The Fundamental Problem of Logical Omniscience

Peter Hawke^{1,2} · Aybüke Özgün^{1,2} · Francesco Berto^{1,2}

Received: 19 June 2019 / Accepted: 12 November 2019 / Published online: 20 December 2019
© The Author(s) 2019

Abstract

We propose a solution to the problem of logical omniscience in what we take to be its fundamental version: as concerning arbitrary agents and the knowledge attitude per se. Our logic of knowledge is a spin-off from a general theory of thick content, whereby the content of a sentence has two components: (i) an intension, taking care of truth conditions; and (ii) a topic, taking care of subject matter. We present a list of plausible logical validities and invalidities for the logic of knowledge per se for arbitrary agents, and isolate three explanatory factors for them: (1) the topic-sensitivity of content; (2) the fragmentation of knowledge states; (3) the defeasibility of knowledge acquisition. We then present a novel dynamic epistemic logic that yields precisely the desired validities and invalidities, for which we provide expressivity and completeness results. We contrast this with related systems and address possible objections.

Keywords Aboutness · Subject matter · Logical omniscience · Dynamic epistemic logic · Thick content

1 Introduction

Formal epistemology tends to bake in logical omniscience. Epistemic logic in the style of [33] models a body of knowledge as a set of possible worlds $\mathcal{K}$ and takes φ as known when φ is true at every world in $\mathcal{K}$. It follows that every logical truth is known

✉ Francesco Berto
fb96@st-andrews.ac.uk

Peter Hawke
P.M.Hawke@uva.nl

Aybüke Özgün
a.ozgun@uva.nl

¹ ILLC, University of Amsterdam, Amsterdam, The Netherlands

² Arché, University of St. Andrews, St. Andrews, UK

by every agent, and if φ is known then so are its logical consequences. Relatedly, a model going back to [11] takes propositional content to be a set of possible worlds. If an agent's doxastic state is then modeled as a mere set of contents, the resulting doxastic logic seems objectionably weak: since $\varphi \wedge \psi$ and φ are generally true at different worlds, this picture allows for agents that believe the former without the latter. Taking belief as closed under entailment provides a simple fix, but renders all logical truths universally believed. Further, Bayesianism models an agent's credences relative to evidence as a probability function. As this standardly assigns 1 to every logical truth, the agent is represented as fully believing all of them, as well as every logical consequence of the evidence [12, Ch.6].

Fans of these formalisms have room to maneuver. One may revise logic or propositional content, or complicate the relationship between attitudes and content [13, 42, 51, 56]. One may interpret the formalisms as being about a perfectly rational Bayesian believer, or about what an ideal agent may infer from its body of knowledge [33, Sect. 2.5].

But deep philosophical issues are nearby. Say that a *problem of logical omniscience* arises for attitude A , relative to a class of agents, when either: (i) it is unclear whether A is closed under deductive consequence at all; or, (ii) given that A is not so closed, it is unclear what restricted closure principles hold for it; or (iii) though it is reasonably clear what principles should hold, it is unclear how to generate them with extant formal tools. (i) and (ii) bear on our basic understanding of propositional attitudes. (i) is widespread: epistemologists debate whether skeptical paradoxes show that even ideal agents cannot know every deductive consequence of their knowledge [15, 16, 32, 65]. Bayesians debate whether even rational credence is closed under logical consequence [59, Sect.5.4]. (ii) is widespread: everyday belief ascriptions seem not to be closed under entailment, but it is less clear what logic, if any, they *do* exhibit [48]. So problems of logical omniscience can be framed for ordinary knowers. Or, for various ideal knowers: with bounded but otherwise perfect cognitive powers [18, Ch. 10]; normatively ideal [59]; or cognitively ideal [35]. Or, for derivative attitudes, such as 'logical commitment given what one knows' [43], or 'knowability in principle given certain evidence or information' [4, 29].

A *fundamental* version of the problem is hit when one focuses on arbitrary agents, abstracting away from cognitive skills and rational commitments; and on the knowledge attitude per se, as opposed to derivative ones. This is the topic of the present paper: the intrinsic logic of the knowledge attitude, based only on the literal content of knowledge ascriptions such as ' a knows that φ ' and plausibly universal features of a body of knowledge, irrespective of what a can or cannot compute, its working memory, or normative constraints.

The fundamental problem has a static aspect and a dynamic aspect. The static: what further facts must an arbitrary agent know if she knows $\varphi_1, \dots, \varphi_n$? The dynamic: how is an arbitrary agent's body of knowledge κ altered by her coming to know φ ? (The dynamic may be less well-known than the static, but has been explored in the literature – we'll get back to this.) For both, we answer by identifying a number of pre-theoretic logical principles, addressing (ii) above. We then use them to motivate a novel logical system, addressing (iii).

Our proposed system capitalizes on two resources. One is a general theory of propositional content we have been developing in various works [2–4, 29, 30], in the ballpark of [65]’s aboutness theory and [23]’s truthmaker semantics. The other is the idea that knowledge is defeasible: adding a piece of knowledge to an existing body of knowledge can lead to loss of knowledge. The latter is a well-known stance [10, 27]. Here’s a primer on the former: treating content as a bare set of possible worlds leaves no hint as to how or why a sentence with that content is true at those worlds. ‘ $2+2=4$ ’ and ‘Equilateral triangles are equiangular’ are true at the same possible worlds (all of them), but differ in content. To grasp that difference is to grasp which aspects of the world bear respectively on their truth. We need to supplement truth conditions with an account of *aboutness*, ‘the relation that meaningful items bear to whatever it is that they are *on* or *of* or that they *address* or *concern*’ [65, p. 1], namely their *topic*, or *subject matter*. Further, what a sentence is about can, intuitively, be properly included in what another one is about. Relatedly, contents stand in mereological relations [65, Section 2.3], [23, Sections 3–5]: they may be merged into wholes which inherit the proper features from the parts [65, Section 3.2]. Thus, the content of an interpreted sentence is fruitfully identified with the *thick proposition* it expresses [65, Sect. 3.3]. This has two components: (i) intension and (ii) topic or subject matter. To know what φ means, one needs to know both what it is *about*, and what it *says* about that topic. φ and ψ express the same thick proposition when (i) they are true at the same worlds, and (ii) they have the same topic. The content of φ includes that of ψ when (i) any φ -world is a ψ -world and (ii) the topic of ψ is part of that of φ . The Boolean connectives are transparent qua subject matter: φ has the same subject matter as $\neg\varphi$ (‘Snow is white’ and ‘Snow is not white’ address the same topic: whether snow is white). $\varphi \wedge \psi$ and $\varphi \vee \psi$ have the same subject matter: a merging of the subject matter of φ and that of ψ . This account of content delivers patterns of logical entailment applicable to any basic propositional attitude, mirroring judgments about indirect speech reports. In particular, it delivers the logic of knowledge ascriptions presented below.

Our approach is reminiscent of [66]’s work on the topic-sensitivity of belief, and of dynamic epistemic logic [60–62]. Yalcin, however, does not present a full-fledged logical system and does not treat propositions as thick. We also depart (Section 4.2) from the characterization of subject matters as partitions of worlds Yalcin borrows from [45]. We follow [60, 61] in thinking that there is a distinctive dynamic aspect to the logical omniscience problem. Authors debating epistemic closure often phrase the latter in terms of competent deduction from known premises, and understand deduction dynamically, as a process [31, 64]: given that one competently deduces ψ from $\varphi_1, \dots, \varphi_n$, all the while knowing the latter, doesn’t one know ψ ? [34], pp. 15–16, distinguishes (what he calls) *pure* and *deductive* epistemic closure principles: the former are such that one just cannot know the premise or antecedent without knowing the conclusion or consequent; the latter (including Hawthornian-Williamsonian closure) require some inferential action to be performed. Van Benthem investigates dynamic closure principles whose general shape has it that $K\varphi$ and $K(\varphi \rightarrow \psi)$ implies $[a]K\psi$, where $[a]$ is a dynamic operator encoding the transformation of the relevant model according to a recipe specified by (inferential) action a . Bjerring and Skipper [7] follow him in thinking that the problem of logical omniscience has

to be addressed dynamically. The treatment of logical omniscience in [60] distinguishes explicit and implicit knowledge [17, 42]. The former is understood statically: it gives the agent's actual knowledge and has negligible logical structure. The latter is a derived attitude, understood dynamically and closed under full classical (monotonic) logical consequence. We argue below that our system gives a richer account of the logical structure of actual knowledge, and avoids over-estimating the fruits of inferential dynamics.

We proceed as follows. Section 2 explains the fundamental problem and critically discusses some attempted responses from the literature. Section 3 presents a list of intuitive logical validities and invalidities for the intrinsic logic of knowledge, taken as defeasible criteria of adequacy for an epistemic logic that solves the fundamental problem. Section 4 spells out our logical system and shows that it meets such criteria. Interesting divergences with traditional dynamic epistemic logic, and with topic-sensitive accounts of doxastic notions, are noted. Section 5 anticipates some objections and offers replies and, where necessary, refinements. A twofold appendix provides proofs for theorems from the philosophical discussion, and results concerning soundness and completeness.

2 Logical Omniscience

2.1 The Goal

Start with three questions:

- (1) *Internal Logic*. What is the internal logic of knowledge? That is: given that an arbitrary agent knows certain facts, what further facts can we conclude she knows?
- (2) *Potential Knowledge*. What can be known on the basis of a body of knowledge? That is: what can an arbitrary agent know via inference from her existing knowledge?
- (3) *Content*. How best to model the contents of knowledge? In particular, how best to model a body of knowledge – the sum total of what an agent knows?

The three are related: the logic of the knowledge attitude depends on the nature of its possible contents. In particular, if contents can stand in mereological relations then an arbitrary agent stands in the knowledge relation to content P when she knows every part of P :¹ [65, Sect. 7.3] calls this 'immanent closure'. In specific cases, it is obvious that one knowledge claim follows from another: one cannot know that Jones is horrendously late without knowing that Jones is late. Such intuitions are *prima facie* data for evaluating a general account.

We intend the answers to (1), (2), (3) to be *purely descriptive* and *universal*. We do not address normative issues, such as: given that an agent knows P , what is she

¹Throughout, P, Q refer to arbitrary propositions, understood to serve as contents for attitudes and meaningful claims; φ, ψ refer to arbitrary interpreted sentences.

thereby permitted to know, or obliged to know? Nor are we asking what a cognitively ideal agent knows given knowledge P , or what a computationally bounded agent thereby knows, inferentially or otherwise. These questions generate interesting issues for which the label ‘problem of logical omniscience’ might be apt [6, 7, 33, 35, 38, 58, 59]. But they are not our fundamental problem. This bears on the intrinsic features of knowledge, and how language reports on them.

Logical tools help with (1), (2), (3). Take a modal propositional language $\mathcal{EL}$ with knowledge modality K and update modality $[]$. Read ‘ $K\varphi$ ’ as ‘the agent knows that φ ’. Read ‘ $[\varphi]\psi$ ’ as ‘ ψ holds after the agent comes to know φ ’.² Suppose we settle on an intuitive epistemic logic, $\mathbb{IEL}$, encompassing a list of plausible validities and invalidities given our intended interpretation of $\mathcal{EL}$. Then (1), (2), (3) compress into a challenge:

- (4) *Goal*. Provide a semantics for $\mathcal{EL}$, such that the compositional meanings of $K\varphi$ and $[\varphi]\psi$ yield $\mathbb{IEL}$.

To solve our fundamental problem is now to meet the Goal.

The solution must be squeezed between two extremes:

- (5) *Lower Bound*: the logic of knowledge *per se* is non-trivial.
 (6) *Upper Bound*: knowledge *per se* is not closed under (classical, monotonic) entailment.

As for *Lower Bound*: given operator $\odot$, a validity is in the *single-premise internal logic* of $\odot$ just in case it is of the form $\odot\varphi \supset \odot\psi$, with $\supset$ the material conditional. The single-premise internal logic of $\odot$ is trivial when $\odot\varphi \supset \odot\psi$ is a validity only if $\varphi = \psi$ (φ and ψ are syntactically identical). The single-premise internal logic of knowledge is *not* trivial. Knowledge distributes over conjunction [64, pp. 276–9]: knowing both Jones and Smith are late entails knowing Jones is late. A satisfactory epistemic logic validates $K(p \wedge q) \supset Kp$, where p and q are atoms.

As for *Upper Bound*: an agent may know something without knowing all of its consequences. Mathematical knowledge on the basis of a conjunction of simple axioms is often elusive: one can know Peano’s Axioms, and these can entail, say, Goldbach’s Conjecture; but one doesn’t know if Goldbach’s Conjecture is true.

As for the dynamics: add update modalities of the form $[\varphi]$ to our language with $\odot$. Suppose the following schema is valid: $[\varphi]\odot\varphi$. Hence, $[\varphi]$ is a ‘bringing it about that $\odot\varphi$ ’ operation. Then, the *single-step internal dynamics* of $\odot$ is the set of validities of the form $[\varphi]\odot\psi$. The internal dynamics is trivial just in case $[\varphi]\odot\psi$ is a validity only if $\varphi = \psi$. The internal dynamics of knowledge is not trivial. If one comes to know that Jones is an expert lawyer, then one comes to know that Jones is a lawyer. Yet for ordinary agents, coming to know φ needn’t result in knowing every consequence of φ : complicated mathematical theorems may elude one even after one comes to know the conjunction of all relevant axioms.

²Studying a language with both K and $[\varphi]$ would be uninteresting if one were convinced that, say, $(K\varphi \wedge K\psi) \supset K\chi$ and $K\varphi \wedge [\psi]K\chi$ are logically equivalent. However, one who holds that knowledge is defeasible denies this equivalence. As we do: we discuss an example shortly.

The Lower and Upper Bound yield desiderata for any framework aiming to achieve the Goal:

- (7) *Lower Desideratum*: the framework yields validities of the form $K\varphi \supset K\psi$ and $[\varphi]K\psi$, where $\varphi \neq \psi$.
- (8) *Upper Desideratum*: there exist φ and ψ where φ entails ψ yet $K\varphi \wedge \neg K\psi$ and $[\varphi]\neg K\psi$ are satisfiable.

2.2 Discussing Some Literature

Here are frameworks that fail to meet some desideratum. On one approach, content is unstructured: an agent's total body of knowledge at actuality @ is modeled as a set of possible worlds. Given that gaining information is narrowing the space of possibilities, the set represents the information the agent has about @ (compare the notion of 'information as range' in [24] and [60]). $K\varphi$ is true at @ exactly when φ is true at every possible world compatible with the agent's information. This, of course, is the classic approach of [33]: epistemic logic is but a normal modal logic. Logical omniscience follows: $K\varphi \supset K\psi$ is valid when φ entails ψ . The Upper Desideratum (8) is missed.

On another approach, content is as structured as a syntax: an agent's total knowledge at @ is modeled as a set of sentences $\mathcal{K}_@$ in a suitable formal language (say, $\mathcal{EL}$ itself).³ (One may then understand the contents of a knowledge attitude as sentences in a language of thought; but see [13] for a critique.) $K\varphi$ is true at @ exactly when $\varphi \in \mathcal{K}_@$. If every set of sentences is a possible body of knowledge, then no significant logical constraints are imposed on the internal logic of knowledge. And any proposed restriction on the class of admissible sets either seems ad hoc, or betrays the leading idea that knowledge contents are as fine-grained as linguistic expressions. The Lower Desideratum (7) is missed.⁴ This quandary has generated a wealth of literature we will not assess here.⁵

A third approach combines the two above [17, 42, 61]. Unstructured content informs an implicit knowledge operator: $I\varphi$ is true at w just in case φ is true at all worlds epistemically accessible from w . Structured content informs an explicit knowledge operator: $E\varphi$ is true at w just in case φ is a member of a given set $\mathcal{K}_w$, containing only sentences true at every world epistemically accessible from w . Such models are sometimes given a finer structure in awareness logics: each w is assigned a set of sentences $\mathcal{A}_w$, thought of as the sentences the agent is *aware of* at w . Then, $E\varphi$ holds at w just in case $\varphi \in \mathcal{A}(w)$ and $I\varphi$ is true at w [17].⁶

³See, for instance, [46] and [40].

⁴One might compare a prominent line in the *impossible worlds* approach, according to which content is taken as a set of worlds, some possible, some impossible. Impossible worlds are treated as 'anarchic', allowing the violation of even basic logical truths. This results in a trivial, or nearly trivial, epistemic logic. See [53] and [52].

⁵For instance, we pass over the impossible worlds approach and non-standard logic approach. See [18, 53, 63], Ch. 9. Other extant approaches are explicitly concerned to model bounded reasoning and bounded rationality. See [18], Ch. 10.

⁶See [55] for a broad overview of this approach and further references.

One may interpret the framework as follows: $E\varphi$ indicates that φ is actually known by the agent; $I\varphi$ indicates that in principle the agent could come to know φ on the basis of her present knowledge, given sufficient conceptual and computational resources. Actual knowledge exhibits little logical structure; potential knowledge displays logical omniscience.

One can enhance this theory with dynamic operators [60]: a ‘realization’ operator $[\#]$ ensures that the truth of both $E\varphi$ and $E(\varphi \supset \psi)$ entails the truth of $[\#\psi]E\psi$ (ψ is explicitly known upon *realizing* that ψ). Technically, $[\#\varphi]\psi$ is true at w just in case the addition of φ to $\mathcal{K}_w$ assures the truth of ψ . Since $E\varphi$ and $E(\varphi \supset \psi)$ need not entail $E\psi$, all this goes some way to representing the power of inference to extend knowledge. A candidate dissolution of the problem of logical omniscience emerges: the Lower Desideratum is met by I ; the Upper Desideratum is met by E . To think that a single logic of knowledge should meet both is to conflate actual and potential knowledge.

