



Hyperintensionality and overfitting

Francesco Berto^{1,2} 

Received: 26 May 2023 / Accepted: 11 March 2024

© The Author(s) 2024

Abstract

A hyperintensional epistemic logic would take the contents which can be known or believed as more fine-grained than sets of possible worlds. I consider one objection to the idea: Williamson's Objection from Overfitting. I propose a hyperintensional account of propositions as sets of worlds enriched with topics: what those propositions, and so the attitudes having them as contents, are about. I show that the account captures the conditions under which sentences express the same content; that it can be pervasively applied in formal and mainstream epistemology; and that it is left unscathed by the objection.

Keywords Aboutness · Subject matters · Topics · Possible worlds semantics · Hyperintensional semantics

1 The hyperintensional revolution

Believers in propositions think that different sentences (or, sentence types) can sometimes say the same thing, or express the same content. We want a semantics to capture propositional content as what is said; and so, to capture same-saying. Call *Standard Possible Worlds Semantics* (SPWS) the view that such contents are sets of possible worlds giving truth conditions; or, equivalently, intensions: functions from possible worlds to extensions—in the relevant case, truth values. Call an operator H *hyperintensional* when $H\varphi$ and $H\psi$ can differ in truth value although φ and ψ express necessarily equivalent contents; a propositional content P , when P can differ from Q although they are necessarily equivalent. Necessary equivalence is usually understood as co-intensionality: truth at the same worlds. A hyperintensional semantics

✉ Francesco Berto
fb96@st-andrews.ac.uk; f.berto@uva.nl

¹ Arché, University of St Andrews, St Andrews, UK

² Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Amsterdam, The Netherlands

would account for some hyperintensional operators by postulating hyperintensional contents, i.e., contents that are individuated more finely than sets of possible worlds.

As propositions are often taken also as the objects of various attitudes, modal-epistemic logic in the Hintikka (1962) tradition represented knowledge and belief as modals—restricted quantifiers over possible worlds—and the things which can be known or believed as SPWS propositions. This was part of the 20th Century’s ‘intensional revolution’: a collective effort to account for a number of concepts (essence, causation, supervenience, conditionality, information) in terms of intensions. Troubles have emerged piecemeal, but have a common source: those notions appear to be hyperintensional. Nowadays there are lots of hyperintensional approaches to content, e.g., structured propositions of various sorts (Soames, 1985, 2010; King, 1996, 2007; King et al., 2014; Duží et al., 2010, 2023), non-normal or impossible worlds semantics (Jago, 2014; Priest, 2016; Berto & Jago, 2019), situation semantics (Barwise & Perry, 1983), truthmaker semantics (Fine, 2017; Fine & Jago, 2018), non-classical logics (Anderson & Belnap, 1975; Anderson et al., 1992; Dunn & Restall, 2002; Standefer, 2023). They differ vastly, but have this much in common: they distinguish contents identified in SPWS. Some talk of a ‘hyperintensional revolution’ (Nolan, 2014).

One early issue in epistemic logic was that the Hintikkan agents are ‘logically omniscient’: they know or believe all logical consequences of what they know or believe. One X s that φ (X being an ascription of knowledge, belief, perhaps some other attitude) at world w when the SPWS proposition that φ is true at all worlds epistemically accessible from w (compatible with one’s evidence at w , one’s belief system there, or whatnot). If entailment is truth preservation at all worlds of all models, one who X s that φ automatically X s all the entailed ψ s. Surely we aren’t like that. Also, all necessary truths turn out to be uninformative. That there are no solutions in positive integers for $x^n + y^n = z^n$ when $n > 2$ would be already known or believed. But it took a proof of over 130 pages to find out that it’s true.

That $2 + 2 = 4$, and that there are no solutions in positive integers for $x^n + y^n = z^n$ when $n > 2$, are true at the same worlds and so, by SPWS, are the same proposition. Perhaps that’s where the problem lies? One may be tempted to adopt some hyperintensional view of content and do epistemic logic systematically on that basis. (E.g., for works recruiting Kit Fine’s truthmaker semantics to do just that, see Krämer (2022); Hawke and Özgün (2023).) Under the label of ‘overfitting’, Sect. 2, introduces Timothy Williamson’s recent arguments against the very idea of being so tempted. Section 3 talks of what I take to be the core issue with SPWS: it offends what Steve Yablo has called ‘our sense of when sentences say the same thing’ (Yablo, 2014, p. 2). The overfitting strategy, it is argued, won’t easily explain this away. Section 4 briefly considers to what extent intensionalists may resort to guises to account for such a sense. Section 5 accounts for it by borrowing from Berto (2022) a little hyperintensional theory, developed by Peter Hawke and myself, where propositions are taken as made of two things: their truth set (the set of worlds where they are true) and their *topic* or *subject matter* (what they are about)—a version of what is coming to be called ‘two-component semantics’ (Yablo, 2014, 2017, 2018; Hawke, 2016; Plebani & Spolaore, 2023; Ferguson, 2023b, c; Berto & Hornischer, 2023; Hawke et al., 2024). Section 6 sketches how the theory can be put to work in formal and mainstream epistemology. Section 7 concludes.

