Knowledge as Issue-Relevant Information

MSc Thesis *(Afstudeerscriptie)*

written by

Wessel Kroon
(born July 6, 1996 in Vught, The Netherlands)

under the supervision of **Dr. Alexandru Baltag**, 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:
June 6, 2024

Members of the Thesis Committee:
Dr. Maria Aloni (Chair)
Dr. Alexandru Baltag (Supervisor)
Dr. Aybüke Özgün
Dr. Malvin Gattinger

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis is a logical study of a notion of knowledge as issue-relevant information. We take issues to be the objects of inquiry. The set of issues pursued by an agent constitutes her epistemic agenda, and only the information that is relevant with respect to the issues on her agenda is processed into knowledge. So, we obtain a notion of knowledge as issue-relevant information. We motivate the study of this notion and compare it to other notions of knowledge involving inquiry that have been proposed within formal epistemology. In particular, we object to earlier work in which a similar notion of issue-sensitive knowledge has been formalized, providing us with criteria for our own framework. We formally flesh out key concepts—information, issues and issue-relevance—and put these together in structures on which knowledge as issue-relevant information can be defined. The laws to which this notion of knowledge adheres are investigated, and contentious epistemic principles from standard epistemic logic are shown to be invalidated in our framework. They are replaced by restricted, weaker principles, that tend to the reasons why the stronger principles are considered contentious. Several phenomena that cannot be modeled in standard epistemic logic, such as paradigm shifts, are shown to be captured by our framework. We define five static epistemic logics that are increasingly expressive and accurately capture knowledge as issue-relevant information, without taking the individual issues on an agent's agenda into account. Subsequently, Dynamic Epistemic Issue Logic is developed, accommodating dynamic updates of an agent's agenda as well as information updates. All logics in this thesis are shown to be sound and complete, as well as decidable. We discuss some limitations of these logics, but argue that they suffice for most purposes. Lastly, we reflect on some adjacent work in formal epistemology, and point out directions for further research.

Acknowledgements

First and foremost, I would like to express my gratitude to my supervisor, Alexandru Baltag. It has been a privilege to work under your supervision. Not only because of your expertise and passion, but also because of the remarkable kindness and patience that you have shown me throughout this project. Our many meetings were always enlightening, and there was not a single one where you did not have something interesting to say (even if it was not always about logic or epistemology). Thank you!

I would also like to thank the members of the Thesis Committee—Maria Aloni, Aybüke Özgün and Malving Gattinger—for reading my thesis and providing insightful comments and questions.

My heartfelt thanks to Wijnand van Woerkom for meticulously proofreading this thesis, and to Joost Hooyman for his assistance in ensuring that all figures have a uniform appearance.

I have met many people at the ILLC who made writing this thesis—and the Master of Logic in general—a more enjoyable experience. While it is impossible to name everyone, I thank all who have contributed to the vibrant atmosphere at the ILLC over the past years. In particular, I would like to thank the “fixed residents” of the MoL room—Paul Seip, Wijnand van Woerkom and Hugh Mee Wong—for enduring countless hours in my presence. Additionally, I would like to thank Bo Flachs and Philemon Huising for enriching my time at the ILLC.

Next, I would like to thank my family and friends for their support. In particular, my parents, for their financial support, but even more for their moral support. I must also thank some friends for checking in with me, allowing me to vent my thoughts, and sometimes providing me with much-needed distractions. Thank you Daniël Everts, Mathijs Geurts, Joost Hooyman, Joost Janssen, Tilman Karsten, Max Meeuwesen, Hugo Peels, Hugo van Winden and everyone from Atlas.

Finally, I would like to thank my love, Bibi Schut, for her unwavering support throughout the entire process. Your steadfast presence and encouragement have been a source of comfort and motivated me to persevere.

Contents

Chapter 1

Introduction

Ever since the publication of the seminal works by Von Wright (1951) and Hintikka (1962), knowledge and information possession are often equated in standard epistemic logic.[1] Later work on dynamic epistemic logic continued this tradition, concerning itself primarily with information updates and deriving the dynamics of knowledge from those.[2]

This need not be a problem in itself. For many applications in fields like AI, computer science and economy, it suffices to study only information (change). Our main concern will not be with such applications, but rather with a more fundamental, philosophical matter: elucidating the properties of knowledge through logic. Naturally, epistemic logic fails to capture ordinary uses of 'knowledge'. Mainly on the grounds that agents within the framework of epistemic logic are endowed with reasoning skills and computational capacities reaching far beyond those of ordinary agents. Therefore, it is clear that epistemic logic is not able to represent the epistemic states of ordinary human agents: epistemic logic was developed in order to model knowledge in an ideal setting, thereby investigating the inherent properties of knowledge.[3] This is the matter that concerns us: the inherent properties of knowledge, not viewed in light of the epistemically flawed world that we inhabit, but in an absolute light devoid of human fallibility. To put it bluntly, we are interested in the basic properties of knowledge that rational agents have in an ideal setting.

Yet, even in an ideal setting, knowledge may depend on more than hard information. In fact, the view that knowledge amounts to information possession is widely contested. Notwithstanding that many notions of knowledge are not phrased in terms of information at all, there are numerous examples. For instance, Williamson (2000) claims that knowledge is information that allows a margin for error. Proponents of the relevant alternatives approach, first put forward by Dretske (1970), maintain that knowledge is information that is immune to error through the elimination of *relevant* alternatives. Yablo (2014) considers knowledge to be topic-sensitive: a topic must be within an agent's epistemic reach before any information about that topic is processed into knowledge.

[1] This is becomes clear when looking at some of the basic resources concerning epistemic logic. Rendsvig and Symons (2021), express knowledge purely in terms of informational indistinguishability of worlds. Fagin et al. (1995, §2.1) express the standard model of knowledge completely in terms of information. And Van Ditmarsch et al. (2015) describe knowing as an *informational* attitude.

[2] Again, this is confirmed in some of the basic resources on dynamic epistemic logic. For instance, see the work by Baltag and Renne (2016) or Baltag, Van Ditmarsch and Moss (2008). The standard textbook on dynamic epistemic logic states that "the book provides a logical approach to change of information" (Van Ditmarsch, Van der Hoek and Kooi, 2008, p. 1).

[3] Hintikka (1962, §2) has made this clear from the start.

Others, such as Schaffer (2007b) and Hintikka (2007), argued that knowledge is question-sensitive; agents process information into knowledge relative to their questions.[4]

Following this line of research, we will take inquiry to be a fundamental activity that guides the information processing of agents. 'Inquiry' should be understood in its etymological sense, referring to the act of querying or requesting information. The idea that knowledge is dependent on inquiry is not novel. It can be traced back to at least as early as Plato and has continued to appear in philosophical work ever since. In general, underlying this idea is the thought that the rational mind does not drift aimlessly in a vast ocean of information, but that it is directed at certain epistemic goals. Inquiry, then, is the activity that charts the course of the rational mind.

One area where this idea has gained significant traction is in the philosophy of science. Since Kuhn's 1962 book *The Structure of Scientific Revolutions*, there has been ample attention for the paradigms that shape scientific practice. A paradigm determines which questions are meaningful and which research directions should be pursued. Thus, an agent's inquiry is shaped by the prevailing paradigm. Information is processed differently in distinct paradigms, leading to disparate epistemic outcomes. An epistemic framework that takes inquiry into account should therefore be able to model phenomena such as paradigm shifts.

This thesis is concerned with the development and study of a logical framework in which not only information possession, but also inquiry is taken into account.

We consider epistemic issues, which we will simply call issues, as the objects of inquiry. The issues pursued by an agent constitute her epistemic agenda, which we will just call her agenda. An agent's agenda shapes and limits her knowledge by restricting the processing of information into knowledge to information that is relevant with respect to the issues on her agenda. In this thesis we formally study the following notion of knowledge:

> An agent knows that P if, and only if, the agent possesses the information that P, and P is relevant with respect to the issues on the agent's agenda.

In a shorter phrasing: knowledge is issue-relevant information. Issues can be viewed as requests for particular pieces of information. The requested pieces of information are the ones that are relevant with respect to that issue. Thus, available information is processed into knowledge only if that information helps an agent to (partially) resolve her issues.

A first attempt at formalizing this notion of knowledge was made by Baltag, Boddy and Smets (2018). One of the aims of this thesis is to improve on their framework, by proposing a formalization that we argue to be more faithful to the underlying philosophical conception.

When viewing knowledge as issue-relevant information, knowledge does not only depend on the information available to an agent, but also on her agenda. Therefore, any dynamic framework capturing this notion of knowledge in a satisfactory manner should not only accommodate information updates, but also agenda updates. The framework presented in this thesis accommodates both.

Throughout this thesis, we restrict ourselves to intensions. Only the content of a proposition is taken into account; its form is ignored. Thus, in formal terms, we identify propositions with the sets of possible worlds in which they are true.

We also assume that rational agents are aware of their agenda, deliberately choosing which issues to include—and which not. This comes down to agents being in full control of their inquisitive activities, which is appropriate in an idealized setting. Furthermore,

[4]Doxastological counterparts of this view can be found in the works of Yalcin (2018) and Hoek (2022; forthcoming).

this assumption fits well with actual scientific practice: scientists are generally aware of the research directions they pursue.

In summary, the goal of this thesis is to logically study knowledge as issue-relevant information in an ideal setting. We restrict ourselves to the intensions of propositions and assume rational agents to be in control of their agenda.

In Chapter 2, we put forward a detailed motivation for studying knowledge as issue-relevant information. We look at philosophical ideas supporting the view that knowledge depends on inquiry. Several examples illustrating the interplay between knowledge and inquiry are given, motivating the incorporation of issues in an epistemic logic. Typical principles from epistemic logic are discussed and it is argued which should hold in a framework incorporating issue-relevance. For the ones that should not hold, we tentatively suggest novel, restricted principles. Finally, we discuss an earlier paper by Baltag, Boddy and Smets (2018) in which a notion of knowledge that depends on an agent's issue(s) is introduced. We argue that this notion of knowledge does not capture the epistemic influence of issues in a satisfactory manner. These observations serve as the starting point for our formal definition of knowledge as issue-relevant information.

In Chapter 3, we start setting forth our logical framework. We start with a formal discussion of the key ingredients of our notion of knowledge: information, issues and issue-relevance. Concurrently, the epistemic issue structures used to formalize our notion of knowledge are introduced. It is explained how agenda updates and the resolution of an issue can be modeled with these structures. Lastly, a formal definition of knowledge as issue-relevant information is given.

In Chapter 4, we investigate the properties to which knowledge and issue-relevance adhere and show how the examples from Chapter 2 can be captured in epistemic issue structures. We investigate the relation between knowledge and truth, at the same time defining a notion of validity in epistemic issue structures. Then the properties of issue-relevance are explored, which are crucial for expressing the laws of knowledge. Next the laws of knowledge in static contexts are investigated, focusing on closure and introspection principles. Thereafter, the properties of the Kripke modalities over the information relation and the issue relation are studied, showing how they relate to the notions of knowledge and issue-relevance. Lastly, the more intricate mechanics of knowledge that involve the actions of agenda update and issue resolution are investigated.

In Chapter 5, we turn to static logics on epistemic issue models that do not take the individual issues comprising an agent's agenda into account. It is argued that a logic based on a language with only a knowledge modality is not expressive enough to capture knowledge as issue-relevant information in a satisfactory manner. Thereafter, several static logics are introduced based on more expressive languages containing modalities for knowledge, issue-relevance, necessity, information and issues, in varying compositions. Each of these logics is proven to be sound and complete, and their mutual relationships are discussed. Lastly, an example is given showing these logics in action.

In Chapter 6, a dynamic logic of epistemic issues is studied. We first introduce Epistemic Issue Logic, which is able to pre-encode agenda updates and issue resolutions. This static logic of epistemic issues is proven to be sound and complete. Thereafter, it is extended into Dynamic Epistemic Issue Logic, which accommodates agenda updates and issue resolutions. Dynamic Epistemic Issue Logic is proven to be sound and complete via reduction axioms. Moreover, decidability is proven for Epistemic Issue Logic, yielding decidability for all logics presented in Chapters 5 and 6 as a corollary. We reflect on the limitations of Dynamic Epistemic Issue Logic and argue that it suffices for practical purposes. This is illustrated by revisiting one of the earlier examples and showing Dynamic Epistemic Issue Logic in action. This chapter concludes the exposition of our logical framework.

In Chapter 7, we compare our framework to other work in formal epistemology and suggest some directions for further research. We start by considering the possibility of a multi-agent extension of our framework and suggest possible notions of group knowledge in such an extension. Subsequently, we discuss epistemic relevance. First we argue that the notion of issue-relevance in our framework differs radically from the notion of relevance in epistemic logics based on relevance logics. Then we turn to the relevant alternatives approach and argue that such an approach may be integrated in our framework, suggesting this as potential further research. Next we look at topic-sensitive approaches to knowledge. Our logic can be considered an intensional counterpart of these hyperintensional logics, and we suggest how our logic can be modified in a sensible manner such that it becomes hyperintensional as well. Lastly, we consider doxastic variants of our notion of knowledge and explain how our framework can be adapted to capture those.

In Chapter 8, a comprehensive conclusion is given.

Lastly, the reader is assumed to be mathematically mature and familiar with dynamic epistemic logic, as it is presented in textbooks such as those by Van Ditmarsch, Van der Hoek and Kooi (2008) or Fagin et al. (1995). Additionally, the reader should be familiar with some standard modal logic, roughly Chapters 1, 2, 4 and 6 in the textbook by Blackburn, De Rijke and Venema (2001).

Chapter 2

Knowledge and Inquiry

In this chapter we motivate our logical study of knowledge as issue-relevant information. In Section 2.1, the idea that knowledge is dependent on inquiry is partially traced throughout philosophical history, emphasising recent developments in formal epistemology. In Section 2.2, the notion of knowledge as issue-relevant information is presented comprehensively and its logical study is motivated. In Section 2.3, we explore epistemic principles commonly found in epistemic logic from a conceptual perspective. We suggest which principles should govern the notion of knowledge as issue-relevant information. In Section 2.4, we outline an earlier proposal by Baltag, Boddy and Smets (2018), who formalized a similar notion of knowledge. In Section 2.5, however, we will argue that their framework does not capture this notion in a satisfactory manner, providing us with desiderata for our own framework. Throughout this chapter—with the exception of Sections 2.4 and 2.5—formalities are avoided and postponed to later chapters.

Finally, before proceeding, it should be mentioned that many of the epistemological problems discussed in this chapter have doxastological counterparts. Some of the cited literature concerns itself with belief rather than knowledge. A careful selection has been made to ensure that the arguments in cited literature are equally applicable in an epistemological setting. Furthermore, this decision is justifiable when knowledge is viewed as a special case of belief.

2.1 Inquiry in epistemology

The idea that knowledge is dependent on inquiry is ancient. In Plato's *Meno* this connection is emphasized when Socrates is confronted with Meno's paradox: "a man cannot search either for what he knows or for what he does not know".[5] Socrates replies to the paradox by arguing that there are epistemic states in which there is neither knowledge nor pure ignorance. In these states, Socrates argues, knowledge may be obtained by just asking questions. He proceeds to demonstrate this by showing that a slave may learn geometrical truths only by asking him a series of questions. As these questions do not provide the slave with new information, Socrates claims that these ideas must have already been within the slave prior to his questioning. Plato thought that because our souls are immortal, we have innate ideas before birth that we may recollect through inquiry. Although these metaphysical ideas have fallen out of favor, we will agree that agents may obtain knowledge through inquiry alone. Not because they recollect knowledge from previous lives, but because the information was already possessed implicitly.

[5] Quoted from a translation of Plato's (1997, p. 880, 80e) *Meno*.

More generally, in Plato's work, inquiry plays a pivotal role when it comes to acquiring philosophical knowledge.[6] Questioning is central to Socrates' method of *elenchus*, along with the critical examination of the answers to these questions. In Plato's dialogues, Socrates outmanoeuvres his opponents by asking the right questions. On top of that, although Socrates himself typically does not claim to have knowledge, he obtains knowledge of his own ignorance through questioning.

Even in Aristotle's work on logic and reasoning, inquiry may have played a more central role than often supposed. Hintikka (2006) argues that this work should be understood in a dialectical context, where reasoning proceeds by question-answer dialogues similar to the Socratic *elenchus*. This is clear in texts like the *Topics*. Yet Hintikka goes further, claiming that even Aristotle's *Analytics*, the conclusion of a logical inference should be understood as the answer to a question.[7]

After the era of ancient Greek philosophy, the idea remained largely dormant for centuries. It resurfaced in different forms intermittently. When Bacon overturns Aristotle's scientific method, questioning remains central to the method that he puts forward as a replacement.[8] In *The Great Instauration*, he compares scientific practice to the questioning of nature.[9] Kant makes a kindred comparison when he states that experiments such as the ones conducted by Galileo, compel nature to answer scientists' questions (Kant, [1787] 1998, p. B xii-xiii).[10]

The importance of inquiry is implicit in pragmatism and can already be discerned in the works of early pragmatists like Peirce, James and Dewey. It is not surprising as a byproduct of the more broad pragmatist thesis that truth is determined by practicality. If truth depends on what is practical for an agent and inquiry is the process by which an agent directs her mind towards what is practical for her, then it is clear that (veridical) knowledge must depend on inquiry. Consequently, most pragmatist conceptions of knowledge are intimately connected to inquiry.

Abduction, as conceived by Peirce, can be interpreted as the formulation of a question that directs attention to certain pieces of information. Thus abduction brings new information into a line of thought via inquiry. Eventually, knowledge is the result of such a line of thought, thus depending on inquiry.[11]

The idea is even more conspicuous in William James' voluntaristic epistemology.

> "We carve out everything, [just as we carve out constellations], to suit our human purposes. For me, this whole 'audience' is one thing, which grows now restless, now attentive. I have no use at present for its individual units, so I don't consider them" (James, [1907] 1979, p. 100).

According to James, reasoning only proceeds with information that has been carved out. Inquiry is the process of carving out information; it carves out the information relevant for the agent. Since inquiry determines which information is considered by an agent in the first place, on James' conception, knowledge depends on inquiry.

[6] An idea that survives to this time, to which Habgood-Coote, Watson and Whitcomb (2022a; 2022b) bear witness.

[7] See, for instance, the work by Hintikka (1999, p. 2).

[8] Bacon ([1620] 1902) argues at length for his new method in *Novum Organum*.

[9] We consulted the following version of the work by Bacon ([1620–1626] 2017).

[10] Kant even tacitly suggests that reason can only produce knowledge of nature that answers a question. "Reason, in order to be taught by nature, must approach nature with its principles in one hand, according to alone the agreement among appearances can count as laws, and, in the other hand, the experiments thought out in accordance with these principles – yet in order to be instructed by nature not like a pupil, who has recited to him whatever the teacher wants to say, but like an appointed judge who compels witnesses to answer the questions he puts to them" (Kant, [1787] 1998, p. B xiii).

[11] Hintikka (1998, p. 524) views abduction as an interrogative step, which can be captured by his own interrogative model of inquiry. Hintikka (1998) treats the relation between his interrogative model of inquiry and Peirce's abduction more thoroughly.

Dewey holds a similar view that is a bit more developed and tacitly suggests the notion of knowledge that we propose. In *How We Think*, Dewey distinguishes five phases of thinking.

> "Upon examination, each instance reveals, more or less clearly, five logically distinct steps: (i) a felt difficulty; (ii) its location and definition; (iii) suggestion of possible solution; (iv) development by reasoning of the bearings of the suggestion; (v) further observation and experiment leading to its acceptance or rejection; that is, the conclusion of belief or disbelief" (Dewey, 1910, p. 72).

The first three phases roughly correspond to the formulation of an issue: an agent identifies which information is required to solve a problem and inquires into that information. This process is akin to Peirce's abduction. In the fourth phase the agent reasons, restricting herself to information that is relevant to the formulated issue. All information that has no bearing on the formulated issue is ignored and can therefore never be the conclusion of such reasoning. The knowledge obtained through this reasoning can therefore rightly be considered information that is relevant with respect to the agent's issue. The fifth phase concerns the actions required to resolve the issue, in case reasoning alone cannot accomplish this. In short, Dewey considers inquiry to be fundamental in reasoning and knowledge acquisition.

After the rise of pragmatism, the idea reappears in Collingwood's logic of question and answer. Collingwood (1940, p. 23) claims that all statements are made in answer to a question. He interprets 'statements' in a broad manner, so that knowledge also becomes the answer to a question (Collingwood, 1940, p. 43). Additionally, agents need not be aware of the questions they ask, allowing inquiry to occur partially unconsciously. Nonetheless, knowledge depends heavily on it. Collingwood (1946) applies these ideas to explain the acquisition of historical knowledge. This framework was later adopted by Gadamer (1975), who also argued that historical knowledge is dependent on inquiry.[12] However, it should be observed that the logic of question and answer does not pass for a logic by today's standards.[13]

In the second half of the twentieth century inquiry became more central in the philosophy of science, albeit implicit sometimes. In his seminal work *The Structure of Scientific Revolutions*, Kuhn (1962) argues that scientific knowledge is not gradually accumulated, but that periods of "normal" science are alternated with periods of revolutionary science. During periods of revolutionary science paradigms change. Paradigms are comprehensive frameworks of assumptions, concepts, and practices that form the basis for understanding, interpreting, and conducting scientific research. It also determines which questions are relevant and meaningful, thereby shaping the process of scientific inquiry. In fact, the questions pursued by scientists are telling for the paradigm in which they are asked. Background assumptions can be viewed as the presuppositions of scientific questions. Additionally, experiments are designed to answer particular questions, as was already observed by Bacon and Kant. A paradigm shift can thus also be interpreted as a change in inquiry.[14] Through Kuhn's work, scientific knowledge, as it is often understood in philosophy of science, became contingent on inquiry.

A more comprehensive historical analysis of the idea that knowledge depends on inquiry is beyond the scope of this thesis. Even so, at this point, the reader should be convinced that on the one hand, the idea is not novel, but on the other hand, it is not a banality that has been presupposed throughout the history of epistemology.

[12] In particular, see the section *The logic of question and answer* (Gadamer, 1975, p. 363-371).

[13] Arguably, it also did not at the time of their publication.

[14] Hintikka (2007, p. 84) goes as far as comparing a body of questions in Collingwood's setting to a Kuhnian paradigm.

The remainder of this section is dedicated to exploring various approaches to knowledge or belief in formal epistemology from recent decades that align with the research direction of this thesis. We return to many of these in Chapter 7, when comparing the logical framework developed in this thesis to adjacent work and exploring potential further research directions.

The relevant alternatives approach forms a more modern strain of knowledge theories that incorporate inquiry. The approach was pioneered by Dretske (1970), who stated: "To know that x is A is to know that x is A within a framework of relevant alternatives, B, C, and D" (Dretske, 1970, p. 1020). Stating it more generally, central to any relevant alternatives theory is the following claim:

> An agent knows that P if, and only if, the agent possesses sufficient information to rule out all relevant alternatives.

The differences between relevant alternatives theories can be found in the manner in which the terms 'relevance', 'alternative' and 'ruling out' are specified (Hawke, 2017, p. 4). For now, however, these differences are not relevant to our discussion.

Thus, a relevant alternatives theorist argues that an agent only has to consider alternatives that are relevant with regard to the desired object of knowledge, that is, the object of knowledge into which an agent is inquiring. In terms of a possible world approach to knowledge: an agent knows that P if, and only if, P is the case in every relevant possible world consistent with the agent's information. So even if an agent's information does not exclude the possibility that some proposition P is false, she may still know P if the possible worlds in which P does not hold are irrelevant. Thus, inquiry on this approach involves a form of presupposition: some alternatives are presupposed to be irrelevant given a line of inquiry.

One instance of a relevant alternatives theory that makes this explicit is Schaffer's (2004; 2006; 2007b) theory of contrastive knowledge. He argues that the knowledge relation is ternary rather than binary: "s knows that p as an answer to question Q" (Schaffer, 2006, p. 235). The view is elegantly captured by the maxim "to know is to know the answer" (Schaffer, 2007b, p. 401), mirroring Collingwood's claim that all knowledge is the answer to a question. The presuppositions of a question, then, determine which alternatives are relevant to know its answer.

An alternative perspective on the relation between knowledge and inquiry is given by Hintikka's (2006) Socratic epistemology and interrogative model of inquiry.[15] Hintikka also attempted to restore the position of inquiry in the study of reasoning and knowledge, taking knowledge acquisition to essentially be a questioning procedure (Hintikka, 2007, p. 17-18). He argues that, consequently, epistemologists should shift their focus from the concept of knowledge to inquiry. Thus, the goals of an agent's inquiry are of central concern; they guide the manner in which agents process information. In Hintikka's words: "all information used in an argument must be brought in as an answer to a question" (Hintikka, 2007, p. 19).

Consequently, attempting to define knowledge itself is a folly according to Hintikka: "The criteria of knowledge concern the conditions on which the results of epistemological inquiry can be relied as a basis of action" (Hintikka, 2007, p. 30). As a consequence, defining knowledge in a general epistemological theory is futile: the criteria for knowledge depend on its application. One consequence of Hintikka's view that knowledge cannot be defined is that the logic corresponding to the interrogative model of inquiry deals primarily with questions, answers and sequences of question-answer pairs (Hintikka, 1999, p. x). It is a logic of knowledge seeking through questioning rather than a logic of knowledge. In a sense, it is the logic of question and answer that Collingwood and Gadamer did not provide. Hintikka (1999, p. 102) argued that these sequences of

[15] For the interrogative model of inquiry, see the works by Hintikka (1981; 1999).

question-answer pairs could be interpreted as question-answer dialogues, matching the Socratic *elenchus* and the dialectic conception of Aristotle's work on reasoning.

The stance that the objects of knowledge or belief are not unstructured propositions is not exclusive to Collingwood, Schaffer and Hintikka. For instance, Hoek (2019) takes the objects of belief to be ordered pairs of propositions and questions, calling these quizpositions.[16] Much like Hintikka, Hoek introduces quizpositions to bridge the gap between the standard account of belief and the goal-directed nature of our beliefs. There exists a myriad of philosophical papers in which the objects of knowledge or belief are propositions that are somehow structured by inquiry or a similar activity.[17]

Nonetheless, unstructured propositions are also still considered to be the objects of knowledge or belief by many, even in settings that incorporate inquiry. For instance, the notion of question-sensitive belief presented by Yalcin (2018) takes unstructured propositions as the objects of belief.[18] On Yalcin's view, the possible worlds considered by agents are typically not maximally specific (Yalcin, 2008, p. 107). Therefore, he argues, we should consider belief states that are coarser than the sets of possible worlds that are usually considered. To achieve this without sacrificing the convenient possible world picture of belief, he proposes to consider resolutions of logical space (Yalcin, 2008, p. 107). They can be considered representations of the question(s) an agent is entertaining or the subject matter(s) in which she is interested (Yalcin, 2018, p. 32). Resolutions structure logical space, foregrounding certain propositions, while backgrounding others. The backgrounded propositions are the ones that escape an agent's attention; the worlds she considers are not specific enough to conceptually distinguish them. Only the foregrounded propositions are candidates for belief. Hence, belief as conceived by Yalcin also depends on inquiry.

The framework of Baltag, Boddy and Smets (2018) bears some resemblance to Yalcin's approach, but in the spirit of Schaffer and Hintikka. The questions an agent entertains govern the conceptual distinctions she makes, shaping and limiting the manner in which she processes information. We comprehensively discuss this framework in Section 2.4.

Within Yalcin's approach, only the intensional content of a proposition is taken into account. However, since recent work by Yablo (2014), an approach going beyond intensions has gained in popularity. In addition to the intension of a proposition, its topic is also taken into account on this approach.[19] The truth conditions of a knowledge ascription, then, do not only consist of a condition taking care of the proposition's intension, but also of a hyperintensional condition that is only true if the topic of the proposition in question is within epistemic reach of the agent.[20] The hyperintensional condition can take on different forms, but discussing these is beyond the scope of this thesis. We call notions of knowledge that take the topics of propositions into account topic-sensitive.[21] Inquiry can be considered as the process that determines the topics in which an agent is interested, so that knowledge can be viewed as dependent on inquiry.

[16]Also see other work of Hoek (2022; forthcoming).

[17]For instance, see the works by Blaauw (2012), Koralus and Mascarenhas (2013), Friedman (2019) and Holguin (2022); but this list is certainly not exhaustive.

[18]Also see the work by Yalcin (2008), who in turn borrows ideas from Stalnaker (1984). Note that Yalcin (2018) takes unstructured propositions to be the objects of belief relative to an agent's belief state.

[19]Topics are also referred to as subject matters in the literature. For the purpose of clarity, we use 'subject matter' when intending a reading that only involves intensions and 'topic' otherwise.

[20]A concept is called hyperintensional if it distinguishes between necessarily equivalent contents (Berto and Nolan, 2023).

[21]Some recent examples of topic-sensitive notions of knowledge are put forward by Yablo (2014, Ch. 7), Berto (2019; 2022), Hawke, Özgün and Berto (2019), Özgün and Berto (2021), Berto and Hawke (2021) and Berto and Özgün (2023).

This wraps up our survey of recent approaches to knowledge that involve inquiry. It furnishes a backdrop against which we can discuss the notion of knowledge as issue-relevant information. We take up this endeavor in the next section.

2.2 Knowledge as issue-relevant information

Epistemic issues, which henceforth we will refer to simply as issues, are the objects of inquiry. Issues are requests for information, highlighting particular aspects of the world. Questions are prototypical issues, and so are subject matters. A question is a request for the information that answers it, while a subject matter can be construed as a request for information pertaining to that particular subject matter. Throughout this thesis, we will often focus on these two particular instances of issues while keeping the more general characterization above in mind.

The slave in Plato's *Meno*, then, obtains knowledge because the questions posed by Socrates compel the slave to add these questions to his epistemic agenda. The slave's epistemic focus is shifted to several issues concerning geometry, causing him to process the information that he already possessed into knowledge. This ancient example motivates the following notion of knowledge:

> An agent knows that P if, and only if, the agent possesses the information that P, and P is relevant with respect to the issues on the agent's agenda.

The core idea behind this notion of knowledge is that information is processed relative to the issues on an agent's epistemic agenda, which we henceforth refer to simply as the agent's agenda. We call information relevant with respect to the issues on an agent's agenda issue-relevant. Information that is not issue-relevant is called issue-irrelevant. Only issue-relevant information is processed into knowledge.

In case of a question, it is clear which information is issue-relevant: those pieces of information that (partially) answer the question. So, when considering questions, the notion above can be rephrased as: an agent knows that P if, and only if, the agent possesses the information that P, and P is a partial answer to the agent's questions. In case of subject matters, it can be rephrased as: an agent knows that P if, and only if, the agent possesses the information that P, and P is about a subject matter in which the agent is interested. In general, information is issue-relevant whenever it (partially) resolves the issues on an agent's agenda.

An agent's agenda limits her knowledge: any information that is not issue-relevant cannot be known. Furthermore, an agent's agenda also shapes her knowledge. Different agendas may give rise to different epistemic states, even if the information possessed by an agent is the same in these cases. For instance, the slave in the *Meno* might have learned very different geometrical truths if Socrates would have asked him different questions. So inquiry plays a fundamental role when taking knowledge to be issue-relevant information.

If issue-relevance is taken into account, situations that cannot be captured in terms of knowledge as information possession can be modeled. Consider the following example, based on an example from Stalnaker (1984, p. 88).[22]

Example 1 (The King of England)**.** In 1700, King William III of England was on the brink of war with France. He wanted to know whether war with France could be avoided. At some point, after receiving adequate intelligence from his subordinates, he knew that war with France could be avoided. However, he did not know that *nuclear* war with France could be avoided.

[22]Note that Stalnaker's example concerns belief, while ours concerns knowledge.

Although the information that war can be avoided entails that nuclear war can be avoided, the King of England did not know the latter. His agenda only contained the question whether war with France could be avoided. The King's issues did not prompt him to consider logical space with a fine enough granularity to conceptually distinguish different types of warfare. So the information that war could be avoided was issue-relevant, whereas the information that nuclear war could be avoided was not. As a consequence, the king had no knowledge about avoiding nuclear war.

The example shows how issues carve logical space, much in the spirit of James: sometimes issues prompt one to only carve out the audience, sometimes to carve out every individual comprising the audience, and sometimes to only carve out some of these individuals. The concepts carved out by the issues on an agent's agenda are the building blocks of her knowledge. It is in this way that an agent's agenda shapes and limits her knowledge.

The carving of logical space, however, goes further than just determining the granularity of an agent's perception. In a sense, it also determines *how* the world is perceived. Suppose the only question on an agent's agenda is whether some object is grue.[23] The knowledge of the agent pertaining to this object will be restricted to knowing it is grue or not grue. If there are issues on the agent's agenda asking whether the object is green, and asking whether the object is blue, then the agent's knowledge of the object's colour will be in terms of green and blue.

So, at least *prima facie*, some issues seem to be better to have on your agenda than others. For instance, it seems generally better to know that an object is green or that it is blue than that it is grue. Lewis (1983) argues that some properties are more natural than others, like green is more natural than grue. So we could put it this way: it is better to have issues that inquire into more natural properties.

Sider (2011) goes even further, arguing that reality has an objective structure. Some concepts carve at the joints of reality's structure, while others do not. According to Sider, "it's better to think and speak in joint-carving terms" (Sider, 2011, p. 61). Since issues carve logical space, it would be better to have issues that carve at the joints of reality's structure. We remain neutral on these metaphysical matters, only pointing out that the epistemic value of different inquiries may differ depending on one's views.

This can be exemplified by looking at phenomena from the history and philosophy of science. In scientific practice, the issues on researchers' agendas sometimes change so that aspects of how the world was perceived before become meaningless.

Example 2 (Relativity of simultaneity)**.** When Einstein introduced his special theory of relativity, our understanding of time and space was revolutionized. Before Einstein, it was commonly assumed that events happening simultaneously for one observer would be simultaneous for all observers. However, Einstein's theory showed that simultaneity is relative and depends on the observer's frame of reference. This rendered the question of whether two distant events truly occur simultaneously meaningless; simultaneity became a relative concept rather than an absolute one.

In terms of our framework: Einstein's revolutionary special theory of relativity forced researchers to retract issues pertaining to the simultaneity of events from their agendas. Any information about the simultaneity of events became issue-irrelevant. Consequently, there no longer was scientific knowledge about simultaneity.

There exist many similar examples. For instance, consider Heisenberg's (1927) paper in which he introduces the uncertainty principle.[24] The formulation of the uncertainty principle in quantum mechanics represents another revolutionary step in science. It

[23] The predicate 'grue' was introduced by Goodman ([1955] 1983): it applies to objects that are observed to be green before a certain time t or observed to be blue after t.

[24] Also see the English translation of Heisenberg's (1983) paper.

states that it is impossible to simultaneously determine both the exact position and momentum of a quantum particle with absolute precision. This principle challenged the classical idea that it was possible, in principle, to precisely measure both the position and momentum of a particle at any given moment. According to some interpretations of quantum mechanics, Heisenberg's uncertainty principle implies inherent limitations to the precision of such measurements. This leads to the conclusion that certain issues, such as the ones pertaining to the exact position and momentum of a particle at a specific moment, are fundamentally meaningless or unknowable.

These examples fit the picture of scientific progress famously described by Kuhn (1962). In general, revolutionary steps that change a paradigm correspond to the retraction of issues. Certain questions, subjects, experiments and methods become meaningless and irrelevant with respect to a researcher's agenda. As a new paradigm is developed to replace the old one, issues are added to the researcher's agenda, leading to new methods, questions and concepts. A complete paradigm shift can thus be described as a process in which issues are both retracted and added to an agenda. The example below neatly captures this.

Example 3 (Galileo's telescope)**.** Galileo Galilei's introduction of the telescope in the seventeenth century led to a significant shift in our understanding of the natural world. Prior to his discoveries, the prevailing belief, influenced by Aristotle, held that direct observation alone could reveal the true nature of the celestial bodies. Consequently, Aristotle and many after him "knew" the moon was perfectly smooth when directly observed. They obtained this "knowledge" by observing the moon with the naked eye.

However, Galileo did not believe that direct observation could reveal the nature of the celestial bodies to us. He believed that our senses could be aided by instruments, like a telescope, to gain new insights into nature. Galileo's telescopic observations revealed previously unseen features on the Moon's surface, such as mountains, valleys, and craters, contradicting the idea of a perfectly smooth lunar sphere. Consequently, Galileo knew that the moon contains craters and mountains when observing it with a telescope.

An Aristotelian agent typically has an issue on her agenda concerning the direct observations that we could make of the moon. The information obtained by any other kind of observation would be issue-irrelevant to those agents. The Aristotelian agent, then, resolved her issue pertaining to the surface of the moon by simply looking at the moon. This provided the issue-relevant information that the moon seemed perfectly smooth when observed by the naked eye, leading to knowledge of the moon's smooth surface. Galileo, however, retracted the issue concerning direct lunar observations from his agenda and added the issue concerning telescope observations of the moon. The information obtained by directly looking at the moon was no longer issue-relevant, instead information gained through telescope observations became issue-relevant. To resolve the new issue, he observed the moon through a telescope, thereafter knowing that the moon contains craters and mountains when observed with a telescope.

The dynamic action of agenda updates that we envision will allow us to model phenomena like paradigm shifts and the revolutionary steps accompanying it. This is not possible in standard dynamic epistemic logic, which further motivates the study of knowledge as issue-relevant information.

The dynamic action of information update that we envision corresponds to the resolution of an issue. It can be considered as an instance of the fifth phase in Dewey's five phases of thinking: it constitutes the action(s) required to resolve an issue. In Galileo's case, looking through his telescope provided him with the information that resolved his issue. In subsequent chapters, we revisit the examples from this section, showing how they can be captured in our framework.

Lastly, we can situate the notion of knowledge as issue-relevant information among the approaches to knowledge discussed in the previous section.

Despite some nomenclatural similarity, the notion of knowledge as issue-relevant information should not be considered a relevant alternatives notion of knowledge. It should rather be understood as a relevant distinctions notion of knowledge.[25] The agents we envision only processes information that (partially) resolve their issues, thus making only conceptual distinctions that are relevant to their inquiry. In Example 1, the conceptual distinction between war or peace was relevant, whereas the conceptual distinction between nuclear war or any other type of war was irrelevant. This does not mean the King of England neglected some alternatives, but only that he did not consider them in full detail. Put differently: the issues we consider do not have presuppositions. Thus the relevant alternatives approach aims primarily at a different aspect of epistemic relevance than we do. We reflect further on this in Sections 4.3 and 7.3.

We diverge from Schaffer's relative alternatives theory of knowledge on more points. For instance, we interpret issues in a broader manner than just as questions. Moreover, although our notion of knowledge is relative to the issues on an agent's agenda, we maintain that the knowledge relation is binary. Unlike Schaffer, we do not take an agent to know propositions as answers to questions—an agent just knows propositions. The issues on an agent's agenda form the background against which information is processed, but ultimately the knowledge relation is between subjects and the propositions they know. This aligns with the intuition that knowledge pertains directly to propositions.

The same observation applies to others who take the objects of knowledge to be propositions that are directed at specific questions.[26] We agree that inquiry plays a fundamental role when it comes to knowledge, but leave this role implicit in knowledge ascriptions.

Recall that Hintikka argued that when inquiry is recognized as fundamental in knowledge acquisition, defining knowledge is a folly. We do not concede this, nor do we deny it. It is clear that our notion of knowledge neglects important aspects of knowledge such as priorities, expectations, reliability of information sources, lack of cognitive resources, etc. At most, we claim that inquiry plays an important role in the attainment of knowledge. Yet it is not the goal of this thesis to present a theory of knowledge. The purpose of our study is to elucidate the logical properties of a particular notion of knowledge that takes inquiry into account.

We also deviate from Hintikka at this point: we intend to provide a logic of knowledge, not of question-answer sequences. Nonetheless, it will become evident that our framework, in its own way, is able to capture sequences of questions and answers.

