak{Q}}}$ and $\sqsubseteq^{\mathfrak{Q}} = \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}}$, we have $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$. Therefore, $(\{0\} \times e, a) \in \mathcal{E}^{\mathfrak{Q}}_{\Box_k}$ is a witness for $\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash \Box_k\psi$.

($\Leftarrow$:) Suppose $\mathfrak{M}^{\mathfrak{Q}}, (0, w) \Vdash \Box_k\psi$. Then there is $e_a \in \mathcal{E}^{\mathfrak{Q}}_k$ s.t. $(0, w) \in e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$ and $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$. Since $(0, w) \in e$ and by definition of $\mathcal{E}^{\mathfrak{Q}}_k$, either one of the following two cases holds:

(1) $e_a = (\{0\} \times f, a)$ for some $f_a \in \mathcal{E}_{\Box_k}$; or
(2) $e_a = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k} \subseteq \mathcal{E}_{\Box_k}$.

In either case, we have $w \in f$ for some $f_a \in \mathcal{E}_{\Box_k}$. Moreover, since $e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$ and by the induction hypothesis, we have in both cases that $f \subseteq [\psi]^{\mathfrak{Q}}$. Finally, since $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$, we have $t^{\mathfrak{Q}}(\psi) \sqsubseteq^{\mathfrak{Q}} a$. Therefore, $f_a \in \mathcal{E}_{\Box_k}$ is a witness for $\mathfrak{Q}, w \Vdash^* \Box_k\psi$.

– Case $\varphi = K_k\psi$:

($\Rightarrow$): Suppose $\mathfrak{Q}, w \Vdash^* K_k\psi$, then there is $e_a \in \mathcal{E}_{K_k}$ s.t. $w \in e \subseteq [\psi]^{\mathfrak{Q}}$ and $t^{\mathfrak{Q}}(\psi) \sqsubseteq^{\mathfrak{Q}} a$. Define $f_a := (\{0,1\} \times e, a)$. We claim that f_a is a witness for $\mathfrak{M}^{\mathfrak{Q}}, (0,w) \Vdash K_k\psi$. By definition of $\mathcal{E}_k^{\mathfrak{Q}}$, $f_a \in \mathcal{E}_k^{\mathfrak{Q}}$. Since $w \in e$, $(0,w) \in f$. Moreover, from $e \subseteq [\psi]^{\mathfrak{Q}}$ and the induction hypothesis it follows that $f \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$. Since $t^{\mathfrak{Q}}(\psi) \sqsubseteq^{\mathfrak{Q}} a$, we have that $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$. It remains to show that $f \cap f' \neq \varnothing$ for all $f'_{b'} \in \mathcal{E}_k^{\mathfrak{Q}}$, which follows from Lemma 3.23. Therefore, $\mathfrak{M}^{\mathfrak{Q}}, (0,w) \Vdash K_k\psi$.

($\Leftarrow$): Suppose $\mathfrak{M}^{\mathfrak{Q}}, (0,w) \Vdash K_k\psi$, then there is $e_a \in \mathcal{E}_k^{\mathfrak{Q}}$ s.t. $(0,w) \in e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$, $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$, and $e \cap e' \neq \varnothing$ for all $e'_{a'} \in \mathcal{E}_k^{\mathfrak{Q}}$. By Lemma 3.23, we have that $e_a = (\{0,1\} \times f, a)$ for some $f_a \in \mathcal{E}_{K_k}$. As $(0,w) \in e$, we have $w \in f$. Moreover, since $e \subseteq [\![\psi]\!]^{\mathfrak{M}^{\mathfrak{Q}}}$, by the induction hypothesis we have that $f \subseteq [\psi]^{\mathfrak{Q}}$. Finally, since $t^{\mathfrak{M}^{\mathfrak{Q}}}(\psi) \sqsubseteq^{\mathfrak{M}^{\mathfrak{Q}}} a$, we have $t^{\mathfrak{Q}}(\psi) \sqsubseteq^{\mathfrak{Q}} a$. Therefore, $f_a \in \mathcal{E}_{K_k}$ is a witness for $\mathfrak{Q}, w \Vdash^* K_k\psi$.

□

Theorem 3.7 (Completeness of HEK)**.** *The logic* HEK *of hyperintensional evidence and knowledge is strongly complete with respect to the class of tsef-models. That is, for any set of formulas* $\Gamma \cup \{\varphi\}$ *of the language* $\mathcal{L}_{[\forall]\Box_k K_k}$*, if* $\Gamma \Vdash \varphi$*, then* φ *is derivable in* HEK *from* Γ*.*

Proof. By Proposition 3.11, it suffices to show that every HEK-consistent set of formulas is satisfiable on some tsef-model. Let G be a HEK-consistent set of formulas. By Lemma 3.10, there is a HEK-MCS Γ s.t. $G \subseteq \Gamma$. Let $\mathfrak{Q}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_{\Box_k}\}_{k\in\mathcal{F}}, \{\mathcal{E}^c_{K_k}\}_{k\in\mathcal{F}}, V^c)$ be the canonical quasi-model for Γ as defined in Definition 3.16. By definition of $\mathfrak{Q}^c$, we have $\Gamma \in W^c$, and by Lemma 3.20, we have $\mathfrak{Q}^c, \Gamma \Vdash^* G$. Now we use the construction in Definition 3.21 to build a structure $\mathfrak{M}^{\mathfrak{Q}^c}$ from the canonical quasi-model $\mathfrak{Q}^c$ for Γ. By Lemma 3.22, $\mathfrak{M}^{\mathfrak{Q}^c}$ is a tsef-model, and by Lemma 3.25, we have $\mathfrak{M}^{\mathfrak{Q}^c}, (0,\Gamma) \Vdash G$. So G is satisfiable on a tsef-model, as required. □

3.2 The belief fragment $\mathcal{L}_{[\forall]B_k}$

In this section, we provide a separate sound and complete axiomatisation for the belief fragment of our language.

3.2.1 Syntax

Definition 3.26 (Syntax of $\mathcal{L}_{[\forall]B_k}$)**.**

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\forall]\varphi \mid B_k\varphi \quad \text{for } k \in \mathcal{F}$$

3.2.2 Semantics

Definition 3.27 (Semantics for $\mathcal{L}_{[\forall]B_k}$ on tsef-models)**.** Given a tsef-model $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$ and a world $w \in W$, the $\models$-semantics for $\mathcal{L}_{[\forall]B_k}$ is defined recursively as follows (we write $[\![\varphi]\!]^{\mathfrak{M}}$ for the truth set $\{w \in W \mid \mathfrak{M}, w \models \varphi\}$ of φ in $\mathfrak{M}$ and shorten $[\![\varphi]\!]^{\mathfrak{M}}$ to $[\![\varphi]\!]$ when $\mathfrak{M}$ is clear from the context):

$$
\begin{array}{lll}
\mathfrak{M}, w \models p & \text{iff} & p \in V(p) \\
\mathfrak{M}, w \models \neg\varphi & \text{iff} & \text{not } \mathfrak{M}, w \models \varphi \\
\mathfrak{M}, w \models \varphi \wedge \psi & \text{iff} & \mathfrak{M}, w \models \varphi \text{ and } \mathfrak{M}, w \models \psi \\
\mathfrak{M}, w \models [\forall]\varphi & \text{iff} & W \subseteq [\![\varphi]\!] \\
\mathfrak{M}, x \models B_k\varphi & \text{iff} & \exists e_a \in \mathcal{E}_k \ (e \subseteq [\![\varphi]\!] \text{ and } t(\varphi) \sqsubseteq a \\
 & & \qquad \text{and } \forall e' \in \mathcal{E}^k : \ e \cap e' \neq \varnothing)
\end{array}
$$

The notions of *satisfaction of a set of formulas*, *validity of a formula*, and *logical semantic consequence* are defined analogously to Definition 3.2.