2 Overfitting

The intensional revolution was resisted by one of the world's greatest philosophers: Quine. The hyperintensional revolution is resisted by one of the world's greatest philosophers: Timothy Williamson. According to him, hyperintensionalists are guilty of *overfitting*, 'the willingness to add extra parameters to an equation until its curve goes almost exactly through all the data points' (Williamson, 2020, p. 264). In science, overfitting is a bad feature of data models:

The flexibility of a model can be roughly measured by the number of its degrees of freedom, of adjustable parameters in the model. [...] By adding more and more degrees of freedom, one can fit just about any data, but in a cheap way which typically brings no insight. The problem with having too many degrees of freedom is not just un informativeness. It is also insensitivity to errors in the data, since the model can accommodate any data, however anomalous. [...] A better methodology is to be very reluctant to add new degrees of freedom, doing it only after potential sources of error in the data have been investigated and shown not to explain the evidence just as well. (Williamson, 2021, pp. 79–80)

Now hyperintensionalists, Williamson has it, base their accounts on intuitive data from language use, speakers' patterns of acceptance and rejection of sentences, which seem to tell against a merely intensional individuation of content: 'a key feature of the hyperintensional revolution is that it is driven by *examples*, especially by apparent counterexamples to intensional principles.' (Williamson, 2021, p. 87) Distinctions are postulated at the level of semantics in an effort to accommodate such data. But the data, Williamson claims, can be shown to be spurious precisely in the cases which were supposed to motivate the need for hyperintensional fine-graining.

Let us see how the objection is unpacked. (Williamson, 2020), a book on indicative and counterfactual conditionals, defends the view that the semantics of the indicative 'if' is given by the material conditional. The counterfactual is then accounted for by combining the material conditional with a 'would' operator, taken as a normal or merely intensional modal. There's a wealth of data providing apparent evidence against the material conditional analysis: competent speakers just don't seem to use the indicative 'if' as we would expect if its content was captured by the material conditional (Edgington, 1995). In the past, materialists tried to deal with the recalcitrant data at the level of pragmatics (Grice, 1989). But Williamson has a new strategy: borrowing terminology from cognitive psychology (e.g., Tversky & Kahneman, 1974), he claims the data are generated by fallible cognitive or epistemic heuristics in place in our assessment of conditionals: "'fast and frugal" or "quick and dirty" ways [of] answering questions which are reliable enough to be useful, but still not perfectly reliable' (Williamson, 2021, p. 83).

Should one 'wonder where to fit such heuristics for assessing sentences into a standard picture of linguistic architecture, with semantics built on syntax and pragmatics built on semantics' (Williamson, 2020, p. 24), he answers that 'the heuristics must come above the semantics, for normally one is in no position to decide between

accepting and rejecting a sentence until one knows what it means.’ (Ibid). However, it comes below the pragmatics:

Pragmatics is the usual first resort for filling the gaps between semantics and language use. However, it is not what is wanted here. In analysing conversational phenomena, pragmatics (legitimately) takes for granted speakers’ capacities to make the very cognitive assessments we are now seeking to understand. At the basic cognitive level, what we seek is a matter of psychology rather than linguistics. (Williamson, 2020, p. 5)

Williamson labels the heuristic governing our primary way of assessing conditionals the Suppositional Rule. It prescribes to take an attitude unconditionally to ‘If φ , then ψ ’ iff one takes that same attitude to the consequent ψ , conditional on the supposition of the antecedent φ (Williamson, 2020, p. 19). E.g., *Accept ‘If φ , then ψ ’ iff you accept ψ on the supposition of φ* . Next, in chapter 3 of his book Williamson goes on to show that the Suppositional Rule is *inconsistent*: when applied to attitudes to logical consequences of hypotheses, it leads to contradictions. When applied to attitudes that admit of degrees, such as credences, it delivers probabilistic paradoxes.¹ The moral is not that we don’t really use the Rule to assess conditionals. Rather, ‘we may have been using an inconsistent rule for “if” all along’ (Williamson, 2020, p. 41).

