Preprint of a paper forthcoming in 'Mind'

Knowability Relative to Information

Francesco Berto[*] and Peter Hawke[†]

May 26, 2018

Abstract

We present a formal semantics for epistemic logic, capturing the notion of *knowability relative to information* (KRI). Like Dretske, we move from the platitude that what an agent can know depends on her (empirical) information. We treat operators of the form $K_A B$ ('B is knowable on the basis of information A') as variably strict quantifiers over worlds with a topic- or aboutness- preservation constraint. Variable strictness models the non-monotonicity of knowledge acquisition while allowing knowledge to be intrinsically stable. Aboutness-preservation models the topic-sensitivity of information, allowing us to invalidate controversial forms of epistemic closure while validating less controversial ones. Thus, unlike the standard modal framework for epistemic logic, KRI accommodates plausible approaches to the Kripke-Harman dogmatism paradox, which bear on non-monotonicity, or on topic-sensitivity. KRI also strikes a better balance between agent idealization and a non-trivial logic of knowledge ascriptions.

Keywords: [Epistemic logic; Information; Aboutness; Epistemic closure; Dogmatism paradox; Logical omniscience.]

1 Introduction

We expect a framework for epistemic logic[1] to perform a balancing act.[2] It should yield sufficient logical structure to justify the use of formal tools. It should allow

[*]University of St Andrews, Department of Philosophy, and University of Amsterdam, Institute for Logic, Language and Computation (ILLC) F.Berto@uva.nl

[†]University of Amsterdam, Institute for Logic, Language and Computation (ILLC) P.M.Hawke@uva.nl.

[1]See Meyer (2001), van Ditmarsch et al. (2008) and van Benthem (2011) for recent introductions.

[2]Many thanks to our anonymous referees for detailed comments that improved the paper substantially. For useful remarks, thanks to Giovanni Ciná, Malvin Gattinger, Davide Grossi, Sonja

the study of a kind of agent that is of genuine interest. There's a well known tension between the desiderata. Emphasis on the former can pull toward modeling idealized agents with unbounded cognitive powers. Emphasis on the latter can pull toward logics that are either too complex and specialized to be candidates for a general framework, or too weak to be of serious interest. What knowledge facts follow from ordinary agent Sarah's knowing that both A and B? Perhaps she has failed to unpack her belief, so she need not know that A. As the contents of Sarah's attitudes are, plausibly, extremely fine-grained, she needn't know that C, where C is logically equivalent to the conjunction of A and B.

Further, we expect a general framework for epistemic logic to maintain a second balancing act, if it is to be useful for philosophers. It should be flexible enough to represent a range of competing positions in philosophical debates, filling the traditional role of logic as a philosophically neutral tool. It should, however, furnish a core epistemic logic capturing substantial, but relatively uncontroversial, aspects of the knowledge concept.

By this measure, standard epistemic logic in the tradition of Hintikka (1962) is remarkably successful. It has the tractability of an unadorned modal logic. It offers a base logic of substance, namely system K. It is expressive enough to embed a natural framework for knowledge update: *public announcement logic* (PAL).[3] It has found widespread use in game theory and computer science (Fagin et al., 1995; van Benthem, 2011; van Ditmarsch et al., 2015). It has proven useful in philosophy as a tool for formalizing theories of knowledge that differ on the issue of introspection, and for framing epistemic paradoxes (Williamson, 2000; van Benthem, 2004; Kvanvig, 2006). Finally, as already observed by Hintikka (1962, Sect. 2), in spite of lacking plausibility as a logic of ordinary knowledge ascriptions, the standard framework can be interpreted in ways that promise some relevance to ordinary agents.

Nevertheless, the framework has shortcomings. With respect to the first balancing act, it is widely viewed as tipping too far in the direction of idealization (Fagin et al., 1995; Humberstone, 2016). With respect to the second balancing act, there is a growing realization that it is not flexible enough to capture key positions in current epistemological debates: far from offering a neutral tool for formalization, it is committed to philosophically controversial theses.

The problem of logical omniscience cuts across these concerns (Stalnaker,

Smets and Shane Steinert-Threlkeld. Special thanks to Alexandru Baltag and Johan van Benthem for detailed and stimulating feedback. This research is published within the project 'The Logic of Conceivability', funded by the European Research Council (ERC CoG), Grant Number 681404.

[3] PAL was introduced by Plaza, in a work that appeared eventually as Plaza (2007). For a general introduction to dynamic epistemic logic, see van Benthem (2011). Standard epistemic logic can embed PAL via reduction axioms, defining dynamic epistemic operators via static ones plus non-epistemic logical vocabulary.

1991). The standard framework has two core features: logical truths are always known; knowledge is closed under known implication. Now, not only do ordinary agents fail to appreciate consequences of their knowledge that they haven't explicitly deduced, let alone cannot conceptualize: it is philosophically controversial whether even fully rational, cognitively ideal agents enjoy logical omniscience. Witness the growing contingent of epistemologists that are sympathetic to 'closure denial' (Dretske, 1970, 2005; Nozick, 1981; Schaffer, 2007; Lawlor, 2013; Yablo, 2014; Holliday, 2015; Hawke, 2016).

The standard picture of knowledge update, upon which public announcement logic is founded, is likewise questionable. On this picture, an agent's knowledge grows monotonically: invariably, more information results in more, or at least no less, knowledge, if we ignore epistemic claims that report on the agent's current body of knowledge.[4] Now, not only are ordinary agents subject to deception, imperfect recall, and irrational aspects of their psychology that can lead to belief updates undermining knowledge: monotonicity is philosophically controversial, again, even for cognitively ideal, fully rational agents – as we will discuss extensively in Sect. 2.3.

Both closure and monotonicity will be core issues for *this* work, which aims at striking a better twofold balance than the standard framework: we introduce a formal semantics for epistemic logic that relaxes the constraints of closure and monotonicity while maintaining both a high degree of simplicity and non-trivial logical properties.

Some idealization is inevitable in the development of a worthwhile epistemic logic. As in Hintikka (1962), we do not aim for a logic that governs ordinary knowledge attributions *per se*. Rather, we intend to capture the notion of *knowability relative to information* (KRI). Our key question is: if her total information is A, what knowledge can a fully rational and computationally unbounded agent base on that information? Thus, we abstract away from certain contingent cognitive handicaps and focus on the quality of the information available to the agent. This echoes a prominent interpretation of the standard framework as a logic of (hard) information (van Benthem, 2011, Ch. 2).

We take inspiration from Dretske (1999).[5] Dretske stresses that knowledge depends on the (empirical) information available to us. We understand information propositionally (one has, or acquires, the information *that* A).[6] The role of

[4]PAL accommodates Moorean phenomena (Holliday and Icard, 2010). Take $p \wedge \neg Kp$. The agent might come to learn this (say, by testimony). But the outcome is not the truth of $K(p \wedge \neg Kp)$, since the update of the agent's knowledge renders $\neg Kp$ false. Update in public announcement logic is monotonic if one restricts attention to non-epistemic claims. This last feature is contentious.

[5]We adopt various basic insights from Dretske (1999). However, we need not be taken to endorse the detailed (probabilistic) theory of information defended by Dretske (1999, Ch.3).

[6]We mention a departure from Dretske's basic commitments (see our Sect. 4): he takes all

incoming information is to narrow down the set of epistemically viable alternatives. We read '$K_A B$' as 'If the total given information were A, then B would be knowable'. Alternatively: 'B can be known on the basis of total information A'. Our focus will be the logic and semantics of knowability ascriptions of the form '$K_A B$'.

Thus, we treat knowability ascriptions as *conditional* claims. Epistemic logic, then, becomes a type of conditional logic. Arguably, this impulse is implicit in the standard framework.[7] We make it explicit. The information-theoretic focus will allow us to address issues of knowledge update in a static system that does not deploy the full machinery of a dynamic logic.[8] Our basic system will invalidate monotonicity: information can grow while knowledge depletes. On the other hand, it will validate *transitivity* as capturing the less controversial sense in which knowledge is 'stable in the face of new information'.

Our formal semantics also combines the possible worlds apparatus with an account of *topics*.[9] The former element allows us to retain many advantages of the dominant model-theoretic approach to epistemic logic. The latter element – a simple mereology of contents, drawing on Berto (2017, 2018) – allows a subtle mix: controversial forms of *epistemic closure* are invalidated, while less controversial ones are validated. Topic-sensitivity can model the limitations of an agent's conceptual apparatus, a crucial source of closure failure in ordinary agents, even logically astute ones. But the topic-sensitivity of knowledge claims, plausibly inherited from an intrinsic topic-sensitivity of information, also provides the most compelling route to closure rejection even for highly idealized agents that have mastery over all concepts, as argued by Yablo (2014, Ch. 7) and Hawke (2016).[10]

Proponents of monotonicity, or of epistemic closure, often emphasize the intuitions that deduction preserves knowledge and that knowledge, as per the venerable Platonic tradition of *epistéme*, rests on conclusive grounds that render it

information to be veridical. Our proposed framework, in contrast, is compatible with there being non-veridical information. Our arguments in support of this framework are consistent with allowing only veridical information, however. On the debate concerning the factivity of information, see Floridi (2015).

[7]The standard framework models agent a's epistemic situation as a set of possible worlds, most straightforwardly understood as a's information or knowledge. Ascriptions $K_a \varphi$ are then naturally understood as capturing what is *knowable on this basis*. Various proposed readings draw out the conditionality. Consider the preferred interpretation in Hintikka (1962): $K_a p$ means roughly 'relative to her knowledge, a is permitted to infer p'. Or consider a purely descriptive interpretation raised in (Hintikka, 1962, Sect. 2.10): 'it follows from what a knows that p'.