3.2.3 Expressivity

The modality B for *fragmented belief* is definable in $\mathcal{L}_{[\forall]B_k}$ as

$$B\varphi := \bigvee_{k\in\mathcal{F}} B_k\varphi.$$

3.2.4 Axiomatisation

Table 3.2 provides a sound and complete axiomatisation of the logic HB of hyperintensional belief over $\mathcal{L}_{[\forall]B_k}$. The proofs for soundness and completeness can be found in Proposition 3.29 and Theorem 3.30, respectively.

Proposition 3.28. *The following are derivable in* HB*:*

(i) $B_k\overline{\varphi} \leftrightarrow \bigwedge_{p\in Var(\varphi)} B_k\overline{p}$

(ii) $B_k\overline{\varphi} \to B_k\overline{\psi}$*, if* $Var(\psi) \subseteq Var(\varphi)$

Proof. Same as in Proposition 3.3, items (i) and (ii). □

The notions of *derivability* of a formula φ in HB, written $\vdash_{\mathsf{HB}}$, and of *consistency* of a (set of) formula(s) of $\mathcal{L}_{[\forall]B_k}$ is defined as previously in Definition 3.4.

(I)	(CPL)	all classical propositional tautologies and modus ponens
(II)	(S5) **axioms and rules for** $[\forall]$**:**	
	$(\mathsf{K}_{[\forall]})$	$([\forall]\varphi \wedge [\forall](\varphi \to \psi)) \to [\forall]\psi$
	$(\mathsf{T}_{[\forall]})$	$[\forall]\varphi \to \varphi$
	$(4_{[\forall]})$	$[\forall]\varphi \to [\forall][\forall]\varphi$
	$(5_{[\forall]})$	$\neg[\forall]\varphi \to [\forall]\neg[\forall]\varphi$
	$(\mathsf{Nec}_{[\forall]})$	from φ, infer $[\forall]\varphi$
(III)	**Axioms for** B_k **with** $k \in \mathcal{F}$**:**	
	(C_{B_k})	$B_k(\varphi \wedge \psi) \leftrightarrow (B_k\varphi \wedge B_k\psi)$
	$(\mathsf{Ax1}_{B_k})$	$B_k\varphi \to B_k\overline{\varphi}$
(IV)	**Axioms connecting** $[\forall]$ **and** B_k**:**	
	$(\mathsf{Ax2}_{B_k})$	$([\forall](\varphi \to \psi) \wedge B_k\varphi \wedge B_k\overline{\psi}) \to B_k\psi$
	(A_{B_k})	$B_k\varphi \to [\forall]B_k\varphi$
	(E_{B_k})	$B_k\varphi \to [\exists]\varphi$

Table 3.2: Sound and complete axiomatisation of the logic HB of hyperintensional belief over $\mathcal{L}_{[\forall]B_k}$

3.2.5 Soundness

Proposition 3.29 (Soundness of HB)**.** *The logic* HB *of hyperintensional belief is sound with respect to the class of tsef-models. That is, for all formulas* $\varphi \in \mathcal{L}_{[\forall]B_k}$ *and all tsef-models* $\mathfrak{M} = (W, \mathcal{T}, \oplus, t, \{\mathcal{E}_k\}_{k\in\mathcal{F}}, V)$, $\vdash_{\mathsf{HB}} \varphi$ *implies* $\mathfrak{M}, w \models \varphi$ *for all* $w \in W$.

Proof. Soundness is proved as usual by showing the validity of the axioms and the preservation of soundness under the inference rules. As we did for the soundness proof of HEK, we omit the axioms and inference rules of classical propositional logic as well as the (S5)-axioms for the global modality $[\forall]$. It is easy to see how the validity of (A_{B_k}) and (E_{B_k}) follows from the semantic definitions. For $(\mathsf{Ax1}_{B_k})$, we refer to the proof in Proposition 3.6 (Soundness of HEK), which is very similar. For (C_{B_k}) and $(\mathsf{Ax2}_{B_k})$, we likewise refer to the proofs there; for the right-to-left direction of (C_{B_k}), simply observe that the witnessing pieces of evidence for $B_k\varphi$ and $B_k\psi$ have non-empty intersection with all evidence in the same frame of mind k; so in particular, they have non-empty intersection with each other. The same remark applies to $(\mathsf{Ax2}_{B_k})$, where the witnessing pieces of evidence for $B_k\varphi$ and $B_k\overline{\psi}$ have non-empty intersection for the same reason. □

3.2.6 Completeness

Compared to HEK, proving completeness for HB is straightforward. In particular, we can directly show completeness with respect to a canonical tsef-model.

The notion of a maximal HB-consistent set is defined as before in Definition 3.8. We denote the set of all HB-consistent sets by $\mathcal{MCS}_{\mathsf{HB}}$. We will make repeated use of the HB-equivalents of Lemma 3.9 (Properties of MCS), Lemma 3.10 (Lindenbaum's Lemma) and Proposition 3.11 (Strong completeness), which we do not state explicitly again.

Theorem 3.30 (Completeness of HB)**.** *The logic of* HB *is strongly complete with respect to the class of tsef-models. That is, for any set of formulas* $\Gamma \cup \{\varphi\}$ *of the language* $\mathcal{L}_{[\forall]B_k}$*, if* $\Gamma \models \varphi$*, then* φ *is derivable in* HB *from* Γ*.*

Definition 3.31. For any HB-MCS Γ, we define

$$\Gamma[\forall] := \{\varphi \in \mathcal{L}_{[\forall]B_k} \mid [\forall]\varphi \in \Gamma\}.$$

Define the binary relation $\sim_{[\forall]}$ on $\mathcal{MCS}_{\mathsf{HB}}$ as follows: for any $\Gamma, \Delta \in \mathcal{MCS}_{\mathsf{HB}}$,

$$\Gamma \sim_{[\forall]} \Delta \quad \text{iff} \quad \Gamma[\forall] \subseteq \Delta.$$

Lemma 3.32. $\sim_{[\forall]}$ *is an equivalence relation.*

Proof. Same as in Lemma 3.15. □

Definition 3.33 (Canonical tsef-model for Γ_0)**.** If Γ_0 is a HB-MCS, then the canonical model for Γ_0 is a tuple $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}_k^c\}_{k\in\mathcal{F}}, V^c)$, where

i) $W^c := \{\Gamma \in \mathcal{MCS}_{\mathsf{HB}} \mid \Gamma_0 \sim_{[\forall]} \Gamma\}$;

ii) $V^c : \text{PROP} \to \mathcal{P}(W^c)$ s.t. $V^c(p) := \widehat{p}$,
where $\widehat{\varphi} := \{\Gamma \in W^c \mid \varphi \in \Gamma\}$ for all $\varphi \in \mathcal{L}_{[\forall]B_k}$;

iii) $\mathcal{T}^c := \mathcal{P}(\text{PROP})$;

iv) $\oplus^c := \cup$ (it follows that $\sqsubseteq^c = \subseteq$);

v) $t^c : \text{PROP} \cup \mathcal{F} \to \mathcal{T}$ s.t.