Topic-sensitive approaches to knowledge can be considered the hyperintensional counterparts of our intensional approach. When interpreting issues as subject matters, the agents we envision can only come to know propositions relevant to the subject matters on their agenda. So, in a sense, our notion of knowledge is sensitive to subject matters. However, issues can only distinguish the subject matters of propositions in terms of their intensions, that is, in terms of possible worlds. Typically, topic-sensitive accounts of knowledge assume that the topic of a proposition is determined by the concepts and objects that it refers to. So on these accounts, the topic of a proposition also depends on the form of a proposition, not only on its content. This explains the hyperintensional nature of topic-sensitive notions of knowledge. In Section 7.4, we will discuss topic-sensitive notions of knowledge in more depth, comparing it to our notion of knowledge as issue-relevant information.

[25] This phrasing is borrowed from Baltag, Boddy and Smets (2018, p. 134).

[26] Such as Hoek (2019), Blaauw (2012), Koralus and Mascarenhas (2013), Friedman (2019) and Holguin (2022).

Our notion of knowledge comes closest to the question-sensitive belief of Yalcin (2018). The issues in our framework resemble his resolutions. However, he does not go beyond sketching models of belief, whereas we will develop a complete logical framework. Furthermore, Yalcin does not consider any actions that change the issues on an agent's agenda or update her information, which we will do.

Having discussed our notion of knowledge in relation to other existing approaches, we turn to the epistemic principles of knowledge as information possession in the next section. We will argue which of these principles should hold for the notion of knowledge as issue-relevant information and suggest novel principles for the ones that should not hold.

2.3 Principles in epistemic logic

One key principle to which agents adhere in standard epistemic logic is the closure of knowledge under deduction. This principle states that whenever a proposition P entails a proposition Q and the agent knows P, the agent also knows Q.[27] Put simply, when knowledge is closed under deduction an agent knows all the logical consequences of her own knowledge. This principle is quite strong. So, naturally, it is not accepted by all philosophers.[28]

For instance, Stalnaker (1991) argues that closure under deduction leaves no room for deductive reasoning since an agent has no use for it. Therefore it distorts any context in which an agent engages in reasoning. Consequently, aspects of an agent's epistemic life that govern her reasoning activities become redundant, such as her goals, priorities and expectations. Because reasoning lies at the heart of rational activity, it is argued that an account of knowledge in which reasoning plays no role may be incomplete, even in an ideal setting (Stalnaker, 1991, p. 428-429). Moreover, it is argued that not all consequences of one's knowledge are worth knowing explicitly, not even for rational agents. In Harman's words: "one should not clutter one's mind with trivialities" (Harman, 1986, p. 12).[29]

Others argue against closure under deduction on the basis of skeptical paradoxes.[30] Generally, they argue as follows. There are certain "skeptical" propositions that are impossible to exclude, like whether you are a brain in a vat, or some other proposition that cannot reliably be excluded given the information at your disposal. The impossibility to exclude these "skeptical" propositions conjoined with closure under deduction, they argue, makes it impossible to know any proposition that entails the denial of those "skeptical" propositions. In terms of the famous example given by Moore (1939): you cannot know you have hands, since the existence of your hands entails that you are not a brain in a vat. Although knowledge is not closed under deduction in our framework, this particular problem will not be tended to directly by our framework. However, the relevant alternatives approach to knowledge does tend to this problem, and we discuss it in relation to our framework in Section 4.3.

We accept arguments in the spirit of Stalnaker (1991); there must be room for aspects of an agent's epistemic life in a formal epistemological framework that go beyond information processing. Therefore knowledge should not be closed under logical consequence.

[27]Alternatively, a similar but non-equivalent principle is: if the agent knows both that P and that P implies Q, then the agent knows Q. For our purposes, however, the difference does not matter: given the principle of necessitation and factivity of knowledge that we will assume, 'P entails Q' and 'the agent knows that P entails Q' turn out to be equivalent.

[28]However, it is also defended by many, such as Stine (1976), Lewis (1996), Hawthorne (2005) and Kripke (2011).

[29]Even from an intensional point of view, a proposition P entails trivialities such as 'either P or Q'.

[30]For instance, Nozick (1981), Dretske (1970; 2005), Yablo (2014), Lawlor (2013), Holliday (2015) and Schaffer (2007a).

Within the framework of standard epistemic logic, assuming the epistemic accessibility relation to be transitive or Euclidean results in agents that are positively or negatively introspective, respectively.[31] A positively introspective agent has knowledge of her own knowledge, while a negatively introspective agent has knowledge of her own ignorance.

In general, negative introspection is rejected because it is very demanding. It requires an agent to consider all facts that she does not know. In Hendricks' words: "The axiom of wisdom or negative introspection is a sort of closed world assumption. A closed world assumption is a forcing assumption if anything is, 'shutting the world down' with the agent, leaving the skeptic nowhere to go" (Hendricks, 2005, p. 87). Hintikka (1962) rejected negative introspection from the beginning, and to our knowledge there exist no serious theories of knowledge that accept it.[32] In conjunction with positive introspection, it gives an agent full access to her epistemic state, a property deemed even too strong for ideally rational agents. For this reason, negative introspection is accepted almost exclusively in situations where it fits the context of a particular application.[33]

Positive introspection is more controversial, since knowing what you know is more plausible than knowing what you do not know. As a result, the literature on this principle is more extensive.[34] Lenzen (1978) gives an overview of some of the early debates on positive introspection.[35] Treating all arguments is beyond the scope of this thesis, but he mentions an argument against positive introspection that is of particular interest to our purposes. The argument stresses that knowing that you know P requires conceptual resources which are not needed to know that P, so that an agent might know P without knowing that she knows P (Lenzen, 1978, p. 74).[36] A distant echo of this argument can be found in our work. Remember that an agent only makes conceptual distinctions that are relevant with respect to her issue. If knowing that P is conceptually distinct from knowing that one knows P, then it might be the case that P is issue-relevant, while knowledge of P is not, suggesting that positive introspection should fail to hold.

Much of the more recent work on positive introspection is based on an objection due to Williamson (2000, Ch. 5), who claims that the principle of positive introspection fails if a margin of error principle for epistemic judgement is taken into account. The validity of this principle, however, is called into doubt by logicians.[37] Furthermore, in an ideal setting, we can assume a margin of error is not called for. Thus, if we wish to challenge the principle of positive introspection in our setting, we should turn to other arguments.

Conveniently, the arguments given against closure under deduction can partly be rehashed to object against positive introspection: if we are to avoid clutter, positive introspection is too strong. The subsequent example illustrates this point.

Example 4 (Exam question)**.** A high school student, Alice, is making a chemistry exam. One of the questions pertains to photosynthesis: 'is the process of photosynthesis an endothermic reaction?' Alice wants to pass her exam, hence she adds the question to her agenda. Moreover, she studied well and therefore possesses the information that photosynthesis is in fact an endothermic reaction. This fact became issue-relevant when

[31] Given reflexivity and transitivity, Euclideanness comes down to symmetry, making the accessibility relation an equivalence relation.

[32] Holliday (2018) claims that negative introspection is universally rejected by epistemologists.

[33] Notwithstanding that some introductory textbooks such as the ones by Fagin et al. (1995), Van Ditmarsch, Van der Hoek and Kooi (2008) and Van Ditmarsch et al. (2015) present systems where agents have both positive and negative introspection as the basic epistemic logic. These books, however, are primarily written with applications in AI, computer science and economy in mind.

[34] In much of the literature concerning positive introspection it is referred to as the KK-principle.

[35] Lenzen himself, Lehrer (1974), Hilpinen (1970) and Hintikka (1962) are some early proponents of positive introspection. It must be noted that Hintikka (1970) seemed to have changed his mind, since later he argues against positive introspection. Other early opponents include Armstrong (1973, §15.1) and Robinson (1971).

[36] Also the work by Annis (1969, p. 168). A similar argument is given by Feldman (1981).

[37] For example, see the works by Mott (1998), Halpern (2004) and Spector (2013).

she added the exam question to her agenda. Thus Alice knows that photosynthesis is an endothermic reaction—it is issue-relevant information. To resolve her issue, this information suffices. She fills in the correct answer and continues the exam. At this point, Alice does not know that she knows the answer. Her knowledge of her knowledge is not issue-relevant: it is not needed to resolve her issue, which is only aimed at correctly answering the exam's questions.

When the exam is finished, a fellow student asks Alice whether she knows the answer to the photosynthesis question. This prompts her to add this question to her agenda, making her own knowledge of photosynthesis' endothermicity issue-relevant. Since she indeed knows that photosynthesis is endothermic, this information is also available to her. Consequently, now she also knows that she knows the answer to the question whether photosynthesis is endothermic.

A rational agent who knows the answer to an exam question has no need to know that she knows that, if she is only interested in doing the exam. Let alone that she needs to know that she knows that she knows the answer, et cetera. The principle of positive introspection would clutter the mind of this agent with unnecessary knowledge of her own knowledge. In contrast, an agent who has to answer a question to which she knows the answer and is wondering whether she can answer the question correctly, has a need to know that she knows the answer. Agents with positive introspection cannot distinguish between these types of propositions. If reasoning about one's own knowledge is to play a role in epistemic logic, the principle of positive introspection must not hold.

The principles discussed above almost fully characterize the notion of knowledge as information possession, in the sense that the axioms corresponding to these properties almost constitute a complete deductive system. For completeness, two less controversial principles are missing.

The first of these is the principle of necessitation, stating that necessary truths are known. In our intensional setting, there is only one necessary truth: the proposition that is the set of all possible worlds; the tautological proposition. For ideal agents this principle is uncontroversial: since they are not able to consider a world in which the tautological proposition does not hold, it can be assumed they know it. This principle only becomes problematic once we go beyond intensional content. Since we do not intend this, it is only reasonable to accept the principle of necessitation.

The second principle is the principle of truth, which we call factivity. It states that any proposition that is known is also true. It is considered one of the essential principles of knowledge and has not been under scrutiny within the context of formal epistemology (Aucher, 2014, p. 5).[38] We also accept it without further ado. Given factivity, positive introspection implies that the statement 'the agent knows that P' is intensionally equivalent to the statement 'the agent knows that she knows that ... she knows that P' for any finite non-zero number of iterations. This strengthens the argument against positive introspection. The intensional content of sentences like 'the agent knows the answer to the exam question' and 'the agent knows that she knows the answer to the exam question' seem to be different, yet positive introspection in tandem with factivity implies that they are always the same. Therefore, given factivity, positive introspection should not hold.

In short, the principle of necessitation and factivity are retained, but there are good reasons to reject closure under deduction and both negative and positive introspection. Dropping these principles, however, leaves a void. An agent might not know *all* consequences of her own knowledge, but surely she knows *some* consequences. Likewise, that

[38]Sequoiah-Grayson and Floridi (2022) discuss the factivity of information more comprehensively.

an agent does not know all she knows does not mean that she does not know she knows anything. Even negative introspection should occasionally occur: consider the example where the agent is asked a question, but this time she does not know the answer. If she is wondering whether she knows the correct answer, it is only natural that she will know that she does not know the answer. So the considerations above call for restricted principles, for otherwise it is completely arbitrary to what extent agents are introspective and which logical consequences of their knowledge they know. Additionally, these restricted principles should address the reasons for rejecting the unrestricted principle. If they do not, the problems identified in the first place are left unsolved.

Under our proposal, information is processed relative to an agent's inquiry. Closure under logical consequence and introspection should be restricted accordingly, taking into account which information is issue-relevant. This gives rise to the following restricted principles.

- Restricted closure principle: if an agent knows that P, and P entails Q, then the agent knows Q whenever Q is issue-relevant.
- Restricted positive introspection: if an agent knows that P, then the agent knows that she knows P whenever her knowledge that P is issue-relevant.
- Restricted negative introspection: if an agent does not know that P, then the agent knows that she does not know P whenever her ignorance of P is issue-relevant.

These restricted principles address the objections discussed above. Firstly, the restricted closure principle ensures that an agent's mind is not cluttered with trivial consequences of her knowledge. If P is known and logically entails Q, then Q is only known if it is issue-relevant. This means that known consequences are never trivialities, because they address the issue(s) on an agent's agenda.

This also leaves space for reasoning within a framework. Because an agent need not know all logical consequences of her knowledge, reasoning could potentially result in new knowledge. Reasoning is instigated by the addition of a new issue to an agent's agenda. If a logical consequence of an agent's knowledge that was previously issue-irrelevant becomes issue-relevant, then the agent reasons towards that logical consequence, acquiring new knowledge. Thus, objections against closure under logical consequence in the spirit of Stalnaker (1991) are also addressed.

Secondly, the restricted positive introspection principle respects the intuition that, in general, the intensional content of 'the agent knows that she knows P' and 'the agent knows P' are different. It leaves the possibility open that an agent knows P without knowing that she knows P. In fact, it ensures that an agent's mind is not cluttered with knowledge about her own knowledge that she does not care for. Only issue-relevant knowledge is known introspectively. Again, this leaves space for reasoning about one's own knowledge since the addition of issues to an agent's agenda may prompt her to gain knowledge of her own knowledge.

Thirdly, the restricted negative introspection principle does not presuppose the world to be closed. There may be propositions that are not known by an agent without her knowing that she does not know them. Thus an agent does not take all facts of the world into consideration. Through inquiry she may consider facts that she previously did not care for epistemically. That way she gains knowledge of her ignorance, like Socrates when he practiced *elenchus*. Thus, once more, the restricted principle leaves space for reasoning, in this case about one's own ignorance.

We have seen that taking knowledge to be information possession brings about some epistemic principles that do not concur with some of our intuitions about knowledge. In particular, closure under deduction and both positive and negative introspection are contentious if epistemic logic is to accurately capture some of the inherent properties of knowledge. When considering knowledge as issue-relevant information, none of these contentious principles are validated, at least on a conceptual level. The principle of necessitation and factivity, however, are untouched.

2.4 An earlier proposal

A formal framework in the spirit of Hintikka and Schaffer is developed by Baltag, Boddy and Smets (2018).[39] We devote an entire section to its discussion for two reasons. Firstly, there are some similarities between our goals and the resulting frameworks. They also intend to restrict the information processing of agents to conceptual distinctions that are relevant with respect to their issues. Moreover, propositions are the objects of knowledge in their framework as well. Secondly, this paper served as one of the starting points for our research. Originally our goal was to enrich their framework by adding actions of issue change, allowing us to model phenomena such as paradigm shifts. However, some of the framework's features dissatisfied us, in particular the notion of relevance with respect to the issues on an agent's agenda. Our efforts to reform the framework shaped this thesis' research for a large part.

Baltag, Boddy and Smets (2018) take the set of fundamental conceptual questions that an agent is actively entertaining as fundamental in her epistemic processes. The set of conceptual questions that an agent entertains is called her interrogative agenda.[40] An agent's interrogative agenda is modeled as a partition of logical space such that each partition cell is a complete answer to her question(s) (Baltag, Boddy and Smets, 2018, p. 138). They call the equivalence relation corresponding to this partition the issue relation. It determines which information is relevant to the agent, structuring her conceptual framework and limiting what she can come to know (Baltag, Boddy and Smets, 2018, p. 134).

Their primary goal is to investigate the interplay between interrogative agendas and information flow in a group, showing that a sensible notion of group knowledge is obtained when taking interrogative agendas into account. We, however, will not focus on multi-agent settings and group knowledge, but on the role interrogative agendas play in the epistemic lives of individual agents. There is plenty to be said about single-agent settings before moving on to multi-agent settings is warranted.

The interrogative epistemic models used by Baltag, Boddy and Smets (2018) restricted to a single-agent setting are based on the epistemic issue models from Van Benthem and Minică (2012), but with an additional condition.

Definition 2.1 (Single-agent interrogative epistemic structure)**.** A single-agent interrogative epistemic structure S is a tuple $(W, \sim, \approx)$, where:

- W a set of possible worlds;
- $\sim$ is a reflexive and transitive relation, the epistemic relation;
- $\approx$ is an equivalence relation, the issue relation.

Additionally, the inclusion $\sim; \approx \subseteq \sim$ holds, where the semicolon denotes relation composition.

Interrogative epistemic models are interrogative epistemic structures augmented with a valuation, which we ignore for now because it will suffice to focus on the underlying

[39]A kindred framework is developed by Boddy (2014), on which we briefly comment in some footnotes.
[40]Boddy (2014) uses the same terminology.

structures. According to Baltag, Boddy and Smets (2018), the issue relation $\approx$ corresponding to an agent's interrogative agenda should be interpreted as a conceptual indistinguishability relation. If $w \approx v$, then w and v correspond to the same world in the agent's "subjective model" (Baltag, Boddy and Smets, 2018, p. 138). The additional constraint $\sim;\approx\subseteq\sim$ ensures that the information possessed by an agent fits the partition induced by her issue(s).

Definition 2.2 (Epistemic states)**.** Given a single-agent interrogative epistemic structure $S = (W, \sim, \approx)$. The epistemic state of the agent at any world $w \in W$ is given by

$$w(\sim) = \{v \in W \mid w \sim v\}.$$

The constraint guarantees that for each world $w \in W$, the epistemic state $w(\sim)$ consists of a union of partition cells corresponding to the issue relation. Knowledge, they assert, corresponds to the Kripke modality over the epistemic relation (Baltag, Boddy and Smets, 2018, p. 140).

Definition 2.3 (Knowledge in interrogative epistemic structures)**.** Let $S = (W, \sim, \approx)$ be an interrogative epistemic structure and $P \subseteq W$ be a proposition. The knowledge modality K is defined as a Kripke modality:

$$K(P) := \{w \in W \mid w(\sim) \subseteq P\}.$$

If $w \in K(P)$, then we say that the agent knows that P in world w.

Baltag, Boddy and Smets (2018) argue that the resulting notion of knowledge adheres to what they call the selective learning principle, which states that "when confronted with information, agents come to know only the information that is relevant for their issues" (Baltag, Boddy and Smets, 2018, p. 144).[41] Underlying this principle is the idea that "all we (can) come to know are answers to our own questions" (Baltag, Boddy and Smets, 2018, p. 134), which is loosely based on Schaffer's (2007) theory of contrastive knowledge. Intuitively, because the epistemic relation is constrained to fit the agent's issue, the conceptual constraints imposed by an agent's interrogative agenda are taken into account.[42]

In sum, two fundamental ideas are put forward. Firstly, the idea that all knowledge is a partial answer to the agent's questions, expressed by the selective learning principle. Secondly, the idea that the issue relation represents a conceptual indistinguishability relation. In the next section we will argue that the framework of Baltag, Boddy and Smets (2018) does not fully comply with these ideas.

2.5 Objections against the earlier proposal

Consider the interrogative epistemic structure in Figure 2.1. Both the information relation and the issue relation are equivalence relations, and $\sim;\approx\subseteq\sim$.[43] According to the semantics for the knowledge modality, both $w \in K(P)$ and $w \in K(Q)$. The proposition P is indeed an answer to the agent's question: worlds in the same partition cell always agree on P's truth value. However, knowledge of Q violates the idea that all knowledge should be an answer to the agent's question. The partition cell on the

[41] The essence of this principle is paraphrased elsewhere as "one cannot find what one is not looking for" (Baltag, Boddy and Smets, 2018, p. 134).

[42] Boddy (2014, p. 24) makes a less bold claim, stating that agents can only know propositions entailed by their issue(s).

[43] Since there is only one agent, the three extra conditions posed by Boddy (2014, p. 31) are also satisfied: $\approx\subseteq\sim$, $\approx;\sim\subseteq\sim;\approx$ and $\approx;\approx\subseteq\approx;\approx$. So it is also a single-agent epistemic group model.

right contains Q-worlds as well as $\neg Q$-worlds. As such, the set of Q-worlds cuts across the partition corresponding to the agent's issue: neither Q nor $\neg Q$ is an answer to the agent's question. Hence Q should not be known by the agent according to the selective learning principle.[44]

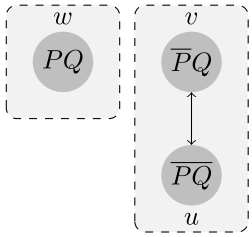

Figure 2.1: A single-agent interrogative epistemic structure S. The issue relation is represented by the dashed areas, in this model the agent's issue corresponds to the binary question whether P is true. The epistemic relation is represented by arrows. Reflexive arrows have been omitted.

So the framework of Baltag, Boddy and Smets (2018) does allow agents to know information that is not relevant to their issues.[45] Consequently, viewing $\approx$ as a conceptual indistinguishability relation seems to go against the proposed notion of knowledge. If the agent knows that Q, then surely she must be able to conceptually distinguish Q-worlds from $\neg Q$-worlds. Yet in Figure 2.1, Q is known at world w and $v \approx u$, with $v \in Q$ and $u \notin Q$. So it seems that $\approx$ does not accurately encode conceptual indistinguishability.

Furthermore, their knowledge modality is a Kripke modality, which means that it satisfies closure under logical consequence. However, in general, the set of conceptually distinguished propositions is not closed under logical consequence. Consider Example 1 again. The King of England conceptually distinguishes avoiding war with France. A logical consequence of avoiding war with France is avoiding *nuclear* war with France, but that is not conceptually distinguished by the King. If we interpret P as 'war can be avoided with France' and Q as 'nuclear war can be avoided with France', then this example fits the structure in Figure 2.1. Proposition P is conceptually distinguished by the King, while Q is not, despite being a logical consequence of P. So in the framework put forward by Baltag, Boddy and Smets (2018), the King knows that nuclear war can be avoided despite lacking the required conceptual resources to distinguish the concept of nuclear war. Since knowledge of a proposition implies that an agent can conceptually distinguish that proposition, the agent should also conceptually distinguish Q. The agent does not, therefore interpreting $\approx$ as a conceptual indistinguishability relation is incoherent.

In short, the models and semantics given by Baltag, Boddy and Smets (2018) do not fully comply with two of the fundamental ideas they put forward: not all knowledge is the answer to a question, and $\approx$ does not fully embody the conceptual indistinguishability relation. The discussion above already suggests what went wrong: the issue relation should be evaluated from a global perspective rather than a local one. Hence a Kripke modality cannot suffice as a knowledge operator.

[44]It also violates Schaffer's maxim "to know is to know the answer", as Q is not an answer to the agent's question.

[45]The structure in Figure 2.1 can also be turned against the framework of Boddy (2014): the proposition Q is semantically entailed by P in the model and can therefore be potentially known. Yet this has a dissatisfying consequence: the agent is able to know that Q is the case, but at the same time cannot distinguish between the Q-world v and the $\neg Q$-world u! This goes against the fundamental idea that $\approx$ represents a conceptual indistinguishability relation. Surely, one must be able to conceptually distinguish truth from falsity with regard to a proposition before that proposition can be known.

In order to determine whether a proposition P is relevant with respect to an issue it needs to be checked whether *all* partition cells internally agree on the truth value of P.[46] In order to uphold the idea that all knowledge is relevant with respect to an agent's issue and that the issue relation $\approx$ encodes which propositions are issue-relevant, the entire model must be taken into account when ascertaining whether a proposition is issue-relevant.

This concludes the chapter on knowledge and inquiry. From the next chapter onwards the logical framework for knowledge as issue-relevant information is build, taking the findings of this chapter into account.

[46]This idea is already expressed by Boddy (2014, p. 26): "Proposition P is an *agent-relevant proposition* for agent a if and only if for all states s and s' the following holds: if P is true at s and a cannot distinguish s from s', then P is true at s'". However, this idea is still only implemented locally.

Chapter 3

Information, Issues and Issue-relevance

In the previous chapter, we set forth our notion of knowledge: an agent knows that P if, and only if, the agent possesses the information that P, and P is issue-relevant. Thus there are three ingredients to our notion of knowledge that we need to flesh out in order to formulate a substantive logical theory: information, issues and issue-relevance. Each of these is treated comprehensively in Sections 3.1, 3.2 and 3.4, respectively. In Section 3.3, the epistemic issue structures that we will employ throughout the thesis are defined. The updates of epistemic issue structures corresponding to the actions of issue addition and retraction are also defined in this section, as well as the update corresponding to the resolution of an issue. In Section 3.5, knowledge as issue-relevant information is defined on epistemic issue structures.

3.1 Information

As mentioned in the previous chapter, we take propositions to be the objects of knowledge. In turn, propositions are defined as sets of possible worlds and the Boolean connectives represent set-theoretic operations on these sets.

Definition 3.1 (Boolean propositions)**.** Given a set of possible worlds W, a proposition P on W is a subset of W. For any pair of propositions P and Q on a set of possible worlds W, the following Boolean propositions are defined:

- $\neg P := W \backslash P$;
- $P \wedge Q := P \cap Q$;
- $P \vee Q := P \cup Q$;
- $P \rightarrow Q := \neg P \vee Q$;
- $\bot := \emptyset$;
- $\top := W$.

If $w \in P$, we say that P is true at world w. The empty proposition $\bot$ is called the inconsistent proposition and the set of all worlds $\top$ is called the tautological proposition.

The set-theoretic operations coincide with our intuition. For instance, $w \in P \wedge Q$ if, and only if, $w \in P$ and $w \in Q$. Likewise, $w \in P \rightarrow Q$ if, and only if, $w \in P$ implies $w \in Q$.

Sticking to tradition, the information possessed by an agent is encoded by an informational indistinguishability relation $\sim$ on the set of possible worlds W: if $w \sim v$ for worlds $w, v \in W$, then the information possessed by the agent does not allow her to distinguish between the states of affairs in worlds w and v.

Definition 3.2 (Information states)**.** Let W be a set of possible worlds and $\sim$ an equivalence relation on W. We call $\sim$ the information relation. The agent's information state at a world $w \in W$ is defined as

$$w(\sim) := \{v \in W \mid w \sim v\}.$$

The agent's information state at w picks out the worlds that are informationally indistinguishable for the agent at w. Since $\sim$ is an equivalence relation, we can view $w(\sim)$ as the equivalence class of $\sim$ that contains w. An agent possesses the information that P at world w if her information state at w is contained in P, that is, whenever $w(\sim) \subseteq P$.

We choose to let $\sim$ be an equivalence relation. Its reflexivity comes down to the agent only possessing true information: one cannot possess the information that P at world w if P is not true at w. This, of course, need not always be the case in reality. False information may be provided, for example by lying agents or faulty measuring instruments. We assume that agents work in an ideal setting, where instruments never fail and agents are provided only truthful information. Thus for our purposes these phenomena are irrelevant and reflexivity is desirable.

The transitivity and Euclideanness of the information relation ensure that information is transparant: every information state contains all information about itself. So, in short, the reflexivity, transitivity and Euclideanness of the information relation ensure that information is well-behaved. Information possession corresponds to knowledge in the modal logic **S5**.[47]

3.2 Issues

Following Yablo (2014, p. 27), we take issues to be systems of differences; patterns of cross-world variation.[48] An issue groups possible worlds in a minimal manner, such that facts at which the issue is directed are constant within each group of worlds. Thus, issues partition logical space.

Definition 3.3 (Issues)**.** Let W be a set of possible worlds. An issue on W is represented by an equivalence relation $\approx$ on W. We call $\approx$ the issue relation. The issue cell of a world $w \in W$ is defined as

$$w(\approx) := \{v \in W \mid w \approx v\}.$$

For instance, the issue concerned with the number of stars groups possible worlds in such a manner that exactly all worlds with an equal number of stars are grouped together.[49] As mentioned in the previous chapter, issues can be interpreted both as questions and as subject matters. We discuss each of these two interpretations below.

[47]The modal logic **S5** corresponds to the class of frames whose relation is an equivalence relation (Blackburn, De Rijke and Venema, 2001, p. 193).

[48]Yablo, as well as others, also refer to issues as subject matters, matters or topics. Throughout this thesis, we stick with 'issues', but we treat questions and subject matters as prototypical issues.

[49]This is a seminal example from Lewis (1988a; 1988b).

There is a rich tradition in which questions are modeled as partitions of logical space.[50] In those models, partition cells correspond to complete answers to a question. In terms of our example: a complete answer to the question 'what is the number of stars?' consists of a definite number. Thus each complete answer corresponds to a subset of the set of possible worlds in which the number of stars is constant. This constant number is the answer to the question represented by that particular partition cell. A partial answer to a question would then consist of a union of some of these partition cells.

When interpreting issues as questions, the agent we envision only distills (partial) answers from the information she has at her disposal. The equivalence relation induced by the partition can then be interpreted as a conceptual indifference relation: if $w \approx v$, then the agent's question(s) do not prompt her to conceptually distinguish w from v.[51] In other words, the differences between w and v are irrelevant to the agent's issue(s). Ultimately, it is this indifference that we are interested in when analysing an agent's information processing.

More recent developments in the semantics of questions deny that questions can be accurately modeled by partitions. Inquisitive semantics has proven itself as one of the most promising approaches when it comes to modeling the meaning of questions.[52] In inquisitive semantics, questions are defined as non-empty, downwards closed sets of sets of possible worlds. Each of these sets of possible worlds represents a body of information that would be sufficient to resolve the question. The maximal elements of such a question are called its alternatives, which encode the complete answers to a question that provide just enough information to answer it.[53] Thus, the framework of inquisitive semantics takes a fundamentally different approach to questions than we do. This is not surprising given our diverging goals: we are not attempting to capture the meaning of questions in terms of which information would resolve them, but rather in terms of which worlds need to be conceptually distinguished to resolve them. This explains why we do not identify questions as downward closed sets of sets of worlds: it would defeat the purpose of capturing which distinctions are conceptually relevant. Thus, to capture conceptual indifference, it is crucial that we only take the alternatives of a question into account.

Moreover, the alternatives of the questions we consider should cover the entire logical space, unlike the questions in inquisitive semantics. Issues construed as questions do not carry any informative content, whereas questions may contain informative content in inquisitve semantics. A question is informative in inquisitive semantics whenever its alternatives do not cover the entire logical space.[54]

On top of that, if we give up the downwards closure and only consider a question's alternatives, representing questions as conceptual indifference relations becomes more complicated. The conceptual indifference relation becomes murky, since different alternatives may overlap in inquisitive semantics. When overlap is allowed, transitivity of the conceptual indifference relation fails. Consequently, a question may not prompt an agent to conceptually distinguish between worlds w and v, and worlds v and u, while it prompts her to conceptually distinguish between worlds w and u, which is nonsensical. This would call for a reinterpretation of the conceptual indifference relation. Thereafter, overlapping alternatives need not be a problem. However, for our purposes, modeling questions as partitions suffices. Therefore we have decided to *not* allow overlap, leaving a more general conception of issues to be addressed in future research.

[50] Key contributions were made by Hamblin (1958), Belnap and Steel (1976), and Groenendijk and Stokhof (1984), among others.

[51] Cf. Hulstijn (1997).

[52] Information on the basics of inquisitive semantics is provided by Ciardelli (2016) and Ciardelli, Groenendijk and Roelofsen (2018).

[53] More details are given by Ciardelli, Groenendijk and Roelofsen (2018, Ch. 2).

[54] Note that questions are also called propositions in inquisitive semantics (Ciardelli, Groenendijk and Roelofsen, 2018, § 2.4).

Issues can also be interpreted as subject matters, as suggested by Lewis (1988a). Then the equivalence relation induced by an issue relates all worlds that are exactly alike with respect to that particular subject matter. For instance, borrowing an example from Lewis (1988a, p. 161), if the subject matter is the seventeenth century, then worlds that agree on all facts pertaining to the seventeenth century are equivalent under that issue relation. The interpretation of the issue relation remains the same as when we interpret issues as questions. If an agent is solely interested in a particular subject matter, then she will be indifferent to worlds that are equivalent with respect to that subject matter.

Thinking of issues as questions has the benefit that we tend to think of inquiry as a questioning procedure, while thinking of issues as subject matters corresponds more closely to the idea that issues highlight certain information that is of interest to the agent.[55] Ultimately, it does not matter. A subject matter X can be glossed as a question 'what is true about X?'. Conversely, a question Y can be glossed as a subject matter 'answers to Y'.[56] Furthermore, under both interpretations issues are represented by equivalence relations that encode conceptual indifference.

3.3 Epistemic issue structures

Considering only one isolated issue is insufficient in the setting we have outlined. We should rather consider the entire body of issues at play. To this end, we distinguish all individual issues that could comprise an agent's agenda.

Definition 3.4 (Basic and compound issues)**.** Let $\mathcal{I}$ be a set of issues on a set of worlds W. An element $x \in \mathcal{I}$ is called a basic issue. For every basic issue $x \in \mathcal{I}$, the corresponding issue relation is denoted by $\approx_x$. A set of basic issues $X \subseteq \mathcal{I}$ is called a compound issue. The agent's agenda is a compound issue $\mathcal{A} \subseteq \mathcal{I}$.

An agent's agenda picks out which basic issues an agent pursues, determining her current issue. Since issues are modeled as equivalence relations induced by a partition, they can be combined in a convenient manner by intersecting them. Suppose we have several issues on a set of possible worlds W, then each of these issues prompts an agent to conceptually distinguish some worlds, while not distinguishing others. The intersection of these indifference relations prompts an agent to conceptually distinguish between all worlds that at least one issue prompts her to distinguish, while not distinguishing all others. This is reasonable: combining issues comes down to combining their potential to make an agent conceptually distinguish between worlds.

Definition 3.5 (Compound issue relations)**.** Let $\mathcal{I}$ be a set of basic issues on a set of worlds W. The issue relation corresponding to a compound issue $X \subseteq \mathcal{I}$ combines all basic issues in X:

$$\approx_X := \bigcap_{x \in X} \approx_x .$$

We stipulate that $\approx_\emptyset := W \times W$. The relation $\approx_\mathcal{A}$ represents the agent's current agenda.

The issue relation corresponding to the empty issue is the universal relation: without issues the agent is not prompted to make any conceptual distinctions. Since the intersection of any number of equivalence relations is itself an equivalence relation, any combination of issues is also an issue. Thus it makes sense to refer to a body of issues as a single issue. Henceforth we can simply refer to the agent's agenda as her current issue. This is also how the issue relation in the single-agent interrogative epistemic structures

[55] Thinking of inquiry as a questioning procedure is appealing to at least some philosophers, such as Hookway (2008) and Friedman (2020), besides the ones already mentioned in Chapter 2.

[56] A similar observation has been made by Habgood-Coote (2022, p. 14).

of Baltag, Boddy and Smets (2018) should be interpreted: as the intersection of the implicit, more basic issues on an agent's interrogative agenda.

Finally, all of the above can be encoded into the following structures.

Definition 3.6 (Epistemic issue structures)**.** Let $\mathcal{I}$ be a set of basic issues on a set of worlds W. An epistemic issue structure S is a tuple $(W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$, where:

- W is a set of worlds;
- $\sim$ is an equivalence relation, the information relation;
- $\approx_x$ is an equivalence relation for every basic issue $x \in \mathcal{I}$;
- $\mathcal{A} \subseteq \mathcal{I}$ is the agent's current agenda.

In general, we use the abbreviation $S = (W, \sim, \approx)$, where:

- W is a set of worlds;
- $\sim$ is the information relation;
- $\approx := \approx_{\mathcal{A}}$ is the equivalence relation corresponding to the agent's agenda.

If $\mathcal{I}$ is a singleton and $\mathcal{A} \neq \emptyset$, then we call S a single issue structure.

Instead of one issue relation, an epistemic issue structure contains an issue relation $\approx_x$ for every basic issue $x \in \mathcal{I}$. The agenda $\mathcal{A}$ then picks out which basic issues the agent is actually pursuing. Hence the current issue of the agent is represented by the compound issue relation $\approx_{\mathcal{A}}$.

The above definition is reminiscent of the structures underlying the epistemic issue Models of Van Benthem and Minică (2012).[57] In fact, our epistemic issue structures can be viewed as generalizations of their structures, hence we have opted to use an identical name. The subclass of single issue structures corresponds to the structures they use. Their issue relation, however, should not be interpreted as encoding conceptual indistinguishability, but as representing the alternatives of a question that an agent wishes to informationally distinguish (Van Benthem and Minică, 2012, p. 2-3). Similarly, the structures underlying the single-agent interrogative epistemic models of Baltag, Boddy and Smets (2018) in which the information relation is an equivalence relation correspond to the subclass of single issue structures that satisfy $\sim;\approx\, \subseteq\, \sim$.

For the purpose of convenience, when the agenda is implicitly understood and we do not consider individual basic issues, we write $S = (W, \sim, \approx)$ to denote epistemic issue structures. In particular, this abbreviation is useful when referring to single issue structures: in these cases $\approx$ unambiguously refers to $\approx_x$, where x denotes the only basic issue in $\mathcal{I}$.

Epistemic issue structures may appear redundantly complex because they contain separate relations for each basic issue. However, this is necessary given that we want to be able to model both the addition and retraction of issues.

Van Benthem and Minică (2012) show how we can model the action of asking a binary question in single issue structures. Asking whether P refines the issue relation so that every issue cell is split into two possibly empty issue cells with P-worlds and $\neg P$-worlds. Conceptually, this action corresponds to adding the question whether P to the agent's agenda. Yet this action only allows us to express the addition of binary questions to an agent's agenda, whereas not all questions on an agent's agenda need to be binary. Wh-questions are in general not binary, while often making up most of

[57] These models were first introduced by Minică (2008).

our actual agendas, so refining a single issue relation using only binary questions is very limiting.

Moreover, besides adding issues to an agenda, we also want to be able to retract issues. As discussed in Chapter 2, sometimes agents lose interest in an issue or it becomes meaningless. Retracting issues is more difficult than adding them, because it requires us to make an issue relation more coarse rather more fine-grained. This difficulty is mirrored in modeling forgetting in dynamic epistemic logic. Agents have perfect recall in dynamic epistemic logic: once information is possessed, it is never forgotten. In our case this translates to: once a conceptual distinction has been made, it cannot be unseen. As a consequence, it is ambiguous how to retract an issue given only a single issue relation.

For instance, consider the structures in Figure 3.1. The issues in the structures on the left and right differ. The proposition $P \vee Q$ forms an issue cell in the left structure, while it does not fit the issue cells in the right structure. The addition of the binary question whether P to the agent's agenda yields the middle structure in both cases. However, given the structure in the middle, it is ambiguous how one would retract the question whether P. Both the left and the right structure are consistent with retracting that question, but yield structures corresponding to different agendas. Furthermore, it might be the case that the question whether P was not on the agent's agenda in the first place, but that it was a byproduct of compounding several of the issues on her agenda. Retracting the question asking whether P then changes the agenda in a more drastic way than intended. In fact, retracting an issue that is not on the agent's agenda should leave it unchanged.

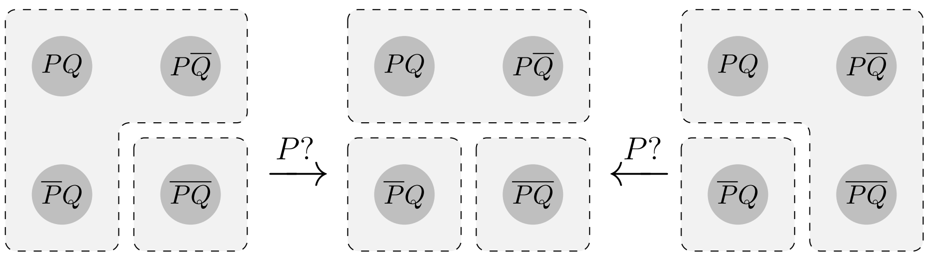

Figure 3.1: Three structures that show an issue on a set of four possible worlds. The structure in the middle is obtained after refining the issues in the structures on the left and right with the question whether P.

Overcoming this ambiguity requires more control over the relation between the individual issues on an agent's agenda and the issue relation that encodes the entirety of these issues. By having distinct issue relations for each basic issue, issues can unambiguously be added and retracted. In epistemic issue structures, updates of the agent's agenda can be modeled in a straightforward manner.