This, however, underestimates the intrinsic logic of actual knowledge. If Smith does not know that Jones is late then it *must* be that Smith does not know that Jones is horrendously late. Inferential relations, on the face of it, cannot explain this: if ψ can *merely* be inferred from φ , there is a possible situation – one where a given agent just fails to make the inference – where φ is known but not ψ . A better explanation: Jones’ being late is part of the content that Jones is horrendously late. By Yabloviaan ‘immanent closure’, to know the latter is to have *already* secured knowledge of the former, no matter how cognitively impaired one (on occasion) is.

Further, the proposal overestimates the intrinsic logic of potential knowledge, on (at least) grounds of the defeasibility of knowledge. An ordinary agent may know φ and know $\varphi \supset \psi$, yet inferring ψ is not a path to knowing ψ , but to losing the knowledge that φ . Beth knows Theorem 1 on the basis of testimony from her trusted tutor. She also believes that Theorem 2 holds, on the basis of testimony from her trusted professor (the professor is mistaken: a rare occurrence). She then comes to know that Theorem 1 and Theorem 2 are inconsistent, on the basis of painstaking mathematical reasoning. She rightly concludes that her tutor might have been mistaken, so rationally suspends her belief in Theorem 1. She thus loses knowledge of Theorem 1. Before she draws her last inference or proves the inconsistency result, there seems no interesting sense in which she potentially or implicitly knows that Theorem 2 is false on the basis of her present knowledge. Notice that we did not appeal to any irrationality on the part of our agent, though examples along this line are easy to concoct.

This motivates a dynamic epistemic logic that better accounts for the delicacies of content and reasoning, towards meeting the Goal. In the next Section we describe the logical features that should be captured.⁷

⁷Our Lower and Upper Bounds above, in a sense, don’t squeeze very tightly. There are proposals, such as that of [37, 38], which resort to an impossible worlds semantics with a sort of logical plausibility metric on worlds, capable of marking a distinction between blatant (e.g., that $2 + 2 = 5$) and subtle (e.g., that Fermat’s Last Theorem fails) impossibilities. This models moderately competent, but not logically omniscient agents capable of reasoning to rule out the former, though often incapable to rule out the latter. The proposal navigates successfully between our two extremes, but the set-up is very different from

3 Intuitive Epistemic Logic

We now propose a list $\mathbb{IEL}$ of intuitive validities and invalidities, generating defeasible criteria of success for an epistemic logic aimed at the Goal. While building the list, we propose three explanations for its items: (1) the *topic-sensitivity* of content; (2) the *fragmentation* of knowledge; (3) the *defeasibility* of its acquisition. One reason logical omniscience is difficult to deal with is its multifarious origin: arbitrary agents fail to be non-omniscient for different, irreducible reasons. We submit that (1)–(3) are three core ones.

The full definition of the formal language for $\mathbb{IEL}$ comes in Section 3.4, after some nuances have been motivated. This will suffice for now: p, q, r are atomic sentence symbols from set $\mathcal{AT}$. The boolean connectives $\neg, \wedge, \vee, \supset$ and $\equiv$ are standard. Read ' $\Box\varphi$ ' as 'It is apriori that φ ' and ' $K\varphi$ ' as 'The agent knows that φ '. Read ' $[\varphi]\psi$ ' as 'After the agent comes to know φ , ψ holds'.⁸ One can come to know φ via receipt of new empirical information, inference, and so on. We model a single agent, so all mention of the agent is suppressed. $\Box$ is a universal modality, quantifying over all possible worlds taken as epistemic scenarios. We model apriori entailment as $\Box(\varphi \supset \psi)$. We use $\checkmark$ to indicate that a schema is intuitively valid, $\times$ that it isn't.

3.1 Topic-Sensitivity

(9) $\times$ **Closure Under Strict Implication:** $\Box(\varphi \supset \psi) \supset (K\varphi \supset K\psi)$

If it is apriori that φ entails ψ , then one knows φ only if one knows ψ . This is invalid: it is apriori that if 113 guests attended the ball, then the number of guests is prime; but one can know the antecedent without knowing the consequent. Borrowing from [57], p. 88: William III might have known that England could avoid war with France, but he needn't have thereby known that England could avoid nuclear war with France. He didn't have the concept *nuclear*, and so wasn't positioned to think about nuclear wars at all. (To bolster this: William III could know and therefore believe that England could avoid war. But, presumably, he didn't have beliefs about nuclear wars, and so didn't believe that nuclear war could be avoided, as [57] notes. As knowledge implies belief, we retrieve our example.)

(9) is connected to two further invalidities:

(10) $\times$ **Apriori Omniscience:** $\Box\varphi \supset K\varphi$

(11) $\times$ **Omniscience Rule:** If φ is valid then so is $K\varphi$

ours: Jago's work models intermediately competent agents that sit strictly between the Lower Bound of complete logical incapacity and the Upper Bound of logical omniscience. It does not give a content-based semantics for ' x knows that φ ' for *arbitrary* agents x , as we claim to do below.

⁸Or, less tersely: 'After the agent comes to know the proposition expressed by φ in the prior discourse context, ψ holds in the posterior discourse context'. The prior context is the context before the proposition in question is added to the agent's body of knowledge; the posterior context is the context after this event. As in standard dynamic epistemic logic, we interpret $[\varphi]\psi$ with care in view of Moorean phenomena: see Section 3.6.

Cantor's Theorem is apriori, but an agent need not know it. William III failed to know that either England could avoid nuclear war with France, or England could not avoid nuclear war with France. Agents can be denied knowledge because they fail to properly grasp the relevant contents. One can grasp claims about the number 113 without being able to grasp claims about primeness; one can grasp claims about war without being able to grasp claims about nuclear war. Thus, one can know a claim about the former subject matter without knowing anything about the latter subject matter, despite that claim entailing something about the latter subject matter.

This is not to say that if one comes to grasp a certain subject matter then one automatically knows every consequence of one's knowledge about that subject matter (at least not when 'grasping' a proposition is understood as merely understanding it, or being positioned to hold an attitude towards it). Presumably, an agent can know there are 113 guests at the ball, fully grasp the subject matter of 'the number of guests is prime' (including the topic of primeness) and still fail to know that number of guests is prime. One can fully grasp 'the number e is irrational' without knowing it to be true (perhaps one is wondering whether it is true). So our conclusions are modest: subject matter is an aspect of content; hence, if ψ is about something that φ is not, then the content of ψ is not included in that of φ ; hence, an agent can know φ without knowing ψ – even if φ entails ψ . In contrast, we note an intuitive validity:

$$(12) \quad \checkmark \text{ Apriori Closure: } \Box(\varphi \supset \psi) \supset (\Box\varphi \supset \Box\psi)$$

This reflects a basic principle in the epistemology of mathematics: if a conjunction of axioms is apriori and it can be determined apriori that those axioms have a certain consequence, then that consequence is itself apriori. As a restricted quantifier over epistemically possible worlds, apriority is a normal modal.

3.2 Thickness, Same-Saying

Section 1 proposed that: content is thick (truth conditions plus topic); the Boolean connectives are topic-transparent; one content includes another when there is both truth and topic preservation; and $K\varphi \supset K\psi$ is valid just in case φ 's content includes ψ 's content. This predicts the following validities and invalidities:

$$(13) \quad \checkmark \text{ Left Simplification: } K(\varphi \wedge \psi) \supset K\varphi$$

$$(14) \quad \checkmark \text{ Right Simplification: } K(\varphi \wedge \psi) \supset K\psi$$

One knows that Jane is a lawyer and a fisherman only if one knows Jane is a lawyer. Knowing the latter is *part of* knowing the former.

$$(15) \quad \times \text{ Negative Addition I: } K\neg\varphi \supset K\neg(\varphi \wedge \psi)$$

$$(16) \quad \times \text{ Negative Addition II: } K\varphi \supset K\neg(\neg\varphi \wedge \psi)$$

William III knew that France will not go to war but didn't know that France will not go to nuclear war. He knew that England will go to war without knowing that England won't both refrain from war and develop a nuclear arsenal. These seem to hold even if we replace knowledge with what knowledge entails: belief. If (15) were valid, then the intentionality of content would be trivialized: in particular, if the logic of knowledge tracks closure under parthood, then everything one knows would be

(partly) about every topic. For example: one knows that Jane is friendly. (15) reflects that part of what one knows is that Jane is not both unfriendly and an admirer of Plutarch. Thus, one's knowledge is partly about Plutarch. (And, 'Plutarch' could be replaced with any referring term.) To say that one does not know anything about Plutarch would be nonsensical on this view, since *everything* one knows is partly about Plutarch.

Our hypothesis also predicts that these patterns are reflected in judgments about indirect speech reports. Mary said that Smith and Jones are both on time; so part of what she said is that Smith is on time. William III said that England will go to war; he did not thereby say that England won't both refrain from war and develop a nuclear arsenal.

3.3 Fragmentation

A simple theory of thick content does not tell the whole non-omniscience story. If it did, the following principle would be valid:

$$(17) \quad \times \text{ Closure Under Known Implication: } K(\varphi \supset \psi) \supset (K\varphi \supset K\psi)$$

If an agent's body of knowledge can always be taken as *one* content, P , and φ and $\varphi \supset \psi$ are both known, then knowledge of φ and $\varphi \supset \psi$ must be included in that of P . Hence, P must both entail ψ and include the topics of φ and ψ in its subject matter. If this is enough for ψ to be part of P , then ψ must also be known, by immanent closure. However, one may know φ and $\varphi \supset \psi$ but fail to put two and two together.⁹ Jones knows that Mary lives in New York, that Fred lives in Boston and that Boston is north of New York. Yet Jones fails to realize the obvious: that Mary will have to travel north to visit Fred [9, p. 199].

A popular response is that Jones' knowledge state must be fragmented across various 'frames of mind' [17, 44, 57, 66]. He knows certain facts about Mary in one frame of mind; certain facts about Fred in another; and certain facts relating Boston to New York in yet another. He may fail to put these bits of knowledge together.¹⁰

Fragmentation also predicts that the following dynamic principles are invalid:

$$(18) \quad \times \text{ MP Update: } [\varphi][\varphi \supset \psi]K\psi$$

$$(19) \quad \times \text{ Reverse MP Update: } [\varphi \supset \psi][\varphi]K\psi$$

One might come to know one premise of modus ponens, then come to know the other, but still fail to know the conclusion by failing to put two-and-two together.

⁹What, intuitively, is it to put 'two and two together'? A weak reading understands this as merely contemplating φ and ψ at the same time, or being aware of one's knowledge of each at the same time. A strong reading understands it as 'merging', partly or fully, the knowledge of φ and ψ into a unified piece of knowledge, as presumably happens when one infers a joint consequence. We are here primarily interested in the latter reading.

¹⁰Classic works on the functioning of human memory [36, 39] provide support for fragmentation. Whether memories are stored or, rather, reconstructed from some sparse partial information base, it seems that our mental hard disk is too big a stock for us to perform an exhaustive research (for the relevant base) each time we aim at retrieving information. Human memory works in a partitioned way, storing in fragments. Intersectorial logical dependencies may pass quietly unnoticed.

Notice that fragmentation alone does not tell the whole non-omniscience story: topic-sensitivity is still needed. If contents were mere sets of possible worlds, and given that a knows φ just in case φ is true at every world in some frame of mind for a , fragmentation would still validate (9), (10) and (11).

3.4 Formal Language

We now properly introduce our formal language, $\mathcal{EL}$. Let $\mathcal{F} = \{1, 2, \dots, n\}$ be a set of n frame symbols. Finitude is needed to keep the corresponding axiomatizations finitary, but does not bear on substantive philosophical points, except for the plausible claim that a finite cognitive agent will not host infinitely many frames of mind. We use i, j, k as metavariables for frame symbols, indexing the agent's different frames of mind and taken as a class of atomic formulas in $\mathcal{AT}$: atom i denotes the content of the agent's i -th frame of mind: ' $i \supset p$ ' reads 'the agent's i -th frame of mind implies p '. We use p, q, r as metavariables ranging over members of $\mathcal{AT} \setminus \mathcal{F}$. We use x, y, z as metavariables ranging over the whole of $\mathcal{AT}$. The BNF for $\mathcal{EL}$:

$$\varphi := x \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid K_i\varphi \mid [i : \varphi]\varphi$$

Read $K_i\varphi$ as 'the agent knows φ in frame of mind i ', and, as a first pass, define

$$K\varphi := \bigvee_{1 \leq i \leq n} K_i\varphi$$

$$[\varphi]\psi := \bigwedge_{1 \leq i \leq n} [i : \varphi]\psi$$

$K\varphi$ says that knowledge tout-court is knowledge in some frame of mind or other. Sentence $[\varphi]\psi$ says that if the agent adds the content of φ (in context) to any of her frames of mind, then ψ holds. That is: her coming to know φ assures that ψ is true. (It does *not* merely express the following weaker claim: if the agent adds the content of φ to *all* of her frames of mind, then ψ holds.) Now for our third source of intuitive (in)validities.

3.5 Defeasibility

If topic-sensitivity and fragmentation were the only sources of non-omniscience, the following would be valid for $i \neq j$:

- (20) $\times$ **Closure Under Merge I:** $K_i(\varphi \supset \psi) \supset (K_j\varphi \supset [i : j]K\psi)$
 (21) $\times$ **Closure Under Merge II:** $K_i(\varphi \supset \psi) \supset (K_j\varphi \supset [j : i]K\psi)$

This says: if in some frame of mind one knows that φ implies ψ , and in another φ , then one will come to know ψ after updating the first frame of mind with the knowledge contained in the second, or vice versa. Such a merge - or a partial merge - presumably results from inference. Hence closure under merge is a version of the proposal [60, Ch.5] that closure ('logical omniscience') holds when regulated by the action of inference.

But Closure under merge is invalid. Consider a variation on a counterexample [41] aims at the weak closure principle in [31, Sect. 2.3]: Edward knows, in different

frames of mind, both p (facts from chemistry) and if p then q (if p , then this homeopathic medicine is too diluted to contain active ingredients). Edward deduces q . He needn't thereby come to know q : he is strongly committed to r ('My mother swears that homeopathy works'), which he sees as weighing against q . Rather than gaining knowledge of q , he loses knowledge of p .

Does the counterexample depend on Edward being irrational? Some think that one cannot know p if one has a defeater for p [1]: the agent has justified beliefs in $d_1, \dots, d_n$, which jointly render p subjectively improbable. If Edward were rational in believing that homeopathy is efficacious on the basis of his total evidence, including his mother's testimony, then he could not have known p in the first place, on the basis of having a defeater.

But can't a perfectly rational agent lose knowledge via competent deduction? In one frame of mind, Mary knows that if there is no mistake on her tax return, then a report to the contrary is mistaken ($p \supset q$). She knows in that frame of mind that her accountant is very reliable on tax matters. She knows that he has reported a mistake on her return. (As it happens, this is the one occasion when he is incorrect.) In a different frame of mind, Mary knows that there is no mistake on her tax return (p). This knowledge is based on the especially fastidious manner in which she completed her return. How should Mary merge? She rightly concludes, given strong reasons to take her accountant as reliable, that she must suspend her belief that there is no error on her return, and so call into question how fastidious her completion really was, or how much she can really trust her fastidious efforts. Thus, she loses the knowledge that there is no mistake, rather than acquiring or preserving the knowledge that report of a mistake must be false.

To generalize, start with two monotonicity principles for epistemic update:

(22) $\times$ **Monotonicity I:** $[\varphi][\psi]K\varphi$

(23) $\times$ **Monotonicity II:** $K\varphi \supset [\psi]K\varphi$

These say that knowledge cannot be lost as more knowledge is acquired. This is obviously false for agents with fallible memories. Other counterexamples have long been noted. Jane knows that she does not know whether p . Presumably, she no longer knows this after coming to know p . Monotonicity failure for epistemic formulas is thus a feature of paradigm systems in the dynamic epistemic logic tradition. However, [28] emphasizes the defeasibility of rational update in belief revision due neither to memory failure nor to epistemic formulas:

(24) $\times$ **Atomic Monotonicity I:** $[p][q]Kp$

(25) $\times$ **Atomic Monotonicity II:** $Kp \supset [q]Kp$

Even standard dynamic epistemic logic systems validate these. But here's a counterexample. Mary knows that her tax return has no mistakes, on the basis of fastidious checking. But after she comes to know that her generally reliable accountant has reported a mistake in the return, she no longer knows that it contains no mistakes. An intuitive assessment: the newly acquired knowledge alters the space of reasons that bear on Mary's belief system, and if she reasonably suspends the belief that p , then she cannot know that p .

Monotonicity fails even within a single frame of mind:

- (26) × **Framed Atomic Monotonicity I:** $[i : p][i : q]K_i p$
 (27) × **Framed Atomic Monotonicity II:** $K_i p \supset [i : q]K_i p$

Our judgment in the previous example seems unaltered if we stipulate that it applies to a single known content. Further, suppose that framed atomic monotonicity is valid. Then the following is, too:

$$K_j q \supset (K_i p \supset [i : j]K_i p)$$

But this is surely invalid. Edward knows, in frame of mind i , certain facts about chemistry that refute homeopathy. In another frame j , he knows that his mother swears by homeopathy. His irrational psychology is such that, were he to put two-and-two together, he would suspend belief and so lose his knowledge of chemistry.