  - $t^c(p) = \{p\}$ for any $p \in \text{PROP}$;
  - $t^c(k) = \{p \in \text{PROP} \mid B_k\overline{p} \in \Gamma_0\}$ for any $k \in \mathcal{F}$; and
  - t^c is extended to $\mathcal{L}_{[\forall]B_k}$ by putting $t^c(\varphi) := Var(\varphi)$;

vi) for each $k \in \mathcal{F}$, $\mathcal{E}_k^c := \{(\widehat{\varphi}, Var(\varphi)) \mid B_k\varphi \in \Gamma_0\}$.

Lemma 3.34 (Existence lemma for $[\forall]$)**.** *Let* Γ_0 *be a* HB-MCS *and* $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}_k^c\}_{k\in\mathcal{F}}, V^c)$ *the canonical tsef-model for* Γ_0*. Then, for all* $\varphi \in \mathcal{L}_{[\forall]B_k}$*, we have*

$$\widehat{\varphi} = \varnothing \textit{ iff } \widehat{[\exists]\varphi} = \varnothing.$$

Proof. Same as in Lemma 3.18. □

Lemma 3.35 (Truth Lemma)**.** *Let* Γ_0 *be a* HB-MCS *and* $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}^c_k\}_{k\in\mathcal{F}}, V^c)$ *the canonical tsef-model for* Γ_0*. Then, for all* $\varphi \in \mathcal{L}_{[\forall]B_k}$ *and all* $\Gamma \in W^c$*, we have*

$$\mathfrak{M}^c, \Gamma \models \varphi \textit{ iff } \varphi \in \Gamma.$$

Proof. By induction on the complexity of φ. The Boolean cases are elementary and omitted. The case $\varphi = [\forall]\psi$ is identical to the one in Lemma 3.20, so we only prove the case for $\varphi = B_k\psi$.

- Case $\varphi = B_k\psi$:

 (⇒:) Suppose $\mathfrak{M}^c, \Gamma \models B_k\psi$. Then there is $(e, a) \in \mathcal{E}^c_k$ s.t. $e \subseteq [\![\psi]\!]$, $t^c(\psi) \sqsubseteq^c a$, and $(e, a) = (\widehat{\chi}, Var(\chi))$ for some $\chi \in \mathcal{L}_{[\forall]B_k}$ s.t. $B_k\chi \in \Gamma_0$. Since $\widehat{\chi} = e \subseteq [\![\psi]\!]$, by the induction hypothesis we have $\widehat{\chi} \subseteq \widehat{\psi}$. Moreover, since $t^c(\psi) \sqsubseteq^c a$, by definition of t^c we have $Var(\psi) \subseteq Var(\chi)$.

 We establish the following three claims:

 (1) $[\forall](\chi \to \psi) \in \Gamma$.
 Proof. Define $D := \Gamma[\forall] \cup \{\chi, \neg\psi\}$. We claim that D is inconsistent: towards a contradiction, suppose D is consistent. Then by Lemma 3.10, there is a HB-MCS Δ s.t. $D \subseteq \Delta$. Since $\Gamma[\forall] \subseteq \Delta$, we have by definition of W^c that $\Delta \in W^c$. Now $\chi \in \Delta$, so by $\widehat{\chi} \subseteq \widehat{\psi}$, we have $\psi \in \Delta$, which contradicts $\neg\psi \in \Delta$.
 So D is inconsistent. Thus, there are $[\forall]\xi_1, ..., [\forall]\xi_n \in \Gamma$ s.t. $\vdash_{\mathsf{HB}} \bigwedge_{i\leq n}[\forall]\xi_i \to (\chi \to \psi)$. By Lemma 3.17, it follows that $\vdash_{\mathsf{HB}} \bigwedge_{i\leq n}[\forall]\xi_i \to [\forall](\chi \to \psi)$. Applying Lemma 3.9 throughout, we have $\bigwedge_{i\leq n}[\forall]\xi_i \in \Gamma$, so by modus ponens, $[\forall](\chi \to \psi) \in \Gamma$.

 (2) $B_k\chi \in \Gamma$.
 Proof. We have $B_k\chi \in \Gamma_0$, so by (A_{B_k}), we have $[\forall]B_k\chi \in \Gamma_0$. Thus, by definition of W^c, we have $B_k\chi \in \Gamma$.

 (3) $B_k\overline{\psi} \in \Gamma$.
 Proof. As shown in (2), we have $B_k\chi \in \Gamma$. Thus, by $(\mathsf{Ax1}_{B_k})$, we have $B_k\overline{\chi} \in \Gamma$. Since we also have $Var(\psi) \subseteq Var(\chi)$, by Proposition 3.28(ii), it follows that $B_k\overline{\psi} \in \Gamma$.

 By claims (1) through (3) and axiom $(\mathsf{Ax2}_{B_k})$, we obtain $B_k\psi \in \Gamma$.

 (⇐:) Suppose $B_k\psi \in \Gamma$. Define $(e, a) := (\widehat{\psi}, Var(\psi))$. We establish the following three claims:

(1) $(e,a) \in \mathcal{E}_k^c$.
Proof. We have $B_k\psi \in \Gamma$, so by (A_{B_k}), modus ponens and Lemma 3.9, it follows that $B_k\psi \in \Gamma_0$. Thus, by definition of $\mathcal{E}_k^c$, we have $(e,a) \in \mathcal{E}_k^c$.

(2) $e \subseteq [\![\psi]\!]$.
Proof. By the induction hypothesis, we have $e = \widehat{\psi} = [\![\psi]\!]$.

(3) $t^c(\psi) \sqsubseteq^c a$.
Proof. By definition of t^c and a, we have $t^c(\psi) = Var(\psi) = a$.

(4) for all $e'_{a'} \in \mathcal{E}_k^c$, $e \cap e' \neq \varnothing$.
Proof. Let $e'_{a'} \in \mathcal{E}_k^c$, then $e' = \widehat{\chi}$ for some $\chi \in \mathcal{L}_{[\forall]B_k}$ s.t. $B_k\chi \in \Gamma_0$. We already established in (1) that $B_k\psi \in \Gamma_0$. Applying Lemma 3.9 throughout, we get from (C_{B_k}) that $B_k(\psi \wedge \chi) \in \Gamma_0$, and by (E_{B_k}) that $[\exists](\psi \wedge \chi) \in \Gamma_0$. So $\widehat{[\exists](\psi \wedge \chi)} \neq \varnothing$, which by Lemma 3.34 implies that $\widehat{\psi \wedge \chi} \neq \varnothing$. So $e \cap e' = \widehat{\psi} \cap \widehat{\chi} = \widehat{\psi \cap \chi} \neq \varnothing$.