[8]This is not to say that our system has the expressive power of a properly dynamic system, nor that pursuing a dynamic variant of our system is without interest.

[9]See Hawke (2017a) for a detailed appraisal of recent work on topicality.

[10]The proposal that knowledge ascriptions are question-sensitive provides another intriguing route (Schaffer, 2007). As Yablo (2014) notes, the two are closely related.

stable. But our framework identifies a closure principle and a stability principle that, we submit, can be accepted by *all* hands in such debates. Monotonicity and closure discontents needn't reject ordinary intuitions – only certain formulations of those intuitions.

Finally: notwithstanding our focus on what is knowable in principle, the KRI framework *can* model important cognitive limitations of an agent. Topic-sensitivity can be used to model the limits on an agent's conceptual resources. Our variably strict operators can also model cognitive systems sensitive to the logical complexity of a piece of information. Sect. 11 offers remarks in this direction.[11]

We proceed as follows. Sect. 2 furnishes preliminaries and presents a version of the standard framework for epistemic logic. We motivate its limitations via a convenient case study: the Kripke-Harman *dogmatism paradox*. As we highlight there, the paradox can be split into sub-paradoxes concerning monotonicity and closure, respectively. Sect. 3 introduces the KRI semantics. Sect. 5-11, then, discuss various principles it validates and invalidates. In particular, Sect. 6 addresses the non-monotonicity of KRI, delivered by the variable strictness of our K_A operators; Sect. 8 and 9 address the failures of forms of logical omniscience and of closure under strict implication, delivered by the topic-sensitivity of K_A. Sect. 11 notes that our binary epistemic operators invalidate principles sometimes (e.g., in Gabbay (1985)) billed as core to conditional logic. We discuss the desirability of meeting these principles in our context. Sect. 12 flags further work and concludes.

2 Preliminaries

2.1 Language

We work with a sentential language $\mathscr{L}$ with a non-empty set $\mathscr{L}_{AT}$ of atomic formulas, p, q, r ($p_1, p_2, ...$), negation $\neg$, conjunction $\wedge$, disjunction $\vee$, a strict conditional $\prec$, a two-place epistemic operator K, round parentheses as auxiliary symbols. We use 'atom' as shorthand for 'atomic formula'. We use A, B, C ($A_1, A_2, ...$), as metavariables for formulas of $\mathscr{L}$. The well-formed formulas are items in $\mathscr{L}_{AT}$ and, if A and B are formulas:

$$\neg A \mid (A \wedge B) \mid (A \vee B) \mid (A \prec B) \mid K_A B$$

[11]This raises a question that we postpone for further work: how does our system compare to extant modifications of epistemic logic for capturing bounded cognition? In particular, it is worth drawing out similarities and contrasts with the tradition that extends the standard framework with a notion of *awareness*, *conceptualization*, *entertainment* or *explicit belief*. See, for instance, Levesque (1984), Fagin and Halpern (1988) and van Benthem (2011, Ch. 5).

Outermost brackets are omitted by default. Expressions of the form 'K_A' work similarly to sententially indexed modals (see Chellas (1975)). We use $\supset$ for the material conditional, defined in the usual manner. In the metalanguage we use variables $w, w_1, w_2, ...$, ranging over possible worlds, x, y, z ($x_1, x_2, ...$), ranging over topics (these will officially enter the stage in Sect. 3), as well as the symbols $\Rightarrow, \Leftrightarrow, \&, or, \sim, \forall, \exists$, with the usual reading. We now look at a standard epistemic logic, semantically presented, for $\mathscr{L}$.

2.2 A Standard Epistemic Logic

The standard approach to (multi-agent) epistemic logic uses the following core ideas. A body of information is modeled as a set of possible worlds. A set of agents is given, and a body of information is associated with each agent at each world. Generally, this is modeled with an agent-relative accessibility relation between worlds. Knowability ascriptions are then interpreted as follows: it is true at world w that a is in a position to know p just in case p is true at every possible world accessible from w by agent a. That is: just in case $\neg p$ is incompatible with the agent's information at w. Public announcement logic adds a natural dynamics to this picture: the receipt of new information is modeled as the intersection between it and the prior body of information (cf. the notion of conditional probability in Bayesian probability theory).

We render this more precisely, in a manner that departs slightly from the usual presentations but lays the groundwork for the KRI framework of Sect. 3. We can eliminate any mention of individual agents (and world-relative accessibilities) without betraying the features we want to emphasize.

A *standard model* for $\mathscr{L}$ is a tuple $\langle W, \Vdash_S \rangle$, where W is a non-empty set of worlds and $\Vdash_S \subseteq W \times \mathscr{L}_{AT}$ is an interpretation of the atomic claims in $\mathscr{L}$. This relates worlds to atoms: we read '$w \Vdash_S p$' as meaning that p is true at w, '$w \nVdash_S p$' as $\sim w \Vdash_S p$. Next, $\Vdash_S$ is extended to all formulas of $\mathscr{L}$ as follows:

- $w \Vdash_S \neg A \Leftrightarrow w \nVdash_S A$
- $w \Vdash_S A \wedge B \Leftrightarrow w \Vdash_S A \;\&\; w \Vdash_S B$
- $w \Vdash_S A \vee B \Leftrightarrow w \Vdash_S A \; or \; w \Vdash_S B$
- $w \Vdash_S A \dashv B \Leftrightarrow \forall w_1 (w_1 \Vdash_S A \Rightarrow w_1 \Vdash_S B)$
- $w \Vdash_S K_A B \Leftrightarrow \forall w_1 (w_1 \Vdash_S A \Rightarrow w_1 \Vdash_S B)$

The interpreted language is quite redundant: $A \dashv B$ and $K_A B$ have the same interpretation. But their meanings will diverge in *our* framework below, and it will prove useful to frame various principles of interest using both linguistic devices.

We define logical consequence in the standard way, as truth preservation at all worlds of all admissible models. With Σ a set of formulas:

- $\Sigma \vDash_S B \Leftrightarrow$ in all standard models $\langle W, \vDash_S \rangle$ and for all $w \in W$: $w \Vdash_S A$ for all $A \in \Sigma \Rightarrow w \Vdash_S B$

We write $A \vDash_S B$ for $\{A\} \vDash_S B$. As a special case, logical validity, $\vDash_S A$, truth at all worlds of all standard models, is $\emptyset \vDash_S A$, entailment by the empty set. One might label the set of all such validities *core standard epistemic logic*.

Given this core framework, one can clarify and contrast more refined epistemic logics by restricting the class of standard models. Each such restriction – a proposed class of admissible models – is a proposal as to which models capture a genuine possibility for an agent's epistemic status, and generates its own set of corresponding validities. Admissibility is key: relative to a core framework, a debate as to which logic is *the* epistemic logic may be framed as a debate over what should count as an admissible model.

The logic induced by the semantics for the extensional operators is just classical propositional logic, $A \multimap B$ and $K_A B$ being notational variants for a strict S5-like conditional, often called 'strict implication'. Key consequences of the standard approach are now easily established:

- (Logical Omniscience) $\vDash_S B \Rightarrow \vDash_S K_A B$ for every A
- (Closure Under Known Implication) $\{K_A B, K_A(B \supset C)\} \vDash_S K_A C$
- (Closure Under Strict Implication) $\{K_A B, B \multimap C\} \vDash_S K_A C$
- (Monotonicity) $\{A \multimap B, K_B C\} \vDash_S K_A C$
- (Transitivity) $\{K_A B, K_B C\} \vDash_S K_A C$

In the current setting, the last three items say the same thing with different symbols.

2.3 Kripke and Harman's Dogmatism Paradox

We now present our case study: a paradox due to Kripke (2011b) which first appeared in Harman (1973, Ch.9, Sect. 2), reporting on a lecture by Kripke. Notably, it applies as much to perfectly ideal agents as to ordinary ones. Appealing replies to the paradox cast doubt either on the above mentioned closure principles, or on monotonicity. Rather than arguing for any reply in particular, we emphasize the *plausibility* of some: whether or not they are best in the final analysis, they deserve to be taken seriously. It is, therefore, desirable to develop a logical framework that allows us to study the theories recommended by such replies.[12]

Suppose that P is true and E is true and R is true, where R is the claim that E is generally a good reason to think that P is false. Let M be the claim that E is misleading information on the question of P. The following seems true:

1. If P is true and E is generally a good reason to think that P is false, then it must be that if E is true then E is misleading information on the question of P. That is: $(P \wedge R) \dashv\!3 (E \supset M)$.

Now suppose that agent a knows that $P \wedge R$ at time t_0, on the basis of information I_1. Using closure under known implication, we may conclude:

2. a is in a position at t_0 to know that: $E \supset M$.

Suppose that a comes to know E at time t_1 on the basis of new information I_2. Presumably, her information is now: $I_1 \wedge I_2$. Using monotonicity, we get:

3. a is in a position at t_1 to know that: $E \supset M$.

Since E also knows E at t_1, we can apply closure under known implication:

4. a is in a position at t_1 to know that M.

If a knows that E is misleading, then presumably she is rational, in the face of E, to continue believing P, ignoring the 'usual implications' of E.

But, as Kripke (2011b) stresses, this result is completely general and therefore coalesces into a principle of dogmatism: knowing agents are immune to rational persuasion with new evidence! This is highly counter-intuitive. It is well known that Kripke first proved certain results in modal logic. Suppose that one comes across a letter, signed by Kripke and addressed to Nozick, in which Kripke confesses to having plagiarized the results. As it happens, the contents of the letter are false (representing a private joke between Kripke and Nozick) but one is unaware of this. Intuitively, the new information – e.g. that such a letter exists – undermines one's rational belief in the claim that Kripke produced the results, and thereby undermines one's knowledge. However, the reasoning from 1 to 4 *seems* to advocate that one can (and should) resist this change in belief, since one knows that the new information is misleading on the question of Kripke's accomplishments. But, intuitively, it is precisely the fact that one does *not* know this that fuels a rational loss of belief.