Now consider the following instance of closure under merge:

$$K_j(p \supset p) \supset (K_i p \supset [i : j]K_i p)$$

There seems an easy path from counterexamples to framed atomic monotonicity to ones for the above schema. One selects a countermodel so that (i) it verifies $K_j(p \supset p)$; (ii) it falsifies $K_i p \supset [i : j]K_i p$ (as in a counter-example to framed atomic monotonicity); (iii) it falsifies $K_j p$; (iv) i and j are the only frames in the model. Closure under merge should be rejected, due to the non-monotonicity of knowledge update.

3.6 Successful Update, Moore Sentences

Neverthelss, $[\varphi]\psi$ must sustain two simple validities:

- (28) ✓ **Atomic Success:** $[p]Kp$
 (29) ✓ **Atomic Consistency Preservation:** $Kp \supset [q]\neg K\neg p$

Learning p implies success: coming to know the proposition expressed by sentence p results in knowing that p . And though coming to know q might lead to losing the knowledge that p , it cannot lead to knowing that $\neg p$: this just is the factivity of knowledge.

One shouldn't endorse these principles beyond atoms, due to Moorean phenomena: take $p \wedge \neg Kp$. One may come to know the proposition expressed by this, say, by testimony. But the acquired knowledge is linguistically elusive: coming to know it does not render the sentence $K(p \wedge \neg Kp)$ true. Indeed, $\neg Kp$ is rendered false by this epistemic event. So $[p \wedge \neg Kp]K(p \wedge \neg Kp)$ is invalid. Our cautious interpretation of $[\varphi]\psi$ in footnote 8 matters: if P is the proposition expressed by $p \wedge \neg Kp$ in context, it is trivially true after coming to know P that one thereby knows P in the resulting context. What is not known is the false proposition expressed by $p \wedge \neg Kp$ in the resulting context, since Kp there reports on the *updated* body of knowledge.

4 A Logic for Thick, Defeasible, Fragmented Knowledge

We now give our formal semantics for language $\mathcal{EL}$ from Section 3.4, meeting the pattern $\mathbb{IEL}$ of validities and invalidities identified in Section 3. In Section 4.2, we draw comparisons with [66].

4.1 Epistemic Update Logic

We work with $\mathcal{EL}$, with frame symbols $\mathcal{F} = \{1, 2, \dots, n\}$.

Definition 1 (Epistemic Update Model) An *epistemic update model* $\mathcal{M}$ is a tuple

$$\langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$$

- W is a non-empty set: *possible worlds*. The power set of W is the set of *intensions* relative to W , denoted $\mathbb{I}$.
- $@ \in W$ is a designated *actual world*. $\mathbb{I}_@$ denotes the set of all intensions that include $@$.
- $\mathbb{T}$ is a non-empty set: *possible topics*.
- $\oplus : \mathbb{T} \times \mathbb{T} \rightarrow \mathbb{T}$ is an idempotent, commutative, associative binary operation: *topic fusion*.
- $\mu : \mathbb{I}_@ \times \mathbb{I}_@ \rightarrow \mathbb{I}_@$ accepts two intensions P and Q , and returns an intension. We call $\mu(P, Q)$ the *update* of P given Q . We stipulate that $\mu(P, Q) \subseteq Q$: an update with information Q leaves every non- Q world eliminated. We also stipulate that $P \cap Q \subseteq \mu(P, Q)$: at best, merging P and Q results in knowing their conjunction. It follows that $@ \in \mu(P, Q)$: intuitively, the update $\mu(P, Q)$ assures knowledge of a content with intension Q , so long as the agent's state is sensitive to the topic of that content.
- $v : \mathcal{AT} \rightarrow \mathbb{I}$ assigns an intension to each atom in the language (including frame symbols): the *truth valuation* function. By stipulation, this assignment obeys:

$$- @ \in v(i) \text{ for all } i \in \mathcal{F}$$

Since knowledge is factive, this reflects our intention that we model the *knowledge* of each frame of mind.

- $t : \mathcal{EL} \rightarrow \mathbb{T}$ assigns a topic to each formula in the language, including frame symbols: the *topic* function. To ensure the topic-transparency of Boolean connectives, the assignment obeys:

$$\begin{aligned} - t(\neg\varphi) &= t(\varphi) \\ - t(\varphi \wedge \psi) &= t(\varphi) \oplus t(\psi) \end{aligned}$$

($\mathbb{T}$, $\oplus$, μ and t obey few constraints. Finer structure would be needed to more deeply explore the nature of subject matter and knowledge update. Our aim, however, is merely to recover $\mathbb{IEL}$. Coarse representations of topic-sensitivity, fragmentation, and defeasibility, will suffice.)

We define *topic inclusion*, $a \leq b$, as $a \oplus b = b$. We call intensions *thin propositions*. A *thick proposition* is an ordered pair $\langle P, a \rangle$ of an intension and a topic (cf.

[65]). Thus, v and t together assign a thick proposition to each atom, including each frame symbol. We write $\kappa(i)$ for the pair $\langle v(i), t(i) \rangle$, and refer to $\kappa(i)$ as the *knowledge state of the agent at @, in her i -th frame of mind*. κ is the agent's *total knowledge state* at @:

$$\kappa := \{\kappa(1), \dots, \kappa(n)\}$$

We define the *satisfaction relation* $\models$ via the following recursion (we define the *intension* of φ with respect to $\mathcal{M}$ as $|\varphi|_{\mathcal{M}} := \{w \in W : \mathcal{M}, w \models \varphi\}$, omit the subscript $\mathcal{M}$ when the model is contextually clear):

Definition 2 ($\models$ -Semantics for $\mathcal{EL}$) Given $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$:

- $\mathcal{M}, w \models x$ iff $w \in v(x)$, where $x \in \mathcal{AT}$
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models (\varphi \wedge \psi)$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models \Box\varphi$ iff $\mathcal{M}, u \models \varphi$ for all $u \in W$
- $\mathcal{M}, w \models K_i\varphi$ iff $t(\varphi) \leq t(i)$ and $\mathcal{M}, u \models \varphi$ for all $u \in v(i)$
- $\mathcal{M}, w \models [i : \varphi]\psi$ iff: if $\mathcal{M}, @ \models \varphi$ then $\mathcal{M}_{\varphi}^i, w \models \psi$, where $\mathcal{M}_{\varphi}^i$ is the model:

$$\langle W, @, \mathbb{T}, \oplus, \mu, v_{\varphi}^i, t_{\varphi}^i \rangle$$

with

- $v_{\varphi}^i(i) = \mu(v(i), |\varphi|)$
- $t_{\varphi}^i(i) = t(i) \oplus t(\varphi)$
- $v_{\varphi}^i(x) = v(x)$ and $t_{\varphi}^i(x) = t(x)$ for all $x \in \mathcal{AT}$ such that $i \neq x$.

There are two natural ways to define truth in a model and validity for this framework. On one hand, we can say that φ is true in $\mathcal{M}$ (and write $\mathcal{M} \models \varphi$) exactly when $\mathcal{M}, @ \models \varphi$. Then we can say that φ is valid just in case it is true in every epistemic update model, and write $\models \varphi$.

Alternatively, we can say that φ is true in $\mathcal{M}$ (and write $\mathcal{M} \models \varphi$) exactly when $\mathcal{M}, w \models \varphi$ for every w in $\mathcal{M}$. Then we can say that φ is valid just in case it is true in every epistemic update model, and write $\models \varphi$.

However, $\models \varphi$ and $\models \varphi$ are in agreement with respect to the principles we canvassed in Section 3.

Theorem 1 Suppose sentence φ is constructed from sentences of the form $\Box\psi$ and $K_i\chi$, using the Boolean connectives and the update modalities. Then $\models \varphi$ iff $\models \varphi$.

Proof Sketch. Call a sentence *global* if, for every model $\mathcal{M}$, φ is either satisfied at every world in $\mathcal{M}$ or at no world in $\mathcal{M}$. It can be shown that sentences of the form $\Box\psi$ and the form $K_i\chi$ are global. Further it can be shown, by induction, that being global is preserved under applications of the boolean connectives and the update modalities. Finally, it can be shown that if φ is global then: $\models \varphi$ iff $\models \varphi$. $\square$

We use EUL to refer to the set of all validities of epistemic update logic: $\{\varphi \in \mathcal{EL} : \models \varphi\}$.

The clause for $K_i\varphi$ reflects our two-component account of content: to know φ is to have a knowledge state that is *both* sensitive to the topic of φ *and* rules out all incompatible worlds.¹¹ Note that, on the above semantics, $K_i\varphi$ is best read as ‘at @, the agent knows that φ in her i -th frame of mind’. On this reading, one expects the truth of $K_i\varphi$ to be invariant between worlds.

Theorem 2 *EUL satisfies IEL: every intuitive validity in Section 3 is in EUL; no intuitive invalidity in Section 3 is in EUL.*

The proofs yielding Theorem 2 come in technical Appendix A.

4.2 Departures from Yalcin

We treat subject matters abstractly, representing only the mereological structure of topics and the transparency of the logical connectives. Otherwise, we pass in silence over what topics *are*. Also, we take subject matter as integral to content: since both knowledge states and interpreted sentences have content, both have subject matter. Knowledge that φ requires that the topic of φ be contained in that of (some fragment of) the agent’s knowledge state.

Our setup is thus doubly divergent from [66], who deploys a specifically Lewisian account of subject matters [45] as partitions of the space of worlds, and treats the content of a claim as merely a set of worlds. On his account, a belief state is a partial function β accepting, if so defined, a subject matter π and returning an intension $\beta(\pi)$, with the constraint that $\beta(\pi)$ is at the *resolution* of π i.e. $\beta(\pi)$ is a union of cells in π . The idea is that $\beta(\pi)$, if defined, represents the agent’s body of (explicit) belief about subject matter π : a thin proposition at the resolution of the subject matter. If $\beta(\pi)$ is not defined, the agent’s belief state is not sensitive to subject matter π and so the agent has no (explicit) beliefs about π . Proposition P is (explicitly) believed relative to the agent’s belief state β just in case there exists a subject matter π such that (i) P is at the resolution of π and (ii) $\beta(\pi)$ is defined and entails P (i.e. every $\beta(\pi)$ world is a P world).

As [66] notes, the partition view of subject matter validates undesirable closure principles. (His Sect. 7 offers intriguing remarks on how to address this.) Let P be any necessary, apriori proposition, e.g., Cantor’s Theorem. This is true at every possible world. Thus, given any partition π , the intension of P is identical to the union of some cells in π (all of them). Thus, on the account of [66], P is believed with respect to any belief state. A form of omniscience follows: agents believe every apriori necessary proposition.

¹¹Compare the clause for ‘analytic implication’ in the conceptivist logic literature (see [19, 22]) and the clause for knowledge ascriptions in the awareness logic literature (see [17]). The version of awareness logic closest to our framework is the one in terms of ‘awareness generated by primitive propositions’ (see [55], p. 81), where an agent is aware of a formula φ just in case it is aware of all of its atomic constituents, taken together. We stress the syntactic features of this approach, not shared by ours: awareness is still given by a construction based on atomic formulas, whereas our topic function assigns topics, nonlinguistic items in the semantics, to formulas, with a recursion on the basic non-epistemic operators (negation and conjunction, thus disjunction).

Further, let $P \& Q$ and P be the non-identical, contingent propositions expressed, respectively, by $p \wedge q$ and p . Let β be a belief state defined on exactly one partition: the division of logical space into the $P \& Q$ worlds and their complement. Suppose the state returns the former set of worlds, given this partition. The content of $p \wedge q$ is thus explicitly believed by an agent with state β , but the content of p is not. So [66] counter-intuitively predicts that one can explicitly believe that Mary arrived at the meeting late and disheveled, without believing that Mary arrived late.

4.3 Further (In)Validities

We note some further validities and invalidities for which intuition seems muddier: it's neither a clear cost nor a clear benefit that our system yields these patterns. Some validities:

- (30) **Double Negation:** $K\varphi \equiv K\neg\neg\varphi$
- (31) **DeMorgan I:** $K\neg(\varphi \wedge \psi) \equiv K(\neg\varphi \vee \neg\psi)$
- (32) **DeMorgan II:** $K\neg(\varphi \vee \psi) \equiv K(\neg\varphi \wedge \neg\psi)$

One knows that Smith and Jones aren't both on time exactly when one knows that either Smith is late or Jones is late. One knows Smith and Jones are both late exactly when one knows that neither is on-time. Compare the following invalidity:

- (33) **Addition:** $K\varphi \supset K(\varphi \vee \psi)$

William III might have known that France will go to war without knowing that either they will go to war or develop a nuclear arsenal. Further, some valid restricted closure principles:

- (34) **Framed Closure:** $K_i(\varphi \supset \psi) \supset (K_i\varphi \supset K_i\psi)$
- (35) **Framed Atomic Rationality:** $[i : p \supset q](K_ip \supset K_iq)$

If one both knows that φ and knows that $\varphi \supset \psi$ within a unified frame of mind, then one knows ψ . If one knows p after coming to know $p \supset q$ within the same frame mind, then one knows q within that same frame of mind.

4.4 Expressive Surprises

$\mathcal{EL}$ is strictly more expressive than the static fragment of the language (call it $\mathcal{EL}_-$) obtained by removing the dynamic operator $[i : \varphi]$. This makes our framework differ crucially from traditional dynamic epistemic logics for which so-called reduction axioms are available. A standard technique for obtaining a complete axiomatization here is by adding a set of reduction axioms to a complete axiomatization of the underlying static logic. These provide a recursive step-by-step translation of every formula containing dynamic operators to a provably equivalent formula without them. Completeness follows from that of the underlying static logic, plus the soundness of the reduction axioms (see, e.g., [62], Sect. 7.4).

However, this only works when the language with dynamic operators is equally expressive as the language without. This is *not* the case for $\mathcal{EL}$. Two examples reveal two different reasons for the lack of expressivity of $\mathcal{EL}_-$. The examples

also motivate a way of extending $\mathcal{EL}_-$ so that the language one gets is sufficiently expressive to (1) still state the validities and invalidities of interest and (2) provide sound and complete axiomatizations for its dynamic extensions by means of reduction axioms. We put the soundness and completeness proofs in Appendix B, but we discuss the examples concerning expressivity in some detail here, for we find them philosophically illuminating.

Some notation: for any $\varphi \in \mathcal{EL}$, let $\mathcal{AT}(\varphi)$ denote the set of all atomic formulas occurring in φ . For a topic fusion operator $\oplus$, let

$$\oplus \mathcal{AT}(\varphi) = t(x_1) \oplus \cdots \oplus t(x_n),$$

where $\mathcal{AT}(\varphi) = \{x_1, \dots, x_n\}$.

Lemma 3 $\mathcal{EL}$ is strictly more expressive than $\mathcal{EL}_-$.

Proof Counterexample 1: This shows the lack of expressivity of $\mathcal{EL}_-$ due to topicality. Consider $\mathcal{M} = \{\langle @ \rangle, @, \mathbb{T}, \oplus, \mu, v, t\}$ and $\mathcal{M}' = \{\langle @ \rangle, @, \mathbb{T}', \oplus', \mu, v, t'\}$, where $v(x) = \{\langle @ \rangle\}$ for all $x \in \mathcal{AT}$, $\mathbb{T} = \{a, b, c, d\}$, $\mathbb{T}' = \{a', b', c', d'\}$, and $\langle \mathbb{T}, \oplus \rangle$ and $\langle \mathbb{T}', \oplus' \rangle$ constitute the join semi-lattices given in Fig. 1a and b, respectively.¹² Moreover, $t(p) = b$, $t'(p) = b'$ for all $p \in \mathcal{AT} \setminus \mathcal{F}$, $t(i) = c$, $t'(i) = c'$ for some $i \in \mathcal{F}$, $t(j) = d$, $t'(j) = d'$ for all $j \in \mathcal{F}$ such that $i \neq j$. For any $\varphi \in \mathcal{EL}$, we let $t(\varphi) = \oplus \mathcal{AT}(\varphi)$ and $t'(\varphi) = \oplus' \mathcal{AT}(\varphi)$. Finally, take some $\mu : \mathbb{I}_@ \times \mathbb{I}_@ \rightarrow \mathbb{I}_@$ that fits the description in Definition 1, such as $\mu(P, Q) = P \cap Q$. (The way we define μ is not important here, since the domains of both models are singletons.) Both $\mathcal{M}$ and $\mathcal{M}'$ are epistemic update models.

We now prove that for all $\varphi \in \mathcal{EL}_-$, we have $\mathcal{M}, @ \models \varphi$ iff $\mathcal{M}', @ \models \varphi$ by induction on the structure of φ ; cases for the elements of $\mathcal{AT}$, the Boolean connectives, and $\Box$ are elementary since they are not concerned with the topic component. Assume inductively that the statement holds for ψ and show it for $\varphi := K_k \psi$.

Case $\varphi := K_k \psi$: Suppose that $\mathcal{M}, @ \models K_k \psi$. This means, by the semantics of K_k , that $t(\psi) \leq t(k)$ and $\mathcal{M}, @ \models \psi$. Note that every frame of mind in both models corresponds to the singleton set $\{\langle @ \rangle\}$. We have two cases:

Case $k = i$: Then, $t(\psi) \leq t(k)$ implies that $t(\psi) = t(i) = c$. This means, by the definition of $t(\psi)$, that $\mathcal{AT}(\psi) = \{i\}$. Therefore, $t'(\psi) = c'$, implying that $t'(\psi) \leq t'(i)$. Also, $\mathcal{M}, @ \models \psi$ implies, by the induction hypothesis, that $\mathcal{M}', @ \models \psi$. Hence, as $i = k$, we obtain that $\mathcal{M}', @ \models K_k \psi$.