Definition 3.7 (Agenda updates)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure. The addition of an issue $X \subseteq \mathcal{I}$ to the agent's agenda yields the updated structure

$$S_{[+X]} := (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A} \cup X).$$

The retraction of an issue $Y \subseteq \mathcal{I}$ from the agent's agenda yields the updated structure

$$S_{[-Y]} := (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A} \backslash Y).$$

What about information updates? Van Benthem and Minică (2012) define the action of announcing a fact. Public announcements are the most basic information updates and therefore the hallmark of dynamic epistemic logics. They can be captured in terms of issue resolution: announcing that P is true is equivalent to resolving the issue whether P, given that P is actually true. Resolution should be interpreted as providing the agent with exactly sufficient information to know that P. So if we take resolution of any issue as information update, then we can capture a broader class of information updates than just announcements.

This also fits the view that inquiry is fundamental to knowledge acquisition. New information is not obtained randomly, but as a means to resolve an issue. It concerns the actions required to resolve an issue, much like the fifth phase in Dewey's description of thinking. Such actions may consist of observation or experiment, "compelling nature" to answer the agent's question, to put it in Kantian terms.

Yet, a complication arises, because the announcement of P can only be captured in terms of issue resolution if there is an issue available that corresponds to the question whether P. Thus, to be able to express the announcement of any proposition P, we either need a sufficiently large number of basic issues so that we can meet this demand or we need to have separate actions of announcement and issue resolution. It will turn out that the former is impossible if we are to have a sound and complete logic, for the set of basic issues $\mathcal{I}$ needs to be finite in order to obtain completeness.[58] The latter might be possible, but as of yet it is an open problem how the resulting logic should be axiomatized. We briefly return to this matter in Section 6.5. For now we opt to only have issue resolution as information updates, for all practical purposes this provides us with sufficient expressivity.

Definition 3.8 (Issue resolution)**.** Let $S = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A})$ be an epistemic issue structure. Resolving an issue $X \subseteq \mathcal{I}$ yields the model

$$S_{[X!]} := (W, \sim \cap \approx_X, \approx_{x \in \mathcal{I}}, \mathcal{A}).$$

Thus the resolution of an issue X refines the agent's information relation such that for every world w the information state $w(\sim)$ is restricted to the issue cell $w(\approx_X)$. This is similar to the resolution action of Van Benthem and Minică (2012), although they have only defined it for a single issue.

3.4 Issue-relevance

At this point we have comprehensively discussed information and issues, also defining the structures that encode both. The structures also allow the modeling of the addition and retraction of issues, as well as issue resolution. In order to formally define knowledge in epistemic issue structures, we are left with defining issue-relevance.

Recall our objections against the framework of Baltag, Boddy and Smets (2018): in order to rightly interpret $\approx$ as a conceptual indistinguishability relation and let all knowledge be an answer to the agent's question(s), the entire issue relation should be taken into account when determining whether a proposition is issue-relevant. Also recall that 'issue-relevant' is shorthand for 'relevant with respect to the issues on an agent's agenda'. For a proposition P to be issue-relevant, then, all current issue cells in an epistemic issue structure must internally agree on the truth value of P. Put differently, P is issue-relevant whenever P does not cut across any cells corresponding to the agent's current issue.

This also matches the characterization of issues from Section 3.2: an issue groups possible worlds in a minimal manner such that the propositions at which it is directed

[58] This problem only arises when the set of possible worlds is infinite, otherwise a finite number of issues suffices to capture all possible public announcements.

are constant within these groups of worlds. Along with our findings from Section 2.5, this leads to the following definition of issue-relevance in epistemic issue structures.

Definition 3.9 (Issue-relevance)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and P a proposition on W. Proposition P is issue-relevant if, and only if, for all $w \in W$, if $w \in P$ then $w(\approx_{\mathcal{A}}) \subseteq P$. If a proposition P is not issue-relevant, we call it issue-irrelevant.

Thus a proposition is issue-relevant if it is closed under the current issue relation. When abbreviating an epistemic issue structure as $S = (W, \sim, \approx)$, a proposition P is issue-relevant if, and only if, for all $w \in W$, if $w \in P$, then $w(\approx) \subseteq P$. The definition states that if P is issue-relevant and P is true at some world w, then P should be true in all worlds in the current issue cell containing w. This ensures that, on a global level, issue-relevant propositions do not cut across any current issue cells.

If we interpret issues as questions, this definition stays faithful to the idea that agents can only know propositions that are (partial) answers to their questions. Furthermore, the issue relation encodes which propositions are conceptually indistinguishable: if P is not issue-relevant, then there exists worlds w and v such that $w \approx v$ with $w \in P$ and $v \notin P$. So P is issue-relevant whenever the agent can always conceptually distinguish between P-worlds and $\neg P$-worlds.

This concludes our exposition of information, issues and issue-relevance, providing us with the required machinery to formally define knowledge.

3.5 Knowledge

A definition of knowledge as issue-relevant information can now be given.

Definition 3.10 (Knowledge)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and P a proposition on W. The knowledge operator K is defined as:

$$K(P) := \{w \in W \mid w(\sim) \subseteq P \text{ and for all } v \in W, \text{ if } v \in P \text{ then } v(\approx_{\mathcal{A}}) \subseteq P\}.$$

If $w \in K(P)$, then we say that the agent knows that P in world w.

So an agent knows that P at some world w if the agent's information state $w(\sim)$ is contained in P and if P fits the partition induced by the current issue relation. That is, whenever the agent possesses the information that P, and P is issue-relevant. Thus, agents truly only come to know information that is relevant with respect to the issues on their agenda; our notion of knowledge satisfies the selective learning principle of Baltag, Boddy and Smets (2018). The definition of knowledge can be phrased in a slightly more familiar manner.

Proposition 3.11 (Truth conditions for knowledge)**.** *Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and let P be a proposition on W. Then $w \in K(P)$ if, and only if, $v \in P$ for all $v \in W$ such that $w \sim v$, and for all $u, v \in W$ such that $u \approx_{\mathcal{A}} v$, $u \in P$ implies $v \in P$.*

Proof. Observe that by definition $w(\sim) \subseteq P$ is equivalent to $v \in P$ for all $v \in W$ such that $w \sim v$. For the second part, assume that $v(\approx_{\mathcal{A}}) \subseteq P$ for all $v \in P$ and take arbitrary $u, v \in W$ such that $u \approx v$. If $u \in P$, it follows that $u(\approx_{\mathcal{A}}) \subseteq P$ by assumption. As $v \in u(\approx_{\mathcal{A}})$, also $v \in P$. The converse is left to the reader. $\square$

Lastly, when abbreviating epistemic issue structures as $S = (W, \sim, \approx)$, we find that $w \in K(P)$ if, and only if, for all $v \in W$ such that $w \sim v$ we have $v \in P$ and for all $u, v \in W$ such that $u \approx v$, $u \in P$ implies $v \in P$. In the next chapter we shall investigate the laws to which this notion of knowledge as issue-relevant information adheres.

Chapter 4

The Laws of Knowledge

The previous chapter detailed how information and issues can be encoded in epistemic issue structures, and how knowledge as issue-relevant information can be defined in these structures. This chapter focuses on the laws to which this notion of knowledge adheres. In Section 4.1, a definition of validity in epistemic issue structures is given. Thereafter, it is shown that knowledge is factive and that the tautological proposition is always known. In Section 4.2, we delve into the properties of issue-relevance, which are crucial for expressing the laws of knowledge. In Section 4.3 and Section 4.4, it is argued that the laws of knowledge as information possession do not hold, but that the restricted principles discussed in Section 2.3 hold instead. In Section 4.5, the properties of the Kripke modalities corresponding to the information relation and the issue relation are explored. It is also shown how issue-relevance can be expressed in terms of Kripke modalities. In Section 4.6, the examples from Section 2.2 are revisited, and dynamic principles of knowledge are formulated.

4.1 Validity, knowledge and truth

In order to study the laws of knowledge, we need a notion of validity.

Definition 4.1 (Validity)**.** A proposition is valid in an epistemic issue structure whenever it equals the set of possible worlds. A proposition is valid if it is valid in every epistemic issue structure.

Recall that we stated in Section 3.3 that we can denote epistemic issue structures in a compressed manner as $(W, \sim, \approx)$ when the dynamic actions and the individual basic issues on the agent's agenda are not of importance. In Sections 4.1–4.5 we will focus on the validity of static principles of knowledge that do not involve individual basic issues, hence we will refer to epistemic issue structures in this compressed manner.

It was argued in Section 3.1 that we assume information to be well-behaved and that therefore all information possessed by an agent is true. All knowledge, which consists of issue-relevant information, is therefore true as well. In addition to all knowledge being true, tautologies are also known: for any epistemic issue structure $S = (W, \sim, \approx)$ and any $w \in W$, we have the inclusions $w(\sim), w(\approx) \subseteq W = \top$.

Proposition 4.2. *The following propositions are valid:*

(i) Factivity of knowledge: $K(P) \to P$*;*

(ii) Knowledge of the tautological proposition: $K\top$*.*

Thus our framework shares factivity of knowledge with standard epistemic logics, assuming that all information available to agents is true. These logics typically require the information relation to be reflexive, which is reflected in the reflexivity constraint on the information relation in epistemic issue structures.

Knowledge of the tautological proposition, (ii), is not surprising. Since we have restricted ourselves to intensions, there is only one tautology. Furthermore, the tautological proposition carries no information. Knowledge of the tautological proposition is therefore not telling.

In conclusion, all known propositions are true, and the tautological proposition is always known.

4.2 Properties of issue-relevance

The definition of knowledge in Section 3.5 consists of two clauses: a proposition is known if an agent can informationally distinguish it, and it is issue-relevant. The former has been studied extensively in the literature; it corresponds to the Kripke modality over the information relation. We briefly return to it in Section 4.5. Issue-relevance, however, is as of yet an opaque notion. It will prove fruitful, in fathoming the laws of knowledge, to distinguish it as a standalone operator. A definition can be derived from Definition 3.9.

Definition 4.3 (Issue-relevance modality)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and let P a proposition on W. If P is issue-relevant in S, then $R(P)$ is valid in S. Since issue-relevance is a global notion, the proposition $R(P)$ can be identified with the tautological proposition $\top = W$ or with the inconsistent proposition $\bot = \emptyset$.

Some properties of issue-relevance immediately follow from the definition above.

Proposition 4.4. *The following propositions are valid:*

(i) Knowledge implies issue-relevance: $K(P) \to R(P)$*;*

(ii) The tautological proposition is issue-relevant: $R(\top)$*;*

(iii) The inconsistent proposition is issue-relevant; $R(\bot)$*.*

The issue-relevance modality allows us to rephrase the truth conditions for knowledge: $w \in K(P)$ if, and only if, $v \in P$ for all $v \in W$ such that $w \sim v$ and $w \in R(P)$. The validity of (i) is an immediate consequence of the truth conditions for knowledge. It confirms that all knowledge is indeed issue-relevant. Additionally, (ii) and (iii) state that the tautological proposition and the inconsistent propositions are always trivially issue-relevant: neither W nor $\emptyset$ can cut across the cells of an issue on W.

The modality R is not normal. It lacks several classical closure properties, including closure under logical consequence.

Proposition 4.5 (Failure of closure for issue-relevance)**.** *The following statements illustrate the failure of classical closure properties for* R*:*

(i) Validity of $P \to Q$ *in an epistemic issue structure* S *does* not *imply validity of* $R(P) \to R(Q)$ *in* S*;*

(ii) Validity of $P \to Q$ *in an epistemic issue structure* S *does* not *imply validity of* $R(Q) \to R(P)$ *in* S*;*

(iii) The proposition $R(P \wedge Q) \to R(P) \wedge R(Q)$ *is* not *valid.*

Proof. For each claim, counterexamples are provided in Figure 4.1. For the first claim, consider the epistemic issue structure S in Figure 4.1 (a). The proposition $P \to Q$ is valid in the structure since $P \subseteq Q$. All worlds within the same issue cell agree on the truth value of P, so $R(P)$ is valid as well. Yet $R(Q) = \bot$, since $v \approx u$ with $v \in Q$ and $u \notin Q$.

For the second claim, consider the epistemic issue structure in Figure 4.1 (b). The proposition $P \to Q$ is valid and $Q = W'$, so $R(Q)$. However, $w' \approx v'$ with $w' \in P$ and $v' \notin P$, so $R(P) = \bot$.

For the third claim, consider the epistemic issue structure in Figure 4.1 (c). Neither $R(P)$ nor $R(Q)$ holds, since both P and Q cut across the issue cell containing two worlds. The conjunction $P \wedge Q$ is false throughout the entire structure, thus $R(P \wedge Q)$ is valid by Proposition 4.4 (iii). $\square$

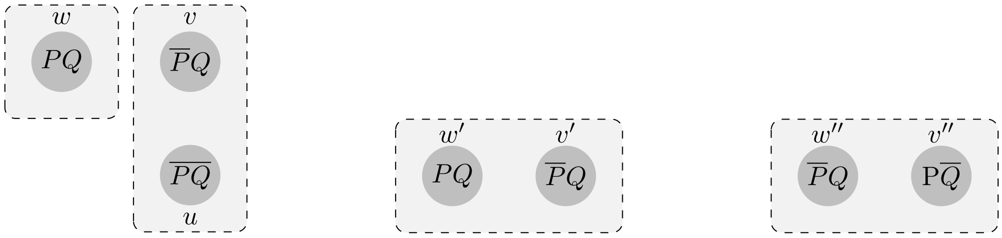

(a) $S = (W, \sim, \approx)$ (b) $S' = (W', \sim', \approx')$ (c) $S'' = (W'', \sim'', \approx'')$

Figure 4.1: Three epistemic issue structures. The dashed areas demarcate the current issue cells. Since the information possessed by the agent does not affect issue-relevance, it can be assumed that the agent is able to informationally distinguish all worlds.

Statements (i) and (ii) express that issue-relevance is neither upwards nor downwards closed, respectively. A direct consequence of (i) and Proposition 4.4 (ii) is that distribution over implication does not hold for R, that is, $R(P \to Q) \to (R(P) \to R(Q))$ is not valid.

Statement (iii) expresses that issue-relevance cannot be distributed over conjunction. The structure in Figure 4.1 (c) exemplifies this using contingent propositions P and $Q = \neg P$ that are each others negations. In this case, their conjunction $P \wedge \neg P = \bot$ is the inconsistent proposition, making it trivially issue-relevant. However, there are conceptually stronger reasons why distribution over conjunction should fail in our intensional setting. In fact, it already follows from (i): if $P \to Q$ is valid in an epistemic issue structure S, then $P \wedge Q = P$. Hence, if $R(P)$ is valid in S while $R(Q)$ is not, then $R(P \wedge Q)$ is valid while $R(Q)$ is not. So if we restrict ourselves to intensions and do not want issue-relevance to be closed under logical consequence, failure of distribution over conjunction is a necessary consequence.

There are even situations in which a conjunction $P \wedge Q$ is non-trivially issue-relevant, while both P and Q are not. The epistemic issue structure in Figure 4.2 illustrates this. An example where this may occur is when studying the quantum entanglement of two particles a and b. Let the proposition P represent 'the spin of particle a is up' and let Q represent 'the spin of particle b is down'. The agent's issue is solely concerned with the joint behaviour of the entangled system, that is, whether $P \wedge Q$ holds or not. The individual propositions P and Q are not independently issue-relevant. Hence, neither $R(P \wedge Q) \to R(P)$ nor $R(P \wedge Q) \to R(Q)$ is valid in this example; the epistemic issue structure in Figure 4.2 confirms this.

Proposition 4.5 shows that issue-relevance is not as well-behaved as normal modal operators. However, despite lacking the above closure properties, there are some closure principles to which issue-relevance adheres.

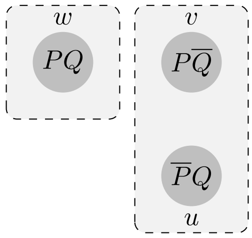

Figure 4.2: An epistemic issue structure $S = (W, \sim, \approx)$ consisting of three worlds. Since the information possessed by the agent does not affect issue-relevance, it can be assumed that the agent is able to informationally distinguish all worlds. The proposition $R(P \wedge Q)$ is valid in S, whereas both $R(P)$ and $R(Q)$ are not.

Proposition 4.6 (Closure principles for issue-relevance)**.** *The following principles are valid:*

(i) Closure under negation: $R(P) \to R(\neg P)$*;*

(ii) Closure under conjunction: $(R(P) \wedge R(Q)) \to R(P \wedge Q)$*;*

(iii) Closure under logical equivalence: if $P \leftrightarrow Q$ *is valid in an epistemic issue structure* S*, then* $R(P) \to R(Q)$ *is valid in* S*.*

Proof. Let $S = (W, \sim, \approx)$ be an epistemic issue structure and pick $w \in W$ arbitrarily.

For (i), assume that $w \in R(P)$. Then every issue cell containing a P-world consists only of P-worlds. By modus tollens it follows that every issue cell containing a $\neg P$-world consists only of $\neg P$-worlds as well, hence $w \in R(\neg P)$.

For (ii), suppose $w \in R(P) \wedge R(Q)$. It follows that P and Q are empty or unions of one or more issue cells. Hence the intersection of P and Q can never cut across an issue cell and must itself be empty or a union of issue cells. Therefore $w \in R(P \wedge Q)$ as well.

For (iii), assume that $P \leftrightarrow Q$ is valid in S. It follows that $P = Q$. It then follows directly from Definition 4.3 that $R(P) = R(Q)$. Hence $R(P) \to R(Q)$. $\square$

Closure under negation states that the negation of any issue-relevant proposition is also issue-relevant. The issue-relevance of a proposition P comes down to the agent being prompted to conceptually distinguish between P-worlds and $\neg P$-worlds. The issue-relevance of $\neg P$ thus comes down to being prompted to conceptually distinguish between $\neg P$-worlds and $\neg\neg P$-worlds, the latter of which are just the P-worlds. So closure under negation is reasonable, and its converse is also valid.

Closure under conjunction, property (ii), states that the conjunction of two issue-relevant propositions is also issue-relevant. If P and Q are issue-relevant the agent conceptually distinguishes both. Therefore, it is unavoidable that a rational agent also conceptually distinguishes $P \wedge Q$. Note that its converse does not hold, as stated in Proposition 4.5 (iii).

Closure under logical equivalence, expressed by (iii), confirms that our notion of issue-relevance is restricted to intensional content.

Now that we have some insight into the behaviour of issue-relevance, we can continue disentangling the laws of knowledge.

4.3 Reasoning and knowledge

Knowledge of a proposition P requires P to be issue-relevant. Proposition 4.5 shows that issue-relevance is not closed under logical consequence and lacks distribution.

Counterparts of these closure principles also fail to hold for knowledge, as is shown by the following example.

Example 5 (The King of England: revisited)**.** We return to the King of England from Example 1, which is captured in the epistemic issue structure in Figure 4.3. The proposition P should be interpreted as 'war can be avoided' and Q as 'nuclear war can be avoided'. That war can be avoided with France is issue-relevant: P does not cut across any issue cells. A logical consequence of avoiding war with France is avoiding *nuclear* war with France, but that is not issue-relevant: $v \approx u$ with $v \in Q$ and $u \notin Q$. Since w is the only world in the information state $w(\sim)$ and $w \in P$, the King possesses the information that P. Hence the King knows that P at the actual world w: $w \in K(P)$. However, since Q is not issue-relevant the King does not know Q; $w \notin K(Q)$. So despite the King knowing that war can be avoided with France, he does not know that nuclear war can be avoided with France.

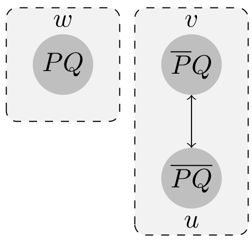

Figure 4.3: An epistemic issue structure S. The agent's current issue corresponds to the binary question whether P is the case. The information relation is indicated by the arrows, reflexive arrows have been omitted. The actual world is denoted by w.

The example shows that knowledge is not closed under logical consequence, inheriting this property from issue-relevance. Observe that Figure 4.3 is identical to Figure 2.1, showcasing the difference between our notion of knowledge and the notion of knowledge in the framework of Baltag, Boddy and Smets (2018).

The key difference between the relevant alternatives approach and our relevant distinctions approach is also shown in this example. It is tempting to try to turn the structure in Figure 4.3 against skeptical paradoxes. The proposition P could be interpreted as 'I have hands' and Q as 'I am *not* a brain in a vat'. Then the agent knows that she has hands, but does not know she is not a brain in a vat. This would be a misinterpretation of the structure, because in the structure the agent also has the information that she is *not* a brain in a vat. Skeptical paradoxes, however, use the fact that we cannot have information about "skeptical" propositions, such as Q. Issue-relevance cannot be used to rule out information or alternatives, whereas this is the fundamental idea in the relevant alternatives approach.

Knowledge does not only inherit failure of closure under logical consequence from issue-relevance, it also inherits failure of distribution over conjunction and implication. Both also fail in the epistemic issue structure in Figure 4.3.

Proposition 4.7 (Failure of closure for knowledge)**.** *The following statements illustrate the failure of classical closure properties:*

(i) Validity of $P \to Q$ in an epistemic issue Structure S does not *imply validity of $K(P) \to K(Q)$ in S;*

(ii) Distribution over conjunction, $K(P \wedge Q) \to K(P) \wedge K(Q)$, is not *valid;*

(iii) Distribution over implication, $K(P \to Q) \to (K(P) \to K(Q))$, is not *valid.*

Thus agents are *not* logically omniscient: they fail to know *all* logical consequences of their knowledge. This is in accordance with the findings in Section 2.3: neither closure under logical consequence nor distribution over implication should hold for knowledge. However, the restricted closure principle that we formulated in Section 2.3 is valid, along with some other weaker closure principles for K.

Proposition 4.8 (Closure principles for knowledge)**.** *The following principles are valid:*

(i) Weak closure under logical consequence: if the proposition $P \to Q$ is valid in S, then $(K(P) \wedge R(Q)) \to K(Q)$ is valid in S;

(ii) Weak distribution over implication: $K(P \to Q) \to ((K(P) \wedge R(Q)) \to K(Q))$;

(iii) Closure under logical equivalence: if $P \leftrightarrow Q$ is valid in S, then $K(P) \leftrightarrow K(Q)$ is valid in S;

(iv) Weak distribution over conjunction: $(K(P \wedge Q) \wedge R(P)) \to K(P)$;

(v) Closure under conjunction: $K(P) \wedge K(Q) \to K(P \wedge Q)$.

Proof. Let $S = (W, \sim, \approx)$ be an epistemic issue structure and pick $w \in W$ arbitrarily.

For (i), suppose that $P \to Q$ is valid in S. By Proposition 4.2 (ii), it follows that $K(P \to Q)$. So this proof reduces to the proof of (ii) given below.

For (ii), suppose $w \in K(P \to Q)$ and $w \in K(P) \wedge R(Q)$. From the former it follows that $w(\sim) \subseteq P \to Q$ and from the latter that $w(\sim) \subseteq P$, thus $w(\sim) \subseteq Q$. Combining this with the right conjunct in the second assumption yields $w \in K(Q)$.

For (iii), suppose $W = P \leftrightarrow Q$. By Proposition 4.6 (iii), $R(P)$ is valid in S if, and only if, $R(Q)$ is valid in S. Moreover, $w(\sim) \subseteq P$ if, and only if, $w(\sim) \subseteq Q$. So $K(P) \leftrightarrow K(Q)$.

For (iv), suppose $w \in K(P \wedge Q) \wedge R(P)$. It follows that $w(\sim) \subseteq P \wedge Q \subseteq P$ and $w \in R(P)$. So $w \in K(P)$.

For (v), suppose $w \in K(P) \wedge K(Q)$, then $w(\sim) \subseteq P, Q$. Hence $w(\sim) \subseteq P \wedge Q$ as well. Furthermore, since R is closed under conjunction it follows that $w \in K(P \wedge Q)$. □

So, in general, logical consequences of known propositions are known if they are issue-relevant. Agents only reason towards issue-relevant consequences of their knowledge. Thus the framework leaves room for reasoning: in principle, agents can obtain new knowledge by reason alone. The weak closure principles above show how vital issue-relevance is when determining which logical consequences of knowledge are also known. In fact, to formulate such principles, the modality R is indispensable.

Although agents are not logically omniscient in its typical sense, they are still exceptionally gifted reasoners: by (i) and (ii), agents reason towards all issue-relevant logical consequences of their knowledge. So agents are logically omniscient with regard to all issue-relevant information. Closure under logical equivalence, (iii), confirms that we have restricted ourselves to intensions. If $P \leftrightarrow Q$ is valid, then $P = Q$, so if only intensions are taken into account $K(P)$ and $K(Q)$ should indeed be equivalent in this case. Finally, since K is not normal, its interaction with conjunction is less straightforward than we are used to. By (iv), individual conjuncts of known conjunctions are only known if they are issue-relevant. Thus, in general, a single proposition of the form $K(P_1 \wedge \cdots \wedge P_n)$ cannot (partially) describe an agent's knowledge base. Instead, a proposition of the form $K(P_1) \wedge \cdots \wedge K(P_n)$ typically carries information about an agent's knowledge base. A proposition of the latter form implies one of the former form by (v), but generally they are not equivalent.

4.4 Knowledge about knowledge

In Section 2.3, we argued that our answer to the questions whether knowledge is positively or negatively introspective is 'no' in either case. We return to the example from that section.

Example 6 (Exam question: revisited)**.** Alice's epistemic situation while making the chemistry exam is captured by the epistemic issue structure S in Figure 4.4 (a). Interpret the proposition P as 'photosynthesis is an endothermic reaction'. Since P is an answer to the exam question that is on her agenda, P should be issue-relevant. Indeed, $R(P)$ is valid in S, since it does not cut across any issue cells. Alice studied well, therefore her information state at the actual world $w_1(\sim)$ is contained in P. So Alice knows that photosynthesis is an endothermic reaction; $w_1 \in K(P)$. However, this does not mean that Alice also knows that she knows that photosynthesis is an endothermic reaction. Since $w_3(\sim) = w_1(\sim) = \{w_1, w_3\} \subseteq P$, also $w_3 \in K(P)$. But because $w_4(\sim) = \{w_4, w_2\} \nsubseteq P$, we have $w_4 \notin K(P)$. Since $w_3 \approx w_4$ with $w_3 \in K(P)$ and $w_4 \notin K(P)$, it follows that $RK(P)$ is not valid in S. Consequently, $KK(P) = \bot$ and in particular $w_1 \notin KK(P)$. So Alice does not know that she knows that photosynthesis is an endothermic reaction, because her agenda does not prompt her to.

The second scenario in which a fellow student asks Alice whether she knows the answer to the photosynthesis question is captured by the epistemic issue structure S' in Figure 4.4 (b). The worlds w_3' and w_4' are not part of the same issue cell, hence $RK(P)$ is valid in S'. Also $w_1'(\sim) = w_3'(\sim) \subseteq K(P)$, so $w_1' \in KK(P)$; Alice knows that she knows that photosynthesis is an endothermic reaction.

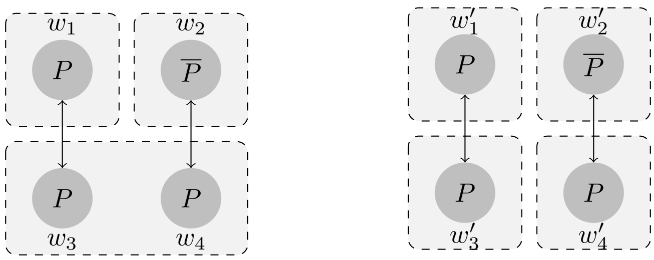

(a) $S = (W, \sim, \approx)$ (b) $S' = (W', \sim', \approx')$

Figure 4.4: Two epistemic issue structures S and S'. The actual worlds are denoted by w_1 and w_1', respectively.

Since $w_1 \notin K(P) \to KK(P)$ in the structure in Figure 4.4 (a), it shows that positive introspection does not hold in general. Moreover, the same structure is also a counterexample against negative introspection: because $w_2(\sim) = \{w_2, w_4\} \nsubseteq P$, we have $w_2 \in \neg K(P)$. Since issue-relevance is closed under negation it follows that $R\neg K(P)$ is not valid, hence $w_2 \notin K\neg K(P)$. Thus $w_2 \notin \neg K(P) \to K\neg K(P)$; negative introspection fails to hold. The failure of classical introspection is summarized by the following proposition.

Proposition 4.9 (Failure of introspection)**.** *The following statements illustrate the failure of full introspection:*

(i) Positive introspection, $K(P) \to KK(P)$, is not *valid;*

(ii) Negative introspection, $\neg K(P) \to K\neg K(P)$, is not *valid.*

Although both positive and negative introspection do not hold in general, there exist valid introspection principles. As with closure, weaker principles can be obtained if one distinguishes between issue-relevant and issue-irrelevant propositions about knowledge or ignorance.

Proposition 4.10 (Introspection principles)**.** *The following principles are valid:*

(i) Weak positive introspection: $(K(P) \wedge RK(P)) \rightarrow KK(P)$;

(ii) Weak negative introspection: $(\neg K(P) \wedge RK(P)) \rightarrow K\neg K(P)$;

(iii) Iterated introspection: $KK(P) \rightarrow KKK(P)$;

(iv) Positive issue-relevance introspection: $R(P) \rightarrow KR(P)$;

(v) Negative issue-relevance introspection: $\neg R(P) \rightarrow K\neg R(P)$.

Proof. Let $S = (W, \sim, \approx)$ be an epistemic issue structure and pick $w \in W$ arbitrarily.

For (i), suppose $w \in K(P) \wedge RK(P)$. The left conjunct gives us $w(\sim) \subseteq P$ and $w \in R(P)$. Since $R(P)$ must be valid in S, $w(\sim) \subseteq R(P)$. Since $\sim$ is an equivalence relation, $w(\sim) = v(\sim)$ for all $v \in w(\sim)$, hence $v(\sim) \subseteq P$ for all $v \in w(\sim)$. So $w(\sim) \subseteq K(P)$. Combining this with the right conjunct yields $w \in KK(P)$.

For (ii), suppose $w \in \neg K(P) \wedge RK(P)$. From the right conjunct and closure of issue-relevance under negation it follows that $w \in R\neg K(P)$, thus it suffices to prove that $w(\sim) \subseteq \neg K(P)$. There are two cases in which $w \in \neg K(P)$: either $w(\sim) \nsubseteq P$ or $w \notin R(P)$. In the former case, there must be a $u \in w(\sim)$ such that $u \notin P$. Suppose for reductio that there exists a $v \in w(\sim)$ such that $v \in K(P)$, then $v(\sim) \subseteq P$ as well. But since $w(\sim) = v(\sim)$, this implies that $u \in P$; contradiction. So $v \in \neg K(P)$ and also $w(\sim) \subseteq \neg K(P)$. In the latter case, necessity of issue-relevance gives $\neg R(P) = W$. It follows that $w(\sim) \subseteq \neg K(P)$, completing the proof.

For (iii), suppose $w \in KK(P)$. Since $w(\sim) = v(\sim)$ for all $v \in w(\sim)$ it follows that $w(\sim) \subseteq KK(P)$, hence it suffices to prove that $w \in RKK(P)$. Assume that $w \notin RKK(P)$ for reductio, then there exist $v, u \in W$ such that $v \approx u$, $v \in KK(P)$ and $u \notin KK(P)$. From $v \in KK(P)$ it follows that $RK(P)$ and $R(P)$ are valid in S. By factivity of knowledge $v \in K(P)$ and $v \in P$. From $v \approx u$ it follows that $u \in K(P)$ and $u \in P$. Since $u \in \neg KK(P)$, there is a $u' \in W$ such that $u \sim u'$ and $u' \in \neg K(P)$. In turn, there is a $u'' \in W$ such that $u' \sim u''$ and $u'' \in \neg P$. By transitivity, $u \sim u''$, but this yields a contradiction: $u \notin K(P)$. Hence $w \in RKK(P)$.

For (iv), assume that $w \in R(P)$. It follows that $R(P) = \top$. From $K\top$ it follows that $w \in KR(P)$.

For (v), assume that $w \in \neg R(P)$. It follows that $R(P) = \bot$. So $\neg R(P) = \top$ and by $K\top$ it follows that $w \in K\neg R(P)$. $\square$

The weakened introspection principles (i) and (ii) state that in order to have introspection with respect to a proposition P, the agent's knowledge of P must be issue-relevant. So instead of requiring all knowledge to be introspective or having introspection as an accidental property, we have it in a very principled manner: an agent has introspection with regard to a proposition if, and only if, the epistemic status of that proposition is issue-relevant. This is in line with the restricted introspection principles formulated in Section 2.3. Moreover, property (iii) shows that once an agent has some introspection, it can be iterated indefinitely. Hence introspection does not come in degrees.

Lastly, properties (iv) and (v) state that agents are always introspective with regard to issue-relevance: they know exactly which propositions are issue-relevant and which are not. This reflects our assumption that rational agents are aware of their agenda and in full control of the issues on it.

4.5 Information and issues

Knowledge is dependent on the information relation as well as the (current) issue relation. The latter governs issue-relevance, for which we have introduced a separate modality R. We can do the same for the information relation, which governs information possession.

Definition 4.11 (Information modality)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and P a proposition on W. The information modality I is defined as:

$$I(P) := \{w \in W \mid w(\sim) \subseteq P\}.$$

If $w \in I(P)$, then we say that the agent possesses the information that P at world w.

The information modality I is simply the Kripke modality over the information relation. Its properties are well known.

Proposition 4.12 (Properties of information possession)**.** *The following principles are valid:*

(i) Distribution over implication: $I(P \to P') \to (I(P) \to I(P'))$*;*

(ii) Closure under logical consequence: if $P \to P'$ *is valid in* S*, then* $I(P) \to I(P')$ *is valid in* S*;*

(iii) Factivity of information: $I(P) \to P$*;*

(iv) Positive introspection of information: $I(P) \to II(P)$*;*

(v) Negative introspection of information: $\neg I(P) \to I\neg I(P)$*.*

These properties are the basic properties of any Kripke modality that corresponds to an equivalence relation.[59] These are expected and only confirm that information is well-behaved, as previously asserted in Section 3.1. Using the information modality, we can formally express that knowledge is issue-relevant information: $K(P) \leftrightarrow I(P) \wedge R(P)$. Its validity can be derived directly from Definitions 3.10, 4.3 and 4.11.

Likewise, we can consider the Kripke modality over the issue relation.

Definition 4.13 (Issue modality)**.** Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure and P a proposition on W. The issue modality Q is defined as:

$$Q(P) := \{w \in W \mid w(\approx) \subseteq P\}.$$

If $w \in Q(P)$, then we say that 'P holds in all worlds that are issue-equivalent to w'.

The phrase "holding in all issue-equivalent worlds" is rather technical. When interpreting issues as questions, '$w \in Q(P)$' can be interpreted as 'the full answer to the current issue at w carries the information that P'.[60] We find that $w \in Q(P)$ whenever the issue cell $w(\approx)$ is contained in P, that is, when P is true throughout the current issue cell of w. If the agent receives the full answer to her current issue, the information relation is restricted to her current issue cells. So if $w \in Q(P)$, it follows that $w \in I(P)$ after fully resolving the agent's issue. This does not mean that P is also known; Example 5 bears witness to this. However, it does mean that the information that P is carried by the full answer to the agent's current issue.

[59] The properties of the modal logic **S5** are given, for instance, by Blackburn, De Rijke and Venema (2001, §4.1).

[60] This interpretation of the issue modality agrees with the original interpretation of Van Benthem and Minică (2012). When construing issues in a more general manner, it can be interpreted as 'P is carried by the information resolving the current issue at world w'.

Since both I and Q are Kripke modalities over an equivalence relation, the issue modality Q also has the properties listed in Proposition 4.12. These properties can be interpreted in a similar manner, stating that the information carried by the full answer to the agent's current issue is well behaved.

Proposition 4.14. *Let $S = (W, \sim, \approx)$ be an epistemic issue structure and let P be a proposition on W. The following statements are equivalent:*

(i) $R(P)$ is valid in S;

(ii) $P \to Q(P)$ is valid in S;

(iii) $Q(P) \vee Q(\neg P)$ is valid in S.

Proof. Let $S = (W, \sim, \approx)$ be an epistemic issue structure. Let P be a proposition on W and take an arbitrary $w \in W$.

$(i) \Rightarrow (ii)$: Suppose that $R(P)$ is valid in S and that $w \in P$. Then by definition of R, $w(\approx) \subseteq P$, thus $w \in Q(P)$.

$(ii) \Rightarrow (iii)$: Suppose that $P \to Q(P)$ is valid in S. We distinguish two cases: $w \in Q(P)$ or $w \notin Q(P)$. In the former case it immediately follows that $w \in Q(P) \vee Q(\neg P)$ as well. In the latter case, it follows that $w \in \neg Q(P)$, which means that $w \in \neg P$ by the contraposition of our assumption. So if a world's issue cell does not exclusively contain P-worlds, then all the worlds in that issue cell are $\neg P$-worlds. Thus $w \in Q(\neg P)$, from which it follows that $w \in Q(P) \vee Q(\neg P)$.

$(iii) \Rightarrow (i)$: Suppose that $Q(P) \vee Q(\neg P)$ is valid in S. Then every issue cell either consists only of P-worlds or only of $\neg P$-worlds, so that P is issue-relevant. $\square$

Thus, issue-relevance can be expressed in terms of validities involving the issue modality.[61]

Throughout the remainder of this chapter, as well as the subsequent two chapters, both the information modality and the issue modality will prove to be invaluable tools. Agenda updates only affect the current issue relation of an agent, whereas issue resolution only affects the information in her possession. The information modality will allow us to distinguish between changes in knowledge due to agenda updates and changes as a result of issue resolution. The normality of the issue modality will be convenient when developing a sound and complete dynamic logic of epistemic issues, as will become apparent in Chapter 6. However, before then, we investigate the dynamic principles of knowledge.

4.6 Dynamic principles of knowledge

We circle back to the example that initiated our journey.[62]

Example 7 (The slave in the *Meno*)**.** Socrates asks Meno's slave how to double the area of a square. The slave does not know and suggests that doubling the length of the sides may double the area of the square. Through a series of questions posed by Socrates,

[61] Van Benthem and Minică (2012, p. 4) use the characterization in Proposition 4.14 (iii), whereas we will mainly use (ii).

[62] The Thesis Committee pointed out that the formalization of this example is flawed. The worlds w_2, w_3 and w_4 represent mathematically *impossible* worlds. Consequently, to make the example work, the information relation needed to be maximally fine-grained in the initial structure S: the agent possesses *all* possible information in the initial structure. However, this goes against the *a priori* nature of mathematical knowledge: mathematical knowledge should be attainable without having any information. Nonetheless, when reinterpreting P and Q as different, contingent propositions, the example still illustrates how an agent may obtain knowledge by merely adding issues to her agenda.

the slave obtains the knowledge that doubling the length of the sides creates a square with an area four times larger than the initial square. Socrates proceeds by asking more questions. Eventually, this leads the slave to know that using the diagonal of the initial square as base yields a square double the size.[63]

Interpret P as ‘doubling the sides of a square makes it four times larger’ and Q as ‘taking the diagonal of a square as base yields a square with double the original area’. Then the initial state of the slave is represented in Figure 4.5 (a), where w_1 is the actual world. Neither P nor Q is issue-relevant, because the slave is not inquiring into these geometrical propositions. Geometrical truths can be known *a priori*, that is, without any information available. Hence all the information required to know P and Q is in possession of the agent; the information relation is simply the identity relation.

When Socrates poses questions, the slave puts the corresponding issue, say x, on his agenda. The first series of questions prompts the slave to make the necessary conceptual distinctions to know P. Thus $\approx_x = \{(w_1, w_2), (w_3, w_4)\}^*$, where R^* denotes the reflexive-symmetric closure of R. After updating his agenda, the slave's epistemic situation is represented by the structure in Figure 4.5 (b). Now $R(P)$ is valid, and because the required information was already possessed by the slave he knows that P: $w_1 \in K(P)$.

Socrates resumes his questioning and the slave updates his agenda by adding the issue y, represented by $\approx_y = \{(w_1, w_3), (w_2, w_4)\}^*$. In the updated structure, shown in Figure 4.5 (c), $R(Q)$ is also valid. Consequently, the slave also knows that taking the diagonal of a square as base yields a square with double the original area. Proving Socrates' point that knowledge can be obtained through mere inquiry.