By claims (1) through (4), $(e,a) \in \mathcal{E}_k^c$ is a witness for $\mathfrak{M}^c, \Gamma \models B_k\psi$.

□

Lemma 3.36. *For any* HB-MCS Γ_0*, the canonical model* $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}_k^c\}_{k\in\mathcal{F}}, V^c)$ *for* Γ_0 *is a tsef-model.*

Proof. Let Γ_0 be a HB-MCS and $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}_k^c\}_{k\in\mathcal{F}}, V^c)$ the canonical model for Γ_0. Then

(i) $W^c \neq \varnothing$ because by $(\mathsf{T}_{[\forall]})$, $\Gamma_0 \in W^c$;

(ii) $\mathcal{T}^c = \mathcal{P}(\text{PROP}) \neq \varnothing$ because $\text{PROP} \neq \varnothing$;

(iii) $\oplus^c = \cup$ is idempotent, commutative and associative;

(iv) $\mathcal{T}^c$ is closed under binary $\oplus^c$ because $\mathcal{P}(\text{PROP})$ is closed under unions;

(v) t^c is clearly well-defined; and moreover, we have $t^c(\varphi) = Var(\varphi) = \bigcup\{\{p\} \mid p \in Var(\varphi)\} = \bigoplus^c\{t^c(p) \mid p \in Var(\varphi)\}$;

(vi) for all $k \in \mathcal{F}$ and $e_a \in \mathcal{E}_k^c$, we have $a \sqsubseteq^c t^c(k)$:

Let $e_a \in \mathcal{E}_k^c$, then (e,a) is of the form $(\widehat{\varphi}, Var(\varphi))$ for some $\varphi \in \mathcal{L}_{[\forall]B_k}$ s.t. $B_k\varphi \in \Gamma_0$. We claim that this implies $\bigwedge_{p\in Var(\varphi)} B_k\overline{p} \in \Gamma_0$: By $(\mathsf{Ax1}_{B_k})$, we have $\vdash_{\mathsf{HB}} B_k\varphi \to B_k\overline{\varphi}$. By Proposition 3.28(i), it follows that $\vdash_{\mathsf{HB}} B_k\varphi \to \bigwedge_{p\in Var(\varphi)} B_k\overline{p}$. Since $B_k\varphi \in \Gamma_0$, by modus ponens and Lemma 3.9, we have $\bigwedge_{p\in Var(\varphi)} B_k\overline{p} \in \Gamma_0$. Thus, $B_k\overline{p} \in \Gamma_0$ for all $p \in Var(\varphi)$, so $a = Var(\varphi) \subseteq \{p \in \text{PROP} \mid B_k\overline{p} \in \Gamma_0\} = t^c(k)$. Therefore, $a \sqsubseteq^c t^c(k)$.

(vii) for each $k \in \mathcal{F}$, $\mathcal{E}_k^c$ is closed under non-empty finite evidence combination, i.e. for all $e_{1,a_1}, ..., e_{n,a_n} \in \mathcal{E}_k^c$, $e_1 \cap \cdots \cap e_n \neq \varnothing$ implies $e_{1,a_1} \sqcap \cdots \sqcap e_{n,a_n} \in \mathcal{E}_k^c$.

Proof. We show that if $e_a, e'_{a'} \in \mathcal{E}_k^c$, then $(e \cap e', a \oplus^c a') \in \mathcal{E}_k^c$. After establishing this, a simple inductive argument suffices to prove the desired result.

Let $e_a, e'_{a'} \in \mathcal{E}_k^c$. Then $(e, a) = (\widehat{\varphi}, Var(\varphi))$ and $(f, b) = (\widehat{\psi}, Var(\psi))$ for some $\varphi, \psi \in \mathcal{L}_{[\forall]B_k}$ s.t. $B_k\varphi \in \Gamma_0$ and $B_k\psi \in \Gamma_0$. It follows from (C_{B_k}) and Lemma 3.9 that $B_k(\varphi \wedge \psi) \in \Gamma_0$. Thus,

$$\begin{aligned}
(e \cap f, a \oplus^c a') &= (\widehat{\varphi} \cap \widehat{\psi}, Var(\varphi) \cup Var(\psi)) && \text{(Definition of } \oplus^c\text{)}\\
&= (\widehat{\varphi \wedge \psi}, Var(\varphi \wedge \psi)) && \text{(Lemma 3.9)}\\
&\in \mathcal{E}_k^c. && \text{(Definition of } \mathcal{E}_k^c\text{)}
\end{aligned}$$

Thus, $\mathfrak{M}^c$ is indeed a tsef-model. □

Theorem 3.30 (Completeness of HB)**.** *The logic of* HB *is strongly complete with respect to the class of tsef-models. That is, for any set of formulas* $\Gamma \cup \{\varphi\}$ *of the language* $\mathcal{L}_{[\forall]B_k}$*, if* $\Gamma \models \varphi$*, then* φ *is derivable in* HB *from* Γ.

Proof. By Proposition 3.11, it suffices to show that every HB-consistent set of formulas is satisfiable on some tsef-model. Let G be a HB-consistent set of formulas. By Lemma 3.10, there is a HB-MCS Γ s.t. $G \subseteq \Gamma$. Let $\mathfrak{M}^c = (W^c, \mathcal{T}^c, \oplus^c, t^c, \{\mathcal{E}_k^c\}_{k \in \mathcal{F}}, V^c)$ be the canonical tsef-model for Γ as defined in Definition 3.33. By definition of $\mathfrak{M}^c$, we have $\Gamma \in W^c$, and by Lemma 3.35, we have $\mathfrak{M}^c, \Gamma \models G$. By Lemma 3.36, $\mathfrak{M}^c$ is a tsef-model. So G is satisfiable on a tsef-model, as required. □

Chapter 4

Comparison with other approaches

In this chapter, we compare our framework with similar existing approaches, some of which inspired our approach. In Section 4.1, we briefly sketch the topological evidence models of Baltag et al. (2016) and Özgün (2017), whose definitions of knowledge and belief strongly inspired our respective formalisations. We point out that our epistemic notions only exploit part of the topological structure of evidence developed there, show that our notion of belief stripped free from its topic-sensitive and fragmented components is strictly stronger than its topological variant, and point to conceptual problems rendering it difficult to combine full evidential topologies with topics. In Section 4.2, we summarise question-based approaches to interpret knowledge and belief, which provide alternatives to our topic-based framework, and we show how some problems pertaining to logical omniscience can inherently not be solved by these approaches. Finally, in Sections 4.3 and 4.4, we compare our framework with awareness logics and justification logics, respectively.

4.1 Topological evidence models

Our approach to defining knowledge and belief based on evidence is strongly inspired by that of Baltag et al. (2016) and Özgün (2017), where van Benthem-Pacuit evidence models are further developed into *topological* evidence models.