Next, the account of the counterfactual presented in Part II of the book agrees with the standard similarity-based SPWS of Stalnaker (1968) and Lewis (1973) in making all *counterpossibles*—counterfactuals with impossible antecedents—trivially true. Various hyperintensional accounts of counterfactuals (Kocurek, 2021, has a beautiful overview) take issue with this, starting from the intuitive data that we are not disposed to accept both of each pair of counterpossibles with the same antecedent but opposite consequents. For a famous example due to Daniel Nolan (1997), take ‘If Hobbes had (secretly) squared the circle, sick children in the mountains of South America at the time would have cared’ and ‘If Hobbes had (secretly) squared the circle, sick children in the mountains of South America at the time would not have cared’: we’ll reject the former and accept the latter. When we counterfactually suppose, *per absurdum*, that Hobbes managed to come up with a square equal in area to a given circle using only ruler and compass, we conclude that sick children of South America could not have cared less in the hypothetical scenario. But in chapter 11 of his book, Williamson argues that the suppositional heuristic for the counterfactual is inconsistent just as well (Sect. 11.4). In chapter 12, he argues that putative counterexamples to the intensionality of counterfactuals proposed by Fine (2012) don’t work, and concludes that no evidence has been provided that counterfactuals are hyperintensional after all.

In (Williamson, 2021), a similar story is told concerning attitudes. One way in which accounts more fine-grained than SPWS have been motivated, at least since

¹ One can check in particular Sects. 3.2–3.4 for the details of Williamson’s arguments. The results are not so surprising when one considers that the (acceptance instance of the) Suppositional Rule is in fact the so-called Ramsey Test for the assessment of indicatives (Ramsey, 1990) Peter Gärdenfors had already shown in the 80s that the Test is inconsistent with natural assumptions on belief revision (Gärdenfors, 1986). The probabilistic paradoxes are akin to the triviality results, due to Lewis (1976) and Hajek (1989), and others, which hit the hypothesis that the probability of an indicative amounts to the probability of the consequent conditional on the antecedent.

Carnap (1947)'s idea of an intensional isomorphism, has to do with contexts created by attitude ascriptions: we seem to sometimes truthfully say that one has an attitude towards the proposition that φ without having it towards an intensionally equivalent ψ : John Doe may believe that $2 + 2 = 4$ without believing Fermat's Last Theorem, etc. But, Williamson claims, Kripke (1979)'s Pierre puzzle should have alerted us to the possibility that our attitude ascriptions are also guided by heuristics which turn out to be occasionally inconsistent:

[Kripke] plausibly suggests that English speakers rely on something like the schema "A normal English speaker who is not reticent will be disposed to sincere reflective assent to 'p' if and only if he believes that p". Plausibly, users of other natural languages rely on analogous schemata. Call this family of schemata *the assent principle*. Combined with the convincing principle that correctly translating a belief-ascription preserves its truth-value, the assent principle generates inconsistency in describing the beliefs of a bilingual speaker in realistically possible circumstances (the famous case of puzzling Pierre). Indeed, as Kripke also explains, the problem arises even in the monolingual case.

Readers of 'A Puzzle about Belief' may interpret Kripke as suggesting that our ordinary concept of belief is incoherent, or something like that. In response, some may attempt to qualify the assent principle in more or less elaborate ways to avoid the contradiction. But there is a simpler possibility. The assent schema may be a normal heuristic for ascribing beliefs. [...] All the complications may be fall-backs we invoke when the basic heuristic fails. If this approach is correct, many of the apparent counterexamples to various theoretical claims presented in the voluminous and inconclusive literature on propositional attitude ascriptions may be errors generated by our implicit reliance on fallible heuristics. The many complicated accounts proposed for the semantics of propositional attitude ascriptions may just be artefacts of overfitting. (Williamson, 2021, p. 84)

The take-home message: we should be wary of 'the danger of giving semantic solutions to epistemic problems' (Williamson, 2020, p. vi), or we'll be dismissing good theory (in particular, SPWS) on the basis of bad data: 'the self-proclaimed hyper-intensional revolution involves multiplying degrees of freedom in order to explain data which may well be unreliable. That looks like a classic case of overfitting.' (Williamson, 2021, p. 93)