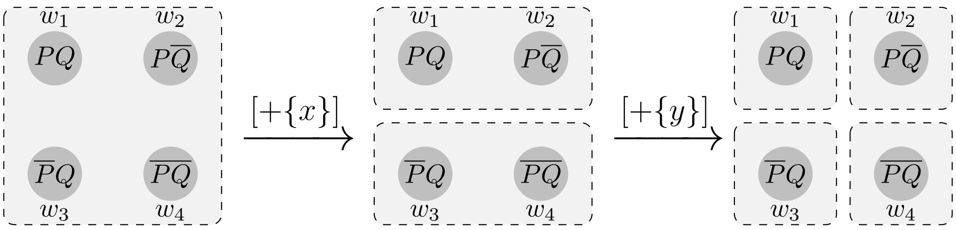

(a) S (b) $S_{[+\{x\}]}$ (c) $S_{[+\{x\}][+[\{y\}]}$

Figure 4.5: From left to right: the dynamic issue-epistemic structures S, $S_{[+\{x\}]}$ and $S_{[+\{x\}][+\{y\}]}$. All information is available to the agent and reflexive arrows have been omitted.

The example of the slave can be generalized to mathematics in general. Obtaining mathematical knowledge is not about obtaining information that rules out possible worlds, but rather about making new conceptual distinctions.[64] The solution of a mathematical problem is often the result of asking deeper, new questions, allowing mathematicians to make new conceptual distinctions. This corresponds to the addition of issues to an agent's agenda. Since standard epistemic logic is only concerned with information possession, it cannot capture the attainment of knowledge as a result of making new conceptual distinctions.

The addition of issues does not only play a role in mathematics, but also in empirical science or more mundane, everyday situations. In an everyday context, new goals

[63] The slave learns more propositions in the actual dialogue. For the sake of clarity, we have distilled two key propositions.

[64] Cf. Carballo (2014). In light of footnote 62, it must be noted that our framework cannot properly capture the attainment of mathematical knowledge as a result of making new conceptual distinctions. However, with some minor modifications our framework can be turned into a hyperintensional framework that could model this. For a sketch of such a framework, see Section 7.4.

correspond to the addition of issues to an agent's agenda. In a scientific context, new research directions and questions correspond to the addition of issues to an agenda. However, as discussed in Section 2.2, the addition of an issue is often preceded by the retraction of an issue.

Example 8 (Relativity of simultaneity: revisited)**.** Before Einstein introduced this theory, simultaneous observations of events a and b were taken to mean that a and b occurred simultaneously. If we interpret P as 'events a and b are observed simultaneously', then Figure 4.6 (a) represents this situation, where w_1 is the actual world. Let x be the basic issue inquiring into the simultaneity of observations regarding events a and b. The corresponding issue relation $\approx_x$ is the identity relation on the set of worlds in S. Since the agent has observed events a and b simultaneously, she possesses the information that P at world w_1. Moreover, she deems such observation meaningful, hence has the issue x on her agenda. Since each issue cell consists of only one world, $R(P)$ is valid. Thus $w_1 \in K(P)$; the agent knows that events a and b are observed simultaneously.

However, after the introduction of the special theory of relativity the simultaneous observation of two events is no longer meaningful, thus x is retracted from the agent's agenda. The updated structure is shown in Figure 4.6 (b). Since the only issue cell in the model contains both a P-world and a $\neg P$-world, $R(P)$ is no longer valid. Consequently, the agent no longer knows that the events a and b seem to occur simultaneously.

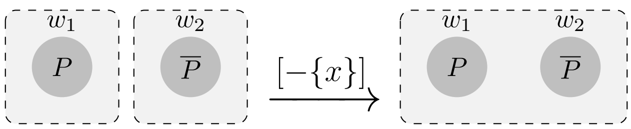

(a) S (b) $S_{[-\{x\}]}$

Figure 4.6: Two epistemic issue structures S and $S_{[-\{x\}]}$. All information is available to the agent and reflexive arrows have been omitted.

The examples above show how knowledge may be gained or lost after adding or retracting issues, respectively. Agenda updates adhere to the following general principles.

Proposition 4.15 (Principles of agenda update)**.** *Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure. We write '$[\pm X]$' when either '$[+X]$' or '$[-X]$' can be inserted. The following propositions are valid.*

- *(i) Agenda updates do not affect information possession: $[\pm X]I(P) \leftrightarrow I(P)$;*
- *(ii) Knowledge is monotonic with respect to issues: $K(P) \to [+X]K(P)$;*
- *(iii) Retraction of issues yields no new knowledge: $[-X]K(P) \to K(P)$;*
- *(iv) Possessed information that is issue-relevant after an agenda update is known after the agenda update: $I(P) \wedge [\pm X]R(P) \to [\pm X]K(P)$.*

Proof. Let $S = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ be an epistemic issue structure.

For (i), observe that agenda updates do not alter the information relation or the set of worlds, so information possession is invariant under agenda updates.

For (ii), it follows from (i) that it suffices to argue that issue-relevant propositions cannot become issue-irrelevant after the addition of an issue. If a proposition P is issue-relevant, it does not cut across the issue cells of the agent's current issue. Since the addition of issues can only refine issue cells, proposition P cannot cut across the updated issue cells.

For (iii), observe that this is a consequence of (ii): if P is known after retracting an issue X, then adding X to agent's agenda yields the initial structure in which P must also be known.

For (iv), it can be observed that this is a consequence of (i) and the fact that $K(P) \leftrightarrow I(P) \wedge R(P)$ is valid. □

Property (i) states that agenda updates do not affect the information that is possessed by an agent; it only changes the manner in which she processes information. Property (ii) and (iii) express that knowledge is monotonic with respect to issues: the addition of issues cannot result in a loss of knowledge. Likewise, any knowledge that an agent has after the retraction of an issue was already in her possession before retracting the issue. Property (iv) characterizes the conditions under which an agent has knowledge after an agenda update: if information that P is possessed by an agent and P is issue-relevant after the agenda update, then P is known after the agenda update.

Besides updating their agendas, agents may resolve issues. Note that issues that are *not* on an agent's agenda may also be resolved. The resolution of an issue comes down to obtaining the information that resolves it, regardless of whether that information is issue-relevant. It conforms to the following principles.

Proposition 4.16 (Principles of issue resolution)**.** *Let* $S = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A})$ *be an epistemic issue structure and let* $X \subseteq \mathcal{I}$*. The following propositions are valid.*

(i) Issue resolution does not affect issue-relevance: $[X!]R(P) \leftrightarrow R(P)$*;*

(ii) Knowledge is monotonic with respect to information: $K(P) \rightarrow [X!]K(P)$*;*

(iii) Issue-relevant information that is possessed after issue resolution is known after issue resolution: $R(P) \wedge [X!]I(P) \rightarrow [X!]K(P)$*.*

Proof. Let $S = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A})$ be an epistemic issue structure.

For (i), observe that issue resolution only alters the information relation in a structure, hence issue-relevance is invariant under agenda updates.

For (ii), it follows from (i) that it suffices to show that information possession is monotonic under issue resolution. Suppose that the information that P is possessed before resolving an issue X. Because $\sim \cap \approx_X \subseteq \sim$, we also have $w(\sim \cap \approx_X) \subseteq w(\sim) \subseteq P$, thus the information that P is also possessed after resolving X.

For (iii), observe that this follows from (i) and the fact that $K(P) \leftrightarrow I(P) \wedge R(P)$ is valid. □

Property (i) states that issue resolution does not affect an agent's issue; it only changes the information she possesses. Property (ii) expresses that learning new information cannot result in the loss of knowledge. Lastly, property (iii) characterizes the conditions under which an agent gains knowledge after resolving an issue: if P is issue-relevant and the information that P is possessed after resolving the issue X, then P is known after resolving X.

Finally, all of the above can be combined into a single example.

Example 9 (Galileo's telescope: revisited)**.** Galileo's situation prior to inventing his telescope is captured by the structure in Figure 4.7 (a), where w_1 denotes the actual world. Interpret P as 'the moon is smooth (when observed by the naked eye)' and Q as 'the moon contains mountains and craters (when observed with a telescope)'. On Galileo's agenda is the issue x, pertaining to information about the moon obtained by direct observation. The corresponding issue relation is $\approx_x = \{(w_1, w_2), (w_3, w_4)\}^*$, where R^* denotes the reflexive-symmetric closure of R. The issue y, which initially is not on Galileo's agenda, pertains to information about the moon obtained by telescope

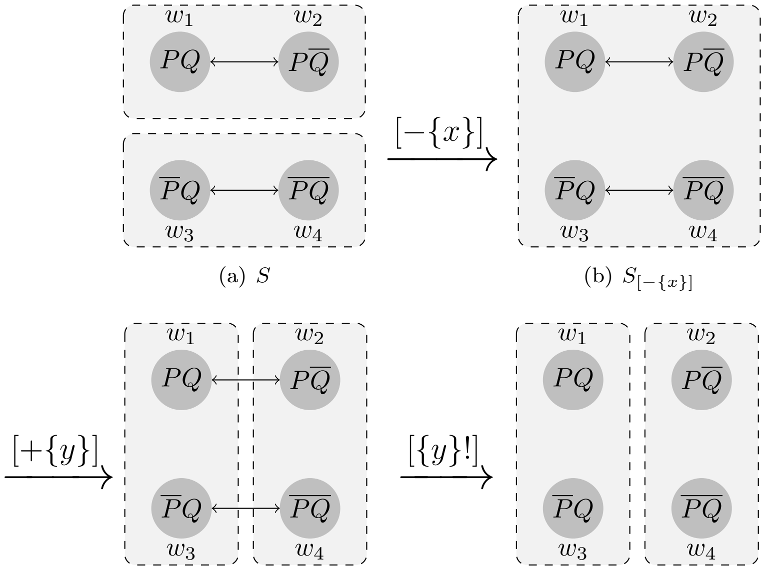

(a) S (b) $S_{[-\{x\}]}$

(c) $S_{[-\{x\}][+\{y\}]}$ (d) $S_{[-\{x\}][+\{y\}][\{y\}!]}$

Figure 4.7: Four epistemic issue structures. The actual world is denoted by w_1.

observations. The corresponding issue relation is $\approx_y = \{(w_1, w_3), (w_2, w_4)\}^*$. The proposition P is issue-relevant and since Galileo has directly observed the moon, he knows P; $w_1 \in K(P)$.

Galileo then doubted that the nature of the moon could be revealed by direct observation. He retracted the issue pertaining to direct observations of the moon from his agenda. The resulting structure is shown in Figure 4.7 (b). In this structure, P is no longer issue-relevant since it cuts across the only issue cell.

Thereafter, Galileo added the issue pertaining to telescope observations of the moon to his agenda. The updated structure is shown in Figure 4.7 (c). In this structure, Q is issue-relevant. However, the information obtained by directly observing the moon is not helpful towards resolving his new issue: Galileo cannot informationally distinguish between Q and $\neg Q$.

Consequently, Galileo resolves the issue y by making telescope observations of the moon. The updated structure is shown in Figure 4.7 (d). In addition to Q being issue-relevant, Galileo now also possesses the information that Q. Thus, he knows that the moon contains mountains and craters when observed with a telescope; $w_1 \in K(Q)$.[65]

This concludes the chapter on the laws of knowledge. We have seen how the examples from Chapter 2 can be captured in epistemic issue structures and the laws to which knowledge adheres. In the two subsequent chapters we define logics on epistemic issue structures.

[65]There is some tension in this example between the interpretations of P and Q and the practice of letting possible worlds only represent *ontic* facts. The propositions P and Q seemingly contradict each other, while being jointly true and false in worlds w_1 and w_4, respectively. This may seem to suggest that they express *phenomenal* facts. However, this is not the intended interpretation: not the agent's senses, but the nature of direct and telescope observations creates the seeming contradiction. In the example, it is taken as an ontic fact that information about the moon obtained by direct observation tells us that the moon is perfectly smooth. Similarly, it is taken as an ontic fact that information about the moon obtained by telescope observations tells us that the moon contains mountains and craters. Thus, in this example, a contradiction is avoided by relativizing information to the method of observation.

Chapter 5

Static Epistemic Logics

In this chapter, we define several sound and complete logics. In Section 5.1, epistemic issue models are defined and it is argued that a logic based on a language containing only a knowledge modality is uninteresting in our framework. In Sections 5.2–5.6, several logics involving knowledge, issue-relevance, necessity, information or issues are developed and shown to be sound and complete. The logic of knowledge and issue-relevance given in Section 5.2 is the least expressive logic, and logics in subsequent sections become progressively more expressive. In Section 5.7, we reflect on the relationship between the logics presented in this chapter and conclude with an example showing the logics in action.

None of the logics introduced in this chapter take dynamic actions or individual basic issues into account. Hence throughout this section we steadily refer to epistemic issue structures in the compressed manner. Furthermore, throughout this chapter we often suppress subscripts when the context permits it. For instance, if L denotes a logic, we may write '$\vDash$' instead of '$\vDash_{\mathsf{L}}$' or '$\mathcal{L}(\Phi)$' instead of '$\mathcal{L}_{\mathsf{L}}(\Phi)$'.

5.1 Knowledge vs. epistemic possibility

We started out by observing that the notion of knowledge as information possession captured by standard epistemic logic is conceptually rather thin. Accordingly, we start the logical study of our framework by considering the language of standard epistemic logic.

Definition 5.1 (Language of epistemic logic)**.** The language $\mathcal{L}_{\mathsf{K}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K\varphi,$$

with $p \in \Phi$.

The Boolean connectives are interpreted as usual. The formula '$K\varphi$' means 'the agent knows that φ'. The language $\mathcal{L}_{\mathsf{K}}(\Phi)$, along with the other languages that we will consider in this and the subsequent chapter, are interpreted on epistemic issue structures augmented with a valuation.

Definition 5.2 (Epistemic issue models)**.** An epistemic issue model M is a tuple $(W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A}, ||\bullet||)$ such that $(W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ is an epistemic issue structure and $||\bullet|| : \Phi \to \mathcal{P}(W)$ is a valuation, assigning atomic propositions to sets of possible worlds. If $(W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A})$ is a single issue structure, then we call M a single issue model.

We use the same abbreviation for epistemic issue structures as before. So when we are not considering dynamics or individual basic issues, we can denote an epistemic issue model as $M = (W, \sim, \approx, || \bullet ||)$. The semantics for $\mathcal{L}_{\mathsf{K}}(\Phi)$ are in consonance with Proposition 3.11.

Definition 5.3 (Semantics for $\mathcal{L}_{\mathsf{K}}(\Phi)$)**.** Let $M = (W, \sim, \approx, || \bullet ||)$ be an epistemic issue model and let $w \in W$. We have the following semantic clauses.

- $M, w \vDash p$ iff $w \in ||p||_M$;
- $M, w \vDash \neg\varphi$ iff $M, w \nvDash \varphi$;
- $M, w \vDash \varphi \wedge \psi$ iff $M, w \vDash \varphi$ and $M, w \vDash \psi$;
- $M, w \vDash K\varphi$ iff for all $u, v \in W$ such that $u \approx v$, $M, u \vDash \varphi$ implies $M, v \vDash \varphi$ and $M, v \vDash \varphi$ for all $v \in W$ such that $w \to v$.

In standard epistemic logic, 'the agent knows that φ' often means the same as 'the agent possesses the information that φ'. Hence '$\neg K\varphi$' expresses 'the agent does not have the information that φ'. Epistemic possibility is the dual of K in standard epistemic logic, often abbreviated as $\hat{K}\varphi := \neg K \neg \varphi$. It expresses that a formula φ is an epistemic possibility if the agent does not have the information that φ is not the case. Thus both knowledge and epistemic possibility can be expressed using only the K modality in standard epistemic logic.

Within our framework, however, this is not the case. The formula '$K\varphi$' expresses that the agent possesses the information that φ, *and* that φ is issue-relevant. Its negation, '$\neg K\varphi$', expresses that *either* the agent does not possesses the information that φ, *or* φ is not issue-relevant. However, it cannot be specified which of these is the case. Consequently, '$\neg K \neg \varphi$' does not express epistemic possibility of φ.[66]

So, because two conditions need to be satisfied in order to have knowledge in our framework, epistemic possibility cannot be expressed as the dual of K. Therefore, the language $\mathcal{L}_{\mathsf{K}}(\Phi)$ is not equipped to fully describe an agent's epistemic state within our framework. As such, a logic of knowledge as issue-relevant information based on the language $\mathcal{L}_{\mathsf{K}}(\Phi)$ is not interesting.

Propositions that are issue-irrelevant are epistemically insignificant for the agent. Hence epistemic possibility of a proposition P should express that P is issue-relevant, but that the agent has insufficient information to exclude P. Thus we can define epistemic possibility semantically as $\hat{K}(P) := \neg K(\neg P) \wedge R(P)$, which is the usual abbreviation of epistemic possibility in conjunction with the issue-relevance of P.

We could augment $\mathcal{L}_{\mathsf{K}}(\Phi)$ with $\hat{K}$ and extend the semantics accordingly to match the expressive power of standard epistemic logic. However, this would be equivalent to simply adding the issue-relevance modality R, because we can define $R(P) := K(P) \vee \hat{K}(\neg P)$: a proposition is issue-relevant exactly when it is known, or its negation is an epistemic possibility. In other words, the modalities $\hat{K}$ and R are interdefinable.

Since issue-relevance is a more primitive notion than epistemic possibility within our approach, we elect to study the logic based on the language containing K and R as primary modalities rather than K and $\hat{K}$. Epistemic possibility of P can then be defined as an abbreviation for not knowing the denial of the issue-relevant proposition P.

In conclusion, within our approach a logic of knowledge based on a language with only a single modal operator lacks expressive power because it cannot express epistemic possibility. Extending the language of such a logic with an epistemic dual is equivalent to adding the issue-relevance modality. The resulting logic is treated in the next section.

[66]A similar problem arises with topic-sensitive notions of knowledge, as discussed by Rossi and Özgün (2023).

5.2 The logic of knowledge and issue-relevance

In the previous section we argued that a language with only a K modality is not enough to obtain a sufficiently strong logic that fully captures knowledge as issue-relevant information. Therefore, we consider the language containing both the knowledge modality K as well as the issue-relevance modality R.

Definition 5.4 (Language of KR)**.** The language $\mathcal{L}_{\mathsf{KR}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K\varphi \mid R\varphi,$$

with $p \in \Phi$.

The Boolean connectives and the modality K are interpreted as before. The formula '$R\varphi$' means 'φ is issue-relevant'.

Definition 5.5 (Semantics of KR)**.** Let $M = (W, \sim, \approx, ||\bullet||)$ be an epistemic issue model and let $w \in W$. In addition to the semantic clauses for $\mathcal{L}_{\mathsf{K}}(\Phi)$, we have:

$$M, w \vDash_{\mathsf{KR}} R\varphi \text{ iff for all } u, v \in W \text{ with } u \approx v,\ M, u \vDash_{\mathsf{KR}} \varphi \text{ implies } M, v \vDash_{\mathsf{KR}} \varphi.$$

Because both K and R lack distribution over conjunction, it will be worthwhile to introduce specific notation for conjunctions of K-formulas and R-formulas. To this end, we define R-conjuncts and K-conjuncts.

Definition 5.6 (R-conjuncts and K-conjuncts)**.** A formula $\rho \in \mathcal{L}_{\mathsf{KR}}(\Phi)$ is a K-conjunct if it is of the form

$$\rho = \left(\bigwedge_{i \in I} K\rho_i\right) \wedge \left(\bigwedge_{j \in J} \neg K\rho_j\right),$$

where I and J are finite index sets and $\rho_k \in \mathcal{L}_{\mathsf{KR}}(\Phi)$ for all $k \in I \cup J$.

A formula $\eta \in \mathcal{L}_{\mathsf{KR}}(\Phi)$ is an R-conjunct if it is of the form

$$\eta = \left(\bigwedge_{i \in I} R\eta_i\right) \wedge \left(\bigwedge_{j \in J} \neg R\eta_j\right),$$

where I and J are finite index sets and $\eta_k \in \mathcal{L}_{\mathsf{KR}}(\Phi)$ for all $k \in I \cup J$.

Observe that K-conjuncts can be interpreted as knowledge bases. They give a partial description of the knowledge and ignorance of an agent. Similarly, R-conjuncts provide a partial description of which formulas are issue-relevant and which are not. Thus an R-conjunct can be interpreted as an issue-relevance base. We will use K-conjuncts and R-conjuncts to express some of the axioms.

Proof system for KR
The proof system for KR consists of:

- The rules and axioms of propositional logic;
- Necessitation for K: from φ infer $K\varphi$;
- Weak closure for K: from $\rho \to \varphi$ infer $(\rho \wedge R\varphi) \to K\varphi$, where ρ is a K-conjunct;
- Factivity for K: $K\varphi \to \varphi$;
- Strong RE for R: from $\eta \to (\varphi \leftrightarrow \psi)$ infer $\eta \to (R\varphi \to R\psi)$, where η is an R-conjunct;

- Closure under negation for R: $R\varphi \to R\neg\varphi$.
- Closure under conjunction for R: $(R\varphi \wedge R\psi) \to R(\varphi \wedge \psi)$;
- Knowledge implies issue-relevance: $K\varphi \to R\varphi$;
- Positive introspection of issue-relevance: $R\varphi \to KR\varphi$;
- Negative introspection of issue-relevance: $\neg R\varphi \to K\neg R\varphi$.

Necessity for K states that theorems are known by the agent. Weak closure for K states that given a knowledge base, any logical consequence of that knowledge base is known if it is issue-relevant. Strong RE for R states that if two formulas are equivalent given an issue-relevance base η, then issue-relevance of one of these formulas implies issue-relevance of the other given that issue-relevance base. Closure under negation and conjunction for R ensure that issue-relevance behaves in a coherent manner. Knowledge implies issue-relevance is a direct consequence of considering knowledge as issue-relevant information: it ensures that all knowledge is issue-relevant. Lastly, positive and negative introspection of issue-relevance express that the agent knows which formulas are issue-relevant and which are not. This reflects the assumption that a rational agent is fully aware of the issues on her agenda.

Proposition 5.7 (Soundness of KR)**.** *The logic* KR *is sound.*

Proof. The validity of all but two axioms is already proven in Chapter 4. We prove the cases of strong RE for R and weak closure for K here. Let $M = (W, \sim, \approx, ||\bullet||)$ be an arbitrary epistemic issue model and let $w \in W$ be arbitrary.

For strong RE for R, assume $M \vDash \eta \to (\varphi \leftrightarrow \psi)$, where η is an R-conjunct. Furthermore, suppose that $M, w \vDash \eta$ and $M, w \vDash R\varphi$. Since issue-relevance and issue-irrelevance are global, $M \vDash \eta$ and $M \vDash R\varphi$ follow. From the former and our initial assumption it follows that $M \vDash \varphi \leftrightarrow \psi$. Combining this with the latter shows that there cannot exist worlds $u, v \in W$ such that $u \approx v$, $M, u \vDash \psi$ and $M, v \nvDash \psi$, therefore $M \vDash R\psi$. It follows that $M \vDash \eta \to (R\varphi \to R\psi)$.

For weak closure for K, suppose $M \vDash \rho \to \varphi$, where $\rho = (\bigwedge_{i \in I} K\rho_i) \wedge (\bigwedge_{j \in J} \neg K\rho_j)$ is a K-conjunct. Furthermore, take an arbitrary $w \in W$ and assume that $M, w \vDash \rho \wedge R\varphi$. From the first conjunct and our initial assumption it follows that $M, w \vDash \varphi$. By the second conjunct and the necessity of issue-relevance it suffices to prove that $M, v \vDash \varphi$ for all $v \in W$ such that $w \sim v$. Suppose for reductio that this is not the case, then there exists a $v \in W$ such that $w \sim v$ and $M, v \nvDash \rho$. Either $M, v \nvDash K\rho_i$ for some $i \in I$ or $M, v \nvDash \neg K\rho_j$ for some $j \in J$. In the former case there exists a $u \in W$ such that $v \sim u$ and $M, u \nvDash \rho_i$, because issue-relevance is global. The transitivity of the information relation then yields a contradiction; $M, w \nvDash \rho$. In the latter case $M, u \vDash \rho_j$ for all $u \in W$ such that $v \sim u$ and $M \vDash R\rho_j$. However, since the information relation is an equivalence relation it follows that $M, u \vDash \rho_j$ for all $u \in W$ such that $w \sim u$, yielding the same contradiction; $M, w \nvDash \rho$. In either case we find that $M, v \vDash \varphi$ for all $v \in W$ such that $w \sim v$, completing the proof. $\square$

Proposition 5.8. *The following theorems and rules are derivable in* KR*:*

(i) Necessitation for R: from φ infer $R\varphi$;

(ii) Issue-relevance is issue-relevant: $RR\varphi$;

(iii) Weak closure under logical consequence for K: from $\varphi \to \psi$ infer $(K\varphi \wedge R\psi) \to K\psi$;

(iv) Weak distribution over implication for K: $K(\varphi \to \psi) \to ((K\varphi \wedge R\psi) \to K\psi)$;

(v) Weakened positive introspection for K: $(K\varphi \wedge RK\varphi) \to KK\varphi$;

(vi) Weakened negative introspection for K: $(\neg K\varphi \wedge RK\varphi) \to K\neg K\varphi$;

(vii) Weak distributivity over conjunction for K: $(K(\varphi \wedge \psi) \wedge R\varphi) \to K\varphi$;

(viii) Closure under conjunction for K: $(K\varphi \wedge K\psi) \to K(\varphi \wedge \psi)$;

(ix) RE for K: from $\varphi \leftrightarrow \psi$ infer $K\varphi \leftrightarrow K\psi$;

(x) Issue-relevance of iterated knowledge: $RK\varphi \to RKK\varphi$;

(xi) Unrestricted introspection: $KK\varphi \to KKK\varphi$.

Proof. Most derivations are easy. Rule (i) follows from necessitation for K and knowledge implies issue-relevance. Theorem (ii) can be proven using positive and negative introspection of issue-relevance and closure under negation for R. Theorems (iii)-(ix) are easily derivable using weak closure for K. The details of these proofs are left to the reader. A proof of theorem (x) is given below.

1. $KK\varphi$	Assumption
2. $KK\varphi \to K\varphi$	Factivity of K
3. $K\varphi$	MP, 1,2
4. $KK\varphi \to RK\varphi$	Knowledge implies issue-relevance
5. $RK\varphi$	MP, 1,4
6. $K\varphi \wedge RK\varphi$	Conjunction introduction, 3, 5
7. $KK\varphi \to (K\varphi \wedge RK\varphi)$	Discharge assumption
8. $K\varphi \to K\varphi$	Tautology
9. $(K\varphi \wedge RK\varphi) \to KK\varphi$	Weak closure for K, 8
10. $(K\varphi \wedge RK\varphi) \leftrightarrow KK\varphi$	Conjunction introduction, 7, 9
11. $((K\varphi \wedge RK\varphi) \leftrightarrow KK\varphi) \to (RK\varphi \to (K\varphi \leftrightarrow KK\varphi)$	Tautology
12. $RK\varphi \to (K\varphi \leftrightarrow KK\varphi)$	MP, 10,11
13. $RK\varphi \to (RK\varphi \to RKK\varphi)$	Strong RE, 12
14. $(RK\varphi \to (RK\varphi \to RKK\varphi)) \to (RK\varphi \to RKK\varphi)$	Tautology
15. $RK\varphi \to RKK\varphi$	MP, 13, 14

Lastly, theorem (xi) can be derived using theorem (x) and weak closure for knowledge. □

The remainder of this section is devoted to proving completeness of KR. Some additional notation and another derivable rule will turn out to be helpful.

Definition 5.9. For any unary connective $\circ$ and set of formulas Ψ, we define

$$\circ\Psi := \{\circ\psi \mid \psi \in \Psi\}.$$

For any finite set of formulas Θ we define

$$\wedge\Theta := \bigwedge_{\theta \in \Theta} \theta.$$

Proposition 5.10. *Let $\Psi \subseteq \mathcal{L}_{\mathsf{KR}}(\Phi)$ be a finite set of formulas and let Λ be the family of sets of formulas that are maximally* KR*-consistent subsets of $\Psi \cup \neg\Psi$. Then the following rule holds for any R-conjunct η:*

$$\begin{array}{l} \textit{from} \vdash_{\mathsf{KR}} (\wedge R\Psi \wedge \eta \wedge (\wedge S)) \to \pm\varphi \textit{ for all } S \in \Lambda \\ \textit{infer} \vdash_{\mathsf{KR}} (\wedge R\Psi \wedge \eta) \to R\varphi, \end{array}$$

where the formula '$\pm\varphi$' should be interpreted as 'either φ or $\neg\varphi$'.

Proof. Let Ψ and Λ be as given above and assume that we have

$$\vdash (\wedge R\Psi \wedge \eta \wedge (\wedge S)) \to \pm\varphi \text{ for all } S \in \Lambda.$$

We split the set Λ into Λ^+ and Λ^- such that:

$$\Lambda^+ := \{S \in \Lambda | \vdash (\wedge R\Psi \wedge \eta \wedge (\wedge S)) \to \varphi\};$$

$$\Lambda^- := \{S \in \Lambda | \vdash (\wedge R\Psi \wedge \eta \wedge (\wedge S)) \to \neg\varphi\}.$$

At least one of these sets must be non-empty. If Λ^+ is non-empty, then

$$\vdash (\wedge R\Psi \wedge \eta) \to \left(\left(\bigvee\nolimits_{S\in\Lambda^+}(\wedge S)\right) \leftrightarrow \varphi\right).$$

The formula $\wedge R\Psi \wedge \eta$ is an R-conjunct. Applying strong RE yields

$$\vdash (\wedge R\Psi \wedge \eta) \to \left(R\left(\bigvee\nolimits_{S\in\Lambda^+}(\wedge S)\right) \to R\varphi\right). \tag{5.1}$$

From $\wedge R\Psi$, closure under negation and closure under conjunction for R, we can derive $R\left(\bigvee_{S\in\Lambda^+}(\wedge S)\right)$. Thus we have

$$\vdash (\wedge R\Psi \wedge \eta) \to \left(R\left(\bigvee\nolimits_{S\in\Lambda^+}(\wedge S)\right)\right). \tag{5.2}$$

Putting together (5.1) and (5.2) yields

$$\vdash (\wedge R\Psi \wedge \eta) \to R\varphi.$$

In case Λ_n^- is the only non-empty set, similar steps can be applied to obtain

$$\vdash (\wedge R\Psi \wedge \eta) \to R\neg\varphi.$$

Applying $\vdash R\neg\varphi \to R\varphi$ then yields the desired result. $\square$

In essence, the rule above states that formulas whose truth value is determined by the truth values of issue-relevant formulas are also issue-relevant. In addition to this rule, we need a notion of a logic's theory.

Definition 5.11 (Theory of a logic)**.** Let L be a logic. A set of formulas Γ is an L-theory if it is L-consistent and any proper extension of Γ is L-inconsistent.

When the logic L is clear from the context, we sometimes write 'theory' instead of 'L-theory'. We also need two standard lemmas: Lindenbaum's Lemma and König's Lemma. We phrase Lindenbaum's Lemma in a more general manner, allowing us to use it in subsequent sections as well.[67]

Lemma 5.12 (Lindenbaum's Lemma)**.** *Let L be any logic based on a countable language. If Γ is an L-consistent set of formulas then there is an L-theory Γ^+ such that $\Gamma \subseteq \Gamma^+$.*

Lemma 5.13 (König's Lemma)**.** *If T is a finitely branching tree with infinitely many vertices, then T has an infinite branch.*

Next, some additional machinery is needed to restrict the canonical model. Since issue-relevance is a global matter, the theories included in a canonical model must all agree on which formulas are issue-relevant—and which are not. Constructing a canonical model in the usual manner will not work: different theories make different R-formulas true, whereas R-formulas need to hold globally in any epistemic issue model. To overcome this problem, we define an equivalence relation on theories and define semi-canonical models relative to theories to ensure that issues are global.

[67]A proof of Lindenbaum's Lemma is given by Blackburn, De Rijke and Venema (2001, p. 197) and a proof of König's Lemma is given by Lévy (1979, p. 298).

Definition 5.14 (Issue equivalence)**.** Two KR-theories Γ and Δ are issue equivalent if for all $\varphi \in \mathcal{L}_{\mathsf{KR}}(\Phi)$, $R\varphi \in \Gamma$ if, and only if, $R\varphi \in \Delta$. The issue equivalence of two KR-theories Γ and Δ is denoted as $\Gamma \equiv_i \Delta$.

Definition 5.15 (Semi-canonical models for KR)**.** Let $\mathcal{I} := \{x\}$ be the set of basic issues. The tuple $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, \mathcal{A}^{\mathsf{KR}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ is called the semi-canonical model relative to the KR-theory Γ_0, where:

- $W^{\mathsf{KR}}_{\Gamma_0} := \{\Gamma \subseteq \mathcal{L}_{\mathsf{KR}}(\Phi) \mid \Gamma \text{ a } \mathsf{KR}\text{-theory and } \Gamma_0 \equiv_i \Gamma\}$;
- $\Gamma \sim^{\mathsf{KR}}_{\Gamma_0} \Delta$ iff for all $\varphi \in \mathcal{L}_{\mathsf{KR}}(\Phi)$, $K\varphi \in \Gamma$ iff $K\varphi \in \Delta$;
- $\Gamma \approx^{\mathsf{KR}}_{\Gamma_0} \Delta$ iff for all $\varphi \in \mathcal{L}_{\mathsf{KR}}(\Phi)$, $R\varphi \in \Gamma_0$ implies $\varphi \in \Gamma$ iff $\varphi \in \Delta$;
- $\mathcal{A}^{\mathsf{KR}}_{\Gamma_0} := \{x\}$;
- $||\bullet||^{\mathsf{KR}}_{\Gamma_0} : \Phi \to \mathcal{P}(W^{\mathsf{KR}}_{\Gamma_0}) : p \mapsto \{\Gamma \mid p \in \Gamma\}$.

Generally, we use the abbreviation $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ when referring to the semi-canonical model relative to Γ_0. Furthermore, often we suppress superscripts and write, for instance, '$\mathcal{M}_{\Gamma_0}$' instead of '$\mathcal{M}^{\mathsf{KR}}_{\Gamma_0}$'.

Recall that epistemic issue models in which the set of basic issues is a singleton and the agenda is non-empty are called single issue models. Both conditions are satisfied in semi-canonical models. Moreover, it is evident that both $\sim^{\mathsf{KR}}_{\Gamma_0}$ and $\approx^{\mathsf{KR}}_{\Gamma_0}$ are equivalence relations. If we interpret $\approx^{\mathsf{KR}}_{\Gamma_0}$ as the issue relation for x, it is clear that $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0}$ is a single issue model.

Observe that if $\Gamma \equiv_i \Delta$, then $\mathcal{M}_\Gamma = \mathcal{M}_\Delta$. So every semi-canonical model corresponds to an equivalence class of $\equiv_i$. We proceed by proving existence lemmas for R and K.

Lemma 5.16 (Existence Lemma for R in KR)**.** *Let Γ_0 be a KR-theory. If $\neg R\varphi \in \Gamma_0$ then there exist $\Delta, \Delta' \in W^{\mathsf{KR}}_{\Gamma_0}$ such that $\Delta \approx^{\mathsf{KR}}_{\Gamma_0} \Delta'$, $\varphi \in \Delta$ and $\neg\varphi \in \Delta'$.*

Proof. Assume Γ_0 is a theory and that $\neg R\varphi \in \Gamma_0$. Consider the sets

$$\Psi := \{\psi \in \mathcal{L}_{\mathsf{KR}}(\Phi) \mid R\psi \in \Gamma_0\} \quad \text{and} \quad \Theta := \{\theta \in \mathcal{L}_{\mathsf{KR}}(\Phi) \mid \neg R\theta \in \Gamma_0\}.$$

The set of formulas Ψ is always countably infinite, hence we can enumerate all its formulas as $\psi_1, \ldots, \psi_n, \ldots$. The set of formulas Θ is either empty or infinite. If it is empty it can be ignored throughout the proof, so assume it is infinite and enumerate its formulas as $\theta_1, \ldots, \theta_n, \ldots$. For any $n \in \mathbb{N}$, let $\Psi_{\leq n}$ and $\Theta_{\leq n}$ denote $\{\psi_1, \ldots, \psi_n\}$ and $\{\theta_1, \ldots, \theta_n\}$, respectively. We claim that for every $n \in \mathbb{N}$, there exists some maximally consistent subset $S_n \subseteq \Psi_{\leq n} \cup \neg\Psi_{\leq n}$ such that both $S_n \cup \{\varphi\} \cup R\Psi_{\leq n} \cup \neg R\Theta_{\leq n}$ and $S_n \cup \{\neg\varphi\} \cup R\Psi_{\leq n} \cup \neg R\Theta_{\leq n}$ are consistent. Assume for reductio that for some $n \in \mathbb{N}$, the desired S_n does not exist and let Λ_n be the family of maximal consistent subsets of $\Psi_{\leq n} \cup \neg\Psi_{\leq n}$. Then φ or $\neg\varphi$ follows from $(\wedge R\Psi_{\leq n}) \wedge (\wedge\neg R\Theta_{\leq n}) \wedge (\wedge S)$ for every $S \in \Lambda_n$. Invoking the rule from Proposition 5.10 yields $\vdash ((\wedge R\Psi_{\leq n}) \wedge (\wedge\neg R\Theta_{\leq n})) \to R\varphi$. As the antecedent is contained in Γ_0, we must also have $R\varphi \in \Gamma_0$, contradicting the assumption. Therefore the desired S_n must exist for all $n \in \mathbb{N}$.

Now construct a tree consisting of all S_n satisfying the consistency constraints, for each $n \in \mathbb{N}$. This tree is infinite because Ψ is infinite and for every $n \in \mathbb{N}$ at least one S_n exists. Let the partial order on these nodes be given by set inclusion. Since every S_n can only be extended into a set S_{n+1} by adding ψ_{n+1} or $\neg\psi_{n+1}$, the tree is finitely branching. Thus by König's Lemma there exists an infinite branch consisting of sets $S_1 \subset \cdots \subset S_n \subset \ldots$. Consider such an infinite branch and define $S := \cup_{n\in\mathbb{N}} S_n$. Either $\psi_n \in S$ or $\neg\psi_n \in S$ for all $n \in \mathbb{N}$, hence extending S with another formula from Ψ must make it inconsistent. So the set of formulas S is a maximal consistent subset of Ψ. Furthermore, both $S \cup \{\varphi\} \cup R\Psi \cup \neg R\theta$ and $S \cup \{\neg\varphi\} \cup R\Psi \cup \neg R\theta$ can proven to be consistent. For the former set of formulas, suppose for reductio that $S \cup \{\varphi\} \cup R\Psi \cup \neg R\theta$ is inconsistent. Then there exist $\chi_1, \dots, \chi_n \in S \cup \{\varphi\} \cup R\Psi \cup \neg R\theta$ such that $\chi_1, \dots, \chi_n \vdash \bot$. However, this implies that there is some $m \in \mathbb{N}$ and node S_m such that $\chi_1, \dots, \chi_n \in S_m \cup \{\varphi\} \cup R\Psi_{\leq m} \cup \neg R\Theta_{\leq m}$, contradicting the earlier result. So $S \cup \{\varphi\} \cup R\Psi \cup \neg R\theta$ must be consistent. An analogous argument proves the consistency of $S \cup \{\neg\varphi\} \cup R\Psi \cup \neg R\theta$.