Roughly, the idea is as follows: an agent's *evidential topology* is the topology generated by her set of evidence $\mathcal{E}$, i.e. the collection of all arbitrary unions of finite intersections of evidence pieces in $\mathcal{E}$. Just like on van Benthem-Pacuit models, a *combined piece of evidence* is simply a non-empty finite intersection of evidence in $\mathcal{E}_0$.

If we strip away the topic and fragmentation components of our evidence models, then a collection $\mathcal{E}$ of combined evidence is simply a family of subsets of W closed under non-empty finite intersections. By adding the empty set to $\mathcal{E}$, we get the *basis* of an evidential topology. Thus, our semantics is built on the basis of evidential topologies.

In topological evidence models, φ is believed iff there exists a *non-empty open subset* of the truth set of φ which is consistent with all combined evidence, or phrased topologically: iff there exists a *dense open subset* of φ. Note the similarity with (the intensional component of) our framework, in which φ is believed iff there exists a *combined piece of evidence* which is consistent with all other combined evidence.

Since every basic subset of an evidential topology which is consistent with all combined evidence is also a dense open subset, belief under (the intensional component of) our semantics implies belief under the topological semantics in Baltag et al. (2016). However, the converse does not hold, so (the intensional reduction of) our notion of belief is strictly stronger than the topological notion. Here is a sketch of a counterexample:

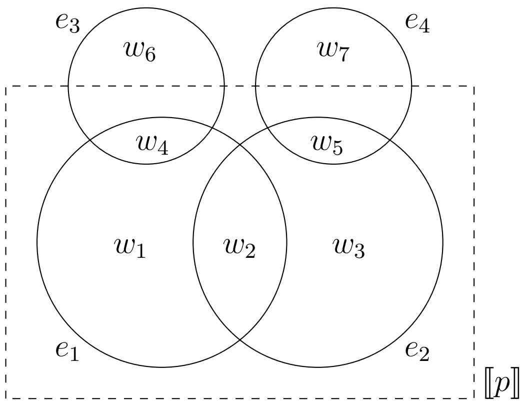

Figure 4.1: Example: The intensional reduction of our belief is strictly stronger than its topological variant

Whereas the union $e_1 \cup e_2$ is a dense open subset of $[\![p]\!]$ and therefore a witness of Bp under the topological semantics, there is no combined piece of evidence consistent with all combined evidence, so Bp does not hold under our semantics.

Intuitively, the conceptual difference between the two semantics is the following: an agent may possess a piece of combined evidence e entailing φ, but which is too "specific" or too "small" to guarantee consistency with all other combined evidence.

Topic-sensitive evidence goes full topology?

The present framework does not make full use of the topological structure of evidence developed in Özgün (2017) and instead uses topics combined with topological bases. This confinement has a reason: the interaction of topics with unions of evidence – from a topological perspective, the missing component in our approach – produce problematic results.

If we want to combine full evidential topologies with topics, we need to introduce a union equivalent of our notion of evidence combination. That is, given two topic-indexed pieces of evidence e_a and f_b, which topic should be assigned to the union $e \cup f$? The obvious choice would be to assign the topic $a \oplus b$ to $e \cup f$, just as we did for evidence combination. But then the following problem arises:

Suppose our agent has two pieces of evidence e_a and f_b, both belonging to the same evidential frame $\mathcal{E}_k$. The agent now performs two subsequent steps of evidence combination: first, e_a and f_b are combined by taking their union, producing $(e \cup f)_{a \oplus b}$. This new piece of evidence is now combined with e_a by means of intersection, yielding $((e \cup f) \cap e)_{(a \oplus b) \oplus a} = e_{a \oplus b}$. Thus, the agent produced a piece of evidence which is intensionally equivalent to e_a, but with topic $a \oplus b$. In general, given any evidential frame $\mathcal{E}_k$ and evidence $e_a \in \mathcal{E}_k$, an agent can *always* produce a new piece of evidence $e_{a'}$, where a' is the *fusion of all topics assigned to evidence in* $\mathcal{E}_k$.

This phenomenon would arguably undermine our objective of modelling evidential relevance. In the fragmentation setting, the topic of all pieces of evidence in one frame effectively reduce to one topic which is attached to that frame. For this reason, we only considered finite intersections as a means of (intensional) evidence combination. We leave the problem of combining full topological evidence models with topics for future research.

4.2 Question-based approaches

We used an abstract notion of *topics* to supply both evidence and propositions with a hyperintensional content component. However, there are alternative notions that could be used instead of topics. In particular, the idea that subject matters correspond to partitions of logical space, often interpreted as *questions*, has sparked a considerable amount of literature (originating with Lewis (1988); more recent developments include Boddy (2014) and Yalcin (2018) and Hoek (2020)). In this section, we briefly delineate the basic conceptual features of question-based accounts of knowledge and belief, sketch how these could be extended to incorporate a notion of question-sensitive evidence, and point out limitations compared to our approach.

A *partition* $\mathcal{P}$ of a set of worlds W is a collection of non-empty subsets of W such that (1) $\mathcal{P}$ covers W (i.e., $\bigcup P = W$); and (2) all subsets in $\mathcal{P}$ are mutually disjoint. We take partitions to represent questions. An *answer* to a question $\mathcal{P}$ is a cell $C \in \mathcal{P}$. A *partial answer* to a question $\mathcal{P}$ is a union of answers to $\mathcal{P}$.

The cells of a partition represent ways the world might be, but they do not settle all possible distinctions. Instead, as phrased by Yalcin (2018), they foreground some distinctions and background others. The foregrounded propositions are those that cut along the lines of the partition, i.e. those that constitute partial answers to the associated question.

Given a set of worlds W, an agent might be interested in obtaining answers to only *some* of all possible questions over W. We call the set of these special questions $\mathcal{A}$ the set of questions *available* to the agent. Intuitively, $\mathcal{A}$ should exhibit some structure. For example, if I ask myself whether $\varphi \wedge \psi$ is true, then I ask myself whether φ is true and whether ψ is true. Thus, the availability of some questions should intuitively imply the availability of others. An example of an implementation is Hoek (2020), but providing the details would go beyond the scope of this comparison.

These fundamental ideas can be applied to develop a question-based notion of evidential relevance as follows: given a set of worlds W, a question $\mathcal{P}$ over W, an (intensional) piece of evidence $e \subseteq W$ and a formula φ, we say that *e is relevant for φ* iff e and the truth set of φ are both partial answers to some question available to the agent.

Note that such a question-sensitive notion of evidence would circumvent some of the problems related to logical omniscience. For example, question-based evidence is not closed under logical entailment: If φ entails ψ and I have a piece of evidence e which is relevant for and entails φ, then φ and e are both partial answers to a question available to me; however, this does not entail that ψ is also a partial answer to a question available to me, so it does not follow that I also know ψ.

However, the questions-based approach is unable to offer a solution to *closure under logical equivalence*: if φ and ψ are logically equivalent and I have relevant evidence e for φ, then φ and e both partially answer some question $\mathcal{P}$ available to me. Since φ and ψ are equivalent, their truth sets are identical, so ψ is also a partial answer to $\mathcal{P}$. Thus, e is also relevant evidence for ψ. This should come as no surprise: logically equivalent formulas are also intensionally equivalent. As partition-based question approaches confine content to possible worlds, they cannot distinguish between such sentences.