3 Saying the same thing

The Williamsonian stance is (consciously) at odds with contemporary research in linguistic semantics. A standard textbook (e.g., Chierchia & McConnell-Ginet, 1990) will tell you at the outset that a key task of semantics is to capture competent speakers' intuitions of synonymy, antonymy, entailment, equivalence, presupposition, etc. The theory is to start from ordinary language use, patterns of assent and dissent, shared judgments on what means what. (Where else could one start from? Self-evident axioms? A priori epiphanies?) It will end up generating analyses and predictions about (further)

use, assent and dissent, etc., which may be corroborated or disconfirmed, in a feedback loop, perhaps on the way to some reflective equilibrium. Everyone agrees that no theory can take on board all our intuitions. For these are bound to be inconsistent; but in the unfortunate case that dialetheism is wrong, no acceptable semantics can be inconsistent. However, as remarked by Rothschild (2021) in his *Mind* review of Williamson's book, Williamson takes the intuitive judgments delivered by our cognitive procedures, not as a guide to the semantics, but as things to be accounted for on their own, possibly in clear disagreement with the endorsed view of content. Conversely, the latter is not to account for patterns of sentence use, if not in the roundabout way mediated by the cognitive procedures.

Two can play the game! Call *underfitting* the over-simplification of models neglecting good data. One should be wary of retaining bad theory in spite of good data. What are the good data? I think they have to do, in the memorable Yabloian words I quoted above, with 'our sense of when sentences say the same thing' (Yablo, 2014, p. 2). Sure, it's our sense: the distinctions supposedly missed by SPWS will be motivated by our fallible judgments—again, where else is one to start from? But some data on (non-)same-saying may not be easily explained away as by-products of fallible heuristics. They may be stable across sentences of different kinds and modal profiles, systematic, and such that accounting for them at the level of content gives explanatory virtues, unmatched by alternative accounts relegating them outside of the semantics.

Different sentences seem to sometimes say different things in spite of being necessary (of the same kind of necessity), or co-necessary, thus co-intensional:

1. $2 + 2 = 4$.
2. $3 + 3 = 6$.
3. Equilateral triangles are equiangular.

These are mathematical necessities. Only one is about equilateral triangles, and made true by what they are like.²

4. Fido is a dog.
5. Kitty is a cat.
6. Water is H_2O .

For a number of essentialists, these are metaphysical necessities. Only one is about water.

7. If it snows, then it snows.
8. Either South Bend is in Indiana, or not.
9. The Liar sentence is not both true and false.

These are (classical) logical necessities (well, with some leeway with the truth predicate). Only one is about the Liar.

10. Clopen sets are sets which are both open and closed.
11. Non-normal modal logics are modal logics weaker than K.

² An anonymous referee of this paper wonders whether abstract types like numbers or geometrical figures may even be in the domain of any possible world, and takes this to be a further issue for the SPWS capacity to tell apart contents of this kind.

These are unrestrictedly necessarily true, if definitions of this kind are. Only one is about clopen sets.

12. Grass is green.
13. That grass is green is true,

These are arguably co-intensional, true at the same worlds. Only one is about a proposition. It is *prima facie* an issue for SPWS that it conflates such intuitively distinct contents. It has been raised several times in the literature. It is, for instance, part of Scott Soames (1985, 1987)'s influential criticisms of the idea that propositional contents are adequately captured as sets of possible worlds.³

Can't the desired distinctions be made within a merely intensional setting? That equilateral triangles are equiangular and that $2 + 2 = 4$ can be told apart, an intensionalist may say, by using extensions and/or mere intensions: those of the subsentential constituents of the relevant sentences. Say, 'equiangular' gets a function from possible worlds to extensions, somehow supposedly embedded in the former proposition, but not in the latter.

But this is a hyperintensional, not merely intensional view of content, probably in the vicinity of structured propositions. That equilateral triangles are equiangular and that $2 + 2 = 4$ keep being true at the same worlds, but now they are distinct contents for they feature distinct constituents. The constituents may be individuated merely intensionally, but the view is incompatible with the claim that contents just are sets of worlds.