Finally, by Lindenbaum's Lemma, we can extend both $S \cup \{\varphi\} \cup R\Psi \cup \neg R\theta$ and $S \cup \{\neg\varphi\} \cup R\Psi \cup \neg R\theta$ into theories Δ and Δ', respectively. Because $R\Psi \cup \neg R\Theta \subseteq \Delta, \Delta'$, both are issue-equivalent to Γ_0 and therefore elements of W_{Γ_0}. These theories also agree on the truth value of all formulas in Ψ, so $\Delta \approx_{\Gamma_0} \Delta'$. Furthermore, $\varphi \in \Delta$ and $\neg\varphi \in \Delta'$. □

Lemma 5.17 (Existence Lemma for K in KR)**.** *Let Γ_0 be a KR-theory and let $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For all $\Gamma \in W^{\mathsf{KR}}_{\Gamma_0}$, if $K\varphi \notin \Gamma$ and $R\varphi \in \Gamma$, then there exists a KR-theory $\Delta \in W^{\mathsf{KR}}_{\Gamma_0}$ such that $\Gamma \sim^{\mathsf{KR}}_{\Gamma_0} \Delta$ and $\varphi \notin \Delta$.*

Proof. Let $\Gamma \in W_{\Gamma_0}$. Suppose that $K\varphi \notin \Gamma$ and $R\varphi \in \Gamma$. We show that there is a theory Δ such that $\Gamma \sim_{\Gamma_0} \Delta$ and $\varphi \notin \Delta$. Let $\psi_1, \dots, \psi_n, \dots$ enumerate all formulas ψ_i such that $K\psi_i \in \Gamma$ and let $\theta_1, \dots, \theta_n, \dots$ enumerate all formulas θ_j such that $\neg K\theta_j \in \Gamma$. For all $n \in \mathbb{N}$, we have

$$\nvdash \left(\left(\bigwedge\nolimits_{i\leq n} K\psi_i\right) \wedge \left(\bigwedge\nolimits_{j\leq n} \neg K\theta_j\right) \wedge \neg K\varphi\right) \to \varphi.$$

If this is not the case, then there exists an $n \in \mathbb{N}$ such that the formula above is a theorem. By weak closure for K it follows that

$$\vdash \left(\left(\bigwedge\nolimits_{i\leq n} K\psi_i\right) \wedge \left(\bigwedge\nolimits_{j\leq n} \neg K\theta_j\right) \wedge \neg K\varphi \wedge R\varphi\right) \to K\varphi.$$

Since the antecedent is contained in Γ, it follows that $K\varphi \in \Gamma$, contradicting the assumption. So $\{K\psi_i \mid i \in \mathbb{N}\} \cup \{\neg K\theta_i \mid i \in \mathbb{N}\} \cup \{\neg K\varphi, R\varphi, \neg\varphi\}$ is consistent and can be extended into a theory Δ. The theory Δ must be issue equivalent to Γ by positive and negative introspection of issue-relevance. Furthermore, by construction, $\Gamma \sim^{\mathsf{KR}}_{\Gamma_0} \Delta$ and $\varphi \notin \Delta$, completing the proof. □

The existence lemmas can be used to prove the Truth Lemma for KR.

Lemma 5.18 (Truth Lemma for KR)**.** *Let $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ be the semi-canonical model relative to some KR-theory Γ_0. For every $\Gamma \in W^{\mathsf{KR}}_{\Gamma_0}$ and $\varphi \in \mathcal{L}_{\mathsf{KR}}(\Phi)$: $\varphi \in \Gamma$ iff $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0}, \Gamma \vDash \varphi$.*

Proof. Let Γ_0 be a theory and let $\mathcal{M}_{\Gamma_0}$ be the semi-canonical model relative to Γ_0. Take an arbitrary $\Gamma \in W_{\Gamma_0}$. The proof is by induction on the structure of formulas.

For $\varphi := p$, the result follows immediately by definition of the valuation $||\bullet||_{\Gamma_0}$.

For $\varphi := \neg\psi$, using the induction hypothesis it follows that $\varphi \in \Gamma$ iff $\psi \notin \Gamma$ iff $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash \psi$ iff $\mathcal{M}_{\Gamma_0}, \Gamma \vDash \varphi$. The case for conjunction is similar.

For $\varphi := R\psi$, suppose $R\psi \in \Gamma$. It suffices to prove that $\{\Delta \in W_{\Gamma_0} \mid \mathcal{M}_{\Gamma_0}, \Delta \vDash \psi\}$ is closed under $\approx_{\Gamma_0}$. This set equals $\{\Delta \in W_{\Gamma_0} \mid \psi \in \Delta\}$ by the induction hypothesis, so it also suffices to show that this set is closed under $\approx_{\Gamma_0}$. Assume $\psi \in \Delta$ and $\Delta \approx_{\Gamma_0} \Delta'$ for arbitrary $\Delta, \Delta' \in W$. Since $R\psi \in \Gamma$, we also have $R\psi \in \Gamma_0$ by definition of W_{Γ_0}. By definition of $\approx_{\Gamma_0}$ it then follows that $\psi \in \Delta'$ as well. So $\{\Delta \in W_{\Gamma_0} \mid \psi \in \Delta\}$ is closed under $\approx_{\Gamma_0}$. It follows that $\mathcal{M}_{\Gamma_0}, \Gamma \vDash R\psi$. Conversely, suppose that $R\psi \notin \Gamma$ for contraposition, which means that $\neg R\psi \in \Gamma_0$. By the Existence Lemma for R, there are theories $\Delta, \Delta' \in W_{\Gamma_0}$ such that $\Delta \approx_{\Gamma_0} \Delta'$, $\psi \in \Delta$ and $\neg\psi \in \Delta'$. By the induction hypothesis it follows that $\mathcal{M}_{\Gamma_0}, \Delta \vDash \psi$ and $\mathcal{M}_{\Gamma_0}, \Delta' \vDash \neg\psi$. Thus $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash R\psi$.

For $\varphi := K\psi$, suppose $K\psi \in \Gamma$. Since knowledge implies issue-relevance it follows that $R\psi \in \Gamma$, which gives $\mathcal{M}_{\Gamma_0}, \Gamma \vDash R\psi$ by the previous clause. Thus it suffices to show that for all Δ such that $\Gamma \sim_{\Gamma_0} \Delta$, it holds that $\mathcal{M}_{\Gamma_0}, \Delta \vDash \psi$. By definition of $\sim_{\Gamma_0}$, we find that $\Gamma \sim_{\Gamma_0} \Delta$ implies $K\psi \in \Delta$. Factivity yields $\psi \in \Delta$, from which we obtain $\mathcal{M}_{\Gamma_0}, \Delta \vDash \psi$ by the induction hypothesis. Conversely, suppose that $K\psi \notin \Gamma$ for contraposition, which means that $\neg K\psi \in \Gamma$. If $R\psi \notin \Gamma$, then it follows from the clause for R that $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash R\psi$, so that $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash K\psi$ as well. If $R\psi \in \Gamma$, then there exists a theory Δ such that $\Gamma \sim_{\Gamma_0} \Delta$ and $\psi \notin \Delta$ by the Existence Lemma for K. By the induction hypothesis it follows that $\mathcal{M}_{\Gamma_0}, \Delta \vDash \neg\psi$ and consequently $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash K\psi$. $\square$

Proving completeness is straightforward given the Truth Lemma above.

Proposition 5.19 (Completeness of KR)**.** *The logic* KR *is strongly complete.*

Proof. Consider an arbitrary KR-consistent set of formulas Γ, by Lindenbaum's Lemma it can be extended into a KR-theory Γ_0. Consider the semi-canonical model $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0}$, which is a single issue model. By the Truth Lemma for KR, Γ_0 is satisfied at Γ_0 in $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0}$. Thus KR is strongly complete with respect to the subclass of single issue models. It follows that KR is strongly complete with respect to epistemic issue models. $\square$

5.3 The logic of knowledge, issue-relevance and necessity

In this section we enhance the logic KR with a universal modality, yielding the more expressive logic KRU. The proof system of KRU is slightly more neat, as some rules of KR can be rewritten as axioms in KRU.

Definition 5.20 (Language of KRU)**.** The language $\mathcal{L}_{\mathsf{KRU}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K\varphi \mid R\varphi \mid U\varphi,$$

with $p \in \Phi$.

The Boolean connectives and the modalities K and R are interpreted in the same manner as before. Moreover, K-conjuncts and R-conjuncts are defined as in Definition 5.6. The formula '$U\varphi$' means 'φ is necessary given the agent's current agenda'.

Definition 5.21 (Semantics of KRU)**.** Let $M = (W, \sim, \approx, || \bullet ||)$ be an epistemic issue model and let $w \in W$. In addition to the semantic clauses of KR, we have the clause:

$$M, w \vDash_{\mathsf{KRU}} U\varphi \text{ iff } M, v \vDash_{\mathsf{KRU}} \varphi \text{ for all } v \in W.$$

The proof system for KRU can be obtained by slightly modifying the proof system for KR.

Proof system for KRU
The proof system for KRU consists of:

- The rules and axioms of propositional logic;
- Necessitation for U: from $\vdash \varphi$ infer $\vdash U\varphi$;

- Kripke's axiom for U: $U(\varphi \to \psi) \to (U\varphi \to U\psi)$;
- Inclusion for K and R: $U\varphi \to K\varphi$ and $U\varphi \to R\varphi$;
- Positive introspection for U: $U\varphi \to UU\varphi$;
- Negative introspection for U: $\neg U\varphi \to U\neg U\varphi$.
- Weak closure for K: $U(\rho \to \varphi) \to U((\rho \wedge R\varphi) \to K\varphi)$, where ρ is a K-conjunct;
- Factivity for K: $K\varphi \to \varphi$;
- Universal RE for R: $U(U\theta \to (\varphi \leftrightarrow \psi)) \to U(U\theta \to (R\varphi \to R\psi))$;
- Closure under negation for R: $R\varphi \to R\neg\varphi$.
- Closure under conjunction for R: $(R\varphi \wedge R\psi) \to R(\varphi \wedge \psi)$;
- Knowledge implies issue-relevance: $K\varphi \to R\varphi$;
- Necessity of issue-relevance: $R\varphi \to UR\varphi$;
- Necessity of issue-irrelevance: $\neg R\varphi \to U\neg R\varphi$;

Axioms already present in the proof system for KR can be interpreted in the same manner. Inclusion for U states that everything that is necessary given the current agenda is known and issue-relevant. Put differently, the tautological proposition is known and issue-relevant. In combination with factivity for knowledge it follows that necessary truths are also true. Along with necessitation, Kripke's axiom and the introspection axioms for U, it ensures that U properly functions as a necessity operator. Universal RE for R expresses that for formulas that are necessarily equivalent given the current agenda, it is necessary given the current agenda that issue-relevance of one implies issue-relevance of the other. Lastly, necessity of issue-relevance and issue-irrelevance mirror our assumption that the agent is fully aware of the issues on her agenda. From the perspective of a rational agent, the issues on her agenda are necessary.

Proposition 5.22 (Soundness of KRU)**.** *The logic* KRU *is sound.*

Proof. Soundness follows almost immediately from Proposition 5.7. The details are left to the reader. □

We turn to completeness. To prove it, a counterpart of Proposition 5.10 is needed.

Proposition 5.23. *Let Ψ be a finite set of formulas in $\mathcal{L}_{\mathsf{KRU}}$ and let Λ be the family of sets of formulas that are maximally consistent subsets of $\Psi \cup \neg\Psi$. Then the following rule holds for any formula χ:*

$$\begin{aligned}&\textit{from } \vdash_{\mathsf{KRU}} (\wedge R\Psi \wedge U\chi \wedge (\wedge S)) \to \pm\varphi \textit{ for all } S \in \Lambda,\\ &\textit{infer } \vdash_{\mathsf{KRU}} (\wedge R\Psi \wedge U\chi) \to R\varphi.\end{aligned}$$

The formula '$\pm\varphi$' should be interpreted as 'either φ or $\neg\varphi$'.

Proof. The proof is similar to the proof of Proposition 5.10. We mention some differences and further details are left to the reader. The set Λ can be split into Λ^+ and Λ^- again, at least one of which must be non-empty. Suppose Λ^+ is non-empty, then

$$\vdash (\wedge R\Psi \wedge U\chi) \to \left(\left(\bigvee_{S\in\Lambda^+}(\wedge S)\right) \leftrightarrow \varphi\right).$$

By necessity of issue-relevance, each conjunct in the antecedent of the formula above is necessarily true, hence it is equivalent to $U(\wedge R\Psi \wedge U\chi)$. After rewriting the formula above, universal RE can be applied. Thereafter the proof is analogous to the proof of Proposition 5.10. $\square$

The completeness proof for KRU is similar to the one for KR, but it deviates at some points. For instance,the canonical model should not be restricted to issue equivalent theories, but to universally equivalent theories.

Definition 5.24 (Universal equivalence)**.** Two KRU-theories Γ and Δ are universally equivalent if for all $\varphi \in \mathcal{L}_{\mathsf{KRU}}(\Phi)$, $U\varphi \in \Gamma$ if, and only if, $U\varphi \in \Delta$. If Γ and Δ are universally equivalent, we denote this as $\Gamma \equiv_u \Delta$.

We define semi-canonical models relative to KRU-theories by restricting the set of worlds to universally equivalent theories.

Definition 5.25 (Semi-canonical models for KRU)**.** Define the semi-canonical model relative to a KRU-theory Γ_0 as $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0} = (W^{\mathsf{KRU}}_{\Gamma_0}, \sim^{\mathsf{KRU}}_{\Gamma_0}, \approx^{\mathsf{KRU}}_{\Gamma_0}, \mathcal{A}^{\mathsf{KRU}}_{\Gamma_0} || \bullet ||^{\mathsf{KRU}}_{\Gamma_0})$ in the same way as $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, \mathcal{A}^{\mathsf{KR}}_{\Gamma_0}, || \bullet ||^{\mathsf{KR}}_{\Gamma_0})$, except that

$$W^{\mathsf{KRU}}_{\Gamma_0} := \{\Gamma \subseteq \mathcal{L}_{\mathsf{KRU}}(\Phi) \mid \Gamma \text{ a } \mathsf{KRU}\text{-theory and } \Gamma_0 \equiv_u \Gamma\}.$$

As before, we sometimes use the abbreviation $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0} = (W^{\mathsf{KRU}}_{\Gamma_0}, \sim^{\mathsf{KRU}}_{\Gamma_0}, \approx^{\mathsf{KRU}}_{\Gamma_0}, || \bullet ||^{\mathsf{KRU}}_{\Gamma_0})$, often suppressing superscripts as well. When interpreting $\approx^{\mathsf{KR}}_{\Gamma_0}$ as the issue relation for the only basic issue, it is clear that $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0}$ is also a single issue model. If the notion of issue equivalence is extended to KRU-theories in the obvious way, the following proposition follows.

Proposition 5.26 (Universal equivalence implies issue equivalence)**.** *Let Γ and Δ be KRU-theories. If $\Gamma \equiv_u \Delta$, then $\Gamma \equiv_i \Delta$.*

Proof. Suppose $\Gamma \equiv_u \Delta$ for two KRU-theories Γ and Δ. Let $R\varphi \in \Gamma$ for some formula $\varphi \in \mathcal{L}_{\mathsf{KRU}}(\Phi)$, by necessity of issue-relevance $UR\varphi \in \Gamma$. By the universal equivalence of Γ and Δ it follows that $UR\varphi \in \Delta$ as well. Factivity of U then yields $R\varphi \in \Delta$, as desired. The proof of the converse is analogous. $\square$

Since $\mathcal{M}^{\mathsf{KRU}}_{\Gamma} = \mathcal{M}^{\mathsf{KRU}}_{\Delta}$ whenever $\Gamma \equiv_u \Delta$, every semi-canonical model for KRU corresponds to an equivalence class of $\equiv_u$. The proposition above then shows that the semi-canonical models for KRU are more fine-grained than the semi-canonical models for KR.

We formulate alternative existence lemmas for R, K and U. The proofs from the previous section can be rehashed for the cases of R and K.

Lemma 5.27 (Existence Lemma for R in KRU)**.** *Let Γ_0 be a KRU-theory. If $\neg R\varphi \in \Gamma_0$ then there exist $\Delta, \Delta' \in W^{\mathsf{KRU}}_{\Gamma_0}$ such that $\Delta \approx^{\mathsf{KRU}}_{\Gamma_0} \Delta'$, $\varphi \in \Delta$ and $\neg\varphi \in \Delta'$.*

Proof. This proof is similar to the proof of Lemma 5.16. We point out key differences and leave the details to the reader. Define the set Ψ as before, but Θ as

$$\Theta := \{\theta \in \mathcal{L}_{\mathsf{KRU}}(\Phi) \mid U\theta \in \Gamma_0\}.$$

Set up the reductio proof for the existence of S_n in the same manner. Observe that a conjunction of universal formulas χ is equivalent to $U\chi$ because U is positively introspective and a normal modality, so that Proposition 5.23 can be applied, proving the existence of S_n. Thereafter a tree similar to the one in the proof of Lemma 5.16 can be constructed and the rest of the proof can be carried out in a similar manner. $\square$

Lemma 5.28 (Existence Lemma for K in KRU)**.** *Let Γ_0 be a KRU-theory and let $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0} = (W^{\mathsf{KRU}}_{\Gamma_0}, \sim^{\mathsf{KRU}}_{\Gamma_0}, \approx^{\mathsf{KRU}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For all $\Gamma \in W^{\mathsf{KRU}}_{\Gamma_0}$, if $K\varphi \notin \Gamma$ and $R\varphi \in \Gamma$, then there exists a KRU-theory $\Delta \in W^{\mathsf{KRU}}_{\Gamma_0}$ such that $\Gamma \sim^{\mathsf{KRU}}_{\Gamma_0} \Delta$ and $\varphi \notin \Delta$.*

Proof. The proof is entirely analogous to the proof of Lemma 5.17. □

Lemma 5.29 (Existence Lemma for U in KRU)**.** *Let Γ_0 be a KRU-theory and let $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0} = (W^{\mathsf{KRU}}_{\Gamma_0}, \sim^{\mathsf{KRU}}_{\Gamma_0}, \approx^{\mathsf{KRU}}_{\Gamma_0}, ||\bullet||^{\mathsf{KR}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For all $\Gamma \in W^{\mathsf{KRU}}_{\Gamma_0}$, if $U\varphi \notin \Gamma$, then there exists a KRU-theory $\Delta \in W^{\mathsf{KRU}}_{\Gamma_0}$ such that $\varphi \notin \Delta$.*

Proof. The proof is similar to the one for standard modal logic. Suppose $U\varphi \notin \Gamma$, we construct a world Δ such that $\Gamma \equiv_u \Delta$ and $\varphi \notin \Delta$. Define $\Delta^- := \{\neg\varphi\} \cup \{\psi \mid U\psi \in \Gamma\}$. If Δ^- is consistent, then by Lindenbaum's Lemma it can be extended into a theory Δ. The theory Δ is universally equivalent to Γ by positive and negative introspection of U, hence $\Delta \in W_{\Gamma_0}$. Furthermore, since $\neg\varphi \in \Delta^-$, we also have $\varphi \notin \Delta$.

It remains to be proven that Δ^- is indeed consistent. Suppose it is not for reductio, then there must be $\psi_1, \dots, \psi_n \in \{\psi \mid U\psi \in \Gamma\}$ such that $\vdash (\psi_1 \wedge \cdots \wedge \psi_n) \to \varphi$. By necessitation for U we obtain $\vdash U((\psi_1 \wedge \cdots \wedge \psi_n) \to \varphi)$. Application of Kripke's axiom then yields $\vdash U(\psi_1 \wedge \cdots \wedge \psi_n) \to U\varphi$. Since U is a normal modality, it follows that $\vdash (U\psi_1 \wedge \cdots \wedge U\psi_n) \to U(\psi_1 \wedge \cdots \wedge \psi_n)$.[68] As $U\psi_1, \dots, U\psi_n \in \Gamma$, we have $U(\psi_1 \wedge \cdots \wedge \psi_n) \in \Gamma$. Applying modus ponens twice yields $U\varphi \in \Gamma$, which is a contradiction. Thus Δ^- is consistent, completing the proof. □

The existence lemmas can be used to prove the Truth Lemma for KRU.

Lemma 5.30 (Truth Lemma for KRU)**.** *Let Γ_0 be a KRU-theory and let $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0} = (W^{\mathsf{KRU}}_{\Gamma_0}, \sim^{\mathsf{KRU}}_{\Gamma_0}, \approx^{\mathsf{KRU}}_{\Gamma_0}, ||\bullet||^{\mathsf{KRU}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For every $\Gamma \in W^{\mathsf{KRU}}_{\Gamma_0}$ and $\varphi \in \mathcal{L}_{\mathsf{KRU}}(\Phi)$: $\varphi \in \Gamma$ iff $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0}, \Gamma \vDash \varphi$.*

Proof. Let Γ_0 be a theory and let $\mathcal{M}_{\Gamma_0}$ be the semi-canonical model relative to Γ_0. Take an arbitrary $\Gamma \in W_{\Gamma_0}$. The proof is by induction on the structure of formulas.

The Boolean cases and the cases for R and K are similar to the ones in the proof of Lemma 5.18.

For $\varphi := U\psi$, suppose $U\psi \in \Gamma$. By construction of $\mathcal{M}_{\Gamma_0}$ every $\Delta \in W_{\Gamma_0}$ contains $U\psi$. Thus by factivity for U it follows that $\psi \in \Delta$ for all $\Delta \in W_{\Gamma_0}$. The induction hypothesis then gives us $\mathcal{M}_{\Gamma_0}, \Delta \vDash \psi$ for all $\Delta \in W_{\Gamma_0}$, thus $M_{\Gamma_0}, \Gamma \vDash U\psi$. Conversely, suppose that $U\psi \notin \Gamma$ for contraposition. By the Existence Lemma for U, there is a $\Delta \in W_{\Gamma_0}$ such that $\psi \notin \Delta$. By the induction hypothesis it follows that $M_{\Gamma_0}, \Delta \nvDash \psi$. Thus ψ is not necessarily true: $M_{\Gamma_0}, \Gamma \nvDash U\psi$. □

Proving completeness is straightforward given the Truth Lemma above.

Proposition 5.31 (Completeness of KRU)**.** *The logic KRU is strongly complete.*

Proof. Consider an arbitrary KRU-consistent set of formulas Γ, by Lindenbaum's Lemma it can be extended into a KRU-theory Γ_0. Consider the semi-canonical model $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0}$, which is a single issue model. By the Truth Lemma for KRU, Γ_0 is satisfied at Γ_0 in $\mathcal{M}^{\mathsf{KRU}}_{\Gamma_0}$. Thus KRU is strongly complete with respect to the subclass of single issue models. It follows that KRU is strongly complete with respect to epistemic issue models. □

[68] This is a property of normal modalities. A proof is given by Blackburn, De Rijke and Venema (2001, p. 35, Example 1.40).

5.4 The logic of information and issue-relevance

There are two conditions for knowledge of a proposition: the agent must possess the information that the proposition is true and the proposition must be issue-relevant. Thus the language $\mathcal{L}_{\mathsf{KR}}$ can also be modified by replacing the knowledge modality K with an information modality I. This results in the logic IR, in which information and issue-relevance are the primary notions.

Definition 5.32 (Language $\mathcal{L}_{\mathsf{IR}}(\Phi)$)**.** The language $\mathcal{L}_{\mathsf{IR}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid I\varphi \mid R\varphi,$$

with $p \in \Phi$. We also have the following abbreviation: $K\varphi := I\varphi \wedge R\varphi$.

The Boolean connectives and the modality R is interpreted in the same manner as before. The formula '$I\varphi$' should be interpreted as 'the agent possesses the information that φ'. The knowledge modality K is no longer a primitive modality, but it is an abbreviation for possessed information that is issue-relevant. So '$K\varphi$' can still be interpreted as 'the agent knows φ'.

Definition 5.33 (Semantics of IR)**.** Let $M = (W, \sim, \approx, || \bullet ||)$ be an epistemic issue model and let $w \in W$. In addition to the semantic clauses given in the preceding sections, we have the clause:

$$M, w \vDash I\varphi \text{ iff } M, v \vDash \varphi \text{ for all } v \in W \text{ such that } w \sim v.$$

So $I\varphi$ is the Kripke modality over the information relation, in agreement with Definition 4.11. Also observe that the truth conditions of a formula $K\varphi$ are the same in IR as in KR. The proof system for IR is given below.

Proof system for IR
The proof system for IR consists of:

- The rules and axioms of propositional logic;
- Necessitation for I and R: from $\vdash \varphi$ infer $\vdash I\varphi$ and $\vdash R\varphi$;
- Kripke's axiom for I: $I(\varphi \to \psi) \to (I\varphi \to I\psi)$;
- Factivity for I: $I\varphi \to \varphi$;
- Positive introspection for I: $I\varphi \to II\varphi$;
- Negative introspection for I: $\neg I\varphi \to I\neg I\varphi$;
- Strong RE for R: from $\vdash \eta \to (\varphi \leftrightarrow \psi)$ infer $\vdash \eta \to (R\varphi \to R\psi)$, where η is an R-conjunct;
- Closure under negation for R: $R\varphi \to R\neg\varphi$;
- Closure under conjunction for R: $(R\varphi \wedge R\psi) \to R(\varphi \wedge \psi)$;
- Information about issue-relevance: $R\varphi \to IR\varphi$;
- Information about issue-irrelevance: $\neg R\varphi \to I\neg R\varphi$.

Axioms already present in the proof system for KR can be interpreted in the same manner. Necessitation for I states that for each theorem the agent possesses the information that it holds. This is trivially so, because theorems do not carry information. Kripke's axiom, factivity and the introspection axioms for I ensure that information possession works in a coherent manner, in line with the discussion in Section 3.1. Information about issue-relevance and issue-irrelevance expresses that the information about an agent's issue is in her possession, reflecting our assumption that rational agents have complete access to the issues on their agenda.

Proposition 5.34 (Soundness of IR)**.** *The logic* IR *is sound.*

Proof. The axioms for I are standard axioms for Kripke modalities corresponding to equivalence relations. The other axioms follow directly from Proposition 5.7 or immediately after application of the trivial theorem $K\varphi \to I\varphi$. $\square$

Again, we will construct semi-canonical models relative to theories, including only those theories that are issue equivalent to ensure that issue-relevance is universal.

Definition 5.35 (Semi-canonical models for IR)**.** Define the semi-canonical model relative to a IR-theory Γ_0 as $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0} = (W^{\mathsf{IR}}_{\Gamma_0}, \sim^{\mathsf{IR}}_{\Gamma_0}, \approx^{\mathsf{IR}}_{\Gamma_0}, \mathcal{A}^{\mathsf{IR}}_{\Gamma_0} || \bullet ||^{\mathsf{IR}}_{\Gamma_0})$ in the same way as $\mathcal{M}^{\mathsf{KR}}_{\Gamma_0} = (W^{\mathsf{KR}}_{\Gamma_0}, \sim^{\mathsf{KR}}_{\Gamma_0}, \approx^{\mathsf{KR}}_{\Gamma_0}, \mathcal{A}^{\mathsf{KR}}_{\Gamma_0}, || \bullet ||^{\mathsf{KR}}_{\Gamma_0})$, except that

$$\Gamma \sim^{\mathsf{IR}}_{\Gamma_0} \Delta \text{ iff for all } \varphi \in \mathcal{L}_{\mathsf{IR}}(\Phi),\ I\varphi \in \Gamma \text{ iff } I\varphi \in \Delta.$$

As before, we sometimes use the abbreviation $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0} = (W^{\mathsf{IR}}_{\Gamma_0}, \sim^{\mathsf{IR}}_{\Gamma_0}, \approx^{\mathsf{IR}}_{\Gamma_0}, || \bullet ||^{\mathsf{IR}}_{\Gamma_0})$, often suppressing superscripts as well. It is evident that $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0}$ is a single issue model if we interpret $\approx^{\mathsf{IR}}_{\Gamma_0}$ as the issue relation for the only basic issue.

Lemma 5.36 (Existence Lemma for I in IR)**.** *Let Γ_0 be an* IR*-theory and let $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0} = (W^{\mathsf{IR}}_{\Gamma_0}, \sim^{\mathsf{IR}}_{\Gamma_0}, \approx^{\mathsf{IR}}_{\Gamma_0}, || \bullet ||^{\mathsf{IR}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For all $\Gamma \in W^{\mathsf{IR}}_{\Gamma_0}$, if $I\varphi \notin \Gamma$, then there exists an* IR*-theory $\Delta \in W^{\mathsf{IR}}_{\Gamma_0}$ such that $\Gamma \sim^{\mathsf{IR}}_{\Gamma_0} \Delta$ and $\varphi \notin \Delta$.*

Lemma 5.37 (Existence Lemma for R in IR)**.** *Let Γ_0 be an* IR*-theory. If $\neg R\varphi \in \Gamma_0$ then there exist $\Delta, \Delta' \in W^{\mathsf{IR}}_{\Gamma_0}$ such that $\Delta \approx^{\mathsf{IR}}_{\Gamma_0} \Delta'$, $\varphi \in \Delta$ and $\neg\varphi \in \Delta'$.*

The proof of the Existence Lemma for I in IR is standard and analogous to the proof of the Existence Lemma for U in KRU (Lemma 5.29). The proof for the Existence Lemma for R in IR is entirely analogous to the proof of its counterpart in KR (Lemma 5.16). We proceed by proving the Truth Lemma for IR.

Proposition 5.38 (Truth Lemma for IR)**.** *Let Γ_0 be an* IR*-theory and let $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0} = (W^{\mathsf{IR}}_{\Gamma_0}, \sim^{\mathsf{IR}}_{\Gamma_0}, \approx^{\mathsf{IR}}_{\Gamma_0}, || \bullet ||^{\mathsf{IR}}_{\Gamma_0})$ be the semi-canonical model relative to Γ_0. For every $\Gamma \in W^{\mathsf{IR}}_{\Gamma_0}$ and $\varphi \in \mathcal{L}_{\mathsf{IR}}(\Phi)$: $\varphi \in \Gamma$ iff $\mathcal{M}^{\mathsf{IR}}_{\Gamma_0}, \Gamma \vDash \varphi$.*

Proof. Let Γ_0 be a theory and let $\mathcal{M}_{\Gamma_0}$ be the semi-canonical model relative to Γ_0. Take an arbitrary $\Gamma \in W_{\Gamma_0}$. The proof is by induction on the structure of formulas.

The Boolean cases and the case for R are similar to the ones in the proof of Lemma 5.18.

For $\varphi := I\psi$, suppose $I\varphi \in \Gamma$. By construction of M_{Γ_0} every Δ with $\Gamma \sim \Delta$ contains $I\varphi$. By factivity for I it follows that $\varphi \in \Delta$ for all these theories. The induction hypothesis then gives us $M_{\Gamma_0}, \Delta \vDash \varphi$, so that $M_{\Gamma_0}, \Gamma \vDash I\varphi$. Conversely, suppose that $I\varphi \notin \Gamma$ for contraposition. By the Existence Lemma for I, there is a theory Δ such that $\Gamma \sim \Delta$ and $\varphi \notin \Delta$. By the induction hypothesis it follows that $M_{\Gamma_0}, \Delta \nvDash \varphi$. Thus the information that φ is not possessed at Γ: $M_{\Gamma_0}, \Gamma \nvDash I\varphi$. $\square$

Completeness follows from the results established above. The reasoning is analogous to the completeness proofs of KR and KRU in the previous sections.

Proposition 5.39 (Completeness of IR)**.** *The logic* IR *is strongly complete.*

5.5 The logic of information, issue-relevance and necessity

The alterations made to KR in the two previous sections can also be combined: a universal modality can be added while replacing the knowledge modality K with the information modality I. We call the resulting logic IRU.

Definition 5.40 (Language $\mathcal{L}_{\mathsf{IRU}}(\Phi)$)**.** The language $\mathcal{L}_{\mathsf{IRU}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid I\varphi \mid R\varphi \mid U\varphi,$$

with $p \in \Phi$. We also have the following abbreviation: $K\varphi := I\varphi \wedge R\varphi$.

The Boolean connectives and modalities are interpreted in the same manner as in the preceding sections. Moreover, the semantic clauses are the same as in previous sections. The proof system for IRU is given below.

Proof system for IRU
The proof system for IRU consists of:

- The rules and axioms of propositional logic;
- Necessitation for U: from $\vdash \varphi$ infer $\vdash U\varphi$;
- Kripke's axiom for I: $I(\varphi \to \psi) \to (I\varphi \to I\psi)$;
- Inclusion for I and R: $U\varphi \to I\varphi$ and $U\varphi \to R\varphi$;
- Factivity for I: $I\varphi \to \varphi$;
- Positive introspection for U and I: $U\varphi \to UU\varphi$ and $I\varphi \to II\varphi$;
- Negative introspection for U and I: $\neg U\varphi \to U\neg U\varphi$ and $\neg I\varphi \to I\neg I\varphi$;
- Universal RE for R: $U(U\theta \to (\varphi \leftrightarrow \psi)) \to U(U\theta \to (R\varphi \to R\psi))$;
- Closure under negation for R: $R\varphi \to R\neg\varphi$;
- Closure under conjunction for R: $(R\varphi \wedge R\psi) \to R(\varphi \wedge \psi)$;
- Necessity of issue-relevance: $R\varphi \to UR\varphi$;
- Necessity of issue-irrelevance: $\neg R\varphi \to U\neg R\varphi$.

All of these rules and axioms can be interpreted in the same manner as in the previous sections. Soundness and completeness follow from the soundness and completeness of KRU and IR.

Proposition 5.41 (Soundness of IRU)**.** *The logic* IRU *is sound.*

Proof. Note that every axiom of IRU is also an axiom of either IR or KRU, both of which are sound. □

We sketch the completeness proof for IRU. Define the semi-canonical model $\mathcal{M}^{\mathsf{IRU}}_{\Gamma_0}$ relative to an IRU-theory Γ_0 by restricting the set of IRU-theories to those that are universally equivalent to Γ_0. Define the relation $\approx^{\mathsf{IRU}}_{\Gamma_0}$ in the same manner as $\approx^{\mathsf{IR}}_{\Gamma_0}$. The proofs for the existence lemmas and the Truth Lemma for IRU are analogous to the proofs in the previous sections. It can then be shown that every IRU-theory Γ_0 can be satisfied in the semi-canonical $\mathcal{M}^{\mathsf{IRU}}_{\Gamma_0}$, proving strong completeness for IRU.

Proposition 5.42 (Completeness of IRU)**.** *The logic* IRU *is strongly complete.*

5.6 The logic of information, issues and necessity

Because R is not a normal modality, all the logics treated thus far are not normal. However, the logic IRU can be augmented to obtain a normal modal logic IQU that captures both knowledge and issue-relevance. The modality R is replaced by an abstract issue modality Q, which is the Kripke modality over the current issue relation.

Definition 5.43 (Language of IQU)**.** The language $\mathcal{L}_{\mathsf{IQU}}(\Phi)$, with Φ a set of atomic propositions, is given in Backus–Naur form by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid I\varphi \mid Q\varphi \mid U\varphi,$$

with $p \in \Phi$. We also have the following abbreviations:

- $R\varphi := U(\varphi \to Q\varphi)$;
- $K\varphi := I\varphi \wedge R\varphi$.

The Boolean cases are interpreted in the same manner as before, and so are the modalities I and U. The formula '$Q\varphi$' can be interpreted as 'φ holds in all issue-equivalent worlds with respect to the agent's current issue'. It only considers the issue relation locally. We have argued that issues can only be sensibly interpreted on a global level, hence the modality Q is not much more than a convenient technical tool. In tandem with the universal modality, however, it can be used to express issue-relevance, as was shown in Proposition 4.14.

The semantics of IQU as well as a simple proof system are provided below.

Definition 5.44 (Semantics of IQU)**.** Let $M = (W, \sim, \approx, || \bullet ||)$ be an epistemic issue model and let $w \in W$. In addition to the semantic clauses given in the preceding sections, we have the clause:

$$M, w \vDash Q\varphi \text{ iff } M, v \vDash \varphi \text{ for all } v \in W \text{ such that } w \approx v.$$

Proof system for IQU
The proof system for IQU consists of:

- The rules and axioms of propositional logic;
- Necessitation for U: from $\vdash \varphi$ infer $\vdash U\varphi$;
- Kripke's axiom for I, Q and U: $I(\varphi \to \psi) \to (I\varphi \to I\psi)$, $Q(\varphi \to \psi) \to (Q\varphi \to Q\psi)$ and $U(\varphi \to \psi) \to (U\varphi \to U\psi)$;
- Inclusion for I and Q: $U\varphi \to I\varphi$ and $U\varphi \to Q\varphi$;
- Factivity for I, Q and U: $I\varphi \to \varphi$, $Q\varphi \to \varphi$ and $U\varphi \to \varphi$;
- Positive introspection for I, Q and U: $I\varphi \to II\varphi$, $Q\varphi \to QQ\varphi$ and $U\varphi \to UU\varphi$;
- Negative introspection for I, Q and U: $\neg I\varphi \to I\neg I\varphi$, $\neg Q\varphi \to Q\neg Q\varphi$ and $\neg U\varphi \to U\neg U\varphi$.

Necessitation for I and Q follow from necessitation for U combined with inclusion. Since we also have Kripke's axiom for I, Q and U, this logic is normal. In fact, all modalities are **S5** modalities corresponding to Kripke modalities over equivalence relations, with the U modality corresponding to the Kripke modality over the universal relation because of inclusion. Consequently the proofs for soundness and completeness are entirely standard.[69]

Proposition 5.45 (Soundness and completeness of IQU)**.** *The logic* IQU *is sound and strongly complete.*

So there is a normal logic that is sound and complete that neatly captures the notion of knowledge as issue-relevant information.

5.7 Synthesis

Every logic treated in this chapter can be considered a fragment of IQU. Figure 5.1 displays the relations between these logics. Their semantics make it evident that for any pair of these logics L and L' such that $\mathcal{L}_{\mathsf{L}}(\Phi) \subseteq \mathcal{L}_{\mathsf{L}'}(\Phi)$, we have $\Gamma \models_{\mathsf{L}} \varphi$ if, and only if, $\Gamma \models_{\mathsf{L}'} \varphi$, for any $\varphi \in \mathcal{L}_{\mathsf{L}}(\Phi)$ and $\Gamma \subseteq \mathcal{L}_{\mathsf{L}}(\Phi)$. Since each logic is sound and strongly complete it follows that $\Gamma \vdash_{\mathsf{L}} \varphi$ if, and only if, $\Gamma \vdash_{\mathsf{L}'} \varphi$, for any $\varphi \in \mathcal{L}_{\mathsf{L}}(\Phi)$ and $\Gamma \subseteq \mathcal{L}_{\mathsf{L}}(\Phi)$.

$$
\begin{tikzcd}
 & \mathsf{KRU} \arrow[dr] & & \\
\mathsf{KR} \arrow[ur] \arrow[dr] & & \mathsf{IRU} \arrow[r] & \mathsf{IQU} \\
 & \mathsf{IR} \arrow[ur] & &
\end{tikzcd}
$$

Figure 5.1: A graph showing the relations between the different logics studied in this chapter. An arrow from a logic L to a logic L' indicates that L is a sublogic of L'. Transitive arrows have been omitted.

Although IQU is the least interesting logic that we studied here from both a conceptual and a technical point of view, it provides a fertile starting point for a dynamic logic of epistemic issues that can capture agenda updates and issue resolution. Naturally, the resolution of an issue refines the information relation, while agenda updates only affect the issue relation. The modalities in IQU are all Kripke modalities over a single relation, thus keeping the information relation and the issue relation separated. As a consequence, the operators of IQU are suitable for pre-encoding both types of dynamics. In the next chapter it is shown how IQU can be augmented in order to obtain a sound and complete logic that takes basic issues into account and accommodates the actions of agenda update and issue resolution.

The chapter is concluded with an example showing the logics of this chapter in action. We revisit an example from Section 2.2 again.