Questions provide an intuitive way of interpreting subject matters, and developing a framework for question-sensitive evidence could be fruitful. Combining a partition-based approach to questions with evidence provides solutions to some of the logical omniscience-related problems; however, since such an approach would still be entirely intensional, it inherently struggles with problems whose solutions require a strictly hyperintensional approach such as ours.

4.3 Awareness logics

Awareness logics (originally developed by Fagin and Halpern (1987)) tackle the problem of logical omniscience by drawing a distinction between *implicit* and *explicit* knowledge. Implicit knowledge corresponds to the classical Hintikkan notion of knowledge, whereas explicit knowledge is a conjunction of two components: φ is explicitly known iff (1) φ is implicitly known and (2) the agent is aware of φ. Logical omniscience is primarily exhibited by implicit, but not by explicit knowledge.

Awareness is typically modelled in terms of *awareness sets*, which are (possibly world-dependent) sets of formulas that specify which formulas the agent is aware off. Thus, the hyperintensional component is implemented in a purely syntactical way. The extent to which awareness sets exhibit structure varies by approach: on the one end of the extreme, any arbitrary set of formulas is admissible as an awareness set; in other approaches, more requirements are imposed.

The extreme variant that allows *any* set of formulas as an awareness set is arguably of limited use in capturing an interesting notion of hyperintensional knowledge. Our goal is to provide epistemic logics that strikes the right balance in terms of epistemic closure principles. This extremely liberal way of determining awareness misses this goal: "explicit attitudes obey no non-trivial logical closure properties" (Berto and Özgün, unpublished, p. 8). For example, $K(\varphi \wedge \psi)$ entails neither $K(\psi \wedge \varphi)$ nor $K\varphi$.

In contrast, our topics-based approach imposes an intuitive mereological structure on the hyperintensional content of both propositions and of evidence, which achieves a better balance in terms of the logical inferences it warrants. For example, $\varphi \wedge \psi$ has the same intensional and hyperintensional content as $\psi \wedge \varphi$; in the most general awareness logic, this is not the case.

In general, "[a] requirement often suggested is that a hyperintensional account of this or that notion shouldn't make it as fine-grained as the syntax of the language one is working with – on pain of giving away the very point of having a semantics for it" (Berto and Nolan 2021, p. 22).

Of course, some more structure can be induced on the awareness set by imposing constraints. For example, in a variant called *awareness generated by primitive propositions* (Ditmarsch et al. (2015, p. 81); a similar approach is pursued in van Ditmarsch and French (2014)), the agent is aware of φ just in case she is aware of all its atomic constituents taken together. This delivers some closure properties such as closure under conjunction.

In general, topic-based approaches to hyperintensional epistemic attitudes such as in Hawke, Özgün, and Berto (2020) and Özgün and Berto (2020) can be mapped to awareness structures. If b represents the greatest topic associated with a concept the agent understands, then the relevant condition for knowledge of φ in these approaches is $t(\varphi) \sqsubseteq b$. The general strategy is to construct an awareness sets that consist of all formulas φ s.t. $t(\varphi) \sqsubseteq b$ holds on the topic model. The approach can be easily generalised to a fragmented setting.

Thus, awareness logic can be seen as a syntax-based generalisation of approaches that use topics to represent conceptual understanding. However, our framework gives a different twist to topics: we use them to formalise the notion of *relevance* of evidence for propositions. There currently is no awareness logics-based approach (in the author's awareness set) that combines evidence with awareness. Such frameworks could be implemented by associating evidence with sets of formulas, each of which represents a sentence for which the evidence is relevant. However, we find that the mereological structure of topics provides an intuitively compelling way of understanding the hyperintensional components of both propositions and of evidence.

4.4 Justification logics

Justification logics (originally developed by Artemov and Fitting (2021)) set out to add a notion of justification to epistemic logic: a proposition P is known in justification logics iff there is a justification for P.

Justifications are represented by syntactic objects called *justification terms*, which together with propositional formulas build up the language of justification logic. If φ is a formula and t is a justification term, then $t : \varphi$ is a formula, which is intended to be read as

t is a justification for φ.

Justification logic makes the *implicit* modalities of standard modal logic *explicit* by unfolding them in terms of justifications. That is, whereas in classical (i.e., Hintikkan) epistemic logic the formula $K\varphi$ (φ is known) is interpreted as truth of φ in all epistemic alternatives, justification logic interprets $K\varphi$ as *there exists a justification for φ*.

Justification logic does not assume that there is a justification for every tautology, i.e. formulas of the form $t : \varphi$ with φ a tautology are not necessarily valid. However, justification terms come with a number of principles which hold in general, one of which is application: If s is a justification for $\varphi \to \psi$ and t is a justification for ψ, then $[s \cdot t]$ is a justification for ψ. As an axiom:

$$(s : (p \to q) \wedge t : p) \to [s \cdot t] : q$$

Note the similarity between the application axiom above and the K axiom of normal modal logic:

$$(\Box(p \to q) \land \Box p) \to \Box q$$

If we interpret $\Box$ as knowledge, then the above axiom turns into the principle of *closure under known implication*. By unfolding $\Box$ into its underlying justifications, justification logic features a kind of justification tracking, which limits epistemic closure under known implication to those instances in which the agent possesses relevant justifications for the involved pieces of knowledge.

Thus, justification logics provide a hyperintensional notion of justification for epistemic logics that gives rise to agents which are not fully logically omniscient. As such, they share the central objectives of our framework. However, there is one important limitation to classical justification logics: they interpret justifications as mathematical proofs. In particular, justifications are assumed to be infallible, indefeasible, and never misleading.

Whereas this kind of justification is certainly important, the resulting logics are of limited use in modelling soft, for example empirical evidence and its associated notions of knowledge and belief, which our framework can capture.

Justification Logic for soft evidence Baltag, Renne, and Smets (2012) provide a framework based on justification logic that aims at representing exactly this kind of soft evidence. Their approach combines the central ideas of justification logic with the distinction between implicit and explicit knowledge drawn in awareness logics (see Section 4.3).

In particular, Baltag, Renne, and Smets (2012) define a notion of defeasible knowledge, which roughly boils down to the following: At a world w, φ is explicitly *known* iff there is a justification term t for φ, t is *available* to the agent, and each formula ψ whose *evidential certificate* c_φ occurs in t is explicitly known. An evidential certificate c_φ for φ is a justification term which is interpreted as a canonical piece of evidence in support of φ.

Therefore, Baltag, Renne, and Smets (2012) capture a notion of defeasible knowledge grounded in soft evidence similar to that developed in this thesis. However, the approach of Baltag, Renne, and Smets (2012) remains rather syntactical; in particular, it is unclear how evidential certificates should be intuitively understood and what evidence corresponds to semantically. In comparison, our framework provides a more semantic and intuitive representation, including the mereological structure of the hyperintensional component of evidence.