How would the heuristics strategy come to the rescue? Williamson (2021) discusses examples used to motivate hyperintensionality in metaphysics, e.g., by proponents of grounding theories. The examples involve operators like 'because', 'it is essential to x that', or ' x brings it about that', such that substitution of necessary equivalents in their scope appears not to be truth-preserving:

14. It is essential to Socrates that he is Socrates.
15. It is essential to Socrates that he is a member of {Socrates}.
16. The proposition that grass is green is true because grass is green.
17. Grass is green because the proposition that grass is green is true.
18. Mary brought it about that John was a contributor.
19. Mary brought it about that John was a self-identical contributor.

We're inclined to accept the first item in each pair and reject the second, although the embedded sentences express co-intensional contents. We are guided by considerations of grounding and explanatory asymmetry, e.g., Socrates' existence and identity grounds the existence and identity of its singleton, but not vice versa; truth is grounded in facts and not vice versa.

But we have the same pattern of inclinations towards these:

20. Vera is a vixen because Vera is a female fox.

³ Soames targets accounts of contents as (mere) sets of truth-supporting circumstances, even more fine-grained than possible worlds. If the criticisms work in general (which is controversial: see Ripley (2012)), they will take down some hyperintensional accounts as well. This is of less importance for us; first, because the focus here is on defending hyperintensionality in general; second, because the toy hyperintensional account proposed below doesn't take propositions merely as sets of truth-supporting circumstances anyway.

21. Vera is a female fox because Vera is a vixen.
22. Richard brought it about that Edward was a king.
23. Richard brought it about that Edward was a male monarch.

However, these differ pairwise only by substitution of synonymous subsentential constituents⁴ they cannot express different propositions. Thus, there is something wrong with our inclinations. This might be taken as casting a shadow on cases (14.)–(19.) too: it may well be that ‘superficial linguistic features can easily deceive us into accepting unsound arguments for hyperintensionality.’ (Williamson, 2021, pp. 90) The diagnosis: in explanatory reasoning, the psychologically salient direction of explanation is from the simpler to the more complex, or from the more to the less familiar: ‘The helpfulness of a (putative) explanation is sensitive to its superficial linguistic form. For explanations are meant to provide *understanding*; how far they do so depends partly on their superficial linguistic features.’ (91) Thus, e.g., ‘The proposition that grass is green’ is a longer sentence than ‘Grass is green’, and grasping its meaning involves having some idea of what a proposition is, hence (16.) looks better than (17.). Sure, ‘female fox’ is two words while ‘vixen’ is only one, but the former two words are more familiar than the latter to the average English speaker, hence (20.) looks better than (21.).

But the hyperintensional stance is an existential claim: some distinct propositional contents, P and Q , are true at the same possible worlds. So for the strategy to work in general against hyperintensionality, it has to be universal. How would it extend to all of our initial list of cases?

The relevant intuitions connect the contents of sentences to what they are about. Aboutness is ‘the relation that meaningful items bear to whatever it is that they are *on* or *of* or that they *address* or *concern*’ (Yablo, 2014, p. 1). What meaningful items are about is their *subject matter* or, as I will also say, their *topic*. Research on topics has been burgeoning in recent decades (Lewis, 1988a, b; Gemes, 1994, 1997; Humberstone, 2008; Fine, 2016, 2017; Hawke, 2018; Moltmann, 2018; Schipper, 2018, 2020; Plebani, 2020; Plebani & Spolaore, 2021). What are topics? Some link them to questions sentences can be taken, in context, as answering to. Lewis (1988b)’s seminal example was *the number of stars*. It maps to the question, ‘What’s the number of stars?’. That splits modal space: two worlds end up in the same cell when they agree on the answer. The splitting gives what ‘There are ten stars’ can be about. Topics are partitions or divisions of modal space. On the other hand, any old sort of thing can serve as a topic: ‘Our topic in this lecture is neural networks’; ‘Today we talk about Margaret Thatcher.’ So other approaches to topic are more object-oriented (Goodman, 1961; Perry, 1989), or state-of-affairs-oriented. A prominent one by Fine (2017, 2020) takes topics as given by states or (exact) truthmakers/falsemakers: ‘There are ten stars’ can be about some situation which makes it true and which, unlike an entire possible world, is wholly relevant and responsible for its truth. Whatever the favored view of topics, there is nothing internalistic about aboutness.