[12] For further discussion of the paradox: Sorensen (1988), Lasonen-Aarnio (2014), Sosa (2017, Ch.10).

This inspires a quandary. Suppose we accept the conclusion of the paradox. Still, our ordinary (purported) claims to knowledge can *obviously* be challenged with new counter-evidence. Thus these claims must be, on reflection, false. Skepticism looms. Alternatively, we need to defy the reasoning that leads to the paradoxical conclusion.

One route for defiance is that of Harman (1973, Ch.9, Sect. 2). It targets *monotonicity*. Consider the step from 2 to 3: if $E \supset M$ is knowable at t_0 then it remains knowable at t_1 if the only change to the agent's psychology is that they have received new information. To abandon monotonicity is to allow that the receipt of I_2 might *reduce* what is knowable. In particular, one might accept counter-instances to monotonicity of the form:

$$(I_1 \wedge I_2) \prec I_1$$

$$K_{I_1}(E \supset M)$$

$$\neg K_{I_1 \wedge I_2}(E \supset M)$$

Another route for defiance is that of Sharon and Spectre (2010, 2017). It targets *epistemic closure*. The paradox relies heavily on the closure of knowledge under strict implication. One may therefore take the paradox as pointing to counter-instances to closure, of the form:

$$K_{I_1}(P \wedge R)$$

$$(P \wedge R) \prec (E \supset M)$$

$$\neg K_{I_1}(E \supset M)$$

Closure has perhaps stronger intuitive plausibility than monotonicity, so it is worth bolstering the appeal of the current act of defiance. Note that Harman's solution, taken by itself, concedes that 2 holds. Thus, at t_0, the agent knows that any counter-evidence to P that she might receive is guaranteed to be misleading. We can accept, with Harman, that if actually presented with new counter-evidence then the agent would be rationally swayed and lose some knowledge.

But a residual paradox remains: it seems that, at t_0, the agent would be rational to do everything she can to *avoid* any possible counter-evidence – especially if she knows that it will hold her under its sway if it appears. As Kripke (2011b) points out, this is an equally repellent form of dogmatism, according to which a rational agent is entitled to actively avoid persons or books or other sources of information that challenge whatever views she takes to constitute her knowledge. Hence, one can appreciate the appeal of restricting closure and thereby allowing for knowing agents that are receptive to counter-argument.[13]

This suggests that the dogmatism paradox encompasses *two* sub-paradoxes: one based on monotonicity, one based on closure. To clarify this, we attend to what we take to be the essential structure of the paradoxical reasoning (notice that this presentation finds no use for Closure Under Known Implication):

5. $P \dashv\!\!\! 3\; \neg(E \wedge \neg P)$ (by classical propositional and modal logic)

6. $K_{I_1}P$ (premise)

7. $K_{I_1}\neg(E \wedge \neg P)$ (by 5, 6 and Closure Under Strict Implication)

8. $K_{I_1 \wedge I_2}E$ (premise)

9. $K_{I_1 \wedge I_2}\neg(E \wedge \neg P)$ (by 7 and Monotonicity)

10. $K_{I_1 \wedge I_2}(E \wedge \neg(E \wedge \neg P))$ (by 8, 9 and Adjunction)

11. $(E \wedge \neg(E \wedge \neg P)) \dashv\!\!\! 3\; P$ (by classical propositional and modal logic)

12. $K_{I_1 \wedge I_2}P$ (by 10, 11 and Closure Under Strict Implication)

To discern the stakes, again interpret E as a claim that inductively supports $\neg P$. We now use $\neg(E \wedge \neg P)$ to capture the idea that E is misleading if E and P are true.[14] 12 captures a significant element of the paradoxical reasoning: new information cannot yield counter-evidence that undermines previous knowledge, since an agent knows that any counter-evidence is misleading (see 10).

But to achieve a paradox using monotonicity, the intervening steps from 7 to 11 are inessential. Our first sub-paradox:

13. $K_{I_1}P$ (premise)

14. $K_{I_1 \wedge I_2}P$ (by monotonicity from 13)

In the abstract, this reasoning is intuitive. Putting aside memory failure, information seems cumulative: new information can only tell one *more* about the world. But the example of losing one's knowledge of the genesis of Kripke's theorem, through the misleading letter, bears directly on the reasoning from 13 to 14, and so on monotonicity directly. We intuitively judge in this *particular* case that knowledge can be lost with the accrual of novel knowledge-producing

[13](Sharon and Spectre, 2010, pp. 310-11) makes a similar point, apparently independently.

[14]This technique for formalizing misleading evidence has proven useful in epistemology: see, for instance, Vogel (2014).

information, since that information undermines formerly rational beliefs (and so knowledge resting on those beliefs).

On the other hand, 5, 6 and 7 provide a closure-based sub-paradox:

15. $P \prec\!\!3\; \neg(E \wedge \neg P)$ (premise)

16. $K_I P$ (premise)

17. $K_I \neg(E \wedge \neg P)$ (by 15, 16 and closure under strict implication)

This is independently puzzling. For emphasis, set E to be I. Then $\neg(I \wedge \neg P)$ says that the agent's total information I is *not* misleading on the question of P. But then 15, 16 and 17 seem to say: if one knows anything, one is positioned to know that one's total information is never misleading. But isn't it objectionably circular to claim that one's total information gives assurance that one's total information is never misleading? We have a version of the classic 'Problem of the Criterion' (Chisholm, 1973; Cohen, 2002).

Thus there is motivation for introducing a framework for epistemic logic with the resources for rejecting *both* monotonicity *and* closure.

3 Semantics for KRI

We now present the KRI semantics for our epistemic language $\mathscr{L}$ from Sect. 2.1. It is informed by three ideas: (1) the content of an interpreted sentence is fruitfully modeled with two components: a *truth set* and a *topic*, or subject matter. Specifically, this is so for a sentence expressing an agent's total information. (2) the topic of information I restricts what is knowable on the basis of I to propositions about that same subject matter. This impacts epistemic closure. (3) total information is a mere upper bound on knowability relative to information: in the best case, an agent knows that I, where I is her total information. But she may not be so lucky: knowledge based on I might defeat knowledge that follows, in the absence of defeaters, from a mere part of I. This impacts monotonicity.

We identify $\mathscr{L}$ with the set of its well-formed formulas. A *frame* for $\mathscr{L}$ is a tuple $\mathfrak{F} = \langle W, \{R_A \mid A \in \mathscr{L}\}, \mathscr{T}, \oplus, t\rangle$, understood as follows. W is a non-empty set of possible worlds. $\{R_A \mid A \in \mathscr{L}\}$ is a set of accessibilities between worlds: each $A \in \mathscr{L}$ has its own $R_A \subseteq W \times W$. Such accessibilities will make our K_A's non-monotonic, addressing one half of the dogmatism paradox. $\mathscr{T}$ is a non-empty set of topics. Abstractly, topics are the situations or distinctions a given bit of information is epistemically relevant for, in a certain context and for the agent involved. Intuitively, the topic of a meaningful sentence or discourse is what it is *about*, a dimension of meaning that (as stressed in such influential works as Yablo (2014)) goes beyond conditions of truth at a possible world: '$2+2=4$' and 'Either Jane is late or she is not' are true at exactly the same possible worlds (all of them). But they differ along the dimension of topic: one is about Jane, the other is not. The notion of topic can naturally explain various hyperintensional aspects of natural language. And the topic-sensitivity of our K_A's will deliver failures of closure that address the other half of the dogmatism paradox.

Mathematics has topology as a sub-topic. Philosophy and mathematics overlap (they have a common sub-topic: logic). The topic *Jane's profession* is included in a larger topic: *Jane*. Thus, topics can have sub-topics, can overlap, and can be included in larger topics. To capture these ideas, we have $\oplus$ as *topic fusion*, a binary operation on $\mathscr{T}$ that combines two topics into, intuitively, the smallest topic of which they are both a part. We take $\oplus$ to satisfy, for all $x,y,z \in \mathscr{T}$:

- (Idempotence) $x \oplus x = x$
- (Commutativity) $x \oplus y = y \oplus x$
- (Associativity) $(x \oplus y) \oplus z = x \oplus (y \oplus z)$

We take on board unrestricted fusion, that is, $\oplus$ is always defined on $\mathscr{T}$: $\forall xy \in \mathscr{T}\ \exists z \in \mathscr{T}(z = x \oplus y)$. We then define *topic parthood*, $\leq$, the usual way: $\forall xy \in \mathscr{T}(x \leq y \Leftrightarrow x \oplus y = y)$. This makes of parthood a partial ordering – for all $x,y,z \in \mathscr{T}$:

- (Reflexivity) $x \leq x$
- (Antisymmetry) $x \leq y\ \&\ y \leq x \Rightarrow x = y$
- (Transitivity) $x \leq y\ \&\ y \leq z \Rightarrow x \leq z$

Thus, $\langle \mathscr{T}, \oplus \rangle$ is a join semilattice, as in Berto (2017, 2018). We can, additionally, stipulate that it be complete: any set of topics $S \subseteq \mathscr{T}$ has a fusion $\oplus S$. As a final technical assumption, we will think of all topics in $\mathscr{T}$ as built via fusions out of the smallest possible topics: namely, *atomic topics*. Atomic topics have no proper parts: $Atom(x) \Leftrightarrow \sim \exists y(y < x)$, with $<$ the strict order defined from $\leq$.