Chapter 5

Conclusions and further research

To sum up, we developed a hyperintensional and fragmented notion of evidence, provided definitions for knowledge and belief based on this conception of evidence, and showed how the resulting logic gives rise to an agent that does not suffer from full logical omniscience. Moreover, we provided sound and complete axiomatisations of a factive evidence and knowledge fragment as well as a belief fragment.

We began with a purely intensional, possible worlds-based representation of evidence in the tradition of van Benthem and Pacuit (2011) and demonstrated how this conception of evidence validates a number of intuitively problematic principles closely related to logical omniscience. To improve on this situation, we took inspiration from Hawke, Özgün, and Berto (2020) and Özgün and Berto (2020) and supplemented the notion of evidence with two complementary components:

On the one hand, we argued that evidence is *hyperintensional*. Specifically, we proposed that the purely intensional requirement e *entails* p is not sufficient for a piece of evidence e to constitute evidence for a proposition p. Instead, we suggested that e additionally needs to be *relevant* for p. To model evidential relevance, we adopted the notion of *topics* from Hawke, Özgün, and Berto (2020) and Özgün and Berto (2020) and developed a way of assigning topics to both evidence and propositions. Having equipped both evidence and propositions with subject matters, we defined the relation of evidential relevance in terms of topic inclusion.

On the other hand, we took an agent's body of evidence to be *fragmented*, i.e. consisting of several bodies of evidence called *evidential frames* that act as "evidential echo chambers": Pieces of evidence can only be combined within the same evidential frame, but not between different frames.

Based on this conception of hyperintensional, fragmented evidence, we defined notions of knowledge and belief. In doing so, we chose an *eviden-*

tialist and *coherentist* approach: First, our notions of knowledge and belief are entirely grounded in evidence. Second, knowledge and belief require a piece of evidence which is consistent with other evidence possessed by the agent.

We argued that knowledge and belief inherit their hyperintensional and fragmented properties from the underlying evidence. Correspondingly, we defined *non-fragmented belief* of φ as possession of a piece of evidence which (a) *entails* φ; (b) is *relevant* for φ; and (c) is *consistent with all other evidence in the same evidential frame*. *Fragmented belief* of φ corresponds to believing φ (in the non-fragmented sense) in some or other evidential frame. *Non-fragmented* and *fragmented knowledge* are defined like the corresponding notions of belief, but with a *factive* witnessing piece of evidence.

We showed that our target notions of fragmented knowledge and belief suffer from the defect of logical omniscience to a significantly reduced extent, and that this is due to the topic-sensitivity and fragmentation of the underlying evidence. For example, topic-sensitivity ensures that agents do not know all validities, and that knowledge is not closed under logical entailment. Fragmentation guarantees that knowledge is not closed under conjunction.

Our technical contributions consist of a sound and complete axiomatisation of the logic of hyperintensional factive evidence and knowledge, whose language is expressive enough to define all fragmented and non-fragmented evidence, knowledge and belief modalities used throughout this thesis. We also provided a sound and complete axiomatisation of the logic of hyperintensional belief.

Our framework shares motivational and conceptual similarities with topological evidence models, question-based approaches to knowledge and belief, awareness logics and justification logics. However, to the best of my knowledge, our approach is the first to provide a full semantic treatment of hyperintensional evidence.

We conclude with mentioning two aspects of logical omniscience that our current framework does not solve and point to existing approaches that could be fruitfully combined with ours in future work.

5.1 Conjunctive parts

Our framework has in common with that of Hawke, Özgün, and Berto (2020) that it validates the following principle:

$$([\forall](\varphi \rightarrow \psi) \wedge K(\varphi \wedge (\psi \vee \neg\psi))) \rightarrow K\psi \tag{5.1}$$

As an example, assume that Goldbach's conjecture G is true and follows as a logical consequence of the conjunction α of the Peano axioms. The above sentence entails that someone who knows α conjoined with $G \vee \neg G$ also knows G, which is implausible.

One response given in Hawke, Özgün, and Berto (2020) traces this problem back to the general way how knowledge of conjunctions is treated. Consider the following sentence:

$$K(\varphi \wedge (\varphi \to \psi)) \to K\psi \tag{5.2}$$

Since both conjuncts φ and $\varphi \to \psi$ occur in the scope of K, they are evaluated in the same evidential frame. Since K_k validates closure under conjunctions, this sentence is valid in our framework. The same explanation is applicable to (5.1) above.

We now briefly sketch the solution proposal in Hawke, Özgün, and Berto (2020). The notion of knowledge in Hawke, Özgün, and Berto (2020) is fragmented. As in our proposal, their knowledge modalities K_k correspond to *frames of mind*, which are conceptually similar to our evidential frames. Fragmented knowledge $K\varphi$ is understood as 'φ is known in some frame of mind'.

Now, according to Hawke, Özgün, and Berto (2020, p. 750), ascriptions of conjunctive knowledge as in 'Jones knows φ and ψ' are generally ambiguous between the following two readings:

1. *Jones knows $\varphi \wedge \psi$ in the same frame of mind;*
2. *Jones knows φ in some frame of mind*
   and *Jones knows ψ in some other frame of mind.*

The suggestion is to modify the definition of knowledge in such a way that it "pulls apart" the conjuncts of a knowledge ascription and evaluates them in a fragmented way. That is, knowledge ascriptions of the form $K(\varphi \wedge \psi)$ are always interpreted according to the second reading given above. Using this modified definition of knowledge breaks the validity of (5.1) and (5.2), because the conjuncts in the scope of K are evaluated in a fragmented manner, and there is no guarantee that the full conjunction is known in any single frame.

In the framework of Hawke, Özgün, and Berto (2020), this fix indeed seems to path the way to a notion of knowledge which is not closed under logical entailment in the sense of (5.1). A similar solution could be applied to our approach, but the implementation would have to be implemented at the level of evidence. While we leave this opportunity for further research, we point out one possible objection to this strategy: a dedicated mathematician who unsuccessfully devoted most of her life to prove Goldbach's conjecture might retort that her failure to know that Goldbach's conjecture is true is not due to the fact that her mind is fragmented and she knows the Peano axioms in one part, and the Law of the Excluded Middle in another; but that it is simply very hard to prove Goldbach's Conjecture. In the next and final section, we turn to a perspective towards the problem of logical omniscience that focuses on the computational cost associated with logical inference.

5.2 Computationally bounded agents

In this thesis, we approach the problem of logical omniscience by making explicit the hyperintensional and fragmented nature of evidence. A fundamtenally different perspective on logical omniscience focuses on the *computational limitations* of epistemic agents who operate in finite space and time.

We consider a brief example of a situation in which computational bounds intuitively restrict an agent's knowledge. Let p be a proposition and φ the following formula:

$$\varphi = \underbrace{\neg\neg\cdots\neg}_{\text{100 negations}}\ p$$

Suppose an agent believes p, and someone asks her whether she believes φ. The intuitive explanation for why it might be difficult for the agent to infer belief of φ from belief of p is that she fails to keep track of the number of negations. Processing complex sentences takes resources, and a lack of these resources can limit the logical inferences one is able to perform.

Our framework does not capture this perspective towards logical omniscience. The reason is that we assume topic transparency of logical connectives: the topic of p is identical to the topic of φ. Since p and φ are also intensionally equivalent, this implies that any piece of evidence for p is also evidence for φ; that belief of p implies belief of φ; and so on.