Case $k \neq i$: Then, $t(\psi) \leq t(k) = d$ means that $t(\psi) = d$. Thus, $\mathcal{AT}(\psi) \subseteq \mathcal{F} \setminus \{i\}$. Then, by the construction of $\mathcal{M}'$, we obtain that $t'(\psi) = \oplus' \mathcal{AT}(\psi) = d'$, hence, $t'(\psi) \leq t'(k) = d'$ for all $k \neq i$. As $\mathcal{M}', @ \models \psi$ (by the induction hypothesis), we again obtain that $\mathcal{M}', @ \models K_k \psi$. The other direction follows similarly.

We then conclude that $\mathcal{M}$ and $\mathcal{M}'$ are modally equivalent with respect to the language $\mathcal{EL}_-$. Now consider the sentence $[i : j]K_i p$ with $j \neq i$. We have $\mathcal{M}, @ \models$

¹²In figures of models, white nodes represent possible worlds, black nodes represent possible topics. Valuation and topic assignment are given by labelling each node with atomic formulas. We omit labelling when a node is assigned every element in $\mathcal{AT}$.

Fig. 1 Models $\mathcal{M}$ and $\mathcal{M}'$ in Counterexample 1

$[i : j]K_i p$: $\mathcal{M}, @ \models j$ but $\mathcal{M}_j^i, @ \not\models K_i p$, since $t_j^i(p) = b \not\leq t_j^i(i) = t(i) \oplus t(j) = a$. On the other hand, $\mathcal{M}', @ \models [i : j]K_i p$ since $\mathcal{M}', @ \models j$ and $(t')_j^i(p) = b' \leq (t')_j^i(i) = t'(i) \oplus t'(j) = a'$. Therefore, $[i : j]K_i p$ can distinguish $\mathcal{M}$ from $\mathcal{M}'$, thus, it cannot be logically equivalent to any formula in $\mathcal{EL}_-$. This proves Lemma 3. $\square$

The counterexample shows that $\mathcal{EL}_-$ is not as expressive as $\mathcal{EL}$ due to topicality. $K_i \varphi$ can express whether the topic of φ is included in the topic of the i th frame of mind. But $K_i \varphi$ cannot express anything concerning a *comparison* between the topics of arbitrary sentences φ and ψ . On the other hand, part of what $[i : \varphi]K_i \psi$ states is that the topic of ψ is included in the topic of the i th frame of mind after having grasped the topic of φ .

This motivates an extension, $\mathcal{EL}_+^+$, of $\mathcal{EL}_-$ with more general *conditional knowledge* operators of the form $K_\varphi \psi$.¹³ Here, φ is not necessarily a frame symbol, but can be an arbitrary sentence in the given language. The semantics are:

$$\mathcal{M}, w \models K_\varphi \psi \text{ iff } t(\psi) \leq t(\varphi) \text{ and } (\forall u \in W)(\text{if } \mathcal{M}, u \models \varphi \text{ then } \mathcal{M}, u \models \psi)$$

The conditional knowledge operator $K_\varphi \psi$ will help us build the topic component of the canonical model in the completeness proof to be given in Appendix B.1.

Counterexample 2: $\mathcal{EL}_+^+$ is still not sufficiently expressive to obtain reduction axioms for the dynamic operator $[i : \varphi]\psi$ since it cannot express the precondition of updates, namely, *truth at the actual state* $@$. To see this, consider the models $\mathcal{M} = \langle \{w_1, w_2, @\}, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ and $\mathcal{M}' = \langle \{w'_1, w'_2, @\}, @, \mathbb{T}, \oplus, \mu, v', t \rangle$, where $\mathbb{T} = \{a\}$, and $t(\varphi) = a$ for all $\varphi \in \mathcal{EL}_+^+$, $v(p) = \{w_1\}$ and $v'(p) = \{@\}$. For all $x \in \mathcal{AT} \setminus \{p\}$, $v(x) = v'(x) = \emptyset$ (see Fig. 2).

We show that for all $\varphi \in \mathcal{EL}_+^+$, $\mathcal{M}, w_2 \models \varphi$ iff $\mathcal{M}', w'_2 \models \varphi$, but, e.g., the dynamic sentence $\neg[i : p]\perp$ can distinguish the two pointed models (where $\perp$ is an abbreviation for some $q \wedge \neg q$). The former is easy to see: since every sentence of the language is assigned to the same topic, the topic component in this particular case does not play any essential role. That is, we have both $\mathcal{M} \models K_\varphi \psi \equiv \Box(\varphi \supset \psi)$ and $\mathcal{M}' \models K_\varphi \psi \equiv \Box(\varphi \supset \psi)$. Moreover, $\mathcal{M}, w_2$ and $\mathcal{M}', w'_2$ cannot be distinguished

¹³Compare system KRI in [4].

Fig. 2 Models $\mathcal{M}$ and $\mathcal{M}'$ in Counterexample 2.

within the basic modal language with the global modality $\Box\varphi$ since they are bisimilar (see, e.g., [8, Chapter 2.2] for a definition of bisimulation for the basic modal language). Therefore, $\mathcal{M}, w_2 \models \varphi$ iff $\mathcal{M}', w'_2 \models \varphi$ for all $\varphi \in \mathcal{EL}^+$.

To show that $\neg[i : p]\perp$ can distinguish the two pointed models, first observe that, for any epistemic update model $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ and $w \in W$,

$$\mathcal{M}, w \models \neg[i : p]\perp \text{ iff } \mathcal{M}, @ \models p.$$

So, the sentence $\neg[i : p]\perp$ really captures the precondition of the corresponding epistemic update with p . It is then not difficult to see that $\mathcal{M}, w_2 \not\models \neg[i : p]\perp$ (since $\mathcal{M}, @ \not\models p$) but $\mathcal{M}', w'_2 \models \neg[i : p]\perp$ (since $\mathcal{M}', @ \models p$).

We therefore need to further extend $\mathcal{EL}^+$ with an operator that expresses ‘ φ is true at the actual state’, that is, the actuality operator $A\varphi$ [14]. The resulting extended static language of knowledge $\mathcal{EL}^*$ is defined recursively by the following grammar:

$$\varphi := x \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid K_\phi\varphi \mid A\varphi,$$

where $x \in \mathcal{AT}$. The dynamic language obtained by adding $[i : \varphi]\psi$ to $\mathcal{EL}^*$ is denoted by $\mathcal{EL}^*$. In Appendix B, we prove soundness and completeness results for $\mathcal{EL}^+$, $\mathcal{EL}^*$ and $\mathcal{EL}^*$. For the dynamic language $\mathcal{EL}^*$, we axiomatize two classes of epistemic update models that represent *maximal* and *minimal* knowledge updates. To anticipate further, the former is modelled similarly to the so-called public announcements [25, 49, 50] and the latter constitutes an example of non-monotonic knowledge update.

5 Objections, Replies, Refinements

5.1 Elusive Logical Knowledge

Our semantics validates:

$$(36) \quad K(p \wedge q) \supset K(\neg p \supset q)$$

$$(37) \quad Kp \supset K\neg \dots \neg p, \text{ where } \neg \dots \neg \text{ is a string of one thousand negations.}$$

It seems clear that ordinary agents can fail to know that $(p \wedge q) \supset (\neg p \supset q)$ and $p \supset \neg \dots \neg p$ are tautologies. One might well describe oneself as knowing that roses are red and violets are blue, but deny knowing that ‘Roses are not red’ materially

implies ‘Violets are blue’. One might well describe oneself as knowing that ‘Roses are red’ is true but hesitate when asked if its thousand-fold negation is true.

Stalnaker [57, Ch.5] came up with a strategy for these cases. (He applied it to the more controversial issue of elusive mathematical knowledge, to which we come below; we only need a milder version.) Ordinary agents can lack basic knowledge about propositional logic, which may influence what knowledge ascriptions they deploy. But, as Stalnaker pointed out, the ignorance in question is plausibly meta-linguistic here, and due to difficulties in parsing the logical form of the relevant sentences. The beginning logic student does not know that certain sentence types must express truths in virtue of their form. She does not know that something she knows (roses are red or violets are blue) can be expressed by a sentence of the form $\neg p \supset q$. If her ignorance is relieved, she has not, after all, learned anything about the colour of flowers: she has learned something about the behavior of certain connectives.

5.2 Elusive Mathematical Knowledge I

One who knows φ , we claim, knows every consequence of φ that preserves subject matter. One might worry that the whole subject matter of a mathematical theory is encapsulated by its axioms. Then one who knows the conjunction of the ZF axioms knows every theorem of ZF set theory.

Whether this has force depends on answers to big questions we cannot address here, due to reasons of space: what is the nature of subject matter? How should we think about the subject matter of mathematical axioms? How do subject matter relationships between quantified statements and their instances work? Consider the claim that $0 \leq S(0)$ entails $S(0) \leq S(S(0))$, where S is the successor function. One view of subject matter [47, 54] has it that the topic of a claim is the set of objects and properties mentioned in that claim, or their associated concepts. Then ‘ $0 \leq S(0)$ ’ has the same subject matter as ‘ $S(0) \leq S(S(0))$ ’: the set with zero, the successor function and the relation of numerical order as members. According to another view [23], the subject matter of φ is a situation-like object: a truthmaker, or fusion of truthmakers. Then the subject matter of ‘ $0 \leq S(0)$ ’ may be represented as a structure of objects and relations: $\langle 0, \text{less}, 1 \rangle$. That of ‘ $S(0) \leq S(S(0))$ ’ may be represented by $\langle 1, \text{less}, 2 \rangle$. If subject matter inclusion amounts to one situation being part of the other, then the subject matter of $S(0) \leq S(S(0))$ is not included in that of $0 \leq S(0)$. A worry about mathematical omniscience does not arise.

Or, take the following set theoretic reasoning:

- (38) **Assumption:** A set exists.
- (39) **Assumption:** For any $\varphi(x)$ in the language of sets and set X , there exists set Y s.t. $x \in Y$ iff $x \in X$ and $\varphi(x)$ is true.
- (40) **Thus:** For any set X , there exists set Y s.t. $x \in Y$ iff $x \in X$ and $(x \in x \wedge x \notin x)$ is true.
- (41) **Thus:** There exists A s.t. $x \in A$ only if $(x \in x \wedge x \notin x)$ is true.
- (42) **Thus:** There exists set A s.t. for every x , $x \notin A$.

Is the subject matter of (40) included in that of (39), and of (41) included in that of the conjunction of (38) and (39)? Clearly, one can know the assumptions without knowing the conclusion. Indeed, one can know (39) without knowing (40).

Thus a proponent of our account of the logic of ordinary knowledge ascription must defend a theory of subject matter according to which the subject matter of a universal claim need not include the subject matter of all of its instances. This seems plausible, though the issue is murky. Suppose Jane is, in fact, in the Netherlands, but Beth has no reason to think so (usually Jane is in Mexico). Presumably, Beth can believe that everyone in the Netherlands is cold without believing Jane is cold.

5.3 Elusive Mathematical Knowledge II

Our semantics validates:

$$(43) \quad (\Box\varphi \wedge K(\varphi \vee \neg\varphi)) \supset K\varphi$$

On the assumption that $\varphi \vee \neg\varphi$ is apriori, this follows inescapably from several key theses:

- A. Content is just thick content i.e. truth conditions plus topic.
- B. P is part of Q iff P entails Q and P 's topic includes Q 's topic.
- C. Immanent closure: $K\varphi$ entails $K\psi$ when ψ 's content is part of φ 's content.
- D. The logical connectives are topic-transparent.
- E. All apriori truths have the same truth conditions.

(43) seems implausible for knowledge per se. Suppose that φ is a complicated tautology. Then (43) entails that φ is known by anyone who knows that φ is either true or false. Assume that Goldbach's Conjecture (GC) is true and apriori. Then (43) entails that GC is known by anyone who knows that GC is either true or false.

The meta-linguistic strategy of Section 5.1 finds some application here. Let φ be a tautology. An advocate of A-E judges that φ and $\varphi \vee \neg\varphi$ have identical content (call it P). Further, she can accept that it is easy to know that the sentence $\varphi \vee \neg\varphi$ is true and, thus, it is often easy to know P . However, she need not accept that, in this case, it is also easy to know that sentence φ is true, or a tautology, or expresses P . The value in identifying φ as a tautology, she says, lies not in learning its austere content P , but in coming to recognize that φ expresses austere content.

However, (43) also demonstrates the limits of the Stalnakerian strategy.¹⁴ It is implausible that 'Goldbach's conjecture' expresses austere content that is easy to know.

So our system needs refinement. B is a natural target. Assuming D, we have identified a plausible counter-example to its right-to-left direction: though GC and $GC \vee \neg GC$ plausibly have the same truth conditions and topic, it seems invariably odd to claim p is part of $p \vee \neg p$. Knowing GC is not part of knowing $GC \vee \neg GC$; to say $GC \vee \neg GC$ is not to (partly) say GC . E is also a natural target. Though GC

¹⁴Compare [20, 21].

is apriori, it is *prima facie* odd to claim that $GC \vee \neg GC$ entails GC . If this entailment fails, a counter-example exists: a circumstance c at which $GC \vee \neg GC$ is true but GC is false. Conclusion: $GC \vee \neg GC$ and GC differ in truth conditions. This involves an important admission: c is a mathematically impossible situation. Hence, this refinement admits ‘impossible worlds’ (though *logically* impossible worlds need not be admitted).

We reserve judgment on the best refinement. At any rate, our system is easily amended to reject E . It treats $\Box$ as the universal necessity modal. A refinement treats $\Box$ as a general S5 necessity modal, by extending our models with a transitive, reflexive, symmetric accessibility relation and interpreting $\Box$ by quantifying over the accessible worlds. In this case, (43) is invalidated when $\Box p$ holds because p is true at every accessible world, $K(p \vee \neg p)$ holds because some frame of mind includes the topic of p , and Kp fails because no frame of mind contains only p worlds.

5.4 Knowledge of Conjunctions

Our semantics validates:

$$(44) \quad (\Box(\varphi \supset \psi) \wedge K(\varphi \wedge (\psi \vee \neg\psi))) \supset K\psi$$

Like (43), (44) is implausible for knowledge per se.¹⁵ Assume that Goldbach’s Conjecture (GC) is true and an apriori consequence of the (conjunction of the) Peano Axioms (α). (44) entails that GC is known by anyone who knows α conjoined with $GC \vee \neg GC$. Like (43), this, at its core, is an inescapable consequence of key theses A-E listed in Section 5.3.

Unlike (43), a refinement based on rejecting E does not seem promising. Here, this amounts to admitting impossible worlds at which $\alpha \wedge (GC \vee \neg GC)$ holds but GC does not. Such a world is not only mathematically impossible, but logically impossible: the Peano Axioms hold, yet one of their logical consequences does not. As such examples are easily multiplied, dropping E heads down the slippery slope toward the extreme liberalism of [51]: for any set of sentences, there exists a world at which exactly those sentences are true. On this view, content is as fine-grained as syntax, contra the recommendations of Section 2.¹⁶

A more promising strategy is to fold the issue of (44) into a more general one concerning conjunctive knowledge. Consider:

$$(45) \quad \textbf{Conjoined MP: } K(p \wedge (p \supset q)) \supset Kq$$

¹⁵Thanks to Alexandru Baltag for pressing us on this point.

¹⁶Compare the logical framework of [53]. Of course, Priestian liberalism has the flexibility to recover any logical desiderata, by merely stipulating a restriction on the class of models that determine logical consequence. The Priestian claims that a content be identified with a set of worlds (i.e. effectively, a set of sets of sentences); that Q is part of P just in case $P \subseteq Q$; and that φ is true at world w just in case φ is in the set of sentences associated with w . Thus, if the full range of Priestian models are admitted, part-hood has no logic. But suppose one wants to preserve simplification: the content of $\varphi \wedge \psi$ must include that of φ , as a matter of logic. The Priestian can achieve this by restricting the *admissible* models to those where: if $\varphi \wedge \psi$ is associated with world w then φ is associated with w . The proposal has an obvious worry: it yields the right logic via mere, ad hoc stipulation. It has little explanatory power: the set of admissible models amounts to a mere *re-description* of the logical desiderata. For a similar critique, see [37] and [5, Sect. 5.3].

Our system validates (45) (contrast (17)). Do ordinary knowledge ascriptions align? Jones knows that Mary lives in New York, knows that Fred lives in Boston and knows that Boston is north of New York, but fails to infer that Mary will have to travel north to visit Fred. Still, we seem comfortable in asserting ‘Jones knows that Mary lives in New York and that Fred lives in Boston and that Boston is north of New York’. Thus, one might worry that, via our treatment of fragmentation, our system departs from intuitive epistemic logic by pulling $K\varphi \wedge K\psi$ and $K(\varphi \wedge \psi)$ apart: ordinary discourse attests that ‘Jones knows φ and ψ ’ is interchangeable with ‘Jones knows φ and Jones knows ψ ’. In particular, one worries that our system errs in validating (45).

To handle such discourse, a proponent of fragmentation must deny that ordinary claim ‘Smith knows φ and ψ ’ invariably reports $K_i(\varphi \wedge \psi)$ i.e. $\varphi \wedge \psi$ is known in a single frame of mind i belonging to Smith. Rather, it may report $K_i\varphi \wedge K_j\psi$ for distinct frames i and j . Hence, ascribing knowledge of a conjunction needn’t communicate that the conjuncts are part of a *unified* piece of knowledge. Unification may be pending, subject to outstanding inferences. This would explain our intuitive distaste for (44): it needn’t follow from ordinary claim ‘Smith knows both the Peano axioms and that either GC holds or it doesn’t’ that Smith has drawn the inferences that would unify her knowledge of the conjuncts, and so yield knowledge of GC .