Example 10 (The King of England: Revisited again)**.** Let $\mathcal{I} = \{x\}$ and let $\Phi = \{p, q\}$. Let x be the issue concerning the question whether war can be avoided. Interpret 'p' as 'war can be avoided' and q as 'nuclear war can be avoided'. We define the single issue model $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$, where:

- $W := \{w, v, u\}$;
- $\sim := \approx_x := \{(v, u), (u, v), (w, w), (v, v), (u, u)\}$;
- $\mathcal{A} := \{x\}$;

[69]These standard proofs are given by Blackburn, De Rijke and Venema (2001, §4.1-3) and Van Ditmarsch, Van der Hoek and Kooi (2008, §7.2), among others.

- $||p|| := \{w\}$ and $||q|| := \{w, v\}$.

Let w denote the actual world. The model is depicted in Figure 5.2 (a), which is similar to the structure in Figure 4.3. In accordance with earlier findings, $M, w \vDash_{\mathsf{KR}} Kp$ and $M, w \vDash_{\mathsf{KR}} \neg Kq$: the King knows that war can be avoided, but he does not know that nuclear war can be avoided. In the more expressive logic KRU, it can be stated that it is necessary that nuclear war can be avoided if war can be avoided; $M, w \vDash U(p \to q)$. Hence the interpretation of the proposition letters makes sense. In IR, which is also more expressive than KR, it can be stated that possessing the information that nuclear war can be avoided is issue-relevant: $M, w \vDash_{\mathsf{IR}} RIq$. Consequently, since information possession is introspective and $M, w \vDash_{\mathsf{IR}} Iq$, it follows that $M, w \vDash_{\mathsf{IR}} KIq$. So the King knows that he possesses the information that nuclear war can be avoided, but he does not know that nuclear war can be avoided. Despite being counterintuitive, this is sensible in our framework. Information may be obtained without it being processed into knowledge, hence an agent may know she possesses information without knowing the content of that information because it is issue-irrelevant. The manner in which logical space is constructed determines whether this occurs.

A slight modification of the model M leads to different results. Consider the single issue model $M' = (W', \sim', \approx'_{x \in \mathcal{I}}, \mathcal{A}', || \bullet ||')$, where:

- $W' := \{w, v, u, v', u'\}$;
- $\sim := \{(v, u), (u, v), (w, w), (v, v), (u, u), (v', v'), (u', u')\}$;
- $\approx_x := \{(v, u), (u, v), (v', u'), (u', v'), (w, w), (v, v), (u, u), (v', v'), (u', u')\}$;
- $\mathcal{A} := \{x\}$;
- $||p||' := \{w\}$ and $||q||' := \{w, v, v'\}$.

The model is depicted in Figure 5.2 (b). The information that nuclear war can be avoided is no longer issue-relevant, because copies of worlds v and u have been added that are informationally distinguishable. So the King no longer knows he possesses the information that nuclear war can be avoided; $M', w \nvDash_{\mathsf{IR}} KIq$. Furthermore, the same formulas of KR are satisfied at w in M and M' and $M', w \vDash_{\mathsf{KRU}} U(p \to q)$ still holds. So the more expressive logics, like IR, can indeed distinguish between models that a less expressive logic like KR cannot distinguish between.

Lastly, in both models it can be seen that the information that nuclear war can be avoided is carried by the information fully answering the King's question whether war can be avoided: $M, w \vDash_{\mathsf{IQU}} Qq$ and $M', w \vDash_{\mathsf{IQU}} Qq$.

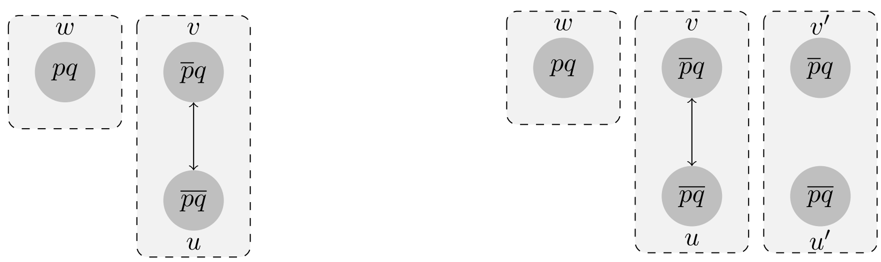

(a) $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$ (b) $M' = (W', \sim', \approx'_{x \in \mathcal{I}}, \mathcal{A}', || \bullet ||')$

Figure 5.2: Two epistemic issue models M and M'. The actual world is denoted by w in both models.

Chapter 6

Dynamic Epistemic Issue Logic

In this chapter, we introduce a dynamic logic of epistemic issues and prove it to be sound, complete and decidable. This dynamic logic takes basic issues into account and accommodates the actions of issue addition, issue retraction and issue resolution. In Section 6.1, we introduce Epistemic Issue Logic, a static logic which we prove to be sound. In Section 6.2, it is proven that that Epistemic Issue Logic is complete. In Section 6.3, Dynamic Epistemic Issue Logic is introduced as an extension of Epistemic Issue Logic and we show that it is both sound and complete. In Section 6.4, we prove that all logics treated in Chapter 5 and the current chapter are decidable. In Section 6.5, we reflect on the findings of this chapter and show Dynamic Epistemic Issue Logic in action, revisiting one of the earlier examples again. As in the previous chapter, throughout this chapter subscripts and superscripts are often suppressed when the context allows it.

6.1 Epistemic Issue Logic

Recall that every $x \in \mathcal{I}$ is a basic issue and that every $X \subseteq \mathcal{I}$ is a compound issue. Also recall that the resolution of an issue $X \subseteq \mathcal{I}$ comes down to providing the agent with exactly sufficient information to informationally distinguish between worlds in different issue cells of the issue X. When interpreting issues as questions, the resolution of an issue X corresponds to receiving the complete answer to X. When interpreting issues as subject matters, the resolution of X corresponds to receiving all information about X.

A logic of epistemic issues should be able to express properties about basic issues. For instance, whether they are on an agent's agenda or which formulas they make issue-relevant. To this end, we introduce the following language for Epistemic Issue Logic (EIL).

Definition 6.1 (Language of EIL)**.** The static language $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ with Φ a set of atomic propositions and $\mathcal{I}$ a finite set of basic issues, is given in Backus–Naur form by:

$$\varphi ::= p \mid Ax \mid \neg\varphi \mid \varphi \wedge \varphi \mid I^X\varphi \mid Q^X\varphi,$$

with $p \in \Phi$, $x \in \mathcal{I}$ and $X \subseteq \mathcal{I}$.

The Boolean formulas are interpreted as usual. The atom 'Ax' is interpreted as 'x is on the agenda'. The modality '$I^X\varphi$' is interpreted as 'the agent possesses the information that φ conditional on resolving X'. The modality '$Q^X\varphi$' is interpreted as 'φ holds in all worlds that are issue-equivalent with respect to the issue X'. The modality I^X pre-encodes the action of resolving an issue, while the atoms Ax are sufficient to pre-encode agenda updates. The modality Q^X does not pre-encode any dynamics: it is simply the issue modality corresponding to the issue X.

Definition 6.2 (Abbreviations in $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$)**.** We have the following abbreviations:

- $A(X) := \left(\bigwedge_{x \in X} Ax\right) \wedge \left(\bigwedge_{x \in \mathcal{I} \setminus X} \neg Ax\right)$ for $X \subseteq \mathcal{I}$;
- $I\varphi := I^{\emptyset}\varphi$;
- $Q\varphi := \bigwedge_{X \subseteq \mathcal{I}} \left(A(X) \to Q^X \varphi\right)$;
- $U\varphi := Q^{\emptyset}\varphi$;
- $R\varphi := U(\varphi \to Q\varphi)$;
- $K\varphi := I\varphi \wedge R\varphi$.

The abbreviation '$A(X)$' expresses 'X is the agent's current agenda': the basic issues comprising X are exactly the basic issues on the agent's agenda. Note that $A(X)$ is always well-defined because $\mathcal{I}$ is required to be finite. The abbreviations '$I\varphi$', '$Q\varphi$', '$U\varphi$', '$R\varphi$' and '$K\varphi$' should be interpreted as in the previous chapter. This is reasonable, since $I^{\emptyset}\varphi$ should be interpreted as 'the agent possesses the information that φ conditional on resolving nothing' and the issue modality '$Q\varphi$' expresses that 'φ holds in all issue-equivalent worlds with respect to the agent's current issue'. The formula '$Q^{\emptyset}\varphi$' expresses 'φ holds in all worlds that are issue-equivalent with respect to the empty issue'. Since no conceptual distinctions need to be made to resolve the empty issue, it only applies to formulas that are necessary given the current agenda. So using '$U\varphi$' as an abbreviation of '$Q^{\emptyset}\varphi$' is conceptually sound. Lastly, issue-relevance and knowledge are the same abbreviations as in the logic IQU, which we studied in Section 5.6. Given the abbreviations above, it follows that $\mathcal{L}_{\mathsf{IQU}}(\Phi) \subseteq \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$.

The semantics for EIL are provided below.

Definition 6.3 (Semantics of EIL)**.** Let $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, ||\bullet||)$ be an epistemic issue model and $w \in W$. In addition to the semantic clauses for propositional letters and Boolean connectives from the previous chapter, we have the following semantic clauses:

- $M, w \vDash_{\mathsf{EIL}} Ax$ iff $x \in \mathcal{A}$;
- $M, w \vDash_{\mathsf{EIL}} I^X\varphi$ iff $M, v \vDash \varphi$ for all $v \in W$ such that $w \sim v$ and $w \approx_X v$;
- $M, w \vDash_{\mathsf{EIL}} Q^X\varphi$ iff $M, v \vDash \varphi$ for all $v \in W$ such that $w \approx_X v$.

Observe that the clause for $I^X\varphi$ can also be written as: $M, v \vDash \varphi$ for all $v \in W$ such that $w(\sim \cap \approx_X)v$. The modalities I, Q and U can be seen to have the same truth conditions in EIL as in IQU. This is obvious for I, and U. For Q, observe that it is the Kripke modality over $\approx_{\mathcal{A}}$, which equals $\approx$ when abbreviating an epistemic issue model as $(W, \sim, \approx, ||\bullet||)$. So EIL properly extends IQU. A proof system for EIL is given below.

Proof system for EIL
A deductive system for EIL consists of:

- The rules and axioms of propositional logic;
- Necessitation for $Q^{\emptyset}$: from $\vdash \varphi$ infer $\vdash Q^{\emptyset}\varphi$;
- Kripke's axiom for I^X and Q^X: $I^X(\varphi \to \psi) \to (I^X\varphi \to I^X\psi)$ and $Q^X(\varphi \to \psi) \to (Q^X\varphi \to Q^X\psi)$;
- Factivity for I^X and Q^X: $I^X\varphi \to \varphi$ and $Q^X\varphi \to \varphi$;
- Positive introspection for I^X and Q^X: $I^X\varphi \to I^X I^X\varphi$ and $Q^X\varphi \to Q^X Q^X\varphi$;

- Negative introspection for I^X and Q^X: $\neg I^X\varphi \to I^X\neg I^X\varphi$ and $\neg Q^X\varphi \to Q^X\neg Q^X\varphi$;
- Resolution: $Q^X\varphi \to I^X\varphi$;
- Necessity of current agenda: $Ax \to Q^{\emptyset}Ax$;
- Monotonicity of conditional information: $I^X\varphi \to I^Y\varphi$ for all $X, Y \subseteq \mathcal{I}$ such that $X \subseteq Y$;
- Monotonicity of issues: $Q^X\varphi \to Q^Y\varphi$ for all $X, Y \subseteq \mathcal{I}$ such that $X \subseteq Y$.

Necessitation for $Q^{\emptyset}$ states that theorems are necessarily true given the current agenda. Kripke's axiom, factivity, positive introspection and negative introspection for I^X ensure that information possession behaves in a coherent manner, even conditional on the resolution of an issue. Likewise, their counterparts for Q^X ensure that information encoded in the issue cells of X behaves as expected. If the information resolving an issue X carries the information that φ, then the information that φ is possessed conditional on resolving X: this is expressed by the resolution axiom. The necessity of current agenda axiom expresses that every basic issue on the agent's agenda is so necessarily given her current agenda. It ensures that the agenda of an agent is global. It reflects our assumption that agents are aware and in full control of their agendas. The monotonicity of conditional information expresses that information possessed conditional on resolving an issue is also possessed conditional on resolving a deeper issue. Likewise, the monotonicity of issues expresses that information carried by the information resolving an issue is also carried by the information resolving a deeper issue.

EIL can proven to be both sound and complete, the latter is taken up in the next section.

Proposition 6.4 (Soundness of EIL)**.** *The logic* EIL *is sound.*

Proof. Most axioms are standard for Kripke modalities over equivalence relations and are therefore easy to prove. The proofs of the other axioms rely on the insight that the intersection of a set of relations is contained in each of the relations in the set. The details are left to the reader. □

6.2 Completeness of Epistemic Issue Logic

In this section, we prove that EIL is strongly complete. The proof can be divided into two steps. First, pseudo models are defined and we argue that epistemic issue models can be viewed as special cases of pseudo models, which we call standard pseudo models. The logic EIL is proven to be sound and complete with respect to pseudo models using a standard completeness-via-canonicity proof. Secondly, it is shown that every pseudo model can be associated with an epistemic issue model and that every pseudo model is a p-morphic image of its associated model. It follows that the same formulas are satisfied in the associated model. Combining both steps yields completeness of EIL with respect to epistemic issue models: By the first step, any EIL-theory can be satisfied in some pseudo model. Then, using the second step, this EIL-theory is also satisfied in the associated epistemic issue model.

We start by defining pseudo models, as well as satisfaction in these models.

Definition 6.5 (Pseudo models)**.** A pseudo model is a tuple $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ such that:

- W is a set of worlds;

- $\sim_X$ is an equivalence relation for every $X \subseteq \mathcal{I}$;
- $\approx_X$ is an equivalence relation for every $X \subseteq \mathcal{I}$;
- $\mathcal{A} \subseteq \mathcal{I}$ is an agenda;
- $|| \bullet || : \Phi \to \mathcal{P}(W)$ is a valuation.

The following conditions are satisfied in pseudo models:

- Anti-monotonicity: For all $X, Y \subseteq \mathcal{I}$, the inclusion $X \subseteq Y$ implies that $\sim_Y \subseteq \sim_X$ and $\approx_Y \subseteq \approx_X$;
- Inclusion: For all $X \subseteq \mathcal{I}$, the inclusion $\sim_X \subseteq \approx_X$ holds;
- Necessity: $\approx_\emptyset$ is the universal relation on W.

If the following conditions also hold, M is called a standard pseudo model:

- Intersection: For all $X, Y \subseteq \mathcal{I}$, $\sim_{X \cup Y} = \sim_X \cap \sim_Y$ and $\approx_{X \cup Y} = \approx_X \cap \approx_Y$;
- Information intersection: For all $X \subseteq \mathcal{I}$, $\sim_X = \sim_\emptyset \cap \approx_X$.

The language $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ can be interpreted on pseudo models using the following satisfaction clauses.

Definition 6.6 (Satisfaction in Pseudo models)**.** Let $M = (W, \sim_{X \subseteq \mathcal{I}}, \approx_{X \subseteq \mathcal{I}}, \mathcal{A}, || \bullet ||)$ be a pseudo model and let $w \in W$ be a world in M. We have the following satisfaction clauses:

- $M, w \vDash p$ iff $w \in ||p||$;
- $M, w \vDash Ax$ iff $x \in \mathcal{A}$;
- $M, w \vDash \neg\varphi$ iff $M, w \nvDash \varphi$;
- $M, w \vDash \varphi \wedge \psi$ iff $M, w \vDash \varphi$ and $M, w \vDash \psi$;
- $M, w \vDash I^X \varphi$ iff $M, v \vDash \varphi$ for all $v \in W$ such that $w \sim_X v$;
- $M, w \vDash Q^X \varphi$ iff $M, v \vDash \varphi$ for all $v \in W$ such that $w \approx_X v$.

Epistemic issue models can be identified with standard pseudo models. Given an epistemic issue model $(W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$, define the relations $\sim_X := \sim \cap \, (\cap_{x \in X} \approx_x)$ and $\approx_X := \cap_{x \in X} \approx_x$. Then $(W, \sim_{X \subseteq \mathcal{I}}, \approx_{X \subseteq \mathcal{I}}, \mathcal{A}, || \bullet ||)$ can be seen to satisfy all the conditions of a standard pseudo model. Conversely, given a standard pseudo model $(W, \sim_{X \subseteq \mathcal{I}}, \approx_{X \subseteq \mathcal{I}}, \mathcal{A}, || \bullet ||)$, define $\sim := \sim_\emptyset$ and $\approx_x := \approx_{\{x\}}$ for every $x \in \mathcal{I}$. Then $(W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$ is an epistemic issue model. Satisfaction of formulas is invariant under these transformations by intersection and information intersection.

Proposition 6.7 (Soundness of EIL relative to pseudo models)**.** *The logic* EIL *is sound with respect to pseudo models.*

Proof. Most axioms can easily shown to be valid because $\sim_X$ and $\approx_X$ are equivalence relations for all $X \subseteq \mathcal{I}$. Resolution follows directly from the inclusion constraint, whereas the monotonicity axioms follow from the anti-monotonicity constraint. The details are left to the reader. □

As in Chapter 5, we need to restrict the canonical model to theories that agree on which formulas are necessarily true. A notion of universal equivalence similar to the earlier notion is defined, thereafter semi-canonical pseudo models can be defined relative to EIL-theories.

Definition 6.8 (Universal equivalence)**.** Two EIL-theories Γ and Δ are universally equivalent if for all $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$, $Q^{\emptyset}\varphi \in \Gamma$ if, and only if, $Q^{\emptyset}\varphi \in \Delta$. If Γ and Δ are universally equivalent, we denote this as $\Gamma \equiv_u \Delta$.

Definition 6.9 (Semi-canonical pseudo model)**.** Define the semi-canonical pseudo model relative to an EIL-theory Γ_0 as $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0} = (W^{\mathsf{EIL}}_{\Gamma_0}, \sim^{\Gamma_0}_X, \approx^{\Gamma_0}_X, \mathcal{A}_{\Gamma_0}, || \bullet ||^{\mathsf{EIL}}_{\Gamma_0})$, with:

- $W^{\mathsf{EIL}}_{\Gamma_0} := \{\Gamma \subseteq \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I}) \mid \Gamma \text{ a L-theory and } \Gamma_0 \equiv_u \Gamma\}$;
- $\Gamma \sim^{\Gamma_0}_X \Delta$ iff for all $\psi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ and $Y \subseteq X$, $I^Y\psi \in \Gamma$ iff $I^Y\psi \in \Delta$;
- $\Gamma \approx^{\Gamma_0}_X \Delta$ iff for all $\psi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ and $Y \subseteq X$, $Q^Y\psi \in \Gamma$ iff $Q^Y\psi \in \Delta$;
- $\mathcal{A}_{\Gamma_0} := \{x \in \mathcal{I} \mid Ax \in \Gamma_0\}$;
- $|| \bullet ||^{L}_{\Gamma_0} : \Phi \to \mathcal{P}(W^{L}_{\Gamma_0}) : p \mapsto \{\Gamma \mid p \in \Gamma\}$.

Proposition 6.10. *Every semi-canonical pseudo model is a pseudo model.*

Proof. Let Γ_0 be an arbitrary EIL-theory and let $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0} = (W^{\mathsf{EIL}}_{\Gamma_0}, \sim^{\Gamma_0}_X, \approx^{\Gamma_0}_X, \mathcal{A}_{\Gamma_0}, || \bullet ||^{\mathsf{EIL}}_{\Gamma_0})$. It is clear that for every $X \subseteq \mathcal{I}$ the relations $\sim^{\Gamma_0}_X$ and $\approx^{\Gamma_0}_X$ are equivalence relations. It remains to be shown that the additional constraints hold.

For anti-monotonicity, let $X, Y \subseteq \mathcal{I}$ such that $X \subseteq Y$ and suppose that $\Gamma \sim^{\Gamma_0}_Y \Delta$. If $Z \subseteq X$, then $Z \subseteq Y$. Thus for all $Z \subseteq X$ and $\psi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ we have $I^Z\psi \in \Gamma$ iff $I^Z\psi \in \Delta$. Hence $\Gamma \sim^{\Gamma_0}_X \Delta$ as well. The argument showing that $\approx^{\Gamma_0}_Y \subseteq \approx^{\Gamma_0}_X$ is analogous.

For inclusion, let $X \subseteq \mathcal{I}$ and suppose that $\Gamma \sim^{\Gamma_0}_X \Delta$. Let $Y \subseteq X$ and suppose that $Q^Y\psi \in \Gamma$. By positive introspection of Q^Y and the resolution axiom we find that $I^YQ^Y\psi \in \Gamma$. By assumption it follows that $I^YQ^Y\psi \in \Delta$. Applying factivity of information possession then yields $Q^Y\psi \in \Delta$. By symmetry $Q^Y\psi \in \Delta$ implies $Q^Y\psi \in \Gamma$, so $\sim^{\Gamma_0}_X \subseteq \approx^{\Gamma_0}_X$.

Lastly, necessity follows by definition of $W^{\mathsf{EIL}}_{\Gamma_0}$: the universal equivalence of all EIL-theories in $W^{\mathsf{EIL}}_{\Gamma_0}$ ensures that $\approx^{\Gamma_0}_{\emptyset}$ is the universal relation. Also note that all EIL-theories in $W^{\mathsf{EIL}}_{\Gamma_0}$ agree on the atom Ax for all $x \in \mathcal{I}$, since agendas are necessary. $\square$

Lemma 6.11 (Pseudo Existence Lemma for Q^X and I^X)**.** *Let Γ_0 be a EIL-theory and let $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0} = (W^{\mathsf{EIL}}_{\Gamma_0}, \sim^{\Gamma_0}_X, \approx^{\Gamma_0}_X, \mathcal{A}_{\Gamma_0}, || \bullet ||^{\mathsf{EIL}}_{\Gamma_0})$ be the semi-canonical pseudo model relative to Γ_0. For all $\Gamma \in W^{\mathsf{EIL}}_{\Gamma_0}$, if $I^X\varphi \notin \Gamma$ ($Q^X\varphi \notin \Gamma$), then there exists an EIL-theory $\Delta \in W^{\mathsf{EIL}}_{\Gamma_0}$ such that $\Gamma \sim^{\Gamma_0}_X \Delta$ ($\Gamma \approx^{\Gamma_0}_X \Delta$) and $\varphi \notin \Delta$.*

Proof. The proof is standard and analogous to the proof of Lemma 5.29. $\square$

Lemma 6.12 (Pseudo Truth lemma)**.** *Let Γ_0 be an EIL-theory and let $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0} = (W^{\mathsf{EIL}}_{\Gamma_0}, \sim^{\Gamma_0}_X, \approx^{\Gamma_0}_X, \mathcal{A}_{\Gamma_0}, ||\bullet||^{\mathsf{EIL}}_{\Gamma_0})$ be the semi-canonical pseudo model relative to Γ_0. For every $\Gamma \in W^{\mathsf{EIL}}_{\Gamma_0}$ and $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$: $\varphi \in \Gamma$ iff $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Gamma \vDash \varphi$.*

Proof. Let Γ_0 be an EIL-theory and let $\mathcal{M}_{\Gamma_0}$ be the semi-canonical pseudo model relative to Γ_0. Take an arbitrary $\Gamma \in W_{\Gamma_0}$. The proof is by induction on the structure of formulas.

For $\varphi := p$ or $\varphi := Ax$, the result follows immediately by the definitions of $|| \bullet ||^{\mathsf{EIL}}_{\Gamma_0}$ and $\mathcal{A}_{\Gamma_0}$, respectively.

For $\varphi := Ax$, the result follows immediately from the definition of $\mathcal{A}_{\Gamma_0}$.

For $\varphi := \neg\psi$, using the induction hypothesis it follows that $\varphi \in \Gamma$ iff $\psi \notin \Gamma$ iff $\mathcal{M}_{\Gamma_0}, \Gamma \nvDash \psi$ iff $\mathcal{M}_{\Gamma_0}, \Gamma \vDash \varphi$. The case for conjunction is similar.

For $\varphi := I^X\psi$, suppose $I^X\psi \in \Gamma$. By construction of $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}$, every Δ with $\Gamma \sim^{\Gamma_0}_X \Delta$ contains $I^X\psi$. By factivity for I^X it follows that $\psi \in \Delta$ for all these theories. The induction hypothesis then gives us $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Delta \vDash \psi$, thus $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Gamma \vDash I^X\psi$. Conversely, suppose that $I^X\psi \notin \Gamma$ for contraposition. By the Pseudo Existence Lemma for I^X, there is an EIL-theory Δ such that $\Gamma \sim_X^{\Gamma_0} \Delta$ and $\psi \notin \Delta$. By the induction hypothesis it follows that $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Delta \nvDash \psi$. Thus $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Gamma \nvDash I^X\psi$. The case for Q^X is similar. □

Proposition 6.13 (Completeness with respect to pseudo models). *The logic* EIL *is strongly complete with respect to pseudo models.*

Proof. Consider an arbitrary EIL-consistent set of formulas Γ, by Lindenbaum's lemma it can be extended into an EIL-theory Γ_0. Consider the semi-canonical model $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}$, which is a pseudo model. It follows that $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Gamma_0 \vDash \Gamma_0$ by Lemma 6.12. Thus EIL is strongly complete with respect to pseudo models. □

Next we turn to transforming pseudo models into epistemic issue models. We employ a technique that is a variant of unraveling.

Definition 6.14 (Paths in pseudo models). Let $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be a pseudo model. A path α in M is a sequence $(w_0, R_0, w_1, R_1, \ldots, R_{n-1}, w_n)$ such that:

- $w_i \in W$ for $0 \leq i \leq n$;
- $R_i \in \{\sim_X \mid X \subseteq \mathcal{I}\} \cup \{\approx_X \mid X \subseteq \mathcal{I}\}$ for $0 \leq i < n$;
- $w_i R_i w_{i+1}$ for $0 \leq i < n$.

For any path $\alpha = (w_0, R_0, \ldots, R_{n-1}, w_n)$ we say that the origin of α is w_0 and for the last element in the path α we write $\overline{\alpha} := w_n$. We define the following relations between paths in M:

- $\alpha \rightarrow_X \beta$ iff $\beta = \alpha \circ (\sim_X, \overline{\beta})$;
- $\alpha \rightsquigarrow_X \beta$ iff $\beta = \alpha \circ (\approx_X, \overline{\beta})$;
- $\alpha \xrightarrow{\sim}_X \beta$ iff $\alpha \rightarrow_Y \beta$ for some $Y \supseteq X$;
- $\alpha \xrightarrow{\approx}_X \beta$ iff $\alpha \rightsquigarrow_Y \beta$ or $\alpha \rightarrow_Y \beta$ for some $Y \supseteq X$.

Observe that $\xrightarrow{\sim}_X \subseteq \xrightarrow{\approx}_X$ for every $X \subseteq \mathcal{I}$. Given any pseudo model, we construct an associated model using a variation of unraveling.[70]

Definition 6.15 (Associated models). Let $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be a pseudo model and let $w_0 \in W$. The associated model of M rooted in w_0 is the tuple $T^M_{w_0} = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$ with:

- T is the set of all paths in M with origin w_0;
- $\sim := (\xrightarrow{\sim}_\emptyset \cup \xleftarrow{\sim}_\emptyset)^*$ is a relation on T;
- $\approx_x := (\xrightarrow{\approx}_{\{x\}} \cup \xleftarrow{\approx}_{\{x\}})^*$ is a relation on T for every $x \in \mathcal{I}$;
- $||p||^T = \{\alpha \in T \mid \overline{\alpha} \in ||p||\}$ is a valuation,

where R^* is the reflexive-transitive closure of R, $\xleftarrow{\sim}_X$ is the inverse of $\xrightarrow{\sim}_X$ and $\xleftarrow{\approx}_X$ is the inverse of $\xrightarrow{\approx}_X$.

[70] Blackburn, De Rijke and Venema (2001, p. 218-219) provide the definition and properties of unraveling.

Let $\alpha \to \beta$ iff $\alpha \to_X \beta$ or $\alpha \rightsquigarrow_X \beta$ for some $X \subseteq \mathcal{I}$. It can be seen that this relation structures the set of all paths in a pseudo model M with origin w_0 into a tree rooted at w_0. It follows that in every associated model $T^M = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$ and $\alpha, \beta \in T$, there is a unique non-redundant path from α to β. That is, there is a unique path between every α and β such that every step in the path is of the form $\alpha_i \to \alpha_{i+1}$ or $\alpha_i \leftarrow \alpha_{i+1}$ and no nodes are repeated on the path.

When the context allows it, we write 'T^M' instead of '$T^M_{w_0}$'. The relations $\sim$ and $\approx_x$ are by definition equivalence relations for every $x \in \mathcal{I}$. Thus associated models are epistemic issue models. As mentioned earlier, this means that we can identify associated models with (standard) pseudo models when taking $\sim_X := \sim \cap (\cap_{x\in X} \approx_x)$ and $\approx_X := \cap_{x\in X} \approx_x$. Thus we can compare any pseudo model M to its associated model T^M as a pseudo model, allowing us to formulate the following lemmas.

Lemma 6.16. *Let $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be a pseudo model and let $T^M_{w_0} = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$ be its associated model for some $w_0 \in W$. Let $X \subseteq \mathcal{I}$ and $\alpha, \beta \in T$. The following statements are true.*

(i) $\alpha \sim_X \beta$ iff the non-redundant path from α to β only consists of steps of the form $\alpha_i \to_{X_i} \alpha_{i+1}$ or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$, where $X_i \supseteq X$;

(ii) $\alpha \approx_X \beta$ iff the non-redundant path from α to β only consists of steps of the form $\alpha_i \rightsquigarrow_{X_i} \alpha_{i+1}$ or $\alpha_i \leftsquigarrow_{X_i} \alpha_{i+1}$ or $\alpha_i \to_{X_i} \alpha_{i+1}$, or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$, such that $X_i \supseteq X$.

Proof. For (i), suppose $\alpha \sim_X \beta$, then $\alpha \sim \beta$ and $\alpha \approx_x \beta$ for every $x \in X$. So the non-redundant path from α to β has the property that every step is of the form $\alpha_i \xrightarrow{\sim}_\emptyset \alpha_{i+1}$ or $\alpha_i \xleftarrow{\sim}_\emptyset \alpha_{i+1}$, *and* every step in the path is of the form $\alpha_i \xrightarrow{\approx}_{\{x\}} \alpha_{i+1}$ or $\alpha_i \xleftarrow{\approx}_{\{x\}} \alpha_{i+1}$ for every $x \in \mathcal{X}$. By definition of $\xrightarrow{\sim}_\emptyset$ and $\xrightarrow{\approx}_{\{x\}}$, each step must be such that: $\alpha_i \to_{X_i} \alpha_{i+1}$ or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$ for some $X_i \subseteq \mathcal{I}$, *and* $\alpha_i \rightsquigarrow_{X_i} \alpha_{i+1}$, or $\alpha_i \leftsquigarrow_{X_i} \alpha_{i+1}$, or $\alpha_i \to_{X_i} \alpha_{i+1}$, or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$ for some $X_i \supseteq X$. It follows that every step must be of the form $\alpha_i \to_{X_i} \alpha_{i+1}$ or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$, where $X_i \supseteq X$. The converse can be obtained by working backwards and applying the definitions.

For (ii), suppose $\alpha \approx_X \beta$, then $\alpha \approx_x \beta$ for all $x \in X$. If $\alpha \approx_x \beta$, then by definition of $\approx_x$ there is a non-redundant path from α to β such that every step in the path is of the form $\alpha_i \xrightarrow{\approx}_{\{x\}} \alpha_{i+1}$ or $\alpha_i \xleftarrow{\approx}_{\{x\}} \alpha_{i+1}$. By definition of $\xrightarrow{\approx}_{\{x\}}$, it follows that each step must be of one of the following forms for some $X_i \ni x$: $\alpha_i \rightsquigarrow_{X_i} \alpha_{i+1}$, or $\alpha_i \leftsquigarrow_{X_i} \alpha_{i+1}$, or $\alpha_i \to_{X_i} \alpha_{i+1}$, or $\alpha_i \leftarrow_{X_i} \alpha_{i+1}$. Since this is the case for all $x \in X$ and the path from α to β is unique, it follows that $X \subseteq X_i$ for every X_i. The converse can be obtained by working backwards and applying the definitions. $\square$

Lemma 6.17. *Let $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be a pseudo model and let $T^M_{w_0} = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$ be its associated model for some $w_0 \in W$. Let $x \subseteq \mathcal{I}$ and $\alpha, \beta \in T$. The following statements are true.*

(i) If $\alpha \xrightarrow{\sim}_X \beta$, then $\overline{\alpha} \sim_X \overline{\beta}$;

(ii) If $\alpha \xrightarrow{\approx}_X \beta$, then $\overline{\alpha} \approx_X \overline{\beta}$.

Proof. For (i), suppose $\alpha \xrightarrow{\sim}_X \beta$, then by definition there exists $Y \supseteq X$ such that $\alpha \to_Y \beta$. Thus $\overline{\alpha} \sim_Y \overline{\beta}$, anti-monotonicity of $\sim_Y$ in the pseudo model M then yields $\overline{\alpha} \sim_X \overline{\beta}$.

For (ii), suppose $\alpha \xrightarrow{\approx}_X \beta$, then by definition there exists $Y \supseteq X$ such that $\alpha \rightsquigarrow_Y \beta$ or $\alpha \to_Y \beta$. In the former case we find that $\overline{\alpha} \approx_Y \overline{\beta}$ and in the latter case that $\overline{\alpha} \sim_Y \overline{\beta}$, which implies $\overline{\alpha} \approx_Y \overline{\beta}$ by inclusion. So in both cases $\overline{\alpha} \approx_Y \overline{\beta}$, the anti-monotonicity of $\approx_Y$ in the pseudo model M then yields $\overline{\alpha} \approx_X \overline{\beta}$. $\square$

Lemma 6.18. *Let $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be a pseudo model and let $T^M_{w_0} = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$ be its associated model for some $w_0 \in W$. Let $x \subseteq \mathcal{I}$ and $\alpha, \beta \in T$. The following statements are true.*

(i) If $\alpha \sim_X \beta$, then $\overline{\alpha} \sim_X \overline{\beta}$;

(ii) If $\alpha \approx_X \beta$, then $\overline{\alpha} \approx_X \overline{\beta}$.

Proof. Both statements are proven by induction on the length of the unique non-redundant path from α to β. The base cases are trivial: for any path of length 0, $\alpha = \beta$, hence since $\sim_X$ and $\approx_X$ are equivalence relations it follows that $\overline{\alpha} \sim_X \overline{\beta}$ and $\overline{\alpha} \approx_X \overline{\beta}$.

For the induction step of (i), suppose that the unique non-redundant path from α to β has length $n+1$ and consider the last step in the path. By Lemma 6.16, this step must be of the form $\alpha_n \rightarrow_{X_n} \alpha_{n+1}$ or $\alpha_n \leftarrow_{X_n} \alpha_{n+1} = \beta$, with $X_n \supseteq X$. So either $\alpha_n \overset{\sim}{\rightarrow}_{X_n} \alpha_{n+1} = \beta$ or $\alpha_n \overset{\sim}{\leftarrow}_{X_n} \alpha_{n+1} = \beta$, both of which imply $\overline{\alpha_n} \sim_X \overline{\beta}$ by Lemma 6.17 (i). The induction hypothesis then gives $\overline{\alpha} \sim_X \overline{\alpha_n}$ and subsequently transitivity of $\sim_X$ yields $\overline{\alpha} \sim_X \overline{\beta}$.

For the induction step of (ii), suppose that the unique non-redundant path from α to β has length $n+1$ and consider the last step in the path. By Lemma 6.16, this step must be of the form $\alpha_n \rightsquigarrow_{X_n} \alpha_{n+1} = \beta$, or $\alpha_n \leftsquigarrow_{X_n} \alpha_{n+1} = \beta$, or $\alpha_n \rightarrow_{X_n} \alpha_{n+1}$ or $\alpha_n \leftarrow_{X_n} \alpha_{n+1} = \beta$, with $X_n \supseteq X$. So either $\alpha_n \overset{\approx}{\rightarrow}_{X_n} \alpha_{n+1} = \beta$ or $\alpha_n \overset{\approx}{\leftarrow}_{X_n} \alpha_{n+1} = \beta$, both of which imply $\overline{\alpha_n} \approx_X \overline{\beta}$ by Lemma 6.17 (ii). The induction hypothesis then gives $\overline{\alpha} \approx_X \overline{\alpha_n}$ and subsequently transitivity of $\approx_X$ yields $\overline{\alpha} \approx_X \overline{\beta}$. $\square$

Since associated models can be interpreted as pseudo models, we can inquire whether a pseudo model M and its associated model T^M are bisimilar. Pseudo models are modal structures augmented with an agenda. The notion of a p-morphism can easily be extended so that it can also be applied to pseudo models: simply add the clause that the agendas of any two pseudo models between which a p-morphism exists must be equal. It is obvious that this entails that satisfaction of atoms $A_x \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ is invariant under p-morphisms. Consequently, by the usual argument, satisfaction is invariant under p-morphisms for all formulas in $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$. It can be proven that every pseudo model M is a p-morphic image of its associated model T^M.

Lemma 6.19. *Every pseudo model $M = (W, \sim_{X\subseteq\mathcal{I}}, \approx_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||)$ is a p-morphic image of its associated model $T^M_{w_0} = (T, \sim^T_{X\subseteq\mathcal{I}}, \approx^T_{X\subseteq\mathcal{I}}, \mathcal{A}, ||\bullet||^T)$, for every $w_0 \in W$.*

Proof. We prove that the mapping $f : T \to W : \alpha \mapsto \overline{\alpha}$ is a p-morphism. Since $\alpha \in ||p||^T$ if, and only if, $\overline{\alpha} \in ||p||$, it follows that atomic proposition letters are preserved under f. Moreover, atoms of the form Ax are also preserved: the agendas in M and T^M are identical.

For the forth condition, assume that $\alpha \sim_X \beta$ ($\alpha \approx_X \beta$) for some $X \subseteq \mathcal{I}$. By Lemma 6.18 (i) ((ii)), $\overline{\alpha} \sim_X \overline{\beta}$ ($\overline{\alpha} \approx_X \overline{\beta}$).

For the back condition, assume that $\overline{\alpha} \sim_X w$ ($\overline{\alpha} \approx_X w$), we need to show that there exists a $\beta \in T$ such that $\alpha \sim_X \beta$ ($\alpha \approx_X \beta$) and $\overline{\beta} = w$. The path $\beta := \alpha \circ (\sim_X, w)$ ($\beta := \alpha \circ (\approx_X, w)$) suffices by definition.

So f is a p-morphism. It is also surjective: for every $w \in W$ there is a reduced path $(w_0, \approx_\emptyset, w) \in T$. Thus the pseudo model M is a bounded morphic image of its associated model T^M. $\square$

So there exists a bisimulation between every pseudo model M and its associated model T^M. Consequently, for all formulas $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ and $\alpha \in T$: $T^M, \alpha \vDash \varphi$ if, and only if, $M, \overline{\alpha} \vDash \varphi$. This corollary can be used to prove the completeness of EIL.

Proposition 6.20 (Completeness of EIL). *The logic* EIL *is strongly complete.*

Proof. Consider an arbitrary EIL-consistent set of formulas Γ, by Lindenbaum's lemma it can be extended into an EIL-theory Γ_0. Let $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}$ be the semi-canonical pseudo model relative to Γ_0. The proof for strong completeness of EIL with respect to pseudo models has shown that $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}, \Gamma_0 \vDash \Gamma_0$. By Lemma 6.19, the associated model $T^{\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}}_{\Gamma_0}$ is an epistemic issue model such that $T^{\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}}_{\Gamma_0}, \Gamma_0 \vDash \Gamma_0$. Thus EIL is strongly complete. □

6.3 Dynamic Epistemic Issue Logic

The logic EIL can be extended into a dynamic logic accommodating the actions of issue addition, issue retraction and issue resolution. We call this logic Dynamic Epistemic Issue Logic (DEIL). The updates of epistemic issue models corresponding to these actions can be derived from Definition 3.7 and Definition 3.8.