Some solutions to this problem focus on representing the computational cost of performing logical inference. For example, Smets and Solaki (2018) develop a framework in which both the cognitive capacity of an agent and the cost associated with different inference rules is captured quantitatively. A complexity-theoretic treatment is given in Artemov and Kuznets (2014).

We suggest that the *bounded computation* perspective towards logical omniscience and our perspective targeting the notion of *evidential relevance* are orthogonal aspects of logical omniscience. Specifically, a piece of evidence can be relevant for two propositions p and q which share the same intensional and hyperintensional content, but an agent may nevertheless fail to realise this because she lacks the required resources to perform the inference. Combining these two perspectives could be a fruitful agenda for future research.

Bibliography

Artemov, Sergei and Melvin Fitting (2021). "Justification Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.

Artemov, Sergei and Roman Kuznets (2014). "Logical Omniscience as Infeasibility". In: *Annals of Pure and Applied Logic*. The Constructive in Logic and Applications 165.1, pp. 6–25. DOI: 10.1016/j.apal.2013.07.003.

Baltag, Alexandru, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets (2016). "Justified Belief and the Topology of Evidence". In: *Logic, Language, Information, and Computation*. Ed. by Jouko Väänänen, Åsa Hirvonen, and Ruy de Queiroz. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 83–103. DOI: 10.1007/978-3-662-52921-8_6.

Baltag, Alexandru, Bryan Renne, and Sonja Smets (2012). "The Logic of Justified Belief Change, Soft Evidence and Defeasible Knowledge". In: *Logic, Language, Information and Computation*. Ed. by Luke Ong and Ruy de Queiroz. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 168–190. DOI: 10.1007/978-3-642-32621-9_13.

Berto, Francesco (2019). "Simple Hyperintensional Belief Revision". In: *Erkenntnis* 84.3, pp. 559–575. DOI: 10.1007/s10670-018-9971-1.

Berto, Francesco and Peter Hawke (2018). "Knowability Relative to Information". In: *Mind* 130.517, pp. 1–33. DOI: 10.1093/mind/fzy045.

Berto, Francesco and Daniel Nolan (2021). "Hyperintensionality". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.

Berto, Francesco and Aybüke Özgün (n.d.). "The Logic of Framing". Unpublished manuscript.

Blackburn, Patrick, Maarten de Rijke, and Yde Venema (2001). *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9781107050884.

Boddy, Rachel (2014). "Epistemic Issues and Group Knowledge". Master's thesis. ILLC, University of Amsterdam.

Builes, David (2020). "A Paradox of Evidential Equivalence". In: *Mind* 129.513, pp. 113–127. DOI: 10.1093/mind/fzz046.

Ditmarsch, Hans van, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Pieter Kooi (2015). *Handbook of Epistemic Logic*. London: College Publications.

Fagin, Ronald and Joseph Y. Halpern (1987). "Belief, Awareness, and Limited Reasoning". In: *Artificial Intelligence* 34.1, pp. 39–76. DOI: 10.1016/0004-3702(87)90003-8.

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi (1995). "Logical Omniscience (Chapter 9)". In: *Reasoning about Knowledge*. Cambridge, MA: MIT Press, pp. 333–389.

Gettier, Edmund L. (1963). "Is Justified True Belief Knowledge?" In: *Analysis* 23.6, pp. 121–123. DOI: 10.2307/3326922.

Hawke, Peter (2018). "Theories of Aboutness". In: *Australasian Journal of Philosophy* 96.4, pp. 697–723. DOI: 10.1080/00048402.2017.1388826.

Hawke, Peter, Aybüke Özgün, and Francesco Berto (2020). "The Fundamental Problem of Logical Omniscience". In: *Journal of Philosophical Logic* 49.4, pp. 727–766. DOI: 10.1007/s10992-019-09536-6.

Hintikka, Jaakko (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press.

Hoek, Daniel (2020). "Questions in Action". Unpublished manuscript.

Ichikawa, Jonathan Jenkins and Matthias Steup (2018). "The Analysis of Knowledge". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2018. Metaphysics Research Lab, Stanford University.

Kelly, Thomas (2016). "Evidence". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

Lehrer, Keith (2018). *Theory Of Knowledge: Second Edition*. London: Routledge.

Lehrer, Keith and Thomas Paxson (1969). "Knowledge: Undefeated Justified True Belief". In: *The Journal of Philosophy* 66.8, pp. 225–237. DOI: 10.2307/2024435.

Lewis, David (1982). "Logic for Equivocators". In: *Noûs* 16.3, pp. 431–441. DOI: 10.2307/2216219.

— (1988). "Relevant Implication". In: *Theoria* 54.3, pp. 161–174. DOI: 10.1111/j.1755-2567.1988.tb00716.x.

Özgün, Aybüke (2017). "Evidence in Epistemic Logic. A Topological Perspective". Doctoral dissertation. University of Amsterdam, University of Lorraine.

Özgün, Aybüke and Francesco Berto (2020). "Dynamic Hyperintensional Belief Revision". In: *The Review of Symbolic Logic*, pp. 1–46. DOI: 10.1017/S1755020319000686.

Smets, Sonja and Anthia Solaki (2018). "The Effort of Reasoning: Modelling the Inference Steps of Boundedly Rational Agents". In: *Logic,* *Language, Information, and Computation*. Ed. by Lawrence S. Moss, Ruy de Queiroz, and Maricarmen Martinez. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 307–324. DOI: 10.1007/978-3-662-57669-4_18.

Solaki, Anthia (2017). "Steps out of Logical Omniscience". Master's thesis. ILLC, University of Amsterdam.

Stalnaker, Robert C. (1984). *Inquiry*. Cambridge, MA: A Bradford Book.

Van Benthem, J. and E. Pacuit (2011). "Dynamic Logics of Evidence-Based Beliefs". In: *Studia Logica* 99.1, pp. 61–92. DOI: 10.1007/s11225-011-9347-x.

Van Benthem, Johan, David Fernández-Duque, and Eric Pacuit (2012). "Evidence Logic: A New Look at Neighborhood Structures". In: *Advances in Modal Logic*. Ed. by T. Bolander, T. Braüner, S. Ghilardi, and L. Moss. Vol. 9. Copenhaguen, Denmark: College Publications, pp. 97–118.

— (2014). "Evidence and Plausibility in Neighborhood Structures". In: *Annals of Pure and Applied Logic*. The Constructive in Logic and Applications 165.1, pp. 106–133. DOI: 10.1016/j.apal.2013.07.007.

Van Ditmarsch, Hans and Tim French (2014). "Semantics for Knowledge and Change of Awareness". In: *Journal of Logic, Language and Information* 23.2, pp. 169–195. DOI: 10.1007/s10849-014-9194-z.

Yablo, Stephen (2014). *Aboutness*. Princeton, NJ: Princeton University Press.

Yalcin, Seth (2018). "Belief as Question-Sensitive". In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47. DOI: 10.1111/phpr.12330.

— (2020). "Fragmented but Rational". Unpublished manuscript.