How to systematize this position? A proponent of our system in Section 4 could bite the bullet, by appealing to the looseness of natural language. A bare claim ‘Smith knows φ and ψ ’, she might posit, is ambiguous between communicating $K_i\varphi \wedge K_j\psi$, for some i and j , and the stronger $K_i(\varphi \wedge \psi)$, for some i . Framed in terms of our proposal in Section 4: ‘Smith knows φ and ψ ’ is ambiguous between $K\varphi \wedge K\psi$ and $K(\varphi \wedge \psi)$.

An appeal to ambiguity will strike some as cheap or otherwise ad hoc. A second strategy rejects B or C and revises the definition of $K\varphi$ so as to invariably express: there exists i and j such that $K_i\varphi \wedge K_j\psi$. We sketch a tentative proposal that piggy-backs on the logical system introduced in Section 4.1. In what follows, φ and ψ are boolean constructions from ordinary proposition symbols. We delay complications raised by frame symbols, epistemic formulas and update formulas for elsewhere. We recursively define φ ’s *basic conjunctive parts*, denoted $P(\varphi)$, with:

$$\begin{aligned} P(p) &= \{p\}. \\ P(\neg p) &= \{\neg p\}. \\ P(\neg\neg\varphi) &= P(\varphi). \\ P(\varphi \wedge \psi) &= P(\varphi) \cup P(\psi). \\ P(\neg(\varphi \wedge \psi)) &= \{\neg(\bigwedge_{\varphi^* \in P(\varphi)} \varphi^* \wedge \bigwedge_{\psi^* \in P(\psi)} \psi^*)\}.^{17} \end{aligned}$$

$P(\varphi)$ is thus a model-independent set of sentences. Given epistemic update model $\mathcal{M}$, $P(\varphi)$ can be associated with a corresponding set of thick propositions: $\mathbf{P}^{\mathcal{M}}(\varphi)$. Conceptually, the basic conjunctive parts of φ are best thought of as the latter set.

We leave it ambiguous whether this proposal rejects B with a more refined account of content part, or rejects C with a more refined account of knowledge closure. It

¹⁷For uniqueness of the right-hand side, we assume that an ordering on the set of formulas has been fixed.

may be checked, by induction, that if $\psi \in P(\varphi)$ then φ 's intension is a subset of ψ 's intension and $t(\psi) \leq t(\varphi)$, relative to $\mathcal{M}$.¹⁸ A sample:

$$\begin{aligned} P(p \wedge q) &= \{p, q\}. \\ P(p \wedge q \wedge \neg r) &= \{p, q, \neg r\}. \\ P(\neg(p \wedge q)) &= \{\neg(p \wedge q)\}. \\ P(\neg(p \wedge q \wedge r)) &= \{\neg(p \wedge q \wedge r)\}. \\ P(\neg(\neg p \wedge \neg q)) &= P(p \vee q) = \{\neg(\neg p \wedge \neg q)\}. \\ P(p \vee \neg p) &= P(\neg(\neg p \wedge \neg \neg p)) = \{\neg(\neg p \wedge p)\}. \end{aligned}$$

Now amend the definition of $K\varphi$ from Section 3.4 as follows (with φ again ranging over boolean complexes of ordinary proposition symbols):

$$K\varphi := \bigwedge_{\psi \in P(\varphi)} \bigvee_{1 \leq i \leq n} K_i \psi$$

$K\varphi$ now says that each basic conjunctive part of φ is known by the agent in some or other frame of mind. The curious reader may check that this does *not* validate (44) and (45). Likewise, it may be checked that the pattern of validities and invalidities from Section 3 is preserved. A pressing issue remains: how to extend the proposal to our full language, including when φ is an epistemic formula?

Acknowledgments Many thanks to our anonymous referees for astute and constructive comments. Special thanks to Alexandru Baltag for stimulating feedback that significantly influenced the contents of Section 5. Versions of the present work were presented at the following conferences and workshops: the fourth meeting of the ‘From Shared Evidence to Group Attitudes’ project, University of Bayreuth, April 18–19, 2018; ‘Reasoning in Social Contexts’ at the Royal Netherlands Academy of Arts and Sciences, Amsterdam, May 31–June 2, 2018; ‘Models of Bounded Reasoning in Individuals and Groups’ at the Lorentz Center, Leiden, July 2–6, 2018; the Super-Special seminar slot at Arché, University of St Andrews, November 29, 2018; and the ‘Workshop on the Occasion of Johan van Benthem’s 70th birthday’ in Amsterdam, June 12, 2019. We thank these audiences for their helpful questions and remarks. This research is published within the project ‘The Logic of Conceivability’, funded by the European Research Council (ERC CoG), Grant Number 681404.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Proof of Theorem 2

We call a sentence ψ of $\mathcal{EL}$ *Boolean* if it does not include occurrences of i , $\Box$, K_i , or $[i : \varphi]$. The following auxiliary lemma will be useful in the proof of Theorem 2:

¹⁸*Proof sketch.* Prove the stronger result that the conjunction of the elements of $P(\varphi)$ is classically equivalent to φ and has the same atoms as φ .

Lemma 4 For all epistemic update models $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$, all $\varphi \in \mathcal{EL}$, $i \in \mathcal{F}$, and all Boolean sentences ψ , we have $|\psi|_{\mathcal{M}} = |\psi|_{\mathcal{M}_\varphi^i}$.

Proof Follows via an easy induction on the structure of ψ . $\square$

Proof of Theorem 2: Invalidities:

In figures of counterexamples, white nodes represent possible worlds, black nodes represent possible topics. Valuation and topic assignment are given by labelling each node with atomic formulae. We omit labelling when a node is assigned every element in $\mathcal{AT}$.

Counterexample for (9)–(11) and (15)–(16): These schemas are invalid due to topicality. Without loss of generality (w.l.o.g.), let $\mathcal{F} = \{i\}$ and consider the model $\mathcal{M}_1 = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ such that $W = \{@\}$, $\mathbb{T} = \{a, b, c\}$ with $b \neq c$, $b < a$, and $c < a$, $v(i) = v(p) = v(q) = \{@\}$, and $t(p) = t(i) = b$ and $t(q) = c$ (Fig. 3). (We do not need to specify μ since it is irrelevant for these schemas.) Then, (9) is invalid since $\mathcal{M}_1, @ \not\models \Box(p \supset q) \supset (Kp \supset Kq)$: $\mathcal{M}_1, @ \models \Box(p \supset q)$ (since $|p| = |q|$) and $\mathcal{M}_1, @ \models K_i p$ (since $v(i) = |p|$ and $t(p) \leq t(i)$). However, $\mathcal{M}_1, @ \not\models K_i q$ since $t(q) \not\leq t(i)$. As $\mathcal{F} = \{i\}$, we obtain that $\mathcal{M}_1, @ \not\models \Box(p \supset q) \supset (Kp \supset Kq)$. Similarly, (10) is invalid since $\mathcal{M}_1, @ \not\models \Box q \supset Kq$, and (11) is invalid since $(q \vee \neg q)$ is valid but $\mathcal{M}_1, @ \not\models K(q \vee \neg q)$ for $t(q) \not\leq t(i)$. Moreover, (15) and (16) are invalid since $\mathcal{M}_1, @ \not\models K\neg p \supset K\neg(\neg p \wedge q)$ and $\mathcal{M}_1, @ \not\models Kp \supset K\neg(\neg p \wedge q)$, respectively: it is easy to verify that $\mathcal{M}_1, @ \models K\neg p$ and $\mathcal{M}_1, @ \models Kp$, but $\mathcal{M}_1, @ \not\models K\neg(\neg p \wedge q)$ as $t(\neg(\neg p \wedge q)) = t(p) \oplus t(q) = a \not\leq t(i)$.

Counterexample for (17)–(19): These schemas are invalid due to fragmentation. W.l.o.g., let $\mathcal{F} = \{i, j\}$ and consider the model $\mathcal{M}_2 = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ such that $W = \{@, w_1, w_2\}$, $v(i) = \{@, w_1\}$, $v(j) = v(p) = \{@, w_2\}$, and $v(q) = \{@\}$, $\mathbb{T} = \{a\}$, and $t(\varphi) = a$ for all $\varphi \in \mathcal{EL}$. (Since every $\varphi \in \mathcal{EL}$ is mapped to the same topic, the topicality constraints will be trivially satisfied.) Then, (17) is invalid since $\mathcal{M}_2, @ \not\models K(p \supset q) \supset (Kp \supset Kq)$: $\mathcal{M}_2, @ \models K_i(p \supset q)$ (since $v(i) = |p \supset q|$) and $\mathcal{M}_2, @ \models K_j p$ (since $v(j) = |p|$), thus, $\mathcal{M}_2, @ \models K(p \supset q)$ and $\mathcal{M}_2, @ \models Kp$. However, $\mathcal{M}_2, @ \not\models Kq$ since $\mathcal{M}_2, @ \not\models K_i q$ and $\mathcal{M}_2, @ \not\models K_j q$ (as $v(i) \not\subseteq |q|$ and $v(j) \not\subseteq |q|$, respectively). For (18) and (19), consider $\mathcal{M}_2$ with $\mu(P, Q) = P \cap Q$ for all $P, Q \in \mathbb{I}_@$. Then, (18) is invalid since $\mathcal{M}_2, @ \models \neg[j : p][i : p \supset q]Kq$:

Fig. 3 Counterexample $\mathcal{M}_1$

Fig. 4 Counterexample $\mathcal{M}_2$

$\mathcal{M}_2, @ \models p$ and $(\mathcal{M}_2)_p^j, @ \models p \supset q$, but, $((\mathcal{M}_2)_p^j)_{p \supset q}^i, @ \not\models K_i q \vee K_j q$ (since $(v_p^j)_{p \supset q}^i(j) = \{ @, w_2 \}$ and $(v_p^j)_{p \supset q}^i(i) = \{ @, w_1 \}$, and $((\mathcal{M}_2)_p^j)_{p \supset q}^i, w_1 \not\models q$ and $((\mathcal{M}_2)_p^j)_{p \supset q}^i, w_2 \not\models q$). Similarly, (19) is invalid since $\mathcal{M}_2, @ \models \neg[i : p \supset q][j : p]Kq$.

Counterexample for (20) and (21): These schemas are invalid due to non-monotonicity of knowledge update. The counterexample $\mathcal{M}_2$ in Fig. 4 with $\mu'(P, Q) = Q$ for all $P, Q \in \mathbb{I}_@$ invalidates (20) and (21): $\mathcal{M}_2, @ \models K_i(p \supset q) \wedge K_j p$, but $\mathcal{M}_2, @ \not\models [i : j]Kq$ and $\mathcal{M}_2, @ \not\models [j : i]Kq$. For $\mathcal{M}_2, @ \not\models [i : j]Kq$, observe that $v_i^j(i) = v_i^j(j) = \{ @, w_2 \} \not\subseteq |q|$, thus, $(\mathcal{M}_2)_j^i \not\models K_i q \vee K_j q$. For $\mathcal{M}_2, @ \not\models [j : i]Kq$, observe that $v_i^j(i) = v_i^j(j) = \{ @, w_1 \} \not\subseteq |q|$, thus, $(\mathcal{M}_2)_i^j \not\models K_i q \vee K_j q$.

Counterexample for (22)–(27): These principles are invalid due to non-monotonicity of knowledge update. It is easy to see that counterexamples invalidating (26) and (27) are also counterexamples for (24) and (25), respectively. Moreover, (24) and (25) are special cases of (22) and (23), respectively. Consider the model $\mathcal{M}_3 = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ such that $W = \{ @, w_1 \}$, $\mathbb{T} = \{ a \}$, $v(i) = v(p) = \{ @ \}$ for all $i \in \mathcal{F}$, and $v(q) = W$, $t(\varphi) = a$ for all $\varphi \in \mathcal{EL}$, and $\mu(P, Q) = Q$ for all $P, Q \in \mathbb{I}_@$ (Fig. 5). As every frame of mind is mapped to the same set of possible worlds, we do not need to consider fragmentation. Similarly since every $\varphi \in \mathcal{EL}$ is mapped to the same topic, the topicality constraints will be trivially satisfied. Then, (26) is invalid since $\mathcal{M}_3, @ \not\models [i : p][i : q]K_i p$: $\mathcal{M}_3, @ \models p$ and $(\mathcal{M}_3)_p^i, @ \models q$ but $((\mathcal{M}_3)_p^i)_q^i, @ \not\models K_i p$ since $(v_p^i)_q^i(i) = W$ and $((\mathcal{M}_3)_p^i)_q^i, w_1 \not\models p$. Similarly, (27) is invalid since $\mathcal{M}_3, @ \not\models K_i p \supset [i : q]K_i p$: $\mathcal{M}_3, @ \models K_i p$ and $\mathcal{M}_3, @ \models q$, but $(\mathcal{M}_3)_q^i, @ \not\models K_i p$ since $v_q^i(i) = W$ and $(\mathcal{M}_3)_q^i, w_1 \not\models p$.

Validities: Let $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ be an epistemic update model and $w \in W$:

(12) $\models \Box(\varphi \supset \psi) \supset (\Box\varphi \supset \Box\psi)$: Suppose $\mathcal{M}, w \models \Box(\varphi \supset \psi)$ and $\mathcal{M}, w \models \Box\varphi$. While the former means that $|\varphi| \subseteq |\psi|$, the latter means $|\varphi| = W$. Therefore, $|\psi| = W$, i.e., $\mathcal{M}, w \models \Box\psi$.

Fig. 5 Counterexample $\mathcal{M}_3$

(13) $\models K(\varphi \wedge \psi) \supset K\varphi$: Suppose $\mathcal{M}, w \models K(\varphi \wedge \psi)$. This means, by the definition of K , that there is an $i \in \mathcal{F}$ such that $\mathcal{M}, w \models K_i(\varphi \wedge \psi)$, i.e., that $v(i) \subseteq |\varphi \wedge \psi|$ and $t(\varphi \wedge \psi) \leq t(i)$. Since $v(i) \subseteq |\varphi \wedge \psi| \subseteq |\varphi|$ and $t(\varphi) \leq t(\varphi) \oplus t(\psi) = t(\varphi \wedge \psi) \leq t(i)$, we obtain $\mathcal{M}, w \models K_i\varphi$. Thus, $\mathcal{M}, w \models K\varphi$. The validity proofs of (14), and (30)–(32) follow similarly.

(28) $\models [p]Kp$: Let $i \in \mathcal{F}$ and suppose $\mathcal{M}, @ \models p$. We then have that $v_p^i(i) = \mu(v(i), |p|) \subseteq |p| = |p|_{\mathcal{M}_p^i}$ (by Lemma 4). Moreover, $t_p^i(i) = t(i) \oplus t(p) \geq t(p) = t_p^i(p)$. Therefore, $\mathcal{M}_p^i, w \models K_i p$, implying that $\mathcal{M}_p^i, w \models Kp$. Thus, $\mathcal{M}, w \models [i : p]Kp$. Since i has been chosen arbitrarily from $\mathcal{F}$, we obtain that $\mathcal{M}, w \models [p]Kp$.

(29) $\models Kp \supset [q]\neg K\neg p$: Suppose $\mathcal{M}, w \models Kp$ and let $i \in \mathcal{F}$ such that $\mathcal{M}, @ \models q$. The former implies that $\mathcal{M}, w \models K_j p$ for some $j \in \mathcal{F}$. In particular, since $@ \in v(k)$ for all $k \in \mathcal{F}$, we have $\mathcal{M}, @ \models p$. Then, by Lemma 4, we obtain that $\mathcal{M}_q^i, @ \models p$. By the definition of μ , we also have that $@ \in v_q^i(k)$ for all $k \in \mathcal{F}$. Therefore, $\mathcal{M}_q^i, w \models \neg K_k \neg p$ for all $k \in \mathcal{F}$, meaning that $\mathcal{M}_q^i, w \models \neg K \neg p$. Hence, $\mathcal{M}, w \models [i : q]\neg K \neg p$. Since i has been chosen arbitrarily from $\mathcal{F}$, we obtain that $\mathcal{M}, w \models [q]\neg K \neg p$, thus, $\mathcal{M}, w \models Kp \supset [q]\neg K \neg p$. $\square$

Appendix B: Soundness and Completeness for $\mathcal{EL}_-^+$, $\mathcal{EL}_-^*$, and $\mathcal{EL}^*$

We provide soundness and completeness results for $\mathcal{EL}_-^+$, $\mathcal{EL}_-^*$, and $\mathcal{EL}^*$. The following syntactic abbreviations will matter: for any $\varphi \in \mathcal{EL}^*$, $\mathcal{AT}(\varphi)$ denotes the set of atomic formulae occurring in φ and we will use ‘ $\overline{\varphi}$ ’ to denote the tautology $\bigwedge_{x \in \mathcal{AT}(\varphi)} (x \vee \neg x)$, following a similar idea in [26].

All our soundness and completeness results are proven with respect to the class of epistemic update models whose topic assignment function is defined in the following way:

- $t : \mathcal{AT} \rightarrow \mathbb{T}$ is a *topic function* assigning a topic to each element in $\mathcal{AT}$. t extends to the whole $\mathcal{EL}^*$ by taking the topic of a sentence φ as the fusion of the elements in $\mathcal{AT}(\varphi)$:

$$t(\varphi) = \oplus \mathcal{AT}(\varphi) = t(x_1) \oplus \cdots \oplus t(x_k)$$

where $\mathcal{AT}(\varphi) = \{x_1, \dots, x_k\}$.