Definition 6.21 (Model updates). Let $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, ||\bullet||)$ be an epistemic issue model.

Adding the issue $X \subseteq \mathcal{I}$ to the agent's agenda yields the model

$$M_{[+X]} := (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A} \cup X, ||\bullet||).$$

Retracting the issue $X \subseteq \mathcal{I}$ from the agent's agenda yields the model

$$M_{[-X]} := (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A} \backslash X, ||\bullet||).$$

Resolving the issue $X \subseteq \mathcal{I}$ yields the updated model

$$M_{[X!]} := (W, \sim \cap \approx_X, \approx_{x \in \mathcal{I}}, \mathcal{A}, ||\bullet||).$$

The language and semantics of EIL are extended accordingly.

Definition 6.22 (Language of DEIL). The dynamic language $\mathcal{L}_{\mathsf{DEIL}}(\Phi, \mathcal{I})$ is obtained by adding the following clauses to the definition of the language $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$:

$$\cdots \mid [+X]\varphi \mid [-X]\varphi \mid [X!]\varphi \quad \text{with } X \subseteq \mathcal{I}.$$

Definition 6.23 (Semantics of DEIL). The semantics for DEIL are obtained by adding the following clauses to the semantics of EIL:

- $M, w \vDash_{\mathsf{DEIL}} [+X]\varphi$ iff $M_{[+X]}, w \vDash \varphi$;
- $M, w \vDash_{\mathsf{DEIL}} [-X]\varphi$ iff $M_{[-X]}, w \vDash \varphi$;
- $M, w \vDash_{\mathsf{DEIL}} [X!]\varphi$ iff $M_{[X!]}, w \vDash \varphi$.

It is easy to see that the following composition principles for the dynamic actions are valid.

Proposition 6.24 (Composition principles). *Let* $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, ||\bullet||)$ *be an epistemic issue model. The following composition principles are valid on* M*:*

(i) $[+X][+Y]\varphi \leftrightarrow [+(X \cup Y)]\varphi$*;*

(ii) $[-X][-Y]\varphi \leftrightarrow [-(X \cup Y)]\varphi$*;*

(iii) $[X!][Y!]\varphi \leftrightarrow [(X \cup Y)!]\varphi$.

It is obvious that an agenda update followed by an issue resolution (or vice versa) cannot be composed into a single action: the former only affects the agent's agenda and the latter only the agent's information relation. Since there are no actions available that affect both the agent's agenda and information relation, these actions cannot be composed. However, we also do *not* have a composition principle for agenda updates in general.

Proposition 6.25 (No composition principle for agenda updates)**.** *There is no composition principle that reduces formulas of the form* $[+X][-Y]\varphi$ *or* $[-X][+Y]\varphi$ *to a formula of the form* $[?]\varphi$*, where '*$[?]$*' denotes a single action.*

Proof. Consider the language $\mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ with $\Phi = \{p, q\}$ and $\mathcal{I} = \{x, y\}$, interpreted on the epistemic issue model $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$ with:

- $W = \{w_1, w_2, w_3\}$;
- $\sim$ the identity relation;
- $\approx_x$ and $\approx_y$ the reflexive-symmetric closures of $\{(w_1, w_2)\}$ and $\{(w_2, w_3)\}$, respectively;
- $\mathcal{A} = \{y\}$;
- $||p|| = \{w_2, w_3\}$ and $||q|| = \{w_1, w_2\}$.

The models M and $M_{[+\{x\}][-\{y\}]}$ are depicted in Figure 6.1. It can be seen that $M, w_2 \vDash Rp \wedge \neg Rq$, while $M_{[+\{x\}][-\{y\}]}, w_2 \vDash \neg Rp \wedge Rq$, i.e. $M, w_2 \vDash [+\{x\}][-\{y\}]\neg Rp \wedge Rq$.

If there exists an action $[?]$ such that $M_{[?]} \vDash \neg Rp \wedge Rq$, then $[?]$ must be an agenda update, because issue resolutions do not affect issue-relevance. However, it was shown in Proposition 4.15 (ii) that issue additions can only make *more* formulas issue-relevant. Similarly, by Proposition 4.15 (iii), issue retractions can only make *fewer* formulas issue-relevant. Thus $[?]$ cannot exist. $\square$

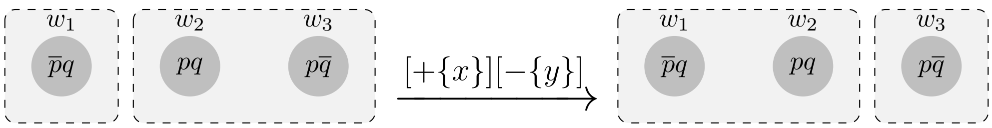

(a) M (b) $M_{[+\{x\}][-\{y\}]}$

Figure 6.1: The dynamic issue-epistemic models M and $M_{[+\{x\}][-\{y\}]}$. All information is available to the agent and reflexive arrows have been omitted.

Intuitively, a sequence of agenda updates corresponds to a single agenda update. Furthermore, sometimes issue addition and issue retraction happen simultaneously. For instance, when a paradigm shift takes place. This limitation can be addressed by introducing an abbreviation for general agenda updates: $[+X{-}Y]\varphi := [+X][-Y]\varphi$, for $X, Y \subseteq \mathcal{I}$ such that $X \cap Y = \emptyset$. The formula '$[+X{-}Y]\varphi$' should be interpreted as 'φ is true after the agenda update that adds the issue X and retracts the issue Y'. Requiring the sets of added and retracted basic issues to be disjoint is sensible: it does not make sense to simultaneously add and retract the same basic issue. When viewing $[+X{-}Y]$ as a single agenda update, the simultaneous addition and retraction of issues can be modeled.

Issue addition and issue retraction can be viewed as special cases of the more general agenda updates: $[+X]\varphi = [+X-\emptyset]\varphi$ and $[-Y]\varphi = [+\emptyset-Y]\varphi$. Accordingly, the agenda update $[+X-Y]$ of a model $M = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A}, ||\bullet||)$ can be taken as:

$$M_{[+X-Y]} := M_{[+X][-Y]}(W, \sim, \approx_{x\in\mathcal{I}}, (\mathcal{A}\cup X)\backslash Y, ||\bullet||).$$

Then it follows that $M, w \vDash [+X-Y]\varphi$ if, and only if $M_{[+X-Y]}, w \vDash \varphi$. It can be proven that there exists a composition principle for agenda updates.

Proposition 6.26 (Composition principle for agenda updates)**.** *Let $M = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be an epistemic issue model. The following composition principle is valid on M:*

$$[+X-Y][+X'-Y']\varphi \leftrightarrow [+(X\cup X')\backslash Y' - (Y\cup Y')\backslash X']\varphi.$$

Proof. Let $M = (W, \sim, \approx_{x\in\mathcal{I}}, \mathcal{A}, ||\bullet||)$ be an epistemic issue model. By definition, $M_{[+X-Y][+X'-Y']} = (W, \sim, \approx_{x\in\mathcal{I}}, (((\mathcal{A}\cup X)\backslash Y)\cup X')\backslash Y', ||\bullet||)$. The sets X and Y and the sets X' and Y' are pairwise disjoint, but otherwise these sets may overlap. If X and X' overlap, then the overlapping basic issues are certainly on the resulting agenda. If X and Y' overlap, then the overlapping basic issues are first added but retracted thereafter, thus they are *not* on the resulting agenda. If Y and X' overlap, then the overlapping basic issues are first retracted from the agenda but added thereafter, thus they are on the resulting agenda. If Y and Y' overlap, then the overlapping basic issues are certainly *not* on the resulting agenda. Thus, effectively, $(X\cup X')\backslash Y'$ is added and $(Y\cup Y')\backslash X'$ is retracted.

It remains to be checked that the sets $(X\cup X')\backslash Y'$ and $(Y\cup Y')\backslash X'$ are disjoint. Suppose for reductio that there is an $x\in\mathcal{I}$ such that $x\in(X\cup X')\backslash Y'$ and $x\in(Y\cup Y')\backslash X'$. From the former it follows that $x\in X\backslash Y'$ or $x\in X'$ and from the latter that $x\in Y\backslash X'$ or $x\in Y'$. If $x\in X\backslash Y'$, then neither $x\in Y\backslash X'$ nor $x\in Y'$ can hold. Likewise, if $x\in X'$, then neither $x\in Y\backslash X'$ nor $x\in Y'$ can hold. So $(X\cup X')\backslash Y'$ and $(Y\cup Y')\backslash X'$ are disjoint. So $M_{[+X-Y][+X'-Y']} = M_{[+(X\cup X')\backslash Y'-(Y\cup Y')\backslash X']}$. Consequently,

$$M, w \vDash [+X-Y][+X'-Y']\varphi$$

if, and only if,

$$M, w \vDash [+(X\cup X')\backslash Y' - (Y\cup Y')\backslash X']\varphi.$$

□

Proposition 6.25 showed that issue addition followed by issue retraction could not be composed into a single action in DEIL. By the proposition above, the formula

$$[+X][-Y]\varphi \leftrightarrow [+(X\backslash Y)-Y]\varphi$$

is valid on epistemic issue models. Hence the simultaneous addition and retraction of issues can be simulated by concatenating an issue addition with an issue retraction.

Agenda updates and issue resolution work independently from each other: agenda updates only affect the agent's agenda, whereas issue resolutions only affect the information relation. Thus agenda updates and issue resolutions commute. As a consequence of Proposition 6.24 and Proposition 6.26, any sequence of actions can be captured by an agenda update followed by an issue resolution.

The proof system of EIL can be extended with reduction axioms to obtain a proof system for DEIL.

Reduction axioms for DEIL
A proof system for DEIL consists of all axioms and rules of EIL supplemented with the following reduction axioms:

- $[+X]p \leftrightarrow p$;

- $[+X]\neg\varphi \leftrightarrow \neg[+X]\varphi$;
- $[+X](\varphi \wedge \psi) \leftrightarrow ([+X]\varphi) \wedge ([+X]\psi)$;
- $[+X]I^Y\varphi \leftrightarrow I^Y[+X]\varphi$;
- $[+X]Q^Y\varphi \leftrightarrow Q^Y[+X]\varphi$;
- $[+X]Ax \leftrightarrow \top$, for $x \in X$;
- $[+X]Ax \leftrightarrow Ax$, for $x \notin X$;
- $[-X]p \leftrightarrow p$;
- $[-X]\neg\varphi \leftrightarrow \neg[-X]\varphi$;
- $[-X](\varphi \wedge \psi) \leftrightarrow ([-X]\varphi) \wedge ([-X]\psi)$;
- $[-X]I^Y\varphi \leftrightarrow I^Y[-X]\varphi$;
- $[-X]Q^Y\varphi \leftrightarrow Q^Y[-X]\varphi$;
- $[-X]Ax \leftrightarrow \bot$, for $x \in X$;
- $[-X]Ax \leftrightarrow Ax$, for $x \notin X$;
- $[X!]p \leftrightarrow p$;
- $[X!]\neg\varphi \leftrightarrow \neg[X!]\varphi$;
- $[X!](\varphi \wedge \psi) \leftrightarrow ([X!]\varphi) \wedge ([X!]\psi)$;
- $[X!]I^Y\varphi \leftrightarrow I^{X\cup Y}[X!]\varphi$;
- $[X!]Q^Y\varphi \leftrightarrow Q^Y[X!]\varphi$;
- $[X!]Ax \leftrightarrow Ax$.

The logic DEIL can proven to be both sound and complete.

Proposition 6.27 (Soundness of DEIL)**.** *The logic* DEIL *is sound.*

Proof. Proving the reduction axioms to be valid is straightforward, the details are left to the reader. The other axioms and rules inherit their validity from EIL. □

The completeness of DEIL can be proven by piggybacking on the completeness result for EIL. The proof relies on two lemmas.

Lemma 6.28. *For any formula* $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ *and any action* $[X]$ *(either* $[X!]$, $[+X]$ *or* $[-X]$*), there exists a formula* $\varphi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ *such that*

$$\vdash_{\mathsf{DEIL}} [X]\varphi \leftrightarrow \varphi'.$$

Proof. We give a proof by induction on the structure of formulas.

For $\varphi := p$ and any action $[X]$, the reduction axiom(s) for proposition letters immediately gives us the desired formula $\varphi' := p$.

For $\varphi := Ax$, we treat each action separately. If the dynamic action is issue addition $[+X]$, then $\varphi' := \top$ suffices if $x \in X$, otherwise $\varphi' := Ax$ suffices. In case of issue retraction $[-X]$, then $\varphi' := \bot$ suffices if $x \in X$, otherwise $\varphi' := Ax$ suffices. In case of issue resolution $[X!]$, $\varphi' = Ax$ suffices. Each case follows directly from the appropriate reduction axiom.

For $\varphi := \neg\psi$ and any action $[X]$, the induction hypothesis gives us a formula ψ' such that $\vdash [X]\psi \leftrightarrow \psi'$. Combining this with the reduction axiom(s) for negation, we find that $\vdash [X]\neg\psi \leftrightarrow \neg\psi'$.

For $\varphi := \psi \wedge \theta$ and any action $[X]$, the induction hypothesis gives us formulas ψ' and θ' such that $\vdash [X]\psi \leftrightarrow \psi'$ and $\vdash [X]\theta \leftrightarrow \theta'$. Combining this with the reduction axiom(s) for conjunction, we find that $\vdash [X]\psi \wedge \theta \leftrightarrow \psi' \wedge \theta'$.

For $\varphi := I^Y\psi$, the induction hypothesis gives us a formula ψ' such that $\vdash [X]\psi \leftrightarrow \psi'$. We distinguish between two cases. If the action is issue addition $[+X]$ or issue retraction $[-X]$, then $\varphi' := I^Y\psi'$ suffices. In case of issue resolution $[X!]$, $\varphi' = I^{X\cup Y}\psi'$ suffices. Each case follows directly from the appropriate reduction axiom.

For $\varphi := Q^Y\psi$ and any action $[X]$, the induction hypothesis gives us a formula ψ' such that $\vdash [X]\psi \leftrightarrow \psi'$. It follows from the appropriate reduction axiom(s) that $\varphi' := Q^Y\psi'$ suffices, completing the induction. □

The following lemma shows that it is DEIL-provable that EIL and DEIL are co-expressive.

Lemma 6.29. *For every formula $\varphi \in \mathcal{L}_{\mathsf{DEIL}}(\Phi, \mathcal{I})$, there exists a formula $\varphi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that*

$$\vdash_{\mathsf{DEIL}} \varphi \leftrightarrow \varphi'.$$

Proof. We give a proof by induction on the structure of formulas.

For $\varphi := p$ and $\varphi := Ax$, it immediately follows that $\varphi' := \varphi$ suffices.

For $\varphi := \neg\psi$, the induction hypothesis gives us a formula $\psi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash \psi \leftrightarrow \psi'$. Hence setting $\varphi' := \neg\psi'$ suffices.

For $\varphi := \psi \wedge \theta$, the induction hypothesis gives us formulas $\psi', \theta' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash \psi \leftrightarrow \psi'$ and $\vdash \theta \leftrightarrow \theta'$. Hence $\varphi' := \psi' \wedge \theta'$ suffices.

For $\varphi := I^Y\psi$, the induction hypothesis gives us a formula $\psi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash \psi \leftrightarrow \psi'$. By necessitation for $Q^\emptyset$, monotonicity of issues and the resolution axiom it follows that $\vdash I^Y\psi \leftrightarrow I^Y\psi'$, hence $\varphi' := I^Y\psi'$ suffices.

For $\varphi := Q^Y\psi$, the induction hypothesis gives us a formula $\psi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash \psi \leftrightarrow \psi'$. By necessitation for $Q^\emptyset$ and monotonicity of issues it follows that $\vdash Q^Y\psi \leftrightarrow Q^Y\psi'$, hence $\varphi' := I^Y\psi'$ suffices.

For $\varphi := [X]\psi$, where $[X]$ is any action, the induction hypothesis gives us a formula $\psi' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash \psi \leftrightarrow \psi'$. It follows that $\vdash [X]\psi \leftrightarrow [X]\psi'$. From Lemma 6.28 it follows that there exists a $\psi'' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ such that $\vdash [X]\psi' \leftrightarrow \psi''$. Hence setting $\varphi' := \psi''$ suffices, completing the induction. □

Proposition 6.30 (Completeness of DEIL)**.** *The logic* DEIL *is strongly complete.*

Proof. Let $\Gamma \subseteq \mathcal{L}_{\mathsf{DEIL}}(\Phi, \mathcal{I})$ be any DEIL-consistent set of formulas. Enumerate the formulas in Γ as $\varphi_1, \varphi_2, \ldots$, which may be finite. By Lemma 6.29, there exist formulas $\varphi'_1, \varphi'_2, \ldots$ such that $\varphi'_i \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ and $\vdash_{\mathsf{DEIL}} \varphi \leftrightarrow \varphi'$. Let Γ' denote the set of these formulas φ'_i. By soundness of DEIL it follows that $\Gamma' \vDash \Gamma$. Since $\Gamma' \subseteq \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ is consistent, there exists an epistemic issue model M containing a world w such that $M, w \vDash \Gamma'$ by completeness of EIL. It follows that $M, w \vDash \Gamma$ as well, hence DEIL is strongly complete. □

6.4 Decidability

In this section we prove that every logic introduced in this thesis is decidable. The decidability of every logic is a corollary of the decidability of EIL. To prove that EIL is decidable, we prove that pseudo models have the finite model property. First, the closure of a formula is defined, which is a finite set of formulas. It is shown that pseudo models can be filtrated through the closure of any formula to obtain finite pseudo models. Satisfaction is invariant for the formulas contained in the closure through which a pseudo model is filtrated, proving that pseudo models have the finite model property. Decidability of EIL follows from the finite model property. Since every logic of Chapter 5 is a sublogic of EIL, they are also decidable. Decidability of DEIL follows from the co-expressiveness of EIL and DEIL: checking the validity of any formula in the language of DEIL is equivalent to checking the validity of some formula in the language of EIL.

We start by defining the closure of a formula.

Definition 6.31 (Closure of a formula)**.** Let $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$, the closure of this formula is the smallest set of formulas $\mathrm{cl}(\varphi)$ such that:

1. $\varphi \in \mathrm{cl}(\varphi)$;
2. If $\psi \in \mathrm{cl}(\varphi)$, then $\mathrm{sub}(\psi) \subseteq \mathrm{cl}(\varphi)$, where $\mathrm{sub}(\psi)$ denotes the set of subformulas of ψ;
3. If ψ is not of the form $\neg\theta$, then $\neg\psi \in \mathrm{cl}(\varphi)$ for all $\psi \in \mathrm{cl}(\varphi)$;
4. $I^X\psi \in \mathrm{cl}(\varphi)$ iff $I^Y\psi \in \mathrm{cl}(\varphi)$ iff $Q^X\psi \in \mathrm{cl}(\varphi)$ iff $Q^Y\psi \in \mathrm{cl}(\varphi)$, for all $X, Y \subseteq \mathcal{I}$.

Observe that for any formula $\varphi \in \mathsf{EIL}(\Phi, \mathcal{I})$, its closure $\mathrm{cl}(\varphi)$ is finite.

Definition 6.32 (Equivalence of theories relative to formulas)**.** Two EIL-theories Γ and Δ are equivalent with respect to a formula $\varphi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ whenever $\Gamma \cap \mathrm{cl}(\varphi) = \Delta \cap \mathrm{cl}(\varphi)$.

Definition 6.33 (Filtrations of semi-canonical pseudo model)**.** Let $\varphi \in \mathsf{EIL}(\Phi, \mathcal{I})$ and let $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0} = (W^{\mathsf{EIL}}_{\Gamma_0}, \sim^{\Gamma_0}_X, \approx^{\Gamma_0}_X, \mathcal{A}_{\Gamma_0}, ||\bullet||^{\mathsf{EIL}}_{\Gamma_0})$ be a semi-canonical pseudo model. The filtration of $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}$ through $\mathrm{cl}(\varphi)$ is $\mathcal{M}^{f(\varphi)}_{\Gamma_0} = (W^{f(\varphi)}_{\Gamma_0}, \sim^{f(\varphi)}_X, \approx^{f(\varphi)}_X, \mathcal{A}^{f(\varphi)}_{\Gamma_0}, ||\bullet||^{f(\varphi)}_{\Gamma_0})$, where:

- $W^{f(\varphi)}_{\Gamma_0} = \{[\Gamma]_{\equiv_\varphi} \mid \Gamma \in W^{\mathsf{EIL}}_{\Gamma_0}\}$, the set of equivalence classes under $\equiv_\varphi$;
- $[\Gamma] \sim^{f(\varphi)}_X [\Delta]$ iff for all $I^X\psi \in \mathrm{cl}(\varphi)$ and $Y \subseteq X$, $I^Y\psi \in \Gamma$ iff $I^Y\psi \in \Delta$, *and* $Q^Y\psi \in \Gamma$ iff $Q^Y\psi \in \Delta$;
- $[\Gamma] \approx^{f(\varphi)}_X [\Delta]$ iff for all $Q^X\psi \in \mathrm{cl}(\varphi)$ and $Y \subseteq X$, $Q^Y\psi \in \Gamma$ iff $Q^Y\psi \in \Delta$;
- $\mathcal{A}^{f(\varphi)}_{\Gamma_0} = \mathcal{A}_{\Gamma_0} \cap \{x \in \mathcal{I} \mid Ax \in \mathrm{cl}(\varphi)\}$;
- $||p||^{f(\varphi)}_{\Gamma_0} = \{[\Gamma] \mid p \in \Gamma\}$ for $p \in \Phi \cap \mathrm{cl}(\varphi)$, otherwise $||p||^{f(\varphi)}_{\Gamma_0} = \emptyset$.

When the context permits it, we often write 'f' instead of '$f(\varphi)$'.

Proposition 6.34. *The filtration $\mathcal{M}^{f(\varphi)}_{\Gamma_0}$ of a semi-canonical pseudo model $\mathcal{M}^{\mathsf{EIL}}_{\Gamma_0}$ is a finite pseudo model.*

Proof. It is evident that $\sim^f_X$ and $\approx^f_X$ are equivalence relations by definition. Furthermore, $\mathcal{A}^f$ is an agenda and $||\bullet||^f_{\Gamma_0}$ is a valuation. Since $\mathrm{cl}(\varphi)$ must be finite, there can only be finitely many equivalence classes under $\equiv_\varphi$, so $W^f_{\Gamma_0}$ is finite. In particular, there can at most be $2^{|\mathrm{cl}(\varphi)|}$ equivalence classes in $W^f_{\Gamma_0}$.

For anti-monotonicity, let $X, Y \subseteq \mathcal{I}$ such that $X \subseteq Y$ and suppose that $[\Gamma] \sim^f_Y [\Delta]$. By the definition of closure, $I^Y\psi, Q^Y\psi \in \mathrm{cl}(\varphi)$ iff $I^X\psi \in \mathrm{cl}(\varphi)$. If $Z \subseteq X$, then $Z \subseteq Y$. Thus for all $Z \subseteq X$ and $I^X\psi \in \mathrm{cl}(\varphi)$ we have $I^Z\psi \in \Gamma$ iff $I^Z\psi \in \Delta$, and $Q^Z\psi \in \Gamma$ iff $Q^Z\psi \in \Delta$. Hence $[\Gamma] \sim^f_X [\Delta]$. The argument showing that $\approx^f_Y \subseteq \approx^f_X$ is analogous.

For inclusion, assume that $[\Gamma] \sim_X^f [\Delta]$. Since $I^X\psi \in \mathrm{cl}(\varphi)$ iff $Q^X\psi \in \mathrm{cl}(\varphi)$, it follows that $[\Gamma] \approx_X^f [\Delta]$ by definition.

For necessity, let $[\Gamma], [\Delta] \in W_{\Gamma_0}^f$ be arbitrary. Let $Q^\emptyset\psi \in \mathrm{cl}(\varphi)$ be arbitrary and suppose $Q^\emptyset\psi \in \Gamma$. Then since all EIL-theories in $W_{\Gamma_0}^{\mathsf{EIL}}$ are universally equivalent, $Q^\emptyset\psi \in \Delta$ as well. It follows that $[\Gamma] \approx_\emptyset^f [\Delta]$. $\square$

Next we show that satisfaction of formulas in the closure of a formula φ is invariant under filtration through that closure, thus showing that EIL has the finite model property with respect to pseudo models.

Lemma 6.35. *Let $\mathcal{M}_{\Gamma_0}^f$ be a filtrated semi-canonical pseudo model that is filtrated through φ. For all $\psi \in \mathrm{cl}(\varphi)$ and all $\Gamma \in W_{\Gamma_0}^{\mathsf{EIL}}$, $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash \psi$ iff $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \vDash \psi$.*

Proof. We give a proof by induction on the structure of formulas.

For $\psi := p$, we have $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash p$ iff $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \vDash p$, by definition of $|| \bullet ||_{\Gamma_0}^f$.

For $\psi := Ax$, we find that for $Ax \in \mathrm{cl}(\varphi)$, $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash Ax$ iff $x \in \mathcal{A}_{\Gamma_0}$ iff $x \in \mathcal{A}_{\Gamma_0}^f$ iff $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \vDash Ax$.

For $\psi := \neg\theta$, by applying the induction hypothesis we find that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash \neg\theta$ iff $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \nvDash \theta$ iff $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \nvDash \theta$ iff $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \vDash \psi$. The proof for conjunction is similar.

For $\psi := I^X\theta$, suppose that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash I^X\theta$ and let $\Delta \in W_{\Gamma_0}^{\mathsf{EIL}}$ with $[\Gamma] \sim_X^f [\Delta]$ be arbitrary. By definition of the filtrated model $\mathcal{M}_{\Gamma_0}^f$, this means that $I^X\theta \in \Gamma$ iff $I^X\theta \in \Delta$, because $I^X\theta \in \mathrm{cl}(\varphi)$ by assumption. Since $I^X\theta$ is satisfied at Γ, the Pseudo Truth Lemma gives $I^X\theta \in \Gamma$. Thus $I^X\theta \in \Delta$, which implies that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Delta \vDash \theta$ by factivity of information possession and the Pseudo Truth Lemma. Application of the induction hypothesis then yields $\mathcal{M}_{\Gamma_0}^f, [\Delta] \vDash \theta$, as desired.

Conversely, suppose that $\mathcal{M}_{\Gamma_0}^f, [\Gamma] \vDash I^X\theta$ and take $\Delta \in W_{\Gamma_0}^{\mathsf{EIL}}$ such that $\Gamma \sim_X^{\mathsf{EIL}} \Delta$ arbitrarily. By definition of semi-canonical pseudo models it follows that for all $\chi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ and $Y \subseteq X$, $I^Y\chi \in \Gamma$ iff $I^Y\chi \in \Delta$. Since this is the case for all $\chi \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$, it is also the case for all formulas in the closure $\mathrm{cl}(\varphi)$. We are left with proving that $Q^Y\chi \in \Gamma$ iff $Q^Y\chi \in \Delta$, for all $I^X\chi \in \mathrm{cl}(\varphi)$ and $Y \subseteq X$. Let $I^X\chi \in \mathrm{cl}(\varphi)$ and suppose that $Q^Y\chi \in \Gamma$ for some $Y \subseteq X$. Positive introspection gives $Q^YQ^Y\chi \in \Gamma$, which means that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma \vDash Q^YQ^Y\chi$ by the Pseudo Truth Lemma. Recall that $\Gamma \sim_X^{\mathsf{EIL}} \Delta$, so we have $\Gamma \approx_X^{\mathsf{EIL}} \Delta$ by inclusion. Thus it follows that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Delta \vDash Q^Y\chi$. Applying the Pseudo Truth Lemma again yields $Q^Y\chi \in \Delta$. By symmetry, it also follows that $Q^Y\chi \in \Delta$ implies $Q^Y\chi \in \Gamma$. Thus $[\Gamma] \sim_X^f [\Delta]$ and it follows that $\mathcal{M}_{\Gamma_0}^f, [\Delta] \vDash \theta$. Application of the induction hypothesis then yields $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Delta \vDash \theta$. The case for Q^X is similar. $\square$

Proposition 6.36 (Decidability of EIL)**.** *The logic* EIL *is decidable.*

Proof. Let $\varphi_0 \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ be arbitrary. It suffices to check all finite pseudo models up to size $2^{|\mathrm{cl}(\varphi_0)|}$ for counterexamples. If φ_0 is not valid in EIL, then by completeness there is an EIL-theory Γ_0 such that $\varphi_0 \notin \Gamma_0$. It follows that $\mathcal{M}_{\Gamma_0}^{\mathsf{EIL}}, \Gamma_0 \nvDash \varphi_0$, and by Lemma 6.35 that $\mathcal{M}_{\Gamma_0}^f, \Gamma_0 \nvDash \varphi_0$. Since $\mathcal{M}_{\Gamma_0}^f$ is a finite pseudo model with size at most $2^{|\mathrm{cl}(\varphi_0)|}$, so this procedure will find a countermodel. If φ_0 is valid, then it is a theorem of EIL. Since EIL is sound with respect to pseudo models, the procedure will terminate without finding a countermodel. Hence EIL is decidable. $\square$

Since the languages of the logics from Chapter 5 are included in the language of EIL and the semantics are the same for these formulas, it follows that all these logics are decidable as well.

Corollary 6.37. *The logics* KR*,* KRU*,* IR*,* IRU *and* IQU *are decidable.*

Lastly, it also follows that the dynamic logic DEIL is decidable.

Corollary 6.38. *The logic* DEIL *is decidable.*

Proof. Let $\varphi_0 \in \mathcal{L}_{\mathsf{DEIL}}(\Phi, \mathcal{I})$. By Lemma 6.28 and Lemma 6.29 there exists a formula $\varphi_0' \in \mathcal{L}_{\mathsf{EIL}}(\Phi, \mathcal{I})$ that is equivalent to φ_0. This formula can be obtained by applying the reduction axioms of DEIL finitely many times. The procedure for determining the validity of φ_0' in EIL then suffices to check whether φ_0 is valid in DEIL. So DEIL is decidable. □

6.5 Synthesis

In Chapter 3 we set out to develop a framework in which the notion of knowledge as issue-relevant information could be captured, as well as agenda updates and information updates. Dynamic Epistemic Issue Logic accomplishes all of this: it enables us to update the agent's agenda by adding or retracting issues, and agents may come to possess new information by resolving issues. Moreover, DEIL contains the logic of knowledge and issue-relevance from Section 5.2 as a sublogic, ensuring that knowledge and issue-relevance behave in accordance with the principles we set out in Chapter 4.

The range of possible information updates, however, is limited. In Section 3.3, we discussed how the public announcement of a proposition P can be viewed as the resolution of the issue whether P. So we can simulate public announcements through issue resolutions. Yet, recall that to express knowledge and issue-relevance in DEIL, the number of basic issues available in its language needs to be finite. So only a finite number of issues can be resolved. Consequently, to capture all possible public announcements, DEIL should be extended with a public announcement operator. As of yet, it is an open problem how such an extension of DEIL can be axiomatized.[71]

Nonetheless, for practical purposes DEIL suffices. We illustrate this by returning to one of the examples seen in earlier chapters.

Example 11 (Galileo's telescope: revisited again)**.** Let $\mathcal{I} = \{x, y\}$ and let $\Phi = \{p, q\}$. Interpret x as the issue pertaining to direct observations of the moon and interpret y as the issue pertaining to telescope observations of the moon. Additionally, interpret 'p' as 'the moon is smooth (when observed by the naked eye)' and 'q' as 'the moon contains mountains and craters (when observed with a telescope)'. Galileo's initial epistemic situation is captured by the epistemic issue model $M = (W, \sim, \approx_{x \in \mathcal{I}}, \mathcal{A}, || \bullet ||)$, where:

- $W := \{w_1, w_2, w_3, w_4\}$;
- $\sim := \approx_x := \{(w_1, w_2), (w_3, w_4)\}^*$;
- $\approx_y := \{(w_1, w_3), (w_2, w_4)\}$;
- $\mathcal{A} = \{x\}$;
- $||p|| := \{w_1, w_2\}$ and $||q|| := \{w_1, w_3\}$,

where R^* denotes the reflexive-symmetric closure of R. Let w denote the actual world. The model M is depicted in Figure 6.2 (a), note that Figure 6.2 is almost identical to Figure 4.7. Initially, Galileo only has the issue pertaining to direct observations of the moon on his agenda. In accordance with earlier findings, $M \vDash Rp$ and $M, w \vDash Ip$: direct observations of the moon are issue-relevant and Galileo has directly observed the moon to be smooth. Hence Galileo knows that the moon is smooth when observed by the naked eye; $M, w \vDash Kp$. After inventing his telescope, Galileo shifted his epistemic

[71] This problem is similar to the open problem regarding the axiomatization of the modal logic of conditional dependence of Baltag and Van Benthem (2021, p. 985).

focus from information obtained by direct observations of the moon to information obtained by looking through his telescope. He retracted the issue x from his agenda, the updated model $M_{[-\{x\}]}$ is depicted in Figure 6.2 (b). The smoothness of the moon is no longer issue-relevant and, consequently, he no longer knows the moon to be smooth: $M \vDash [-\{x\}]\neg Rp$ and $M, w \vDash [-\{x\}]\neg Kp$. Galileo replaced the issue x by y, making information obtained by telescope observations of the moon issue-relevant. The updated model $M_{[+\{y\}-\{x\}]}$ is depicted in Figure 6.2 (c), it can be seen that $M \vDash [+\{y\}-\{x\}]Rq$. However, as Galileo has only observed the moon with his naked eye, after changing his agenda he does not possess the information that, when observed with a telescope, the moon contains mountains and craters; $M \vDash [+y-\{x\}]\neg Iq$. To resolve the issue y on his agenda, Galileo observes the moon through his telescope. The updated model $M_{[+\{y\}-\{x\}][\{y\}!]}$ is depicted in Figure 6.2 (d). Now, Galileo possesses the information that the moon contains mountains and craters; $M, w \vDash [+\{y\}-\{x\}][\{y\}!]Iq$. As this information is also issue-relevant, he now knows that the moon contains mountains and craters; $M, w \vDash [+\{y\}-\{x\}][\{y\}!]Kq$.

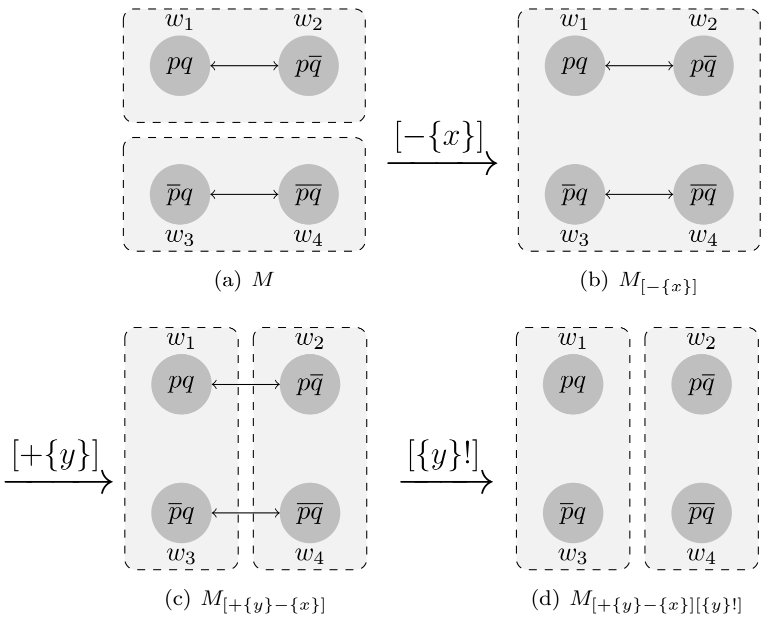

Figure 6.2: Four epistemic issue models. The actual world is denoted by w_1.

The example above shows that DEIL can model paradigm shifts in a satisfactory manner. In fact, the logic DEIL is strong enough to model situations in which a finite number of issues is at play. For most, if not all, practical purposes this is sufficient.

Lastly, we can compare the logics of this chapter with the logics from Chapter 5. Figure 6.3 displays the relations between all the logics defined in this thesis. Every logic from Chapter 5 is contained in the logics EIL and DEIL. Moreover, even though EIL and DEIL are co-expressive, the language of DEIL contains sentences that are not part of the language of EIL.

For any pair of these logics L and L' such that the language of L is contained in the language of L', $\Gamma \vDash_{\mathsf{L}} \varphi$ if, and only if, $\Gamma \vDash_{\mathsf{L}'} \varphi$, for any $\varphi \in \mathcal{L}_{\mathsf{L}}(\Phi, \mathcal{I})$ and $\Gamma \subseteq \mathcal{L}_{\mathsf{L}}(\Phi, \mathcal{I})$. Since each logic is sound and strongly complete it also follows that $\Gamma \vdash_{\mathsf{L}} \varphi$ if, and only if, $\Gamma \vdash_{\mathsf{L}'} \varphi$, for any $\varphi \in \mathcal{L}_{\mathsf{L}}(\Phi, \mathcal{I})$ and $\Gamma \subseteq \mathcal{L}_{\mathsf{L}}(\Phi, \mathcal{I})$.

$$
\begin{tikzcd}
 & \mathsf{KRU} \arrow[dr] & & & & \\
\mathsf{KR} \arrow[ur] \arrow[dr] & & \mathsf{IRU} \arrow[r] & \mathsf{IQU} \arrow[r] & \mathsf{EIL} \arrow[r] & \mathsf{DEIL} \\
 & \mathsf{IR} \arrow[ur] & & & &
\end{tikzcd}
$$

Figure 6.3: A graph showing the relations between the different logics studied in this thesis. An arrow from a logic L to a logic L' indicates that L is a sublogic of L'. Transitive arrows have been omitted.

This concludes our logical study of knowledge as issue-relevant information. We have seen that our framework captures this notion of knowledge as it was introduced in Section 2.2. It addresses the shortcoming of knowledge as information possession, invalidating the principles that were argued to be contentious in Section 2.3. Moreover, the restricted principles put forward in their place are all valid in our framework. Several sound and complete static logics of knowledge, issue-relevance, necessity, information and issues have been defined, elucidating the principles governing the notion of knowledge as issue-relevant information. Moreover, Dynamic Epistemic Issue Logic extends these logics and accommodates the actions of issue addition, issue retraction and issue resolution that were envisioned in Section 2.2. The motivating examples from Chapter 2 can all be formalized in a satisfactory manner. Lastly, all of the logics in this thesis are shown to be decidable.

In the next chapter, we explore some of the connections between our framework and different approaches to knowledge in which inquiry is taken into account. We also suggest directions for further research.

Chapter 7

Related Work and Further Research

In this chapter we compare our logical framework to adjacent work in formal epistemology. Alongside this, we suggest directions for further research. In Section 7.1, we return to the work of Baltag, Boddy and Smets (2018) and consider the possibility of a multi-agent extension of our framework. Potential notions of group knowledge in this extended framework are discussed as well. In Section 7.2, we briefly consider the connection between issue-relevance and relevance in epistemic logics that are based on relevance logics. In Section 7.3, the differences between our framework and frameworks in the spirit of the relevant alternatives approach are considered. We point out how our framework could be enhanced to integrate relevant alternatives, while also warning for pitfalls associated with such an integration. In Section 7.4, knowledge as issue-relevant information is compared to topic-sensitive notions of knowledge, and we suggest how our framework can be adjusted to obtain a topic-sensitive notion of knowledge. In Section 7.5, the parallels between knowledge as issue-relevant information and Yalcin's notion of question-sensitive belief are examined. Furthermore, we describe how our framework can be adapted in order to accommodate belief as issue-relevant information.