$\langle \mathscr{T}, \oplus \rangle$ is needed to assign topics to formulas of $\mathscr{L}$, as follows. Our t in $\mathfrak{F}$ above is a function $t : \mathscr{L}_{AT} \to \mathscr{T}$, such that if $p \in \mathscr{L}_{AT}$, then $t(p) \in \{x \in \mathscr{T} | Atom(x)\}$: atomic formulas have atomic topics (this is an idealization: grammatically simple sentences of everyday language can involve intuitively complex topics; but it will streamline our discussion). Next, t is extended to the whole of $\mathscr{L}$ by taking a formula as having as topic the fusion of what its atomic formulas are about. If the set of atomic formulas in a formula $A \in \mathscr{L}$ is $\mathfrak{A}\mathfrak{t}A = \{p_1, ..., p_n\}$, then:

- $t(A) = \oplus\mathfrak{A}\mathfrak{t}A = t(p_1) \oplus ... \oplus t(p_n)$

Topical hyperintensionality is less fine-grained than the syntax of our language. By induction on the construction of formulas, not only $t(A) = t(\neg\neg A)$ (remember Frege on how Double Negation is *Sinn*-preserving), but also $t(A) = t(\neg A)$: the topic of a formula is that of its negation ('Snow is white' is about snow's whiteness, just as 'Snow is not white'). And not only $t(A \wedge B) = t(B \wedge A)$, but also, e.g., $t(A \wedge B) = t(A) \oplus t(B) = t(A \vee B)$. These identities are taken as requirements for a good theory of aboutness- or content-inclusion in, e.g., Yablo (2014), Fine (2015), and Hawke (2017a).

A frame becomes a *model* $\mathfrak{M} = \langle W, \{R_A \mid A \in \mathscr{L}\}, \mathscr{T}, \oplus, t, \Vdash\rangle$ when one adds an interpretation $\Vdash \subseteq W \times \mathscr{L}_{AT}$. This relates worlds to atomic formulas: again, we read '$w \Vdash p$' as meaning that p is true at w, '$w \nVdash p$' as $\sim w \Vdash p$. Next, $\Vdash$ is extended to all formulas of $\mathscr{L}$ as follows:

- (S$\neg$) $w \Vdash \neg A \Leftrightarrow w \nVdash A$
- (S$\wedge$) $w \Vdash A \wedge B \Leftrightarrow w \Vdash A \ \& \ w \Vdash B$
- (S$\vee$) $w \Vdash A \vee B \Leftrightarrow w \Vdash A \ or \ w \Vdash B$
- (S$\rightarrow$) $w \Vdash A \rightarrow B \Leftrightarrow \forall w_1(w_1 \Vdash A \Rightarrow w_1 \Vdash B)$
- (SK) $w \Vdash K_A B \Leftrightarrow \forall w_1(wR_A w_1 \Rightarrow w_1 \Vdash B) \ \& \ t(B) \leq t(A)$[15]

Read '$wR_A w_1$' as: 'relative to w: w_1 is epistemically accessible on the basis of total information that A' (or: 'relative to w: w_1 is not ruled out by knowledge that can be based on the total information that A'). Accessibilities, thus, are indexed to information: different informational inputs will commit the agent to different epistemic possibilities.

Following Lewis (1973)'s worlds semantics for counterfactuals, (SK) can be equivalently expressed using set-selection functions. Each $A \in \mathscr{L}$ comes with a function $f_A : W \rightarrow \mathscr{P}(W)$, taking as input the world where the information is had by the agent, and giving as output the set of epistemically accessible worlds, $f_A(w) = \{w_1 \in W | wR_A w_1\}$. Let $|A|$ denote A's *truth set*: $\{w \in W | w \Vdash A\}$. Then we can rephrase (SK), equivalently, as:

- (SK) $w \Vdash K_A B \Leftrightarrow f_A(w) \subseteq |B| \ \& \ t(B) \leq t(A)$

Set-selection functions also tersely express a natural Basic Constraint on the semantics – that for all $A \in \mathscr{L}$ and $w \in W$:

[15] Compare *analytic implication* in the *conceptivist* literature: see Ferguson (2014) for an overview.

- (BC) $|A| \subseteq f_A(w)$[16]

BC says that A-worlds are always A-selected: no world in which A is true is ruled out by knowledge based on total information A. Besides being intrinsically plausible, this will come handy to prove some simple results below. From now on, we will only consider models satisfying BC. Notice what BC does *not* ensure: that every world in which A is false is ruled out by knowledge based on A.[17]

We again define logical consequence as truth preservation at all worlds of all models. With Σ a set of formulas:

- $\Sigma \vDash B \Leftrightarrow$ in all models $\mathfrak{M} = \langle W, \{R_A \mid A \in \mathscr{L}\}, \mathscr{T}, \oplus, t, \Vdash\rangle$ and for all $w \in W$: $w \Vdash A$ for all $A \in \Sigma \Rightarrow w \Vdash B$

We keep having $A \vDash B$ as shorthand for $\{A\} \vDash B$ and $\vDash A$ as shorthand for $\emptyset \vDash A$. The set of validities induced by this framework gives a *core logic* for KRI, from which more refined theories are set by restricting the class of admissible models. We start by commenting on the key clause of the semantics: (SK).

4 Knowability, Information

(SK) asks two things for $K_A B$ to come out true: (1) it embeds a *truth-conditional* requirement – that B be true throughout a selected set of worlds compatible with the information that A; (2) it embeds a *topicality* requirement: that B be fully on topic with respect to A. Knowability is, then, determined by the available information twice over: once via the worlds it makes epistemically accessible, and once via the topic it concerns.

This complies with insights about informativeness, and its relation to knowledge. Consider the picture emerging from Dretske (1999). Knowledge depends on information: to learn that Beth's grandmother is ill, one requires information to that effect. Information should not be conflated with meaning: if I am passed

[16] Caution could tempt one towards a weaker basic constraint. For instance: for all $w \in W$, if $w \in |A|$ then $w \in f_A(w)$ (Chellas, 1975, p. 42). This yields a strictly weaker logic than BC (e.g. Transitivity is lost). This strikes us as unnecessarily cautious. Thanks to an anonymous referee of this Journal for pressing us on this.

[17] One anonymous referee of this Journal rightly asked, what of further constraints on R_A or f_A, e.g., making of our R_A's equivalence relations, or transitive ones, etc.? In the standard framework of epistemic logic, these are linked to the debate about the validity of principles like the KK principle or Positive Introspection (if one knows that A, does one know that one knows that A?). We take no commitment on these, given our aim, stated at the start, of providing a fairly neutral epistemic logic, and the fact that such principles are controversial already in the standard Hintikkan framework. We will, however, discuss the plausibility of one further constraint involving both f_A and topics in our Sect. 11 below.

a note that reads 'Beth's grandmother is ill', written by someone who chose that sentence using a random device, then that sentence is meaningful, but carries no information about the state of health of Beth's grandmother. Even if the sentence is true, I cannot *learn* anything about Beth's grandmother from it.

Nevertheless, information may be regarded as semantic (Floridi, 2015) to the extent that, firstly, it eliminates possibilities, just as the truth of a meaningful sentence is, in general, compatible with some possibilities and not others; and, secondly, it is about something, just as a meaningful sentence has a subject matter that it addresses.[18] Abstractly, an information source divides logical space into a partition of possibilities, and selects between them (definitively, if it is noise-free). What the information licenses as true is captured by the selection. What the information is about is captured by the distinctions that mark the borders of the partition. If the information source (say, a voice on the telephone) reports on the health of Beth's grandmother (call her 'X'), then it divides logical space, roughly, into cells such as: X is fit and hearty; X is under the weather; X has been hospitalized. It need not discriminate between issues such as: is X the grandmother of Sue, or is it Y? Nor need it carry the information that $2+2=4$, despite this being strictly implied by any true claim. Nor need it carry information about the source itself: it needn't report that the telephone connection is noise-free, for instance.

Dretske takes information to be veridical: "*false* information and *mis*-information are not kinds of information – any more than decoy ducks and rubber ducks are kinds of ducks" (Dretske, 1999, p. 45). KRI semantics makes no such assumption. We will be neutral on whether information is always true, or can be false: no invalidity we prove depends on the existence of false information; and no validity depends on the assumption that information must be true.

We now unpack some logical validities and invalidities in the semantics. These will highlight how KRI fares with respect to the issues presented for the standard epistemic framework in Sect. 2.

5 Factivity, Conjunction, Paradox

Our first logical validity is:

(Factivity) $\{K_A B, A\} \vDash B$

This is immediately guaranteed by our Basic Condition.[19] This validity expresses the factivity of KRI: when B is knowable relative to the information that

[18] These aspects have long been recognized, though emphasized in distinct traditions: compare *information-as-range* and *information-as-correlation* in van Benthem and Martinez (2008).

[19] *Proof*: let $w \Vdash A$ and $w \Vdash K_A B$. By the former, $w \in |A|$ so BC applies: $w \in f_A(w)$. Then by the latter and (*SK*), $w \Vdash B$.

A, and A *is* true, B must be true as well.

It is easy to show that our framework does *not* validate a different factivity principle:

$$K_A B \nvDash A$$

This is crucial for accommodating theorists that allow non-veridical information. However, it is not an endorsement of the possibility of false information. Recall that our intuitive reading of $K_A B$ has a subjunctive flavor: 'If the total given information were A, then B would be knowable'. Now if information is necessarily veridical, then one should also accept: 'If the total given information were A, then A would be true'. But one need *not* accept that the intuitive reading of $K_A B$ entails 'A is true': the subjunctive might be true, intuitively, because receiving A positions one to know B at all (nearby) worlds where A is true.