$\leq$ again denotes *topic inclusion* as defined in Section 4.1. The topic of a complex sentence φ , defined from its primitive components in $\mathcal{AT}(\varphi)$, makes all the logical connectives and modal operators in $\mathcal{EL}^*$ topic-transparent, that is,

- $t(\Box\varphi) = t(A\varphi) = t(\neg\varphi) = t(\varphi)$
- $t(K_\psi\varphi) = t(\varphi \wedge \psi) = t(\varphi \vee \psi) = t(\varphi \supset \psi) = t(\varphi \equiv \psi) = t(\varphi) \oplus t(\psi)$
- $t([i : \varphi]\psi) = (t(i) \oplus t(\varphi)) \oplus t(\psi)$

Topic-transparency of the truth-functional connectives has already been argued for in Section 1 and formally captured in Definition 1. The only further constraint we impose here is the topic-transparency of the modal operators $K_\psi\varphi$, $\Box\varphi$, $A\varphi$, and $[i : \varphi]\psi$. This is admittedly an idealization, but for technical purposes, in this paper

we can live with this idealization, as Theorem 2 still holds with respect to epistemic update models with the above constraint. In the remainder of this appendix, all models are implicitly assumed to obey this constraint on t .

We first provide a sound and complete axiomatization for $\mathcal{EL}^*_-$ (Appendix B.1). The completeness result for $\mathcal{EL}^+_+$ follows similarly, so we omit many details and only point out the differences (Appendix B.1.1). The completeness for $\mathcal{EL}^*$ will follow from the completeness of $\mathcal{EL}^*_-$ via a set of sound reduction axioms (Appendix B.2).

B.1 The (static) Logic of Knowledge Over $\mathcal{EL}^*_-$

Since $\mathcal{EL}^*_-$ does not have the dynamic operator, the update function μ does not play any role in its interpretation in epistemic update models. We therefore opt for simplicity and interpret $\mathcal{EL}^*_-$ in what we call *epistemic models*, $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, v, t \rangle$, obtained by removing μ from epistemic update models.

To recap, given an epistemic model $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, v, t \rangle$ and $w \in W$, we define the *satisfaction relation* $\models$ for the atomic formulae, Booleans, and $\Box\varphi$ as in Definition 2; for $K_\psi\varphi$ and $A\varphi$ we have:

$$\begin{aligned} \mathcal{M}, w \models K_\psi\varphi & \text{ iff } t(\psi) \leq t(\varphi) \text{ and } (\forall u \in W) (\text{ if } \mathcal{M}, u \models \varphi \text{ then } \mathcal{M}, u \models \psi) \\ \mathcal{M}, w \models A\varphi & \text{ iff } \mathcal{M}, @ \models \varphi. \end{aligned}$$

Truth in a model and *validity* are defined as before (see Section 4.1). Soundness and completeness are defined in a standard way with respect to the global notion of validity denoted by $\models$.

Lemma 5 *Let $\langle W, @, \mathbb{T}, \oplus, v, t \rangle$ be an epistemic model. Then, for any $\varphi \in \mathcal{EL}^*_-$, $w \in W$, and update function μ (as described in Definition 1), we have*

$$\langle W, @, \mathbb{T}, \oplus, v, t \rangle \models \varphi \text{ iff } \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle, w \models \varphi.$$

Proof Follows via an easy induction on the structure of φ as μ does not play any role in the interpretation of the sentences in $\mathcal{EL}^*_-$. $\square$

Observation 6

- (1) For all $\varphi \in \mathcal{EL}^*_-$, $\models \bar{\varphi}$.
- (2) $\models K_\varphi\psi \equiv (K_\varphi\bar{\psi} \wedge \Box(\varphi \supset \psi))$.¹⁹

Table 1 presents a sound and complete axiomatization EL^* of epistemic logic over $\mathcal{EL}^*_-$. The notion of derivation, denoted by $\vdash_{\text{EL}^*}$, in EL^* is defined as usual. Thus, $\vdash_{\text{EL}^*} \varphi$ means φ is a theorem of EL^* . For any set of formulas $\Gamma \subseteq \mathcal{EL}^*_-$ and any $\varphi \in \mathcal{EL}^*_-$, we write $\Gamma \vdash_{\text{EL}^*} \varphi$ if there exists finitely many formulas $\varphi_1, \dots, \varphi_n \in \Gamma$ such that $\vdash_{\text{EL}^*} (\varphi_1 \wedge \dots \wedge \varphi_n) \supset \varphi$.

¹⁹Observation 6. (2) shows that $K_\varphi\psi$ is the analytic implication of [22].

Table 1 Axiomatization EL^* (over $\mathcal{EL}^*$)

(CPL)	all classical propositional tautologies and Modus Ponens
(S5 $\Box$)	S5 axioms and rules for $\Box$
	(I) Axioms for A :
(K $_A$)	$A(\varphi \supset \psi) \supset (A\varphi \supset A\psi)$
(Dua $_A$)	$A\varphi \equiv \neg A\neg\varphi$
	(II) Axioms for A and $\Box$:
(Ax1)	$\Box\varphi \supset A\varphi$
(Ax2)	$A\varphi \supset \Box A\varphi$
	(III) Axioms for i and A :
(Ax3)	$A(1 \wedge \dots \wedge n)$
	(IV) Axioms connecting K and $\Box$:
(Ax4)	$K_\varphi\psi \equiv (K_\varphi\overline{\psi} \wedge \Box(\varphi \supset \psi))$
(Ax5)	$K_\varphi\psi \supset \Box K_\varphi\psi$
	(V) Axioms for K :
(Ax6)	$(K_\varphi\overline{\psi} \wedge K_\psi\overline{\chi}) \supset K_\varphi\overline{\chi}$
(Ax7)	$(K_\varphi\overline{\psi} \wedge K_\varphi\overline{\chi}) \supset K_\varphi(\overline{\psi} \wedge \overline{\chi})$
(Ax8)	$K_\varphi\overline{\psi}$, if $\mathcal{AT}(\psi) \subseteq \mathcal{AT}(\varphi)$

Theorem 7 *The following are derivable from EL^* :*

- (1) *from $\vdash_{\text{EL}^*} \varphi \equiv \psi$, $\vdash_{\text{EL}^*} K_\varphi\overline{\psi}$, and $\vdash_{\text{EL}^*} K_\psi\overline{\varphi}$, infer $\vdash_{\text{EL}^*} K_\chi\varphi \equiv K_\chi\psi$ and $\vdash_{\text{EL}^*} K_\varphi\chi \equiv K_\psi\chi$,*
- (2) *(Nec $_A$) from $\vdash_{\text{EL}^*} \varphi$ infer $\vdash_{\text{EL}^*} A\varphi$,*
- (3) *$\vdash_{\text{EL}^*} A(\varphi \wedge \psi) \equiv A\varphi \wedge A\psi$,*

Proof (1) follows from Ax4, Ax6, and S5 $\Box$. Item (2) follows from necessitation for $\Box$ and Ax1, and (3) is derived from K $_A$ and (Nec $_A$) in a standard way. $\square$

Soundness of EL^* is a matter of routine validity check, so we skip its proof. The rest of this section is devoted to the completeness proof of EL^* , which is presented in full detail.

We say that Γ is EL^* -consistent if $\Gamma \not\vdash_{\text{EL}^*} \perp$, and EL^* -inconsistent otherwise. We omit the tag EL^* and say (in)consistent when the logic is contextually clear. A sentence φ is consistent with Γ if $\Gamma \cup \{\varphi\}$ is consistent (or, equivalently, if $\Gamma \not\vdash_{\text{EL}^*} \neg\varphi$). Finally, a set of formulas Γ is a *maximally consistent set* (or, in short, *mcs*) if it is consistent and any set of formulas properly containing Γ is inconsistent [8].²⁰

²⁰Notions of derivation, (in)consistent, and maximally consistent sets for the systems studied in Appendices B.1.1 and B.2 are defined similarly.

Lemma 8 For every mcs Γ of EL^* and $\varphi, \psi \in \mathcal{EL}^*$, the following hold:

- (1) $\Gamma \vdash_{\text{EL}^*} \varphi$ iff $\varphi \in \Gamma$,
- (2) if $\varphi \in \Gamma$ and $\varphi \supset \psi \in \Gamma$, then $\psi \in \Gamma$,
- (3) if $\vdash_{\text{EL}^*} \varphi$ then $\varphi \in \Gamma$,
- (4) $\varphi \in \Gamma$ and $\psi \in \Gamma$ iff $\varphi \wedge \psi \in \Gamma$,
- (5) $\varphi \in \Gamma$ iff $\neg\varphi \notin \Gamma$.

Proof Standard. □

Lemma 9 (Lindenbaum's Lemma) Every EL^* -consistent set can be extended to a maximally consistent one.

Proof Standard. □

Let $\mathcal{X}$ be the set of all maximally consistent sets of EL^* . Define $\sim_\square$ on $\mathcal{X}$ as

$$\Gamma \sim_\square \Delta \text{ iff } \forall \varphi \in \mathcal{EL}^* (\square\varphi \in \Gamma \text{ implies } \varphi \in \Delta).$$

It is standard to prove that $\sim_\square$ is an equivalence relation, as $\square$ is an S5 operator. To define the canonical model, we need some auxiliary definitions and lemmas.

Lemma 10 For all $\Gamma \in \mathcal{X}$, $@_\Gamma = \{\varphi \in \mathcal{EL}^* : A\varphi \in \Gamma\}$ is a maximally consistent set such that $\Gamma \sim_\square @_\Gamma$ and $1 \wedge \dots \wedge n \in @_\Gamma$.

Proof Suppose that $@_\Gamma = \{\varphi \in \mathcal{EL}^* : A\varphi \in \Gamma\}$ is not consistent. This means that there are finitely many $\varphi_1, \dots, \varphi_k \in @_\Gamma$ such that $\vdash_{\text{EL}^*} (\varphi_1 \wedge \dots \wedge \varphi_k) \supset \perp$. This implies, in particular, that $\vdash_{\text{EL}^*} (\varphi_1 \wedge \dots \wedge \varphi_{k-1}) \supset \neg\varphi_k$. Then, by K_A , Theorem 7.(2) and 7.(3), we have $\vdash_{\text{EL}^*} (A\varphi_1 \wedge \dots \wedge A\varphi_{k-1}) \supset A\neg\varphi_k$. As $A\varphi_1, \dots, A\varphi_{k-1} \in \Gamma$, we have $A\neg\varphi_k \in \Gamma$. This means, by Dua_A , that $\neg A\varphi_k \in \Gamma$, contradicting consistency of Γ . Therefore, $@_\Gamma$ is consistent. Now suppose that $@_\Gamma$ is not maximal. This means that there is a $\psi \in \mathcal{EL}^*$ such that $\psi \notin @_\Gamma$ and $@_\Gamma \cup \{\psi\}$ is consistent. But, as $\psi \notin @_\Gamma$, we have $A\psi \notin \Gamma$, thus, by Lemma 8.(5) and Dua_A , $A\neg\psi \in \Gamma$. This means that $\neg\psi \in @_\Gamma$, which makes $@_\Gamma \cup \{\psi\}$ inconsistent. Therefore, $@_\Gamma$ is maximal too. By $\text{Ax}3$ and Lemma 8.(3), we have $A(1 \wedge \dots \wedge n) \in \Gamma$, therefore, $1 \wedge \dots \wedge n \in @_\Gamma$. Now take an arbitrary $\varphi \in \mathcal{EL}^*$ and suppose that $\square\varphi \in \Gamma$. Then, by $\text{Ax}1$ and Lemmas 8.(2) and 8.(3), $A\varphi \in \Gamma$. This means that $\varphi \in @_\Gamma$. Therefore, we also obtain that $\Gamma \sim_\square @_\Gamma$. □

For $\Gamma \in \mathcal{X}$, let $\approx_\Gamma \subseteq \mathcal{EL}^* \times \mathcal{EL}^*$ such that

$$\varphi \approx_\Gamma \psi \text{ iff } K_\varphi \overline{\psi}, K_\psi \overline{\varphi} \in \Gamma.$$

In the following proofs, we make repeated use of Lemma 8 in a standard way as in the proof of Lemma 10 and omit mention of it.

Lemma 11 For all $\Gamma \in \mathcal{X}$, $\approx_\Gamma$ is an equivalence relation. Moreover, for all $\Gamma, \Delta \in \mathcal{X}$ such that $\Gamma \sim_\square \Delta$, we have $\approx_\Gamma = \approx_\Delta$.

Proof Let $\Gamma \in \mathcal{X}$ and $\varphi, \psi, \chi \in \mathcal{EL}^*_-$.

- reflexivity: By Ax8, we have $\vdash_{\text{EL}^*} K_\varphi \bar{\varphi}$, thus, $\varphi \approx_\Gamma \varphi$.
- symmetry: Suppose $\varphi \approx_\Gamma \psi$. This means, by the defn. of $\approx_\Gamma$, that $K_\varphi \bar{\psi}, K_\psi \bar{\varphi} \in \Gamma$. Therefore, $K_\psi \bar{\varphi}, K_\varphi \bar{\psi} \in \Gamma$, i.e., $\psi \approx_\Gamma \varphi$.
- transitivity: Suppose $\varphi \approx_\Gamma \psi$ and $\psi \approx_\Gamma \chi$. This means that (a) $K_\psi \bar{\varphi} \in \Gamma$, (b) $K_\varphi \bar{\psi} \in \Gamma$, (c) $K_\psi \bar{\chi} \in \Gamma$, and (d) $K_\chi \bar{\psi} \in \Gamma$. Then, by Ax6, (b), and (c), $K_\varphi \bar{\chi} \in \Gamma$. Similarly, by Ax6, (a), and (d), $K_\chi \bar{\varphi} \in \Gamma$. Therefore, $\varphi \approx_\Gamma \chi$.

For the last part, let $\Gamma, \Delta \in \mathcal{X}$ such that $\Gamma \sim_\square \Delta$. Suppose $\varphi \approx_\Gamma \psi$. This means that $K_\varphi \bar{\psi}, K_\psi \bar{\varphi} \in \Gamma$. Then, by Ax5, we obtain that $\square K_\varphi \bar{\psi}, \square K_\psi \bar{\varphi} \in \Gamma$. As $\Gamma \sim_\square \Delta$, we conclude that $K_\varphi \bar{\psi}, K_\psi \bar{\varphi} \in \Delta$, i.e., $\varphi \approx_\Delta \psi$. For the other direction, use the symmetry of $\sim_\square$. $\square$

Let $[\varphi]_\Gamma = \{\psi \in \mathcal{EL}^*_- : \varphi \approx_\Gamma \psi\}$, i.e., $[\varphi]_\Gamma$ is the equivalence class of φ with respect to $\approx_\Gamma$.

Definition 3 (Canonical Model for Γ_0) Let Γ_0 be a mcs of EL^* . The canonical model for Γ_0 is the tuple $\mathcal{M}^c = \langle W^c, @_{\Gamma_0}, \mathbb{T}^c, \oplus^c, v^c, t^c \rangle$, where

- $W^c = \{\Gamma \in \mathcal{X} : \Gamma_0 \sim_\square \Gamma\}$,
- $@_{\Gamma_0} = \{\varphi \in \mathcal{EL}^*_- : A\varphi \in \Gamma_0\}$,
- $\mathbb{T}^c = \{[\varphi]_{\Gamma_0} : \varphi \in \mathcal{EL}^*_-\}$ (we omit the subscript Γ_0 when it is clear from the context),
- $\oplus^c : \mathbb{T}^c \times \mathbb{T}^c \rightarrow \mathbb{T}^c$ such that $[\varphi] \oplus^c [\psi] = [\varphi \wedge \psi]$,
- $t^c : \mathcal{EL}^*_- \rightarrow \mathbb{T}^c$ such that, for all $x \in \mathcal{AT}$, $t^c(x) = [x]$ and $t^c(\varphi) = \oplus^c \mathcal{AT}(\varphi)$,
- $v^c : \mathcal{AT} \rightarrow \mathcal{P}(W^c)$ such that $v^c(x) = \{\Gamma \in W^c : x \in \Gamma\}$.

The topic inclusion relation $\leq^c$ on the canonical model is defined in the usual way.

Lemma 12 Given a mcs Γ_0 of EL^* and the canonical model $\mathcal{M}^c = \langle W^c, @_{\Gamma_0}, \mathbb{T}^c, \oplus^c, v^c, t^c \rangle$ for Γ_0 as described above, we have

- (1) $@_{\Gamma_0} \in W^c$,
- (2) $i \in @_{\Gamma_0}$ for all $i \in \mathcal{F}$,
- (3) for all $\varphi \in \mathcal{EL}^*_-$ and $\Gamma \in W^c$, if $\varphi \in @_{\Gamma_0}$ then $A\varphi \in \Gamma$,
- (4) for all $\Gamma \in W^c$, $@_\Gamma = @_{\Gamma_0}$.

Proof

- (1) By Lemma 10.
- (2) By Lemmas 10 and 8.(4).
- (3) Let $\varphi \in \mathcal{EL}^*_-$ such that $\varphi \in @_{\Gamma_0}$, i.e., $A\varphi \in \Gamma_0$. Then, by Ax2, we have $\square A\varphi \in \Gamma_0$. Hence, by the definition of W^c , $A\varphi \in \Gamma$ for all $\Gamma \in W^c$.
- (4) $@_{\Gamma_0} \subseteq @_\Gamma$ follows directly from item (3). For $@_\Gamma \subseteq @_{\Gamma_0}$, suppose $\varphi \in @_\Gamma$, i.e., that $A\varphi \in \Gamma$. Then, by Ax2, we have $\square A\varphi \in \Gamma$. As $\Gamma \in W^c$, i.e., $\Gamma \sim_\square \Gamma_0$, we obtain that $A\varphi \in \Gamma_0$, i.e., $\varphi \in @_{\Gamma_0}$. $\square$

Lemma 13 *The canonical model $\mathcal{M}^c = \langle W^c, @_{\Gamma_0}, \mathbb{T}^c, \oplus^c, v^c, t^c \rangle$ is an epistemic model.*

Proof

- $@_{\Gamma_0} \in W^c$ and $@_{\Gamma_0} \in v^c(i)$ for all $i \in \mathcal{F}$: The former is shown in Lemma 12.(1) and the latter is a consequence of Lemma 12.(2) and the definition of v^c .
- $\oplus^c$ is idempotent, associative, and commutative: All these properties of $\oplus^c$ follow from its definition and Ax8. $\square$

In order to prove an appropriate Truth Lemma, we need the following auxiliary lemmas.