Appeals to our sensitivity to superficial linguistic structure, syntactic simplicity or complexity, or familiarity with some terms rather than others, can’t plausibly work

⁴ An anonymous referee of this paper mentions that they think a male monarch needn’t be a king. So they’d object to the Williamsonian view that in (22.) and (23.) one is substituting synonyms.

for all of the above cases. ‘ $2 + 2 = 4$ ’ and ‘ $3 + 3 = 6$ ’ have the same simple syntactic structure; so do ‘Fido is a dog’ and ‘Kitty is a cat’. Neither includes words less familiar than the other: ‘vixen’ may sound less familiar than ‘female fox’ to the average English speaker, but it would be bizarre to claim that people are generally more familiar with the numeral ‘2’ than with the numeral ‘3’, or with ‘dog’ than with ‘cat’.

Nor do the cases above obviously involve considerations of explanatory perspicuity. They have been presented unembedded, not in contexts involving operators tied to explanatory reasoning. It’s not even directly about patterns of assent and dissent, or rational requirements constraining them. Qua competent speakers of English (in possession of the relevant, truthful information) we may accept all of the contents expressed by sentences (1.)–(13.). We may also be rationally committed to accepting them all—because they are true, necessarily true, logically entailed by other truths, necessarily equivalent to them and to each other, or whatnot. Still we will take them as expressing distinct contents, for they are about different things.

One may bring in an operator ‘says the same as’ and come up with widespread, but mutually inconsistent folk judgments on sentences of the form ‘That φ says the same as that ψ ’. One can come up with inconsistent intuitions involving more or less any operator. But how are we to overrule our judgments of non-same-saying in all of the above cases? If it were to be applied to them, Williamsonian talk of ‘fast and frugal heuristics generating illusions’, taken from cognitive psychology, may start to look misleading. The Kahneman-Tversky illusions, such as the conjunction fallacy (people are prone to judge it more likely that Linda is a bank teller and active in the feminist movement than that Linda is a bank teller: Tversky and Kahneman (1974)) are normally overcome by the experimental subjects, after they have understood the relevant explanation. It’s not that easy to concoct an explanation that will make competent speakers accept that, when it is said that $2 + 2 = 4$ and then that equilateral triangles are equiangular, the same thing has been said twice, setting aside superficial linguistic structure or explanatory salience.

4 Guises and disguises

So one may raise an Objection from Underfitting against intensionalists who stick to SPWS propositions in the face of certain resilient intuitions. Intensionalists will probably agree that an explanation is called for, but they’ll be happy to locate the explanatory material outside of semantic content.

One way may be to resort to guises or modes of presentation, taken as constraints on mental representations at work in pragmatics, or perhaps at the cognitive level intermediate between semantics and pragmatics Williamson has called our attention to. Guises may be legitimate devices for a number of purposes. They may be useful to hyperintensionalists as well: (neo-)Russellians on propositions like Salmon (1986) and Soames (1987) use them although Russellian structuralism makes for a hyperintensional individuation of content, way more fine-grained than SPWS. But can guises accommodate all good intuitions of aboutness and same-saying? The historical debate on guises quickly became subtle, but here’s a reconstruction in broad strokes.

Guises were not supposed to be constituents of content. While using them in the business of addressing Frege's puzzle, Salmon (1986) distinguished the information 'semantically encoded' in a sentence from the 'pragmatically imparted' one. People were (and, are) impressed by arguments for the direct reference theory of names and by Kripke's criticisms of descriptivism, thus reluctant to add components to the semantics of names besides their denotation—especially components that would look like Fregean descriptive senses. Guises were introduced to be activated in opaque intentional contexts. The issue was that, even on a hyperintensional Russellian account of propositions, that Hesperus is Hesperus is the same proposition as that Hesperus is Phosphorus, whereas *prima facie* it appears that we can truthfully say that the ancients believed the former, not the latter.

But intuitions of aboutness and same-saying are shaky here, in a way they are not for at least some of our examples (1.)–(13.) above. Perhaps 'Hesperus is Hesperus' and 'Hesperus is Phosphorus' are about the same thing, namely planet Venus, or whatever single topic is suitably associated to it. It's difficult to argue that this is true of ' $2 + 2 = 4$ ', ' $3 + 3 = 6$ ' and 'Equilateral triangles are equiangular'. Does the difference in what these say reduce to our ways of mentally representing the same things? If so, which things?