The next validities show that KRI is closed with respect to conjunction introduction and elimination:

(Simplification) $K_A(B \wedge C) \vDash K_A B \qquad K_A(B \wedge C) \vDash K_A C$ [20]

The 'tracking' notion of knowledge due to Nozick (1981) does not necessitate that one who knows a conjunction is positioned to know the conjuncts. According to Kripke (2011a), this is a damning defect for Nozick's approach. KRI is free from such a defect.[21]

The companion of Simplification is:

(Adjunction) $\{K_A B, K_A C\} \vDash K_A(B \wedge C)$[22]

If, given information A, both B and C are knowable, then $B \wedge C$ is knowable, too. Despite its intuitive plausibility, there is a case for viewing this validity as a drawback. Consider the preface paradox, due to Makinson (1965). An author has written a particularly well-researched book. In fact, every claim in the book is an instance of knowledge for her. Nevertheless, with appropriate epistemic modesty, her preface admits that she cannot guarantee that her long book is error-free as a whole. One might conclude that the author is not positioned to know

[20] *Proof*: we do the first one (for the second, replace B with C appropriately). Let $w \Vdash K_A(B \wedge C)$. By (SK), for all w_1 such that $wR_A w_1$, $w_1 \Vdash B \wedge C$, thus by (S$\wedge$), $w_1 \Vdash B$. Also, $t(B \wedge C) = t(B) \oplus t(C) \leq t(A)$, thus $t(B) \leq t(A)$. Then, by (SK) again, $w \Vdash K_A B$.

[21] See Hawke (2016) for further frameworks for epistemic logic that validate simplification without validating closure under strict implication.

[22] *Proof*: let $w \Vdash K_A B$ and $w \Vdash K_A C$, that is, by (SK): for all w_1 such that $wR_A w_1$, $w_1 \Vdash B$ and $w_1 \Vdash C$, so by (S$\wedge$) $w_1 \Vdash B \wedge C$. Also, $t(B) \leq t(A)$ and $t(C) \leq t(A)$, thus $t(B) \oplus t(C) = t(B \wedge C) \leq t(A)$. Then, by (S$K$) again, $w \Vdash K_A(B \wedge C)$.

the conjunction of every claim in the book. One plausible reaction is that we have identified a counter-instance to Adjunction. The current system – like the standard framework – cannot accommodate a theory of knowability that embraces this reaction.

Nevertheless – unlike the standard framework – KRI can accommodate subtler forms of the debate. A principle akin to Adjunction is invalid in the semantics:

(Unwise Adjunction) $\{K_{A_1}B, K_{A_2}C\} \nvDash K_{A_1 \wedge A_2}(B \wedge C)$

(This invalidity is closely related to that of Monotonicity, which we address next.) Now consider a second take on the core issue of the preface phenomenon: for every claim in her book, the author has presumably received some information that renders that claim knowable relative to *that exact* information. However, might it be that her total information does not render the book's conjoined content knowable? By denying that information and its resultant knowledge can always be simultaneously adjoined, a proponent of KRI can answer in the affirmative.

6 Non-monotonicity, Transitivity, Stability

KRI is non-monotonic in the following sense:

$K_A B \nvDash K_{A \wedge C} B$ [23]

Topicality is preserved here, for in general if $t(B) \leq t(A)$, then also $t(B) \leq t(A) \oplus t(C) = t(A \wedge C)$. Our epistemic operators, however, are of variable strictness: $f_A(w)$ can differ from $f_{A \wedge C}(w)$.

Here's a toy example drawing on Hawthorne (2004), p. 71.[24] Assume that information is veridical. At the actual world @, agent a reads in *The Times* that Manchester United won. We use M to name the proposition that Manchester United won and T to name the proposition that *The Times* reported that M. *The Times* is a trusted and reliable source, that offers a correct report. Hence, a is informed that $M \wedge T$ and thereby comes to know $M \wedge T$. We can model this with

[23] *Counterexample*: let $W = \{w, w_1\}$, w R_p-accesses nothing, $wR_{p \wedge r}w_1$, $w_1 \nVdash q$, $t(p) = t(q) = t(r)$. Then $w \Vdash K_p q$, but $w \nVdash K_{p \wedge r} q$.

[24] Hawthorne's example is similar, but developed with a different purpose: to serve as a puzzle about closure. As he acknowledges, the puzzle is essentially the closure-based sub-paradox of the dogmatism paradox. Hawthorne's verdict is that closure can be preserved: knowing that A puts one in a position to know that any evidence against A is misleading, but the latter is 'junk knowledge' that is 'destroyed' if new evidence is actually received, as in Sorensen (1988) (see Sharon and Spectre (2010, Sect.4) for push-back). This last part indicates that Hawthorne (2004) advocates monotonicity-rejection. Our development of his example is in this spirit.

a set-selection function: $f_{M\wedge T}(@) = |M| \cap |T|$, with $@ \in |M| \cap |T|$. Hence: $@ \Vdash K_{M\wedge T}M$.

But a reads *The Globe* next, which reports that Manchester United lost. This, unbeknownst to a, is a rare instance of a misprint in *The Globe*, which is itself trusted and reliable. Hence, *The Globe* is uninformative about the game's outcome (i.e. on the question of M). Nevertheless, glancing at the report yields some new information for a: proposition G, that *The Globe* reported a loss. Intuitively, receiving this new information undermines a's knowledge that M. In particular, she should rationally suspend judgment on this claim. We can model this as follows: $f_{M\wedge T\wedge G}(@) = |T| \cap |G|$, with $@ \in |T| \cap |G|$ and $|M| \cap |T| \cap |G| \subsetneq |T| \cap |G|$. Note that this accords with constraint BC, since $|M| \cap |T| \cap |G| \subseteq f_{M\wedge T\wedge G}(@)$. Thus, the information $M \wedge T \wedge G$ leaves only $T \wedge G$-worlds epistemically accessible, but allows for some $\neg M$-worlds. Hence: $@ \Vdash K_{M\wedge T\wedge G}G$ but $@ \Vdash \neg K_{M\wedge T\wedge G}M$.

What if one allows for false information? Such a theorist might describe the situation differently: since both *The Globe* and *The Times* are reliable and trusted, they both furnish information on the question of M. However, they conflict: yielding M and $\neg M$, respectively. The total information is thus $T \wedge G \wedge M \wedge \neg M$. Presumably, knowledge of M cannot be achieved here, since the conflicting pieces of information cancel each other out. Hence: $@ \Vdash K_{M\wedge T}M$ but $@ \Vdash \neg K_{T\wedge G\wedge M\wedge\neg M}M$. This is modeled with $f_{T\wedge G\wedge M\wedge\neg M}(@) = |T| \cap |G|$.

On the other hand (thanks to BC) KRI respects:

(Transitivity) $\{K_A B, K_B C\} \vDash K_A C$ [25]

Knowledge is *stable*: old knowledge cannot be lost as new knowledge is accumulated. The intuitive case for monotonicity is that it captures the core idea of the stability of knowledge. KRI suggests a different hypothesis: knowledge is stable is that it respects Transitivity. Suppose C is known on the basis of information B. And suppose that one's information is refined insofar as new information A is received upon which knowledge of B can be based. Transitivity says that C is still knowable: no knowledge is lost in the update from B to A.[26]

We illustrate with a version of the dogmatism paradox that hinges on Transitivity. Suppose:

$K_{P\wedge R}(E \supset M)$ and $K_{E\wedge P\wedge R}(P \wedge R)$

[25] *Proof.* Assume that $w \Vdash K_A B$ and $w \Vdash K_B C$. Thus: $\forall w_1(wR_A w_1 \Rightarrow w_1 \Vdash B)$ & $t(B) \leq t(A)$ and $\forall w_2(wR_B w_2 \Rightarrow w_2 \Vdash C)$ & $t(C) \leq t(B)$. Then $t(C) \leq t(B) \leq t(A)$. Further: by BC, we have that $|B| \subseteq f_B(w)$ and, by (SK), that $f_B(w) \subseteq |C|$. Thus, $|B| \subseteq |C|$. Now, by (SK) again, we have that $f_A(w) \subseteq |B|$. Hence, $f_A(w) \subseteq |C|$.

[26] This echoes the *Xerox Principle* endorsed by Dretske (1999, p. 57): if A carries the information that B, and B carries the information that C, then A carries the information that C.

with P, E, R, M as in Sect. 2.3. That is: suppose that the joint information that P is true and that E supports $\neg P$ renders it knowable that E is misleading if true; and that refining the information to $E \wedge P \wedge R$ renders it jointly knowable that P and that E supports $\neg P$. In this case, an advocate of Transitivity must accept that an agent with the refined information is positioned to know that E is misleading if true:

$$K_{E \wedge P \wedge R}(E \supset M)$$

Once again, when generalized, this seems an objectionable conclusion. However, defiance in the style of Harman (1973) is here best interpreted as doubt about the truth of $K_{E \wedge P \wedge R}(P \wedge R)$. That an agent has received, in total, the information that $E \wedge P \wedge R$ need not position her to know P: her resultant knowledge that E is true and E supports $\neg P$ defeats rational belief in P. Defiance in the style of Sharon and Spectre (2010) is here best interpreted as doubt about the truth of $K_{P \wedge R}(E \supset M)$. That an agent has received, in total, the information that $P \wedge R$ cannot, in general, position her to know that E is misleading if true. Thus, standard responses to the paradox provide little motivation for rejecting Transitivity.