7.1 Multi-agent settings and group knowledge

In Section 2.5, we criticized the framework for issue-sensitive knowledge given by Baltag, Boddy and Smets (2018). In particular, we argued against their formalization of issue-relevance and replaced it with our own notion. Subsequently, we proceeded by developing a framework capturing knowledge as issue-relevant information for single agents. However, the primary goal of Baltag, Boddy and Smets (2018) was to develop a notion of group knowledge that is weaker than common knowledge, but stronger than distributed knowledge. They attempted this by investigating information flow in groups of agents who only conceptually distinguish information relevant with respect to their own interrogative agendas. A natural next step would thus be to extend our setting to include multiple agents and investigate group knowledge in this setting. In particular, it would be interesting to see whether the results of Baltag, Boddy and Smets (2018) can be partially recovered.

In standard epistemic logic, generalizing the single agent setting to include multiple agents is straightforward: every agent receives its own epistemic relation, encoding the knowledge in her possession. Generalizing our framework in a similar manner will be more difficult: assigning an information relation and an agenda to each agent would not be enough. Since issue-relevance is defined to be necessary given the agents' current agenda, each agent would know exactly which propositions are issue-relevant for arbitrary other agents. We call this the principle of public agendas. This principle may be admissible in some contexts. For instance, when considering scientific knowledge in an ideal setting: in general, researchers are not only aware of their own research agenda, but also of their colleagues' agendas. Hence in the ideal scientific community, it is issue-relevant for every agent what her colleagues find issue-relevant. Consequently, since issue-relevance is necessary given the agents' agendas, agents know what every other agent finds issue-relevant. This allows ideally rational scientists to effectively share information with their colleagues.

In most contexts, however, the principle of public agendas is too strong. Introspection with respect to one's own agenda is sensible for rational agents; it comes down to being aware of one's own goals and issues. Yet when agents have access to the agendas of all other agents, this can be interpreted as mindreading. In principle, agents should only be able to access each others' agenda through sharing. Failure of the principle of public agendas is especially important when trying to study how the flow of information is restricted in groups. If information processing is shaped and limited by the agenda of an agent, then it only makes sense that this impedes the unrestricted flow of information when agendas are not public. Specifically, if we are to recover some of the results of Baltag, Boddy and Smets (2018), then the principle of public agendas should not hold.

One manner in which our framework can be extended while avoiding public agendas is by assigning an agenda to each agent in each world. Let G be a finite set of agents, W the set of possible worlds and $\mathcal{I}$ the set of basic issues, then the agenda map $\mathcal{A} : G \times W \to \mathcal{P}(\mathcal{I})$ would map each pair consisting of an agent and world to a set of basic issues that constitute the agent's agenda at that world. To ensure that information about the issues on one's own agenda is consistent, we should require that $\mathcal{A}(a, w) = \mathcal{A}(a, v)$ for all $w, v \in W$ and $a \in G$ such that $w \sim_a v$, where $\sim_a$ denotes the information relation of agent a. Since agendas may differ within issue cells, it follows that issue-relevance is no longer necessary given the agents' current agendas. Consequently, it no longer needs to be the case that for every proposition it is issue-relevant for every agent whether that proposition is issue-relevant for any other agent. Thus, the agenda map prohibits agendas from being public.

Introducing such an agenda map, however, would increase the complexity of our models exponentially. Moreover, introspection with respect to one's own agenda would no longer hold, since $RR(P)$ would no longer be valid. As a consequence, the rationality of agents is threatened. So a challenge that could be taken on in the future is to develop a multi-agent extension of our framework in which agents are rational, agendas are *not* public and models are *not* overly complex.

Given an appropriate multi-agent extension of our framework, we may turn to notions of group knowledge. Some potential notions of group knowledge are suggested below.

(i) Distributed knowledge with consensus: P is distributed knowledge with consensus if the information that P is possessed after pooling all agents' information, and P is issue-relevant for each agent in the group;

(ii) Distributed knowledge with pooled agendas: P is distributed knowledge with pooled agendas if the information that P is possessed after pooling all agents' information, and P is issue-relevant with respect to the union of all agents' agendas;

(iii) Knowledge by pooling agendas: let there be a single body of information available to all agents. A group knows that P if P is carried by the available body of information, and P is issue-relevant with respect to the union of all agents' agendas;

(iv) Common knowledge: P is common knowledge at world w iff $K_{a_1}K_{a_2}\cdots K_{a_n}(P)$ is true at world w for every $n \in \mathbb{N}$ and $a_i \in G$ for $1 \leq i \leq n$.

Distributed knowledge with consensus can be used to model phenomena like scientific consensus: researchers share all their information and the pooled information that (partially) resolves the issues on which all researchers agree constitute scientific knowledge on which there is consensus. On the other hand, distributed knowledge with pooled agendas may apply in situations in which agents cooperate: besides their information, they also share their issues and goals. The group then collectively tries to resolve the issues of every agent in the group. Knowledge obtained by pooling agendas, (iii), is a special instance of (ii) in which the information relations of all agents in the group are identical. It is already captured by epistemic issue models: interpret $\mathcal{I}$ as a set of agents and each $\approx_x$ as the current issue relation for each agent $x \in \mathcal{I}$. The information relation $\sim$ represents the body of information available to all agents. The modality R then captures issue-relevance with respect to the union of the agendas of agents in the group $\mathcal{A} \subseteq \mathcal{I}$. Thus the modality K captures possessed information that is issue-relevant with respect to the union of the agendas of agents in $\mathcal{A}$, and so K corresponds to knowledge obtained by pooling agendas. Lastly, common knowledge, (iv), is defined in the usual manner. Note, however, that this notion is very strong: the proposition P not only needs to hold in any world that can be accessed by a path over the union of every agent's information relation, but for every agent $a \in G$ the proposition $K_{a_1}K_{a_2}\cdots K_{a_n}(P)$ also needs to be issue-relevant, for every $n \in \mathbb{N}$ and $a_i \in G$ for $1 \leq i \leq n$. It is doubtful that investigating such a strong notion of group knowledge is worthwhile.

The four notions above could serve as starting points for further research. However, each of them relies on pooling information and agendas separately. In contrast, group knowledge as conceived by Baltag, Boddy and Smets (2018) relies on agents sharing all they *know*. So another research direction would be to see whether we can mimic this process in a multi-agent extension of our framework. Firstly, information should not be pooled directly, but only after the information of each agent is "rounded off" to fit their current agenda. Secondly, it should be decided which information shared with the group constitutes group knowledge; an analogue of issue-relevance is needed for group knowledge. Two options for this are already given above: we may look at propositions relevant with respect to the union or intersection of all agents' agendas. In any case, much work has to be done before the notion of group knowledge of Baltag, Boddy and Smets (2018) can be studied in an extension of our framework.

In summary, finding a multi-agent extension of our framework and exploring different notions of group knowledge present interesting paths for further research. Extending our framework is not trivial if overly complex models and the principle of public agendas are to be avoided, while maintaining agents' rationality. We suggested some notions of group knowledge that may be worthwhile to explore, ranging from notions of distributed knowledge, to common knowledge. Yet these notions rely on sharing information rather than sharing knowledge, which would be the ultimate goal if we are to recover some of the results of Baltag, Boddy and Smets (2018).

7.2 Relevance and truth

The reader should be wary when relating the notion of issue-relevance to earlier work on relevance in formal epistemology. In particular when it comes to epistemic logics based on a relevance logic, such as in the following examples. Wansing (2002) introduces a relevant epistemic logic to dissolve Fitch's knowability paradox.[72] Bilkova et al. (2010)

[72] Brogaard and Salerno (2019) provide a thorough introduction to the knowability paradox.

replace normal modal logic by a weaker relevance logic to bypass closure under deduction, among other strong properties of epistemic logic based on **S5**. A relevance logic of questions and answers is developed by Punčochář (2020).

Relevance logics focus on the concept of relevance in reasoning and inference. In particular, they require the antecedent and consequent of an implication to be be relevantly related in order for an implication to hold.[73] This avoids certain paradoxes of strict and material implication.[74] However, the principle of bivalence, which states that every proposition has exactly one truth value, is violated by such logics. Thus these logics are non-classical.

This non-classicality is a direct consequence of the fundamental idea underlying relevance logics: that relevance plays a role in truth itself. The logical framework that we have set forth, however, is classical. Our notion of issue-relevance does not restrict closure under deduction by requiring consequences to be sufficiently relevant to the known antecedents *in themselves*. We only require consequences to be relevant with respect to the agent's issue(s), hence implications without relevantly related antecedents and consequents may be true and known in the framework we have put forward. Thus, issue-relevance should not be confused with the notion of relevance employed in relevance logics.

7.3 The relevant alternatives approach

In Section 2.2, we juxtaposed our notion of knowledge as issue-relevant information with relevant alternatives notions of knowledge. We argued that our notion of knowledge should be considered a relevant *distinctions* notion rather than a relevant *alternatives* notion: the issues in our framework only prompt agents to make conceptual distinctions, not to exclude worlds deemed irrelevant. Epistemic logics in the spirit of the relevant alternatives approach typically contain a mechanism that allows an agent to know P without P being true in all worlds that an agent cannot informationally distinguish.[75]

We have seen that our framework is able to capture a wide range of examples, but the relevant alternatives approach has its own merits. For instance, relevant alternatives theorists such as Schaffer (2006) block closure under logical consequence by allowing worlds to be excluded as irrelevant. Therefore, as we explained in Section 4.3, skeptical paradoxes can be solved within the relevant alternatives approach, while this cannot be accomplished in our framework. Moreover, although our framework can model scientific phenomena like paradigm shifts, it is not able to accurately model the background assumptions in a paradigm. This is possible within the relevant alternatives approach: an inquiry that presupposes P causes the agent to neglect all $\neg P$-worlds. So, there is something to gain by integrating the relevant alternatives approach in our framework.[76]

Prima facie, if we assume that inquiry determines which alternatives are relevant, there is a manner to integrate the relevant alternatives approach in our framework. Instead of partitions of logical space, redefine epistemic issues as partitions of a subset of logical space. If an issue x partitions a subset $V \subseteq W$, then V is the set of relevant worlds with respect to the issue x. Then the subset that is partitioned determines which possible worlds are relevant, whereas the manner in which this subset is divided into issue cells determines which conceptual distinctions are relevant. Compound issues can still be defined as intersections of basic issue relations. After redefining issues in this manner, the framework and logics of this thesis can be redeveloped. This would yield a formal framework of knowledge as issue-relevant information, in which both relevant

[73] Priest (2008, Ch. 9–10) offers an introduction to relevance logics.
[74] Mares (2022) gives a comprehensive introduction to these paradoxes.
[75] For instance, see the logics put forward by Holliday (2012), or Xu and Chen (2018).
[76] An overview of other merits of relevant alternatives notions of knowledge is given by Hawke (2016).

distinctions and relevant alternatives are taken into account. In addition to the examples treated in this thesis, examples involving skeptical paradoxes and (scientific) background assumptions could also be formalized in this adapted framework.

Yet, implementing this modification may not be so simple. Firstly, the presuppositions of issues would not only influence which possible worlds are relevant, but also which distinctions are. For example, when an issue excludes all possible worlds in which an agent is a brain in a vat, it becomes necessary (by presupposition) that the agent is *not* a brain in a vat if that issue is on her agenda. Consequently, being a brain in a vat also becomes issue-relevant. This may not be problematic, as it makes sense that an issue presupposing you are not a brain in a vat prompts the agent to deem this issue-relevant. However, the informational content of issues may also affect which propositions are issue-relevant in a less trivial manner. It may be that presupposing P to be irrelevant makes another proposition Q issue-relevant, since excluding all P-worlds from issue-cells may make Q fit the agent's issue. Thus, presuppositions obscure why some propositions are issue-relevant.

Secondly, on this approach, an agent's agenda could rule out more worlds than would be desirable. If it turns out that you are a brain in a vat, then the actual world is ruled out if you know that you have hands. So knowledge would, on first sight, no longer be factive. Moreover, it might be the case that all worlds are ruled out by unsuitable agendas. Suppose an agent has two questions on her agenda: 'is mom home or is dad home?' and 'why is nobody home?'. The first question presupposes that either her mom is home or her dad is home, whereas the second question presupposes that nobody is home. In tandem, they rule out all possible worlds, breaking the model.

Thirdly, not all relevant alternatives theories are explained in terms of inquiry and issues. For instance, some relevant alternatives theories take knowledge to be context-sensitive. According to these theories, it is the context that determines which alternatives are relevant.[77] To overcome this, the agenda in our framework should be reinterpreted so that it becomes the set of basic issues matching the current context.

We should briefly mention another possible approach. When assuming that issues also have informational content, they become more like the issues in inquisitive semantics. Recall that in Section 3.2 we explained that there is no technical reason to disallow issue cells to overlap: when allowing overlap, the issue relation only needs to be reinterpreted as something other than a conceptual indistinguishability relation. So we could also redefine issues as sets of maximal sets of worlds that cover a part of logical space. Observe that these issues still differ from the ones in inquisitive semantics, since they are not downwards closed. However, an issue in this sense can be identified with the set of alternatives of an issue in inquisitive semantics. Hence, issues can then rightly be interpreted as a set of relevant alternatives. Redefining issues in this manner and subsequently redeveloping our logical framework would yield an epistemic logic of relevant alternatives.

There is no guarantee that either approach yields a logical framework that fully encompasses any established relevant alternatives theory and its merits. Further research is required in order to ensure that any integration of relevant alternatives into our framework properly formalizes the core ideas of the corresponding relevant alternatives theory.

So, while our notion of knowledge emphasizes conceptual distinctions prompted by epistemic issues, the relevant alternatives approach focuses primarily on which worlds are ruled out by issues. The relevant alternatives approach may be integrated in our framework by redefining issues as partitions of subsets of logical space or as sets of possibly overlapping alternatives that cover part of logical space. Implementing either of these modifications may pose challenges, such as the risk of the informational content

[77]For instance, Lewis (1996) argues that the set of relevant alternatives is context-dependent.

of issues affecting which conceptual distinctions should be made and overly restrictive agendas ruling out the actual world or all possible worlds. Furthermore, there is no guarantee that either of these approaches fully embodies a relevant alternatives theory and its merits. Nonetheless, a successful implementation could give us a better logical understanding of epistemic relevance.

7.4 Topic-sensitive notions of knowledge

Issues can be interpreted as subject matters. Hence the notion of knowledge studied in this thesis can justifiably be viewed as sensitive to subject matters. However, our agents are only able to distinguish the subject matters of formulas in terms of their intension. Therefore knowledge and issue-relevance are closed under logical equivalence: agents cannot distinguish between intensionally equivalent formulas. However, recall that topic-sensitive notions of knowledge, as we discussed them in Sections 2.1 and 2.2, are hyperintensional: there is a hyperintensional condition that φ needs to satisfy, besides the agent possessing the information that φ, before the agent knows that φ. This is mirrored in our notion of knowledge: in addition to the usual condition that φ must be informationally possessed, we impose the extra condition that φ is issue-relevant. Since 'issue-relevant' can be interpreted as 'relevant with respect to the subject matters on the agent's agenda', our approach can be considered an intensional sibling of the topic-sensitive notions of knowledge.

Like our notion of knowledge as issue-relevant information, topic-sensitive notions of knowledge are not closed under logical consequence. Hence, topic-sensitive notions also block the logical omniscience of agents. However, our agents do know all tautological formulas. This property, it is argued, is too strong for non-ideally rational agents.[78] The additional hyperintensional condition sees to it that the principle of necessitation is blocked for knowledge: not all tautological propositions are known under topic-sensitive notions of knowledge. So, in general, topic-sensitive notions of knowledge are introduced in an attempt to model non-ideal agents incapable of making some distinctions. In contrast, we model ideal agents that rationally *choose* to not make some conceptual distinctions. Hence our approach diverges from topic-sensitive notions of knowledge on this matter.

Nonetheless, our notion of issue-relevance can be adjusted so that it becomes hyperintensional. Say that a formula φ is hyperintensionally issue-relevant if not only the set of worlds corresponding to φ fits the agent's current issue, but also every set of worlds corresponding to a subformula of φ. In terms of our notion of issue-relevance: φ is hyperintensionally issue-relevant if, and only if, every subformula of φ is issue-relevant. It is evident that for a Boolean formula φ this is equivalent to every proposition letter occurring in φ being issue-relevant, since issue-relevance is closed under negation and conjunction. For formulas involving the modality K this is not the case, otherwise positive and negative introspection would have been valid in our framework.

Hyperintensional issue-relevance respects the intuition that the topic of a sentence should be downwards closed syntactically. If an agent is interested in the subject matter of the sentence φ, then, intuitively, all subformulas of φ must be about a subject matter that is contained in the subject matter of φ. For instance, if $\varphi = p \vee \neg p$, then φ is issue-relevant, but p and $\neg p$ need not be; they may cut across one of the agent's current issue cells. Hence φ need not be hyperintensionally issue-relevant: the principle of necessitation does not hold for hyperintensional issue-relevance.

[78] See, for instance, the arguments given by Hawke, Özgün and Berto (2019), or Berto and Hawke (2021).

Call a formula φ hyperintensionally known if the agent possesses the information that φ, and φ is hyperintensionally issue-relevant. Hyperintensionality is then inherited from hyperintensional issue-relevance, so hyperintensional knowing is rightly called hyperintensional. Consequently, an agent may hyperintensionally know simple tautologies such as '$2 + 2 = 4$', while not hyperintensionally knowing more complex tautologies such as 'every differentiable function is continuous'. Therefore, under the notion of hyperintensional knowledge, agents are no longer ideally rational. Furthermore, failure of closure under logical consequence and failure of both positive and negative introspection are inherited from issue-relevance.

Yet, there is a complication: if we only make the minimal changes sketched above, agents are still fully aware of their agendas. Atoms of the form 'Ax' are always true or false in an entire epistemic issue model, hence always hyperintensionally issue-relevant. It may be argued that non-ideal agents should not always be aware of the issues on their agenda. Hence, if we want to avoid this, we need to ensure that the agent's agenda is not fully accessible to her. This problem is similar to the problem of public agendas in multi-agent settings, discussed in Section 7.1.

In summary, topic-sensitive notions of knowledge can be considered the hyperintensional siblings of our intensional notion of knowledge. They also invalidate closure of knowledge under logical consequence. However, the agents in our framework are ideal, whereas topic-sensitive notions of knowledge are typically introduced to model non-ideal agents. This difference is caused by the additional hyperintensional condition put on knowledge. One direction for further research would thus be to replace issue-relevance by a hyperintensional notion, such as the one given above. The resulting notion of knowledge is topic-sensitive, but further modifications are required if we want to correctly model non-ideal agents.

7.5 Logics of belief

In Section 2.2 we mentioned that, to a great extent, our notion of knowledge as issue-relevant information resembles the notion of question-sensitive belief of Yalcin (2018). Our issues correspond to his resolutions, both are partitions of logical space. The propositions that he calls 'foregrounded by a resolution' are what we call 'issue-relevant'.[79] An agent believes that P whenever the agent possesses the information that P, and P is foregrounded by the agent's resolution. So, this mirrors the truth conditions for knowledge as issue-relevant information.

Yalcin introduces question-sensitive belief as a means to address philosophical problems in the philosophy of content and, in particular, the problem of logical omniscience. To this end, he only sketches doxastic models and considers belief from a semantic perspective. A logic of this notion of question-sensitive belief is lacking as of yet. A direction for further research would thus be to make modifications to our framework in order to obtain a logic of Yalcin's question-sensitive belief. This logic may contribute to a better understanding of the philosophical questions treated by Yalcin.

Making the required modifications should be easy. Our issues can be left as they are, whereas the information relation should be adjusted in order to have the properties that correspond to belief. Yalcin does not specify what these properties should be, so we have some leeway here. We could require the information relation, now interpreted as a doxastic relation, to be serial and transitive.[80] Or we could follow the more standard approach of Fagin et al. (1995) and also require the information relation to be Euclidean.

[79] He also considers a more local notion that would correspond more closely to the notion of issue-relevance considered by Baltag, Boddy and Smets (2018).

[80] As suggested, for example, by Hendricks and Rendsvig (2018), who claim that this would correspond to the notion of belief put forward by Hintikka (1962).

Ultimately, it does not matter, as it is well known how logics on modal structures with these properties can be axiomatized. Furthermore, the proofs in Chapters 5 and 6 should carry over after making minor adjustments.

Counterparts of our epistemic framework in which the information relation has different properties are not condemned to be mere supplements to Yalcin's philosophical work; they can be full-fledged dynamic logics of belief or knowledge in themselves. For instance, we can relax the assumption that all information provided to agents is true and drop the reflexivity constraint on the information relation. The resulting logic is interesting in itself, as it allows us to also model situations in which measuring instruments are faulty or agents are lied to. Furthermore, a notion of belief as issue-relevant information can potentially address the doxastological counterparts of the problems discussed in Chapter 2, which motivated us to consider knowledge as issue-relevant information in the first place.[81]

[81] In fact, Example 1 was originally introduced by Stalnaker (1984) as a problem concerning belief.

Chapter 8

Conclusion

We started this thesis by observing that many philosophers and logicians have grown dissatisfied with equating knowledge and information possession—even in ideal settings. Several proposals for notions of knowledge that go beyond mere information possession have been proposed and studied by others, and we set out to follow this line of research.

In our view, inquiry is a fundamental activity when it comes to knowledge acquisition. We saw that this idea is not novel and can be discerned throughout the history of philosophy, but that it is not a banality commonly presupposed by philosophers. Additionally, we saw that inquiry still plays a role in contemporary formal epistemology.

We have taken issues to be the objects of inquiry, where issues are interpreted as requests for information. They prompt agents to conceptually distinguish between worlds that disagree on the truth of a proposition that is relevant to them. Questions and subject matters were identified as prototypical issues: a question prompts an agent to conceptually distinguish its answers, while a subject matter prompts agents to conceptually distinguish information about that subject matter. We defined the agenda of an agent as the set of issues that she pursues. We then defined knowledge as information that is possessed by an agent, and that is relevant with respect to the issues on the agent's agenda: knowledge as issue-relevant information.

This notion of knowledge is similar to the one introduced by Baltag, Boddy and Smets (2018). However, we argued that their approach does not formalize issue-relevance in a satisfactory manner: their agents can come to know propositions that are not answers to their questions, and the issue relation does not represent conceptual indistinguishability. In response, we proposed that the entire issue relation should be taken into account when determining whether a proposition is issue-relevant.

Subsequently, we fleshed out the key notions necessary to formulate a substantive logical theory: information, issues and issue-relevance. Information and issues were encoded in epistemic issue structures, on which issue-relevance and knowledge were defined. Moreover, we defined updates of epistemic issue structures corresponding to the actions of issue addition, issue retraction and issue resolution.

Contentious principles common in standard epistemic logic, such as closure under logical consequence, positive introspection and negative introspection were shown to be invalidated in our framework. We argued that this was desirable, since this allows room for reasoning, avoids clutter and respects the intuition that the intensional content of 'knowing that' and 'knowing that one knows that' differ. Instead of these traditional principles, our framework validates restricted, more realistic laws of knowledge: for instance, closure under logical consequence and introspection are restricted depending on the issue-relevance of the proposition in question. As a consequence, all the motivating examples could be formalized neatly.

We introduced several logics on epistemic issue models. The logics KR, KRU, IR, IRU and IQU captured our notion of knowledge as issue-relevant information, but only did so in a static setting without taking the individual basic issues on the agent's agenda into account. Each of these logics was proven to be sound and complete, as well as decidable. The static logic EIL contained all of these logics as sublogics, while also taking individual basic issues into account. Furthermore, its dynamic extension DEIL was seen to accommodate both agenda and information updates. Both these logics were shown to be sound, complete and decidable.

The logics in this thesis are limited to settings with a finite number of basic issues, putting some limitations on the range of possible issues and information updates. However, they can model all of the examples that were introduced in Chapter 2. Furthermore, we argued that, despite this limitation, these logics suffice for practical purposes. Thus we developed a logical framework that captures knowledge as issue-relevant information in a satisfactory manner.

There still is more work that can be done. First and foremost, we could extend our framework to a multi-agent setting, and explore different notions of group knowledge in this extended framework. Secondly, the framework can potentially be adapted to obtain a framework that integrates the relevant alternatives approach. Thirdly, a hyperintensional notion of issue-relevance can be defined, which would bring about a topic-sensitive notion of knowledge. Interesting results may be obtained when investigating the logic corresponding to this hyperintensional notion of knowledge. Finally, by making slight adjustments to the logics in this thesis we can obtain logics for belief as issue-relevant information.

In conclusion, in this thesis we have successfully developed a logical framework that formalizes knowledge as issue-relevant information. It challenges the notion of knowledge as information possession and invalidates epistemic principles that hold in standard epistemic logic. As a consequence, it allows us to model phenomena that cannot be modeled in standard dynamic epistemic logic, such as paradigm shifts. Several sound and complete logics capturing knowledge as issue-relevant information were defined, some of which are also able to neatly capture the actions of agenda update and issue resolution. Lastly, the framework presented in this thesis provides fruitful ground for further research.

Bibliography

Annis, David (1969). “A Note on Lehrer’s Proof That Knowledge Entails Belief”. In: *Analysis* 29.6, pp. 207–208.

Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge University Press.

Aucher, Guillaume (2014). “Principles of Knowledge, Belief and Conditional Belief”. In: *Interdisciplinary Works in Logic, Epistemology, Psychology and Linguistics*. Ed. by Manuel Rebuschi, Martine Batt, Gerhard Heinzmann, Franck Lihoreau, Michel Musiol and Alain Trognon. Springer International Publishing, pp. 97–134.

Bacon, Francis ([1620] 1902). *Novum Organum*. Ed. by Joseph Devey. P.F. Collier.

— ([1620–1626] 2017). *New Atlantis and the Great Instauration*. Ed. by Jerry Weinberger. Second Edition. Chichester, West, Sussex, UK: Wiley-Blackwell.

Baltag, Alexandru and Johan van Benthem (2021). “A Simple Logic of Functional Dependence”. In: *Journal of Philosophical Logic* 50.5, pp. 939–1005.

Baltag, Alexandru, Rachel Boddy and Sonja Smets (2018). “Group Knowledge in Interrogative Epistemology”. In: *Jaakko Hintikka on Knowledge and Game Theoretical Semantics*. Ed. by Hans van Ditmarsch and Gabriel Sandu. Springer, pp. 131–164.

Baltag, Alexandru, Hans van Ditmarsch and Larry S. Moss (2008). “Epistemic Logic and Information Update”. In: *Philosophy of Information*. Ed. by Pieter Adriaans and Johan van Benthem. MIT Press.

Baltag, Alexandru and Bryan Renne (2016). “Dynamic Epistemic Logic”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

Belnap, Nuel D. and Thomas B. Steel (1976). *The Logic of Questions and Answers*. New Haven: Yale University Press.

Berto, Francesco (2019). “Simple Hyperintensional Belief Revision”. In: *Erkenntnis* 84.3, pp. 559–575.

— (2022). *Topics of Thought. The Logic of Knowledge, Belief, Imagination*. Oxford: Oxford University Press.

Berto, Francesco and Daniel Nolan (2023). “Hyperintensionality”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Winter 2023. Metaphysics Research Lab, Stanford University.

Berto, Francesco and Aybüke Özgün (2023). “The Logic of Framing Effects”. In: *Journal of Philosophical Logic* 52.3, pp. 939–962.

Berto, Franz and Peter Hawke (2021). “Knowability Relative to Information”. In: *Mind* 130.517, pp. 1–33.

Bilkova, Marta, Ondrej Majer, Michal Pelis and Greg Restall (2010). “Relevant Agents”. In: *Advances in Modal Logic*. Vol. 8, pp. 22–38.

Blaauw, Martijn (2012). “Contrastive Belief”. In: *Contrastivism in Philosophy: New Perspectives*. Ed. by Martijn Blaauw. Routledge.

Blackburn, Patrick, Maarten de Rijke and Yde Venema (2001). *Modal Logic*. New York: Cambridge University Press.

Boddy, Rachel (2014). “Epistemic Issues and Group Knowledge”. MA thesis. University of Amsterdam.

Brogaard, Berit and Joe Salerno (2019). “Fitch’s Paradox of Knowability”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University.

Carballo, Alejandro Pérez (2014). “Structuring Logical Space”. In: *Philosophy and Phenomenological Research* 92.2, pp. 460–491.

Ciardelli, Ivano (2016). “Questions in Logic”. PhD thesis. University of Amsterdam.

Ciardelli, Ivano, Jeroen Groenendijk and Floris Roelofsen (2018). *Inquisitive semantics*. Oxford University Press.

Collingwood, R. G. (1940). *An Essay on Metaphysics*. Oxford: Clarendon Press.

— (1946). *The Idea of History*. Oxford: Clarendon Press.

Dewey, John (1910). *How We Think*. London and Boston: D.C. Heath.

Dretske, Fred I. (1970). “Epistemic Operators”. In: *Journal of Philosophy* 67.24, pp. 1007–1023.

— (2005). “The Case Against Closure”. In: *Contemporary Debates in Epistemology*. Ed. by Matthias Steup and Ernest Sosa. Blackwell, pp. 13–25.

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses and Moshe Vardi (1995). *Reasoning About Knowledge*. MIT Press.

Feldman, Richard (1981). “Fallibilism and Knowing That One Knows”. In: *Philosophical Review* 90.2, pp. 266–282.

Friedman, Jane (2019). “Inquiry and Belief”. In: *Noûs* 53.2, pp. 296–315.

— (2020). “The Epistemic and the Zetetic”. In: *Philosophical Review* 129.4, pp. 501–536.

Gadamer, Hans-Georg (1975). *Truth and Method*. Ed. by Joel Weinsheimer and Donald G. Marshall. Second, Revised Edition. New York: Continuum.

Goodman, Nelson ([1955] 1983). *Fact, Fiction, and Forecast*. Fourth Edition. Cambridge, Massachusetts: Harvard University Press.

Groenendijk, Jeroen and Martin Stokhof (1984). “dies on the Semantics of Questions and the Pragmatics of Answers”. PhD thesis. University of Amsterdam.

Habgood-Coote, Joshua (2022). “Knowing More (About Questions)”. In: *Synthese* 200.1, pp. 1–23.

Habgood-Coote, Joshua, Lani Watson and Dennis Whitcomb (2022a). “Can a Good Philosophical Contribution Be Made Just by Asking a Question?” In: *Metaphilosophy* 54.1, pp. 54–54.

— (2022b). “Commentary on “Can a Good Philosophical Contribution Be Made Just by Asking a Question?”” In: *Metaphilosophy* 54.1, pp. 55–60.

Halpern, Joseph Y. (2004). “Intransitivity and vagueness”. In: *Principles of Knowledge Representation and Reasoning: Proceedings of KR2004*. Ed. by Didier Dubois, Christopher Welty and Mary-Anne Williams. AAAI Press.

Hamblin, C. L. (1958). “Questions”. In: *Australasian Journal of Philosophy* 36.3, pp. 159–168.

Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. Cambridge, MA, USA: MIT Press.

Hawke, Peter (2016). “Relevant Alternatives in Epistemology and Logic”. In: *Epistemology, Knowledge and the Impact of Interaction*. Ed. by Ángel Nepomuceno Fernández, Olga Pombo Martins and Juan Redmond. Springer Verlag.

— (2017). “The Problem of Epistemic Relevance”. PhD thesis. Stanford University.

Hawke, Peter, Aybüke Özgün and Francesco Berto (2019). “The Fundamental Problem of Logical Omniscience”. In: *Journal of Philosophical Logic* 49.4, pp. 727–766.

Hawthorne, John (2005). “The Case for Closure”. In: *Contemporary Debates in Epistemology*. Ed. by Matthias Steup and Ernest Sosa. Blackwell, pp. 26–43.

Heisenberg, Werner (1927). “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik”. In: *Zeitschrift für Physik* 43, pp. 172–198.
— (1983). “The Physical Content of Quantum Kinematics and Mechanics”. In: *Quantum Theory and Measurement*. Princeton University Press, pp. 62–84.
Hendricks, Vincent and Rasmus Rendsvig (2018). “Hintikka’s Knowledge and Belief in Flux”. In: *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics*. Ed. by Hans van Ditmarsch and Gabriel Sandu. Springer, pp. 317–337.
Hendricks, Vincent F. (2005). *Mainstream and Formal Epistemology*. New York: Cambridge University Press.
Hilpinen, Risto (1970). “Knowing That One Knows and the Classical Definition of Knowledge”. In: *Synthese* 21.2, pp. 109–132.
Hintikka, Jaakko (1962). *Knowledge and Belief*. Ithaca, N.Y., Cornell University Press.
— (1970). “’Knowing That One Knows’ Reviewed”. In: *Synthese* 21.2, pp. 141–162.
— (1981). “On the Logic of an Interrogative Model of Scientific Inquiry”. In: *Synthese* 47.1, pp. 69–83.
— (1998). “What Is Abduction? The Fundamental Problem of Contemporary Epistemology”. In: *Transactions of the Charles S. Peirce Society* 34.3.
— (1999). *Inquiry as Inquiry: A Logic of Scientific Discovery*. Kluwer Academic.
— (2006). *Analyses of Aristotle*. Springer.
— (2007). *Socratic Epistemology: Explorations of Knowledge-Seeking by Questioning*. New York: Cambridge University Press.
Hoek, Daniel (2019). “The Web of Questions: Inquisitive Decision Theory and the Bounds of Rationality”. PhD thesis. New York University.
— (2022). “Questions in Action”. In: *Journal of Philosophy* 119.3, pp. 113–143.
— (forthcoming). “Minimal Rationality and the Web of Questions”. In: *Unstructured Content*. Ed. by Dirk Kindermann, Peter van Elswyk and Andy Egan. Oxford University Press.
Holguin, Ben (2022). “Thinking, Guessing, and Believing”. In: *Philosophers’ Imprint* 22.1, pp. 1–34.
Holliday, Wesley H. (2012). “Epistemic Logic, Relevant Alternatives, and the Dynamics of Context”. In: *New Directions in Logic, Language and Computation*. Ed. by Daniel Lassiter and Marija Slavkovik. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 109–129.
— (2015). “Fallibilism and Multiple Paths to Knowledge”. In: *Oxford Studies in Epistemology* 5, pp. 97–144.
— (2018). “Epistemic Logic and Epistemology”. In: *Handbook of Formal Philosophy*. Ed. by Vincent F. Hendricks and Sven Ove Hansson. Springer, pp. 351–369.
Hookway, Christopher (2008). “Questions, Epistemology, and Inquiries”. In: *Grazer Philosophische Studien* 77.1, pp. 1–21.
Hulstijn, J. (1997). “Structured Information States: Raising and Resolving Issues”. In: *Proceedings of Mundial ’97, Formal Semantics and Pragmatic of Dialogue*. Ed. by G. Jager and A. Benz, pp. 99–117.
James, William ([1907] 1979). *Pragmatism*. Cambridge, MA: Harvard University Press.
Kant, Immanuel ([1787] 1998). *Critique of Pure Reason*. Ed. by Paul Guyer and Allen W. Wood. The Cambridge Edition of the Works of Immanuel Kant. Cambridge University Press.
Koralus, Philipp and Salvador Mascarenhas (2013). “The Erotetic Theory of Reasoning: Bridges Between Formal Semantics and the Psychology of Deductive Inference”. In: *Philosophical Perspectives* 27.1, pp. 312–365.
Kripke, Saul A. (2011). “Nozick on Knowledge”. In: *Philosophical Troubles. Collected Papers Vol I*. Ed. by Saul A. Kripke. Oxford University Press.

Kuhn, Thomas Samuel (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Lawlor, Krista (2013). *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.

Lehrer, Keith (1974). *Knowledge*. Clarendon Press.

Lenzen, Wolfgang (1978). *Recent Work in Epistemic Logic*. Acta philosophica Fennica. North-Holland Publishing Company.

Lévy, Azriel (1979). *Basic Set Theory*. Springer-Verlag.

Lewis, David (1983). "New Work for a Theory of Universals". In: *Australasian Journal of Philosophy* 61.4, pp. 343–377.

— (1988a). "Relevant Implication". In: *Theoria* 54.3, pp. 161–174.

— (1988b). "Statements Partly About Observation". In: *Philosophical Papers* 17.1, pp. 1–31.

— (1996). "Elusive Knowledge". In: *Australasian Journal of Philosophy* 74.4, pp. 549–567.

Mares, Edwin (2022). "Relevance Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2022. Metaphysics Research Lab, Stanford University.

Minică, Ştefan (2008). "Dynamic Logic of Questions". PhD thesis. University of Amsterdam.

Moore, George Edward (1939). "Proof of an External World". In: *Proceedings of the British Academy* 25.5, pp. 273–300.

Mott, Peter (1998). "Margins for Error and the Sorites Paradox". In: *The Philosophical Quarterly* 48.193, pp. 494–504.

Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Özgün, Aybüke and Francesco Berto (2021). "Dynamic Hyperintensional Belief Revision". In: *The Review of Symbolic Logic* 3, pp. 766–811.

Plato (1997). *Plato: Complete Works*. Ed. by J. M. Cooper. Hackett.

Priest, Graham (2008). *An Introduction to Non-Classical Logic: From If to Is*. New York: Cambridge University Press.

Punčochář, Vít (2020). "A Relevant Logic of Questions". In: *Journal of Philosophical Logic* 49.5, pp. 905–939.

Rendsvig, Rasmus and John Symons (2021). "Epistemic Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2021. Metaphysics Research Lab, Stanford University.

Robinson, Richard (1971). "The Concept of Knowledge". In: *Mind* 80.317, pp. 17–28.

Rossi, Niccolò and Aybüke Özgün (2023). "A Hyperintensional Approach to Positive Epistemic Possibility". In: *Synthese* 202.44, pp. 1–29.

Schaffer, Jonathan (2004). "From Contextualism to Contrastivism". In: *Philosophical Studies* 119.1-2, pp. 73–104.

— (2006). "Contrastive Knowledge". In: *Oxford Studies in Epistemology* 1.

— (2007a). "Closure, Contrast, and Answer". In: *Philosophical Studies* 133.2, pp. 233–255.

— (2007b). "Knowing the Answer". In: *Philosophy and Phenomenological Research* 75.2, pp. 383–403.

Sequoiah-Grayson, Sebastian and Luciano Floridi (2022). "Semantic Conceptions of Information". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2022. Metaphysics Research Lab, Stanford University.

Sider, Theodore (2011). *Writing the Book of the World*. Oxford, England: Oxford University Press.

Spector, David (2013). "Margin for error semantics and signal perception". In: *Synthese* 190.15, pp. 3247–3263.

Stalnaker, Robert C. (1984). *Inquiry.* Cambridge University Press.
— (1991). "The Problem of Logical Omniscience, I". In: *Synthese* 89.3, pp. 425–440.
Stine, G. C. (1976). "Skepticism, Relevant Alternatives, and Deductive Closure". In: *Philosophical Studies* 29.4, pp. 249–261.
Van Benthem, Johan and Ştefan Minică (2012). "Toward a Dynamic Logic of Questions". In: *Journal of Philosophical Logic* 41.4, pp. 633–669.
Van Ditmarsch, Hans, Joseph Y. Halpern, Wiebe van der Hoek and Barteld Kooi, eds. (2015). *Handbook of Epistemic Logic.* College Publications.
Van Ditmarsch, Hans, Wiebe Van der Hoek and Barteld Kooi (2008). *Dynamic epistemic logic.* Springer.
Von Wright, Georg Henrik (1951). *An Essay in Modal Logic.* Amsterdam, Netherlands: North-Holland Publishing Company.
Wansing, Heinrich (2002). "Diamonds Are a Philosopher's Best Friends". In: *Journal of Philosophical Logic* 31.6, pp. 591–612.
Williamson, Timothy (2000). *Knowledge and its Limits.* New York: Oxford University Press.
Xu, Zhaoqing and Bo Chen (2018). "Epistemic Logic with Evidence and Relevant Alternatives". In: *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics.* Ed. by Hans van Ditmarsch and Gabriel Sandu. Springer, pp. 535–555.
Yablo, Stephen (2014). *Aboutness.* Princeton University Press.
Yalcin, Seth (2008). "Modality and Inquiry". PhD thesis. Massachusetts Institute of Technology.
— (2018). "Belief as Question-Sensitive". In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47.