Lemma 14 *For all $\varphi \in \mathcal{EL}^*$, $t^c(\varphi) = [\varphi]$.*

Proof Let $\varphi \in \mathcal{EL}^*$.

$$\begin{aligned}
 t^c(\varphi) &= \oplus^c \mathcal{AT}(\varphi) && \text{(by the definition of } t^c(\varphi)) \\
 &= t^c(x_1) \oplus^c \dots \oplus^c t^c(x_k) && \text{(for } \mathcal{AT}(\varphi) = \{x_1, \dots, x_k\}) \\
 &= [x_1] \oplus^c \dots \oplus^c [x_k] && \text{(by the definition of } t^c(x)) \\
 &= [x_1 \wedge \dots \wedge x_k] && \text{(by the definition of } \oplus^c) \\
 &= [\varphi] && \text{(Ax8)}
 \end{aligned}$$

$\square$

Lemma 15 *For all $\Gamma \in W^c$ and $\varphi, \psi \in \mathcal{EL}^*$, $t^c(\psi) \leq^c t^c(\varphi)$ iff $K_\varphi \overline{\psi} \in \Gamma$.*

Proof Let $\Gamma \in W^c$, and $\varphi, \psi \in \mathcal{EL}^*$, and observe that

$$\begin{aligned}
 t^c(\psi) \leq^c t^c(\varphi) &\text{ iff } t^c(\psi) \oplus^c t^c(\varphi) = t^c(\varphi) && \text{(by the definition of } \leq^c) \\
 &\text{ iff } [\psi] \oplus^c [\varphi] = [\varphi] && \text{(Lemma14)} \\
 &\text{ iff } [\psi \wedge \varphi] = [\varphi] && \text{(by the definition of } \oplus^c) \\
 &\text{ iff } K_{\psi \wedge \varphi} \overline{\varphi}, K_\varphi \overline{(\psi \wedge \varphi)} \in \Gamma_0 && \text{(by the definition of } [\varphi]) \\
 &\text{ iff } K_{\psi \wedge \varphi} \overline{\varphi}, K_\varphi \overline{(\psi \wedge \varphi)} \in \Gamma && \text{(Lemma11)}
 \end{aligned}$$

Suppose $t^c(\psi) \leq^c t^c(\varphi)$. Then, by the observation above, we have $K_{\psi \wedge \varphi} \overline{\varphi}, K_\varphi \overline{(\psi \wedge \varphi)} \in \Gamma$. By Ax8, we also have that $K_{\psi \wedge \varphi} \overline{\psi} \in \Gamma$. Then, $K_{\psi \wedge \varphi} \overline{\psi} \wedge K_\varphi \overline{(\psi \wedge \varphi)} \in \Gamma$. Thus, by Ax6, we obtain that $K_\varphi \overline{\psi} \in \Gamma$. For the other direction, suppose that $K_\varphi \overline{\psi} \in \Gamma$. By Ax8, we also have $K_\varphi \overline{\varphi} \in \Gamma$. Therefore, by Ax7, we obtain that $K_\varphi \overline{(\psi \wedge \varphi)} \in \Gamma$. By Ax8, we already have that $K_{\psi \wedge \varphi} \overline{\varphi} \in \Gamma$. Therefore, following the same steps above, we obtain that $t^c(\psi) \leq^c t^c(\varphi)$. $\square$

Lemma 16 (Existence Lemma for $\square$) *For every mcs $\Gamma \in \mathcal{X}$ and $\varphi \in \mathcal{EL}^*$, if $\square\varphi \notin \Gamma$ then there is a mcs Δ such that $\Gamma \sim_\square \Delta$ and $\varphi \notin \Delta$.*

Proof Standard, see, e.g., [8, Lemma 4.20]. $\square$

Lemma 17 (Truth Lemma) *Let Γ_0 be a mcs of EL^* and $\mathcal{M}^c = \langle W^c, @_{\Gamma_0}, \mathbb{T}^c, \oplus^c, v^c, t^c \rangle$ be the canonical model for Γ_0 . Then, for all $\Gamma \in W^c$ and $\varphi \in \mathcal{EL}^*_-$,*

$$\mathcal{M}^c, \Gamma \models \varphi \text{ iff } \varphi \in \Gamma.$$

Proof The proof is by induction on the structure of φ . The cases for the atomic formulae, Booleans, and $\Box\psi$ are standard, where the case for $\Box\psi$ uses Lemma 16. Toward showing the cases for $\varphi := K_\psi\chi$ and $\varphi := A\psi$, suppose inductively that the statement holds for ψ and χ .

Case $\varphi := K_\psi\chi$:

$$\begin{aligned} \mathcal{M}^c, \Gamma \models K_\psi\chi &\text{ iff } t^c(\chi) \leq^c t^c(\psi) \text{ and } \forall \Delta \in W^c (\text{ if } \mathcal{M}^c, \Delta \models \psi \text{ then } \mathcal{M}^c, \Delta \models \chi) \\ &\quad \text{(by the semantics)} \\ &\text{ iff } K_\psi\bar{\chi} \in \Gamma \text{ and } \forall \Delta \in W^c (\text{ if } \mathcal{M}^c, \Delta \models \psi \text{ then } \mathcal{M}^c, \Delta \models \chi) \quad \text{(Lemma 15)} \\ &\text{ iff } K_\psi\bar{\chi} \in \Gamma \text{ and } \forall \Delta \in W^c (\text{ if } \psi \in \Delta \text{ then } \chi \in \Delta) \quad \text{(by the induction hypothesis)} \\ &\text{ iff } K_\psi\bar{\chi} \in \Gamma \text{ and } \forall \Delta \in W^c (\psi \supset \chi \in \Delta) \quad \text{(Lemma 8)} \\ &\text{ iff } K_\psi\bar{\chi} \in \Gamma \text{ and } \Box(\psi \supset \chi) \in \Gamma \quad \text{(Lemma 16, the definition of } W^c) \\ &\text{ iff } (K_\psi\bar{\chi} \wedge \Box(\psi \supset \chi)) \in \Gamma \quad \text{(Lemma 8)} \\ &\text{ iff } K_\psi\chi \in \Gamma \quad \text{(Ax4 and Lemma 8)} \end{aligned}$$

Case $\varphi := A\psi$:

$$\begin{aligned} \mathcal{M}^c, \Gamma \models A\psi &\text{ iff } \mathcal{M}^c, @_{\Gamma_0} \models \psi \quad \text{(by the semantics)} \\ &\text{ iff } \varphi \in @_{\Gamma_0} \quad \text{(IH, Lemma 12.(1))} \\ &\text{ iff } A\varphi \in \Gamma \quad \text{(Lemma 12.(4))} \end{aligned}$$

□

Corollary 18 *EL^* is a sound and complete axiomatization of $\mathcal{EL}^*_-$ with respect to the class of epistemic models.*

Proof We prove only completeness since soundness is a matter of routine validity check. Let $\models_{\text{EL}^*} \varphi$. This means that $\{\neg\varphi\}$ is consistent, and by Lemma 9, can be extended to a mcs Γ_0 . Then, by Lemma 17, we obtain that $\mathcal{M}^c, \Gamma_0 \not\models \varphi$, where $\mathcal{M}^c$ is the canonical model for Γ_0 . □

B.1.1 The (static) Logic of Knowledge Over $\mathcal{EL}^+_+$

We now provide a sound and complete axiomatization for the fragment $\mathcal{EL}^+_+$ of $\mathcal{EL}^*_+$ without the actuality operator.

Theorem 19 *EL^+ given in Table 2 is a sound and complete axiomatization of $\mathcal{EL}^+_+$ with respect to the class of epistemic models.*

The proof of Theorem 19 follows similarly to the proof of Corollary 18 except that we need to replace Lemma 10 by the following lemma.

Lemma 20 *For all $\Gamma \in \mathcal{X}$, there exists a $@_\Gamma \in \mathcal{X}$ such that $\Gamma \sim_\Box @_\Gamma$ and $1 \wedge \dots \wedge n \in @_\Gamma$.*

Table 2 Axiomatization EL^+ (over $\mathcal{EL}^+_{\square}$)

(CPL)	all classical propositional tautologies and Modus Ponens
(S5 $_{\square}$)	S5 axioms and rules for $\square$
	(I) Axioms for i and $\square$:
(Ax1)	$\Diamond(1 \wedge \dots \wedge n)$
	(II) Axioms connecting K and $\square$:
(Ax2)	$K_{\varphi}\psi \equiv (K_{\varphi}\overline{\psi} \wedge \square(\varphi \supset \psi))$
(Ax3)	$K_{\varphi}\psi \supset \square K_{\varphi}\psi$
	(I) Axioms for K :
(Ax4)	$(K_{\varphi}\overline{\psi} \wedge K_{\varphi}\overline{\chi}) \supset K_{\varphi}\overline{\chi}$
(Ax5)	$(K_{\varphi}\overline{\psi} \wedge K_{\varphi}\overline{\chi}) \supset K_{\varphi}(\overline{\psi} \wedge \overline{\chi})$
(Ax6)	$K_{\varphi}\overline{\psi}$, if $\mathcal{AT}(\psi) \subseteq \mathcal{AT}(\varphi)$

Proof Let $\Gamma \in \mathcal{K}$. Consider the set $\{\varphi \in \mathcal{EL}^+_{\square} : \square\varphi \in \Gamma\} \cup \{1 \wedge \dots \wedge n\}$ and suppose it is inconsistent. This implies that there are finitely many $\varphi_1, \dots, \varphi_k \in \{\varphi \in \mathcal{EL}^+_{\square} : \square\varphi \in \Gamma\}$ such that $\vdash_{\text{EL}^+} (\varphi_1 \wedge \dots \wedge \varphi_k) \supset \neg(1 \wedge \dots \wedge n)$. Then, by S5 $_{\square}$, we obtain $\vdash_{\text{EL}^+} (\square\varphi_1 \wedge \dots \wedge \square\varphi_k) \supset \square\neg(1 \wedge \dots \wedge n)$. As each $\square\varphi_i \in \Gamma$ for $1 \leq i \leq k$, by Lemma 8, we have $\square\neg(1 \wedge \dots \wedge n) \in \Gamma$. This means that $\neg\Diamond(1 \wedge \dots \wedge n) \in \Gamma$, contradicting consistency of Γ due to Ax1 and Lemma 8.(3). Therefore, by Lemma 9, $\{\varphi \in \mathcal{EL}^+_{\square} : \square\varphi \in \Gamma\} \cup \{1 \wedge \dots \wedge n\}$ extends to a $\mathcal{EL}^+_{\square}$ -maximally consistent set $@_{\Gamma}$ such that $1 \wedge \dots \wedge n \in @_{\Gamma}$. As $\{\varphi \in \mathcal{EL}^+_{\square} : \square\varphi \in \Gamma\} \subseteq @_{\Gamma}$, we also have $\Gamma \sim_{\square} @_{\Gamma}$. $\square$

Lemma 20 guarantees the existence of an appropriate actual world $@_{\Gamma_0}$ in the canonical model for a maximally EL^+ -consistent set Γ_0 .

B.2 The Logic of Knowledge Update Over $\mathcal{EL}^*$

This section presents soundness and completeness results for the dynamic language $\mathcal{EL}^*$ of epistemic update. We axiomatize two classes of epistemic update models that represent *maximal* and *minimal* knowledge updates, respectively. These are the two extreme cases that an epistemic update function μ (given in Definition 1) can interpret: the former is modelled similarly to the so-called public announcements [25, 49, 50], and the latter represents an agent who learns the new piece of information without merging its intension with her prior information state.

Definition 4 (Maximal Epistemic Update Model) A *maximal epistemic update model* is a tuple $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ where $\langle W, @, \mathbb{T}, \oplus, v, t \rangle$ is an epistemic model and $\mu : \mathbb{I}_{@} \times \mathbb{I}_{@} \rightarrow \mathbb{I}_{@}$ is an update function such that $\mu(P, Q) = P \cap Q$.

Theorem 21 A sound and complete axiomatization $\text{EUL}_{\max}^*$ of $\mathcal{EL}^*$ with respect to the class of maximal epistemic update models is obtained by adding to EL^* the set of axioms and rules in Table 3.

Definition 5 (Minimal Epistemic Update Model) A minimal epistemic update model is a tuple $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$ where $\langle W, @, \mathbb{T}, \oplus, v, t \rangle$ is an epistemic model and $\mu : \mathbb{I}_@ \times \mathbb{I}_@ \rightarrow \mathbb{I}_@$ is an update function such that $\mu(P, Q) = Q$.

Theorem 22 A sound and complete axiomatization $\text{EUL}_{\min}^*$ of $\mathcal{EL}^*$ with respect to the class of minimal epistemic update models is obtained by replacing axiom (R_i) in Table 3 by

$$[i : \varphi]i \equiv A\varphi \supset (\varphi \wedge (i \vee \neg i)).$$

B.2.1 Proofs of Theorems 21 and 22

The proofs of Theorems 21 and 22 are by the so-called standard DEL-technique *completeness via reduction* (or *translation*), as briefly explained in Section 4.4. For a detailed presentation of completeness by reduction, we refer the reader to [62, Chapter 7.4].

Definition 6 (Complexity measure for $\mathcal{EL}^*$) The complexity $c(\varphi)$ of a formula $\varphi \in \mathcal{EL}^*$ is a natural number recursively defined as

$$\begin{aligned} c(i) &= c(p) = 1 \\ c(\neg\varphi) &= c(\Box\varphi) = c(A\varphi) = c(\varphi) + 1 \\ c(\varphi \wedge \psi) &= 1 + \max\{c(\varphi), c(\psi)\} \\ c(K_\varphi\psi) &= 4 + c(\varphi) + c(\psi) \\ c([i : \varphi]\psi) &= (3 + |\mathcal{AT}(\varphi)| + c(\varphi)) \cdot (|\mathcal{AT}(\psi)| + c(\psi)), \end{aligned}$$

where $|\mathcal{AT}(\varphi)|$ is the number of elements in $\mathcal{AT}(\varphi)$.

Table 3 Reduction axioms and inference rules for $[i : \varphi]$ for maximal epistemic update models

(Nec $_{\Box}$)	from ψ , infer $[i : \varphi]\psi$	
(R $_i$)	$[i : \varphi]i \equiv A\varphi \supset (i \wedge \varphi)$	
(R $_{x \neq i}$)	$[i : \varphi]x \equiv A\varphi \supset (x \wedge (i \vee \neg i))$	if $x \neq i$
(R $_{\neg}$)	$[i : \varphi]\neg\psi \equiv A\varphi \supset \neg[i : \varphi]\psi$	
(R $_{\wedge}$)	$[i : \varphi](\psi \wedge \chi) \equiv [i : \varphi]\psi \wedge [i : \varphi]\chi$	
(R $_{\Box}$)	$[i : \varphi]\Box\psi \equiv A\varphi \supset \Box[i : \varphi]\psi$	
(R $_A$)	$[i : \varphi]A\psi \equiv A\varphi \supset A[i : \varphi]\psi$	
(R $_K1$)	$[i : \varphi]K_\psi\chi \equiv A\varphi \supset (K_\psi\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi))$	if $i \notin \mathcal{AT}(\psi) \cup \mathcal{AT}(\chi)$
(R $_K2$)	$[i : \varphi]K_\psi\chi \equiv A\varphi \supset (K_\psi\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi))$	if $i \in \mathcal{AT}(\psi) \setminus \mathcal{AT}(\chi)$
(R $_K3$)	$[i : \varphi]K_\psi\chi \equiv A\varphi \supset (K_\psi(\bar{\chi} \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi))$	if $i \in \mathcal{AT}(\chi) \setminus \mathcal{AT}(\psi)$
(R $_K4$)	$[i : \varphi]K_\psi\chi \equiv A\varphi \supset (K_\psi\bar{\chi} \wedge (\chi \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi))$	if $i \in \mathcal{AT}(\psi) \cap \mathcal{AT}(\chi)$

Lemma 23 For all $\varphi, \psi, \chi \in \mathcal{EL}^*$ and $i, x \in \mathcal{AT}$,

- (1) $c(\varphi) > c(\psi)$ if ψ is a proper subformula of φ ,
- (2) $c([i : \varphi]i) > c(A\varphi \supset (i \wedge \varphi))$,
- (3) $c([i : \varphi]x) > c(A\varphi \supset (x \wedge (i \vee \neg i)))$, for $i \neq x$
- (4) $c([i : \varphi]\neg\psi) > c(A\varphi \supset \neg[i : \varphi]\psi)$,
- (5) $c([i : \varphi](\psi \wedge \chi)) > c([i : \varphi]\psi \wedge [i : \varphi]\chi)$,
- (6) $c([i : \varphi]\Box\psi) > c(A\varphi \supset \Box[i : \varphi]\psi)$,
- (7) $c([i : \varphi]A\psi) > c(A\varphi \supset A[i : \varphi]\psi)$,
- (8) $c([i : \varphi]K_\psi\chi) > c(A\varphi \supset (K_\psi\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi)))$,
- (9) $c([i : \varphi]K_\psi\chi) > c(A\varphi \supset (K_{\psi \wedge \varphi}\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi)))$,
- (10) $c([i : \varphi]K_\psi\chi) > c(A\varphi \supset (K_\psi(\chi \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi)))$,
- (11) $c([i : \varphi]K_\psi\chi) > c(A\varphi \supset (K_{\psi \wedge \varphi}(\chi \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi)))$.