An advocate of inductive knowledge might be suspicious of Transitivity.[27] Let S be the (true) claim that smoke is rising above the treeline, along with background information on the frequent correlation between smoke and wildfire. Let F be the (true) claim that there is a raging wildfire in the forest. Let C be the claim that there is an inhabited cabin in the vicinity, with a chimney leading from its fireplace. S, we suppose, provides inductive knowledge of F, in the absence of defeaters. Further, we suppose that C is exactly such a defeater. Hence, an alleged counter-example to Transitivity:

$$K_{S \wedge C}S \text{ and } K_S F, \text{ but } \neg K_{S \wedge C}F$$

That is: to receive the information that there is smoke positions one to know there is (smoke and) fire, unless defeating information is also received.

We reject this counter-example: the above formalization seems a poor representation of the scenario at issue. That smoke signals fire is analogous to a voice on a telephone signaling that Beth's grandmother is ill, the headline of *The Times* signaling that Manchester United won, or Koplik spots signaling that a patient has measles. The former situation carries information about the latter. Coming to know that there is fire on the basis of smoke is like coming to know grandma is ill from a telephone call: the information that F is thereby transmitted, in a manner conducive to knowledge. To subsequently learn of the cabin is to lose knowledge

[27] Thanks to Alexandru Baltag for highlighting the issue of inductive knowledge.

of F *despite* having received the information that F, just as one loses the knowledge that grandma is ill when given a reason to doubt the testimony of the speaker, or doubt the quality of the telephone line.[28] Such thinking is central in philosophical theories of information: the idea that information about a situation may flow to a receiver via a second situation - a *carrier* - is prominent in Dretske (1999, Ch.5), Skyrms (2010, Ch.3) and situation theory (Barwise and Etchemendy, 1987; Barwise and Seligman, 1995; van Benthem and Martinez, 2008; Seligman, 2014). Consider Skyrms (2010, p. 44):

> At this point some philosophers will say "You might as well say that Smoke carries information about fire." Well, doesn't it? Don't fossils carry information about past life forms? Doesn't the cosmic background radiation carry information about the early stages of the universe? *The world is full of information.* [Skyrms' emphasis]

A better formalization of the above scenario, therefore, does *not* bear on Transitivity:

$K_{S\wedge F\wedge C}S$ and $K_{S\wedge F}F$, but $\neg K_{S\wedge F\wedge C}F$

To receive the information that there is smoke is to receive the information that there is fire, positioning one to know there is (smoke and) fire, unless defeating information is also received.

(If skeptical that smoke carries the information that there is a wildfire for agents that know of the cabin, one might prefer: $K_{S\wedge C}S$ and $K_{S\wedge F}F$, but $\neg K_{S\wedge C}F$.)

7 Disjunction, Paradox

The following principle fails in our basic system:

(Addition) $K_A B \nvDash K_A(B\vee C)$ [29]

This inference fails for the right reason, according to theorists such as Yablo (2014), who endorse the topic-sensitivity of knowability: although $A \vDash A\vee B$, disjunction can bring in alien topics.

[28] Here *evidence* and *information* seem to pull apart. F, let's say, becomes part of one's information when one sees (and correctly interprets) the smoke. However, F does not seem to be part of one's evidence: rather, knowledge that F seems inferentially based one's evidence e.g. the appearance as of smoke.

[29] *Counterexample*: let $W=\{w,w_1\}$, wR_pw_1, $w_1 \Vdash q$, $t(p)=t(q)\neq t(r)$. Then $t(q)\leq t(p)$, so by (SK), $w\Vdash K_pq$. But $t(q\vee r)=t(q)\oplus t(r)\nleq t(p)$, thus $w\nVdash K_p(q\vee r)$.

This is easily motivated as a feature of agents that lack unlimited *conceptual* tools, even when they have unlimited deductive powers. If an agent knows that $2+2=4$ but they do not possess the concept of an astronaut, it is at best misleading to claim that their information positions them to know that either $2+2=4$ or Neil Armstrong was an astronaut.

But topic-sensitivity grounds a plausible rationale for rejecting unrestricted closure even for ideal agents with a full repertoire of concepts. A topic is closely associated with a set of distinctions, issues, or questions (Lewis, 1988; Yablo, 2014; Hawke, 2017a). To say that knowability is topic-sensitive is just to say that what is knowable on a certain body of information depends on what that information is *about*: what distinctions it speaks to; what issues it resolves, or, leaves open. Now, the most compelling counter-examples to unrestricted closure can be understood as counter-examples to Addition, rooted in an enrichment of topic or subject matter. $A \vee \neg B$ is equivalent to $\neg(\neg A \wedge B)$ twice over, that is, both in a truth-conditional sense and qua topic. Then the validity of Addition would commit one to:

$$K_A B \vDash K_A \neg(\neg B \wedge C)$$

But various cases impress philosophers as counter-examples to this principle – at least those that resist radical skepticism or Moorean dogmatism.[30] Knowing that one has hands (based on ordinary information) does not put one in a position to know that one is not a handless brain-in-a-vat (Cohen, 1988). Knowing that the wall before one is red (based on the visual information of it looking red) does not put one in a position to deny that the wall is not red but subject to trick lighting (Cohen, 2002). Knowing that the animal in the zebra enclosure is a zebra (based on the visual information that it looks like a zebra) does not put one in a position to know that the animal is not a cleverly disguised non-zebra (Dretske, 1970, 2005). Or, back to our previous example: knowing that Kripke produced result X in modal logic (based on testimony in the classroom) does not put one in a position to deny the veridicality of a letter signed by Kripke confessing that he is a fraud.

A different disjunction-involving issue, also adequately modeled in the semantics, has to do with the fact that KRI can be nonprime due to indeterminacy in the available information. Your information is sufficient for you to come to know that Mary is either left- or right-handed (you have seen that she is a normally endowed human being, etc.), but insufficient to establish which one it is. So we need, and we get:

[30]For further discussion, see Hawke (2016). For an opposing verdict, see Roush (2010) for a nuanced defense of the validity of the above principle. For push-back, see Avnur et al. (2011) and Hawke (2017b, Sect. 3.4.5).

$K_A(B \vee C) \nvDash K_A B \vee K_A C$ [31]

Here, too, the inference fails for the right reason. Topicality is there, but the different worlds one has access to will fill in the unspecified details in different ways. There can be worlds where B but not C, and worlds where C but not B, and both can be compatible with what information one has.

8 Omniscience

KRI invalidates the rule of Logical Omniscience from 2.2.

$\vDash B \vee \neg B \not\Rightarrow \vDash K_A(B \vee \neg B)$

Topic-preservation fails: $t(B)$ need not be included in $t(A)$, and so $t(B \vee \neg B)$ need not be included in $t(A)$. This is a happy outcome if the goal to reason about agents lacking the total conceptual repertoire.

On the other hand, as we envision the KRI semantics as governing agents that are idealized in the sense of being logically astute and fully rational, one would expect some version of logical omniscience to be captured by the system. It is straightforward to see that KRI validates a principle of omniscience with a topicality constraint:

(Topical Omniscience) If $\vDash B$ and for all models $t(B) \leq t(A)$, $\vDash K_A B$

For instance:

$\vDash K_{A \wedge B}(B \vee \neg B)$

A logically astute agent will always be in the position to know a logical truth, once she is provided with information allowing her access to the concepts involved in it.

[31] *Counterexample*: let $W = \{w, w_1, w_2\}$, $wR_p w_1$, $wR_p w_2$, $w_1 \Vdash q$ but $w_1 \nVdash r$, $w_2 \Vdash r$ but $w_2 \nVdash q$, $t(p) = t(q) = t(r)$. Then by (S$\vee$), $w_1 \Vdash q \vee r$ and $w_2 \Vdash q \vee r$, so for all w_x such that $wR_p w_x$, $w_x \Vdash q \vee r$. Also, $t(q \vee r) = t(q) \oplus t(r) \leq t(p)$, thus by (S$K$), $w \Vdash K_p(q \vee r)$. However, $w \nVdash K_p q$ and $w \nVdash K_r r$ for both q and r fail at some R_p-accessible world. Thus by (S$\vee$), $w \nVdash K_p q \vee K_p r$.]

9 Closure under (Known) Implication

Another invalidity displays *the* essential form of closure failure for KRI, namely that of Closure Under Strict Implication from 2.2:

$$\{K_A B, B \prec C\} \nvDash K_A C$$ [32]

Although all the B-worlds are C-worlds, thus all the A-selected B-worlds are C-worlds, the strict implication is not topic-preserving: A may be information about the topic of B yet not be information about the topic of C. Thus, given A, one can come to know B but not C even if there just is no way for B to be true while C is not.

The idea applies nicely to Cartesian skepticism (Dretske, 1970): one's ordinary empirical information, delivered via sensory perception, may put one in the position to know one has hands. Having hands is incompatible with the possibility that one is a bodiless brain in a vat whose phenomenal experience is systematically misleading. Yet it might seem implausible that ordinary empirical information puts one in a position to rule out a brain-in-vat scenario.

Crucially, KRI not only invalidates Closure Under Strict Implication, but seems to invalidate the right instances of the principle. Looking again at the results of Sect. 5 and 7: $K_A(B \wedge C)$ ensures $K_A B$, but $K_A B$ does not ensure $K_A(B \vee C)$. While the former appears indisputable, it is far from clear that knowability is closed under the introduction of arbitrary disjuncts: intuitively, the received information may not be about the topic of the alien disjunct.