Even when guises are used only to explain why we accept some attitude ascriptions and reject others, it may be taken as a mandatory feature of guise theory that it account for some compositional phenomena, in particular involving embeddings. But this easily makes guises look like Fregean senses in disguise. A critique along these lines can be found in a famous review of Salmon's book by Forbes (1987). An objection to Forbes can be found in Branquinho (1990): Forbes' Fregean theory and Salmon's Russellian theory disagree in their assignment of truth values to attitude ascriptions. However, Graham Oppy (1992) argued, I think successfully, that the structural problem remains:

[I]t is as obvious that there must be a compositional theory involving Salmonian guises which issues in an assignment of assertability-values to sentence-context pairs as it is that there must be a compositional meaning theory for languages which issues in an assignment of truth-values to sentence-context pairs. (How else could we account for the fact that speakers can recognise the assertability-values of novel sentences? How else could speakers have the ability to produce and understand a potentially infinite range of sentences with attached assertability-values?) Moreover, it is equally clear that this theory will have exactly the same structure as the neo-Fregean theory of Fregean propositions. That is, it is clear that what Salmon's theory does is to shift some of the structure which is found in the Fregean theory from semantics to pragmatics. (Oppy, 1992, p. 4)

Now compositionality may be insufficient for falling on the side of content: there may well be compositionality at the epistemic and/or pragmatic levels. Still, I suspect something similar to what happened in the debate on guises may happen when someone attempts to protect SPWS by systematically explaining away all the putative counterexamples as mere cognitive differences alien to content. To be credible, the

explanation will have to start looking like a disguised theory of content: as semantics, under another name.

Williamson (2020)'s account of the indicative as material conditional plus heuristics is not quite there. Endorsing the simplest semantics for the indicative 'if', he offloads the explanatory work semanticists expect from a semantics to the epistemic heuristics. The latter can only look simple to the extent that the account of the relevant epistemic procedures is underdeveloped.⁵ I conjecture that, once fully developed, it may look at least as complicated as the rival, more complex semantics for the conditional he criticizes, and structurally like something many semanticists would call a semantics, but for the epistemic relabeling.

In the next section, I'll introduce an overfitting-free, hyperintensional account of content (a distillate of ideas from Berto (2022)) for a simple propositional language, capturing our robust judgments of (non-)same-saying. In the section after that I'll sketch how it can be pervasively put to use in epistemic logic.

5 A Little Hyperintensionality...

One who says 'Midori is an accountant' addresses a certain topic: one talks about Midori's job, what Midori does, or just Midori. And one says that things are such-and-so with respect to that topic. What one says is true just in case Midori's job is or includes being an accountant. We may then understand a proposition P as a pair, $P = \langle W_P, T_P \rangle$, and so as made of two components: (1) W_P is the truth set of P : the set of Worlds where it's true. (2) T_P is the Topic of P : what it's about, or directed to. W_P is just our old SPWS proposition ('thin proposition', as Yablo (2014) has it). The whole P also features a topic (a 'thick' proposition, or 'directed' if one likes: a proposition that points at a subject matter).

The literature on subject matters generally agrees on the space of topics having a natural mereological structure (Yablo, 2014; Humberstone, 2008; Yablo, 2014; Fine, 2016). Topics can have proper parts; distinct topics may have common parts. *Mathematics* includes *arithmetic*. *Mathematics* and *philosophy* share subject matter, having (certain parts of) *logic* in common. Correspondingly, what a proposition is about can overlap with, or be properly included in, what another one is about.

This also gives natural ideas of same-saying, saying more, etc. We claim that φ says at least as much as ψ (what ψ says is part of what φ says), when (1) φ entails ψ and (2) what φ is about includes what ψ is about. (As per Yablo's motto: 'Content-inclusion is implication plus subject matter inclusion', (Yablo, 2014, p. 15).) We claim that φ and ψ say the same when they are (1) co-intensional and (2) about the same things.

⁵ 'The distance between usage and semantics in Williamson's account is a sign, not of developed science, but rather of an underdeveloped theory. [...] What Williamson presents as a good prediction of his theory, as opposed to a bad prediction of the material conditional view, is mostly a demonstration of the wide range of data that can be accommodated by exploiting the underspecification of heuristics. This problem is compounded by the fact that Williamson also takes instances in which conditionals seem to act like material conditionals to vindicate his overall account. Given the freedom Williamson's account allows in explaining linguistic judgments, it is no surprise that the survey of empirical problems for the material conditional view in chapters 7 and 8 is triumphalist.' (Rothschild, 2013, pp. 11–12)