Proof Follows by easy calculations using Definition 6. And, obviously, (8)–(10) follow from (11). $\square$

Definition 7 (Translation $f : \mathcal{EL}^* \rightarrow \mathcal{EL}_-^*$) The translation $f : \mathcal{EL}^* \rightarrow \mathcal{EL}_-^*$ is defined as follows:

$$\begin{aligned}
 f(x) &= x \\
 f(\neg\varphi) &= \neg f(\varphi) \\
 f(\Box\varphi) &= \Box f(\varphi) \\
 f(\varphi \wedge \psi) &= f(\varphi) \wedge f(\psi) \\
 f(K_\varphi\psi) &= K_{f(\varphi)}f(\psi) \\
 f([i : \varphi]i) &= f(A\varphi \supset (i \wedge \varphi)) \\
 f([i : \varphi]x) &= f(A\varphi \supset (x \wedge (i \vee \neg i))), \text{ for } x \neq i \\
 f([i : \varphi]\neg\psi) &= f(A\varphi \supset \neg[i : \varphi]\psi) \\
 f([i : \varphi](\psi \wedge \chi)) &= f([i : \varphi]\psi \wedge [i : \varphi]\chi), \\
 f([i : \varphi]\Box\psi) &= f(A\varphi \supset \Box[i : \varphi]\psi) \\
 f([i : \varphi]A\psi) &= f(A\varphi \supset A[i : \varphi]\psi) \\
 f([i : \varphi]K_\psi\chi) &= f(A\varphi \supset (K_\psi\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi))), \text{ if } i \notin \mathcal{AT}(\psi) \cup \mathcal{AT}(\chi) \\
 &= f(A\varphi \supset (K_{\psi \wedge \varphi}\bar{\chi} \wedge \Box[i : \varphi](\psi \supset \chi))), \text{ if } i \in \mathcal{AT}(\psi) \setminus \mathcal{AT}(\chi) \\
 &= f(A\varphi \supset (K_\psi(\chi \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi))), \text{ if } i \in \mathcal{AT}(\chi) \setminus \mathcal{AT}(\psi) \\
 &= f(A\varphi \supset (K_{\psi \wedge \varphi}(\chi \wedge \varphi) \wedge \Box[i : \varphi](\psi \supset \chi))), \text{ if } i \in \mathcal{AT}(\chi) \cap \mathcal{AT}(\psi) \\
 f([i : \varphi][j : \psi]\chi) &= f([i : \varphi]f([j : \psi]\chi)).
 \end{aligned}$$

We need the following lemma in order to be able to use the derived replacement of equivalents rule given in Theorem 7.(1) in the completeness proofs of $\text{EUL}_{\max}^*$ and $\text{EUL}_{\min}^*$. For this lemma to go through, it is crucial that the reduction axioms R_i and $R_{x \neq i}$ have occurrences of each element in $\mathcal{AT}(\varphi)$ and i on the right-hand-side of the equivalences, where φ is the sentence inside the dynamic operator and i is the updated frame of mind.

Lemma 24 For all $\varphi \in \mathcal{EL}^*$, $\mathcal{AT}(\varphi) = \mathcal{AT}(f(\varphi))$.

Proof The proof follows by an easy c -induction on the structure of φ and uses Lemma 23. Note that the case for $\varphi := [i : \psi]\chi$ requires subinduction on χ . $\square$

Lemma 25 For all $\varphi \in \mathcal{EL}^*$, $\vdash_{\text{EUL}_{\max}^*} \varphi \equiv f(\varphi)$.

Proof The proof follows by c -induction on the structure of φ and uses Lemma 23. Cases for the atomic formulae, the Boolean connectives, and $\square$ are elementary. Here we only show the cases for $\varphi := K_\psi \chi$ and $\varphi := [i : \psi]\chi$, where the latter requires subinduction on χ . Suppose inductively that $\vdash_{\text{EUL}_{\max}^*} \psi \equiv f(\psi)$, for all ψ with $c(\psi) < c(\varphi)$.

Case $\varphi := K_\psi \chi$

By Lemma 23.(1) and the induction hypothesis (IH), we have $\vdash_{\text{EUL}_{\max}^*} \psi \equiv f(\psi)$. Moreover, by Lemma 24, we have $\mathcal{AT}(\psi) = \mathcal{AT}(f(\psi))$. Therefore, by Ax8 in Table 1, we obtain $\vdash_{\text{EUL}_{\max}^*} K_\psi f(\psi)$ and $\vdash_{\text{EUL}_{\max}^*} K_{f(\psi)} \bar{\psi}$. Then, by Theorem 7.(1), we obtain $\vdash_{\text{EUL}_{\max}^*} K_\psi \chi \equiv K_{f(\psi)} \chi$. Similarly, we also have $\vdash_{\text{EUL}_{\max}^*} \chi \equiv f(\chi)$ and $\mathcal{AT}(\chi) = \mathcal{AT}(f(\chi))$, thus, $\vdash_{\text{EUL}_{\max}^*} K_\chi f(\chi)$ and $\vdash_{\text{EUL}_{\max}^*} K_{f(\chi)} \bar{\chi}$. Then, by Theorem 7.(1) again, we obtain $\vdash_{\text{EUL}_{\max}^*} K_{f(\psi)} \chi \equiv K_{f(\psi)} f(\chi)$. Therefore, by CPL, we conclude that $\vdash_{\text{EUL}_{\max}^*} K_\psi \chi \equiv K_{f(\psi)} f(\chi)$, with $K_{f(\psi)} f(\chi) = f(K_\psi \chi)$ by Definition 7.

Case $\varphi := [i : \psi]\chi$: we prove only the cases $\chi := x$ with $x \neq i$, and $\chi := [j : \theta]\alpha$. All the other cases follow similarly by using the corresponding reduction axioms, Lemma 23, and Definition 7.

Subcase $\chi := x$, for $x \neq i$

1. $\vdash_{\text{EUL}_{\max}^*} [i : \psi]x \equiv (A\varphi \supset (x \wedge (i \vee \neg i)))$ $R_{x \neq i}$
2. $\vdash_{\text{EUL}_{\max}^*} (A\varphi \supset (x \wedge (i \vee \neg i))) \equiv f(A\varphi \supset (x \wedge (i \vee \neg i)))$ Lemma 23.(3), IH

And, $f(A\varphi \supset (x \wedge (i \vee \neg i))) = f([i : \psi]x)$ by Definition 7.

Subcase $\chi := [j : \theta]\alpha$

By Lemma 23.(1) and IH, we know that $\vdash [j : \theta]\alpha \equiv f([j : \theta]\alpha)$

1. $\vdash_{\text{EUL}_{\max}^*} [j : \theta]\alpha \equiv f([j : \theta]\alpha)$ Lemma 23.(1), IH
2. $\vdash_{\text{EUL}_{\max}^*} [i : \psi]([j : \theta]\alpha \equiv f([j : \theta]\alpha))$ Nec $\square$
3. $\vdash_{\text{EUL}_{\max}^*} [i : \psi][j : \theta]\alpha \equiv [i : \psi]f([j : \theta]\alpha)$ R $\neg$, R $\wedge$
4. $\vdash_{\text{EUL}_{\max}^*} [i : \psi]f([j : \theta]\alpha) \equiv f([i : \psi]f([j : \theta]\alpha))$ IH
5. $\vdash_{\text{EUL}_{\max}^*} [i : \psi][j : \theta]\alpha \equiv f([i : \psi]f([j : \theta]\alpha))$ 3, 4, CPL

And, $f([i : \psi]f([j : \theta]\alpha)) = f([i : \psi][j : \theta]\alpha)$ by Definition 7. $\square$

Proof of Theorem 21 Let $\varphi \in \mathcal{EL}^*$ such that $\not\vdash_{\text{EUL}_{\max}^*} \varphi$. By Lemma 25, there is $\psi \in \mathcal{EL}_{\max}^*$ such that $\vdash_{\text{EUL}_{\max}^*} \varphi \equiv \psi$. Therefore, $\not\vdash_{\text{EUL}_{\max}^*} \psi$. As $\psi \in \mathcal{EL}_{\max}^*$ and $\text{EL}^* \subseteq \text{EUL}_{\max}^*$, we also have $\not\vdash_{\text{EL}^*} \psi$. Then, by Corollary 18, there is an epistemic model $\mathcal{M} = \langle W, @, \mathbb{T}, \oplus, v, t \rangle$ and $w \in W$ such that $\mathcal{M}, w \not\models \psi$. Observe that the tuple $\langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle$, where $\mu : \mathbb{I}_{@} \times \mathbb{I}_{@} \rightarrow \mathbb{I}_{@}$ such that $\mu(P, Q) = P \cap Q$, is a maximal epistemic update model. Then, by Lemma 5, we have $\langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle, w \not\models \psi$. Therefore, by the soundness of $\text{EUL}_{\max}^*$ with respect to maximal epistemic update models, we obtain that $\langle W, @, \mathbb{T}, \oplus, \mu, v, t \rangle, w \not\models \varphi$. $\square$

Proof of Theorem 22 Follows the same steps as the proof of Theorem 21. We need to replace every occurrence of $[i : \varphi]i \equiv A\varphi \supset (i \wedge \varphi)$ by $[i : \varphi]i \equiv A\varphi \supset (\varphi \wedge (i \vee \neg i))$. $\square$

References

1. Bergmann, M. (2006). *Justification without awareness*. Oxford: Clarendon Press.
2. Berto, F. (2018). Aboutness in imagination. *Philosophical Studies*, 175(8), 1871–1886.
3. Berto, F. (2019). Simple hyperintensional belief revision. *Erkenntnis*, 84(3), 559–575.
4. Berto, F., & Hawke, P. (2018). Knowability relative to information. *Mind*, Forthcoming.
5. Berto, F., & Jago, M. (2019). *Impossible worlds*. Oxford: Oxford University Press.
6. Bjerring, J.C. (2012). Impossible worlds and logical omniscience: an impossibility result. *Synthese*, 190, 2505–2024.
7. Bjerring, J.C., & Skipper, M. (2019). A dynamic solution to the problem of logical omniscience. *Journal of Philosophical Logic*, 48(3), 501–521.
8. Blackburn, P., de Rijke, M., Venema, Y. (2001). *Modal logic, volume 53 of Cambridge tracts in theoretical computer science*. Cambridge: Cambridge University Press.
9. Braddon-Mitchell, D., & Jackson, F. (2007). *The philosophy of mind and cognition: an introduction*, 2nd edn. Oxford: Blackwell Publishers.
10. Brown, J. (2018). *Fallibilism: evidence and knowledge*. Oxford: Oxford University Press.
11. Carnap, R. (1947). *Meaning and necessity*. Chicago University Press.
12. Christensen, D. (2004). *Putting logic in its place: formal constraints on rational belief*. Oxford: Oxford University Press.
13. Cresswell, M.J. (1985). *Structured meanings: the semantics of propositional attitudes*. Cambridge: MIT Press.
14. Crossley, J., & Humberstone, L. (1977). The logic of “actuality”. *Reports on Mathematical Logic*, 8, 11–29.
15. Dretske, F.I. (1970). Epistemic operators. *Journal of Philosophy*, 67(24), 1007–1023.
16. Dretske, F.I. (2005). The case against closure. In Steup, Sosa (Eds.) *Contemporary debates in epistemology*. 1st edn.: Wiley Blackwell, 13–26.
17. Fagin, R., & Halpern, J. (1988). Belief, awareness and limited reasoning. *Artificial Intelligence*, 34, 39–76.
18. Fagin, R., Halpern, J., Moses, Y., Vardi, M. (1995). *Reasoning about knowledge*. Cambridge: MIT Press.
19. Ferguson, T.M. (2014). A computational interpretation of conceptivism. *Journal of Applied Non-Classical Logics*, 24(4), 333–367.
20. Field, H. (1986). Critical notice: Stalnaker, Robert “Inquiry” *Philosophy of Science*, 53(3), 425–448.
21. Field, H. (1986). Stalnaker on intentionality: on Robert Stalnaker’s “Inquiry” *Pacific Philosophical Quarterly*, 67(2), 98–112.
22. Fine, K. (1986). Analytic implication. *Notre Dame Journal of Formal Logic*, 27, 169–79.
23. Fine, K. (2016). Angellic content. *Journal of Philosophical Logic*, 45(2), 199–226.
24. Floridi, L. (2015). Semantic conceptions of information. The Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/information-semantic/>.
25. Gerbrandy, J., & Groeneveld, W. (1997). Reasoning about information change. *Journal of Logic, Language and Information*, 6(2), 147–169. ISSN 1572-9583.
26. Giordani, A. (2019). Axiomatizing the logic of imagination. *Studia Logica*, 107(4), 639–657.
27. Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
28. Harman, G. (1986). *Change in view*. Cambridge: MIT Press.
29. Hawke, P. (2016). Questions, topics and restricted closure. *Philosophical Studies*, 73(10), 2759–2784.
30. Hawke, P. (2018). Theories of Aboutness. *Australasian Journal of Philosophy*, 96(4), 697–723.
31. Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
32. Hawthorne, J. (2005). The case for closure. In Steup, Sosa (Eds.) *Contemporary debates in epistemology*. 1st edn.: Wiley Blackwell, 26–43.

33. Hintikka, J. (1962). *Knowledge and belief. An Introduction to the logic of the two notions*. Ithaca: Cornell University Press.
34. Holliday, W. (2012). *Knowing what follows: Epistemic closure and epistemic logic*. PhD Dissertation, Stanford.
35. Holliday, W.H. (2015). Fallibilism and multiple paths to knowledge. In T. Gendler, & J. Hawthorne (Eds.) *Oxford studies in epistemology*, (Vol. 5 pp. 97–144): Oxford University Press.
36. Howe, M. (1970). *Introduction to human memory*. New York: Harper and Row.
37. Jago, M. (2007). Hintikka and Cresswell on logical omniscience. *Logic and Logical Philosophy*, 15, 325–54.
38. Jago, M. (2014). *The impossible. An essay on hyperintensionality*. Oxford: Oxford University Press.
39. Katzley, A. (1975). *Human memory: Structures and processes*. San Francisco: W.H Freeman.
40. Konolige, K. (1986). *A deduction model of belief*. San Mateo: Morgan Kaufman.
41. Lawlor, K. (2005). Living without closure. *Grazer Philosophische Studien*, 69(1), 25–49.
42. Levesque, H. (1984). A logic of implicit and explicit belief. *National conference on AI, AAAI-84*, 198–202.
43. Levi, I. (1991). *The fixation of belief and its undoing*. Cambridge: Cambridge UP.
44. Lewis, D. (1982). Logic for equivocators. *Nous*, 16, 431–441.
45. Lewis, D. (1988). Relevant implication. *Theoria*, 54(3), 161–174.
46. Moore, R., & Hendrix, G. (1982). Computational models of belief and the semantics of belief claims. In P. S. Peters, & E. Saarinen (Eds.) *Processes, beliefs and questions* (pp. 107–127). Dordrecht: Reidel.
47. Parry, W.T. (1989). Analytic implication: its history, justification and varieties. In Norman, & Sylvan (Eds.) *Directions in relevant logic* (pp. 101–118): Kluwer Academic Publishers.
48. Partee, B.H. (1973). The semantics of belief-sentences. In Hintikka, K.J.J., Moravcsik, J.M.E., Suppes, P. (Eds.) *Approaches to natural language: proceedings of the 1970 Stanford workshop on grammar and semantics* (pp. 309–336): D. Reidel Publishing Company.
49. Plaza, J. (1989). Logics of public communications. In *Proceedings of the 4th international symposium on methodologies for intelligent systems* (pp. 201–216). Reprinted as [50].
50. Plaza, J. (2007). Logics of public communications. *Synthese*, 158(2), 165–179.
51. Priest, G. (2005). *Towards non-being. The logic and metaphysics of intentionality*. Oxford: Oxford University Press.
52. Priest, G. (2016). Thinking the impossible. *Philosophical Studies*, 173, 2649–2662.
53. Rantala, V. (1982). Impossible world semantics and logical omniscience. *Acta Philosophica Fennica*, 35, 106–115.
54. Ryle, G. (1933). About. *Analysis*, 1(1), 10–12.
55. Schipper, B. (2015). Awareness. In H. van Ditmarsch, J. Halpern, W. van der Hoek, B. Kooi (Eds.) *Handbook of epistemic logic* (pp. 79–146). London: College Publications.
56. Sedlar, I. (2015). Substructural epistemic logics. *Journal of Applied Non-Classical Logics*, 25, 256–285.
57. Stalnaker, R. (1984). *Inquiry*. Cambridge: MIT Press.
58. Stalnaker, R. (1991). The problem of logical omniscience, i. *Synthese*, 89(3), 425–440.
59. Titelbaum, M. (2014). *Quitting certainties: a Bayesian framework modeling degrees of belief*. Oxford: Oxford University Press.
60. van Benthem, J.F.A.K. (2011). *Logical dynamics of information and interaction*. Cambridge: Cambridge University Press.
61. van Benthem, J.F.A.K., & Velázquez-Quesada, F. (2009). Inference, promotion, and the dynamics of awareness. *Knowledge, Rationality and Action*, 177, 5–27.
62. van Ditmarsch, H., van der Hoek, W., Kooi, B. (2008). *Dynamic epistemic logic*. Dordrecht: Springer.
63. Wansing, H. (1990). A general possible worlds framework for reasoning about knowledge and belief. *Studia Logica*, 49, 523–39.
64. Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
65. Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.
66. Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1), 23–47.