Or suppose that one rejects unrestricted closure on the basis that various epistemic paradoxes (the dogmatism paradox, the Cartesian paradox) are best interpreted as counter-instances. One should then hope to invalidate any instance of Closure Under Strict Implication that can be used to construct such a paradox. Now the semantics provides the following (easily, via failure of topic-preservation):

$$\{K_A B, B \prec C\} \nvDash K_A(B \wedge C)$$

If one accepts that $K_A B$ and $B \prec C$ always ensures $K_A(B \wedge C)$ then various paradoxes can be constructed. Suppose that $K_A(P \wedge R)$ and $(P \wedge R) \prec (E \supset M)$, where P, E, R, M are as in the Kripke-Harman dogmatism paradox. Then we could conclude: $K_A((P \wedge R) \wedge (E \supset M))$. In other words, if it is known both that P and that E is generally a reason to reject P, then we could draw the dogmatic

[32] *Counterexample*: let $W = \{w, w_1\}$, $wR_p w_1$, $w \nVdash q$, $w_1 \Vdash q$, $w_1 \Vdash r$, $t(p) = t(q) \neq t(r)$. Then $f_p(w) \subseteq |q|$ and $t(q) \leq t(p)$, thus by (SK), $w \Vdash K_p q$. Also, $|q| \subseteq |r|$, thus by (S-$\prec$), $w \Vdash q \prec r$. But although $f_p(w) \subseteq |r|$, $t(r) \nleq t(p)$, thus $w \nVdash K_p r$.

conclusion that it is knowable that P, that E is generally a reason to reject P *and* that if E were true then E would be misleading evidence. This dogmatic conclusion seems no better than that in the original puzzle.

On the other hand, closure under known material implication does hold – and for good reasons. In the current setting, call this principle *Closure Over Known Implication and Topic*:

(COOKIT) $\{K_A B, K_A(B \supset C)\} \vDash K_A C$ [33]

COOKIT should hold. Here, both B and $B \supset C$ are fully on-topic with respect to the information that A. Also, relative to that information, it is knowable both that B and that if B is true, C is. Then the agent is in a position to know that C, relative to the same information A (the final proviso is essential: given the non-monotonic features of K highlighted above, the inference may fail if the index for the available information is allowed to change across the involved formulas). If, for instance, your information puts you in the position to know both that Peano's postulates are true and that if these are then Goldbach's conjecture is, then you will also be in the position to know Goldbach's conjecture.

Authoritative closure sympathizers tend to cite the powerful intuition that the conclusion of a deductive argument from known premises must result in knowledge: see, for instance, Williamson (2000, p.118); Hawthorne (2004, Sect.1.5); Kripke (2011a, p.200). After all, this is the basis for the entire enterprise of mathematics: few want to deny the epistemic sanctity of mathematical results. This is often translated into a conviction in closure under strict implication – at least, if we restrict attention to computationally unbounded and fully rational agents, for, the rationale goes, the truth of $B \multimap C$ is best understood in the setting of epistemic logic as an a priori truth of some kind.

Now a proponent of KRI need not deny the intuition that deduction is a sanctified means for extending knowledge. She can dispute, however, that closure under strict implication best captures this intuition, given apparent counter-examples that can be extracted from epistemic paradoxes. Instead, she posits COOKIT as the uncontroversial core of the intuition.

10 Closure, A Priority

Does acceptance of COOKIT court trouble with regards to epistemic paradox? It might be proposed, for instance, that the dogmatism paradox can be reconstructed

[33] *Proof*: let $w \Vdash K_A B$ and $w \Vdash K_A(B \supset C)$. By the former and (S$K$), for all w_1 such that $wR_A w_1$, $w_1 \Vdash B$, and $t(B) \leq t(A)$. By the latter and (SK) again, for all w_1 such that $wR_A w_1$, $w_1 \Vdash B \supset C$. Thus for all w_1 such that $wR_A w_1$, $w_1 \Vdash C$. Also, $t(B \supset C) = t(B) \oplus t(C) \leq t(A)$, thus $t(C) \leq t(A)$. Thus by (SK), $w \Vdash K_A C$.

for an adherent of COOKIT. The story is told as follows: suppose that a has the information that $P \wedge R$ at time t_0. Further, since it is knowable a priori that $(P \wedge R) \dashv\!\!\exists\ (E \supset M)$, it is also knowable a priori $(P \wedge R) \supset (E \supset M)$, and hence knowable on the basis of $P \wedge R$ that $(P \wedge R) \supset (E \supset M)$. But then COOKIT yields that a is in a position to know, on the basis of $P \wedge R$, that E must be misleading if true.

This reasoning betrays a confusion. A proponent of KRI need not accept that if it is knowable a priori that $(P \wedge R) \supset (E \supset M)$, then it is knowable on the basis of $P \wedge R$ that $(P \wedge R) \supset (E \supset M)$. In general, she need not accept that if A is knowable a priori then A is knowable on the basis of *every* body of information B. This is not licensed by the intuitive reading of 'on the basis of' that has been exploited. It is knowable a priori that $2+2=4$. It would be odd to conclude that $2+2=4$ can be known on the basis of the news that Beth's grandmother is ill.

This clarifies that the semantics embeds an absolute notion of a priority: what can be known without any empirical information by a computationally unbounded, fully rational agent with access to the full repertoire of concepts. Contrast the notion of relative a priority: what can be known without empirical information, given a fixed, possibly incomplete, universe of concepts. Let $\top$ denote one's favorite tautology. Then we read $\top \dashv\!\!\exists\ A$ as 'A is a priori'. For A is knowable a priori exactly when conceptual limitations are forgotten and A is true at every possible world (this is only plausible, of course, if worlds are understood as basic epistemic possibilities, possibly in contrast to basic metaphysical possibilities).

With this in mind, a proponent of the KRI semantics can judge the case from the dogmatism paradox as follows. First, the case is best described as:

$$K_{P\wedge R}(P \wedge R)$$

$$\neg K_{P\wedge R}((P \wedge R) \dashv\!\!\exists\ (E \supset M))$$

$$\top \dashv\!\!\exists\ ((P \wedge R) \dashv\!\!\exists\ (E \supset M))$$

$$\neg K_{P\wedge R}(E \supset M)$$

That is, though $P \wedge R$ is known on the basis of the empirical information $P \wedge R$, it is not knowable on this basis that $(P \wedge R) \dashv\!\!\exists\ (E \supset M)$. Rather, this fact is known a priori. In particular, this knowledge is based on concepts that go beyond those that comprise the topic of $P \wedge R$. In this case, in accord with COOKIT, the proponent of our system can deny that a is positioned by her empirical information to know that $E \supset M$.

11 Minimal Conditional Logic

Three final principles KRI does not validate:

(Reflexivity) $\not\models K_A A$

(Cautious Transitivity) $\{K_A B, K_{A\wedge B} C\} \not\models K_A C$

(Cautious Monotonicity) $\{K_A B, K_A C\} \not\models K_{A\wedge B} C$ [34]

Gabbay (1985) proposes these as a minimal foundation for a logic of non-monotonic derivations. In particular, they hold appeal as a base logic of *ceteris paribus* conditionals. The KRI semantics piggy-backs on a non-monotonic conditional logic. Thus, we consider whether there are *prima facie* motivations for rejecting these validities in a minimal logic for KRI.

Start with Reflexivity. Consider the following line of reasoning:

1. $K_B C$ (Assumption)
2. $K_{A\wedge B}(A \wedge B)$ (by Reflexivity)
3. $K_{A\wedge B} B$ (by Simplification)
4. $K_{A\wedge B} C$ (by Transitivity)
5. If $K_B C$ then $K_{A\wedge B} C$ (by discarding 1)

Thus, if Reflexivity is conjoined with background principles we found to be independently good, we validate Monotonicity – in the form the dogmatism paradox calls into question. One who accepts Reflexivity must either reject a Harman-like response to the dogmatism paradox, or bear the cost of rejecting Simplification or Transitivity.

Counter-examples to Reflexivity can arguably be furnished. If a theorist allows non-veridical information, counter-examples are obvious: if an agent's total information I has a false part, then factivity assures that the agent does not know that I. Plausible counter-examples exist even when restricting to veridical information: examples in Sect. 6 can be adapted to this effect. Here is another: suppose that Mary watches Barack Obama deliver his state of the union address, from a

[34] Let p, q be atomic formulas (in all of the following, topic-assignments don't matter). First, a countermodel to Reflexivity. Let $W = \{w_1, w_2\}$, let $|p| = \{w_1\}$ and let $f_p(w_1) = W$. It follows that $w_1 \not\Vdash K_p p$. Second, a countermodel to Cautious Monotonicity. Let $W = \{w_1, w_2\}$. Let $|p| = W$ and $|q| = \{w_1\}$. Let $f_p(w_1) = |p|$ and $f_{p\wedge p}(w_1) = |q|$. It follows that $w_1 \Vdash K_p p \wedge K_{p\wedge p} q \wedge \neg K_p q$. Finally, a countermodel to Cautious Transitivity. Let $W = \{w_1, w_2\}$. Let $|p| = \{w_1\}$. Let $f_p(w_1) = |p|$ and $f_{p\wedge p}(w_1) = W$. It follows that $w_1 \Vdash K_p p \wedge \neg K_{p\wedge p} p$.

front row seat, hearing distinctly that his first topic is trade. A week later, Mary's memory of the speech remains vivid. Presumably, her senses informed her that his first topic was trade, she thereby came to know it, and she now preserves this knowledge via memory. However, an epistemic peer then claims that Obama's first topic was gun control, reminding Mary that her memory can be unreliable. Given this, it can be rational for Mary to suspend (or weaken) her belief that the first topic was trade, losing her knowledge. Nevertheless, it remains true, in an important sense, that Mary has the information that the first topic is trade (T): she received that information through a perceptual event that, at the time, was conducive to knowledge. The event and its interpretation remain vividly stored in her memory. So, if $I \wedge T$ is Mary's total information: $\neg K_{I\wedge T}(I \wedge T)$.[35]

Turn to Cautious Transitivity and Cautious Monotonicity. The semantics invalidates these because set-selection functions operate on formulas and, in the base framework, few constraints regulate how sets are selected for different formulas. BC requires the sets selected for p and for $p \wedge p$, for instance, to both contain every p-world, but otherwise, no constraint is imposed. Models are allowed where, for instance, $|p| \subsetneq W$, $f_p(w) = |p|$ and $f_{p\wedge p}(w) = W$.

Thus, the question as to whether Cautious Transitivity and Cautious Monotonicity should be treated as logical truths is bound up with substantive issues: does a piece of information have a logical structure, and in particular one that mirrors the syntax of a sentence with which it is expressed? If so, to what extent should an epistemic logic accommodate agents whose cognition is sensitive to syntax? One might wish to model agents whose capacity to extract knowledge from information tracks the complexity of the information's structure. This impulse is waged against an insistence that $f_{A\wedge B}$ always selects the same set as f_A when, say, $|A| = |A \wedge B|$.

One possible view has it that information is unstructured. Or one might accept that information is structured, but hold that this structure should be ignored when dealing with idealized agents, as in the KRI setting. In this case, since p and $p \wedge p$ have the same topic and truth set, they should be treated as equivalent. With this in mind, consider the class of models that satisfy a Twice Over Equivalence principle:

(TOE) $|A| = |B|$ & $t(A) = t(B) \;\Rightarrow\; f_A(w) = f_B(w)$ for all w

TOE says that if A and B are equivalent twice over (both true at the same worlds and about the same same topic), then their set-selection functions always output the same values. Models complying with TOE filter out a number of syntactic differences concerning the way information is presented, but still allow, via

[35] Compare the famous position of Sellars (1997) that 'the given' is a myth: even basic perceptual evidence - and the knowledge directly based thereon - is subject to revision and defeat.

differences in topicality, hyperintensional distinctions involving pieces of information with coincident truth sets. It is easy to check that Cautious Monotonicity and Cautious Transitivity are validated if we impose this restriction on the admissible models. It may also be confirmed that, however, Monotonicity, in full generality, is *not* validated by this restricted class. Compliance with TOE allows that if

$$|A \wedge B| \subsetneq |A \wedge (B \vee \neg B)|$$

then the set selected for $A \wedge B$ need not be a subset of that selected for $A \wedge (B \vee \neg B)$, despite these sentences sharing a topic and the former entailing the latter.

12 Conclusion and Further Work

This paper only presented a first exploration of KRI – a general epistemic logic framework which seems to us both formally simple and capable of properly dealing with a number of issues in mainstream epistemology. A first direction of development would consist, of course, in coming up with a proof system, sound and complete with respect to the semantics. One second direction may come, as hinted above, from making the framework dynamic in the sense of dynamic epistemic logic, thereby capturing the process of knowledge update on the basis of newly acquired information by means of model-transformations.

A third direction may be to relate and compare our semantics with recent work in aboutness and truthmaker theory mentioned above, such as Yablo (2014) and Fine (2014, 2015). So far these theories have not been developed having epistemic notions in sight (although Yablo's Chapter 7 does get into the relations between knowledge and aboutness, in particular, in connection to epistemic closure). While Yablo retains a possible worlds apparatus, characterizing subject matters – what sentences are about – as divisions of the space of worlds, Fine is not friendly to the notion of world and works with a space of truthmakers which can be fused into further truthmakers. We have followed an intermediate path, combining possible worlds with a mereology of topics. Comparing these different approaches, possibly in order to assess their relative merits, makes for further interesting work.

References

Y. Avnur, A. Brueckner, and C. T. Buford. No Closure on Skepticism. *Pacific Philosophical Quarterly*, 92:439–447, 2011.

J. Barwise and J. Etchemendy. *The Liar: An Essay on Truth and Circularity*. Oxford University Press, 1987.

J. Barwise and J. Seligman. *Information Flow, The Logic of Distributed Systems*. Cambridge University Press, 1995.

F. Berto. Aboutness in imagination. *Philosophical Studies*, 2017. doi:10.1007/s11098-017-0937-y. URL `https://doi.org/10.1007/s11098-017-0937-y`.

F. Berto. Simple hyperintensional belief revision. *Erkenntnis*, 2018. doi:10.1007/s10670-018-9971-1. URL `https://doi.org/10.1007/s10670-018-9971-1`.

B. Chellas. Basic conditional logic. *Journal of Philosophical Logic*, 4:133–53, 1975.

R. Chisholm. *The Problem of the Criterion*. Marquette University Press, 1973.

S. Cohen. How to Be a Fallibilist. *Philosophical Perspectives*, 2:91–123, 1988.

S. Cohen. Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2):309–329, 2002.

F. Dretske. Epistemic Operators. *Journal of Philosophy*, 67(24):1007–1023, 1970.

F. Dretske. *Knowledge and the Flow of Information*. CLSI Publications, 1999.

F. Dretske. The Case Against Closure. In Steup, Turri, and Sosa, editors, *Contemporary Debates in Epistemology*. Wiley Blackwell, first edition, 2005.

R. Fagin and J. Halpern. Belief, awareness and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, 1995.

T.M. Ferguson. A computational interpretation of conceptivism. *Journal of Applied Non-Classical Logics*, 24(4):333–367, 2014.

K. Fine. Truthmaker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43:549–77, 2014.

K. Fine. Angellic content. *Journal of Philosophical Logic*, On line first, 2015.

L. Floridi. Semantic conceptions of information. *The Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/information-semantic/, 2015.

D. Gabbay. *Theoretical Foundations for Non-monotonic Reasoning in Expert Systems*. Springer, Berlin, 1985.

G. Harman. *Thought*. Princeton University Press, 1973.

P. Hawke. Questions, topics and restricted closure. *Philosophical Studies*, 73(10): 2759–2784, 2016.

P. Hawke. Theories of aboutness. *Australasian Journal of Philosophy*, 2017a. doi:10.1080/00048402.2017.1388826. URL `https://doi.org/10.1080/00048402.2017.1388826`.

P. Hawke. *The Problem of Epistemic Relevance*. PhD thesis, Stanford University, 2017b. URL `https://www.illc.uva.nl/Research/Publications/Dissertations/DS-2017-06.text.pdf`. Updated ILLC Version.

J. Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.

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

W. Holliday. Fallibilism and Multiple Paths to Knowledge. In T. Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 5, pages 97–144. Oxford University Press, 2015.

W. Holliday and T. Icard. Moorean phenomena in epistemic logic. In L. Beklemishev, V. Goranko, and V. B. Shehtman, editors, *Advances in Modal Logic*, volume 8, pages 178–199. College Publications, 2010.

L. Humberstone. *Philosophical Applications of Modal Logic*. College Publications, 2016.

S. Kripke. Nozick on Knowledge. In *Philosophical Troubles: Collected Papers, Volume 1*. Oxford University Press, 2011a.

S. Kripke. On Two Paradoxes of Knowledge. In *Philosophical Troubles: Collected Papers, Volume 1*. Oxford University Press, 2011b.

J. Kvanvig. *The Knowability Paradox*. Oxford University Press, 2006.

M. Lasonen-Aarnio. The Dogmatism Puzzle. *Australasian Journal of Philosophy*, 92(3):417–432, 2014.

K. Lawlor. *Assurance: An Austinian View of Knowledge Claims*. Oxford University Press, 2013.

H. Levesque. A logic of implicit and explicit belief. *National Conference on AI*, AAAI-84:198–202, 1984.

D. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.

D. Lewis. Relevant Implication. *Theoria*, 54(3):161–174, 1988.

D. Makinson. The Paradox of the Preface. *Analysis*, 25(6):205–207, 1965.

J.J. Meyer. Epistemic logic. In L. Goble, editor, *The Blackwell Guide to Philosophical Logic*, pages 182–202. Blackwell, Oxford, 2001.

R. Nozick. *Philosophical Explanations*. Harvard University Press, 1981.

J. Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.

S. Roush. Closure on Skepticism. *Journal of Philosophy*, 107(5):243–256, 2010.

J. Schaffer. Closure, Contrast and Answer. *Philosophical Studies*, 133(2):233–255, 2007.

J. Seligman. Situation theory reconsidered. In A. Baltag and S. Smets, editors, *Johan van Benthem on Logic and Information Dynamics*, volume 5 of *Outstanding Contributions to Logic*, pages 895–932. Springer, 2014.

W. Sellars. *Empiricism and the Philosophy of Mind*. Harvard University Press, 1997.

A. Sharon and L. Spectre. Dogmatism repuzzled. *Philosophical Studies*, 148(2): 307–321, 2010.

A. Sharon and L. Spectre. Evidence and the openness of knowledge. *Philosophical Studies*, 174:1001–1037, 2017.

B. Skyrms. *Signals*. Oxford University Press, 2010.

R. Sorensen. Dogmatism, junk knowledge, and conditionals. *Philosophical Quarterly*, 38:433–454, 1988.

E. Sosa. *Epistemology*. Princeton Foundations of Contemporary Philosophy. Princeton University Press, 2017.

R. Stalnaker. The problem of logical omniscience, i. *Synthese*, 89(3):425–440, 1991.

J. van Benthem. What One May Come to Know. *Analysis*, 64:95–105, 2004.

J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.

J. van Benthem and M. Martinez. The Stories of Logic and Information. In P. Adriaans and J. F. A. K. van Benthem, editors, *Handbook of the Philosophy of Information*, pages 217–280. Elsevier Science Publishers, 2008.

H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, Dordrecht, 2008.

H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi, editors. *Hanbook of Epistemic Logic*. College Publications, London, 2015.

J. Vogel. $e\&\neg h$. In D. Dodd and E. Zardini, editors, *Scepticism and Perceptual Justification*, pages 87–107. Oxford University Press, 2014.

T. Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000.

S. Yablo. *Aboutness*. Princeton University Press, Princeton, 2014.