Next, the truth-functional logical vocabulary should be topic-transparent: it must add no subject matter of its own. There's some, though not universal, agreement on this: Yablo (2014) makes a forceful case for negation; Hawke (2018) and Fine (2020) make forceful cases for all the truth-functional connectives. The topic of $\neg\varphi$ should be the same as that of φ : 'Grass is not green' is exactly about what 'Grass is green' is about. It certainly doesn't address the topic of *negation*. Conjunction and disjunction merge topics: 'Carlos is short and handsome', 'Carlos is short or handsome' are about the height and looks of Carlos. The topics of $\varphi \wedge \psi$ is that of $\varphi \vee \psi$: the fusion of the topic of φ and that of ψ .⁶

Here's a simple hyperintensional semantics for a plain sentential language, using such ideas to capture same-saying. It is taken from ch. 2 of Berto (2022). The language \mathcal{L} has a countable set \mathcal{L}_{AT} of atoms, p, q, r (p_1, p_2, \dots), negation \neg , conjunction \wedge , disjunction \vee , the box of necessity \Box , two-place operators \approx and \supseteq , round parentheses as auxiliary symbols $(,)$. We use $\varphi, \psi, \chi, \dots$, as metavariables for formulas of \mathcal{L} . The well-formed formulas are the atoms and, if φ and ψ are well-formed, so are the following: $\neg\varphi \mid \Box\varphi \mid (\varphi \wedge \psi) \mid (\varphi \vee \psi) \mid (\varphi \approx \psi) \mid (\varphi \supseteq \psi)$. Outermost brackets are usually omitted. We identify \mathcal{L} with the set of its well-formed formulas. Read ' $\varphi \approx \psi$ ' as saying that φ and ψ have the same topic, ' $\varphi \supseteq \psi$ ' as

⁶ On request by an anonymous referee, in this long footnote I'll rehearse from the literature some arguments for the claim that the truth-functional connectives are topic-transparent (see, e.g. the aforementioned Yablo (2014); Hawke (2016); Fine (2020), but also Plebani and Spolaore (2021, 2023); Hawke et al. (2024)).

Well, why would one think otherwise, to begin with? The idea is inspired by the *Tractatus*' 'fundamental insight', namely that 'the "logical constants" do not represent' ([89], 4.0312). (Remember 5.44: 'if there was an object called " \sim ", then " $\sim\sim p$ " would have to say something other than " p ". For the one proposition would then treat of \sim , the other would not.'). And one may worry that, by following Wittgenstein, one is pushing talk purporting to be about the meaning of the logical vocabulary in the realm of pseudo-talk aiming at the ineffable.

But we needn't follow Wittgenstein that far. Surely we can talk about negation—it's only that we don't do it by negating sentences. Just in the same sense in which 'Grass is green' and 'Grass is not green' can be about *grass*, *what grass is like*, or *features of grass*, etc., so can 'Negation is a connective' and 'Negation is not a connective' be about *negation*, *what negation is like*, or *features of negation*, etc. But 'Grass is not green' can't be about *that* just because 'not' shows up in it.

One who thinks otherwise is probably taking what a sentence can be about—what topics it can be concerned with, what issues it can address, what questions it can be taken as an answer to in a conversational context—as plainly fixed by the contents (perhaps the reference or denotation) of its subsentential constituents; which, in general, it isn't ('Midori loves Akiko' and 'Akiko loves Midori' have constituents with the same meanings—Midori, Akiko, loving—but can be about different topics). It may well be that negation is an operator from truth values to truth values, and the 'not' in 'Grass is not green' means or denotes that, but of course that's not what 'Grass is not green' is about. 'Negation is an operator from truth values to truth values': this can be about negation. And in 'Negation is not an operator from truth values to truth values', it's 'negation', not 'not', that allows the sentence to be about negation. 'Negation is an operator from truth values to truth values' is about that same topic: it says the opposite, about *it*.

Besides, there's no easy candidate for a topic that's systematically introduced in discourse by a negation, a conjunction, or a disjunction. For there to be such a thing, one would need discourse contexts where, e.g., 'Grass is green' is on-topic, but 'Grass is not green' isn't, or vice versa, because the topicality component added by 'not' makes a difference—and it's hard to come up with such contexts. Surely at most one will be true in any context (unless dialetheism); but that is handled by truth conditions, not by topics. Similarly, there's no easily found context where 'Midori is an accountant' and 'Akiko is an accountant' are both on-topic but their conjunction isn't. But if 'and' added topic, there should be such contexts (Berto, 2022, Chap. 2). Surely there will be contexts where some φ and ψ are both acceptable for they're found probable enough, while their conjunction isn't; but probability of truth, too, is the business of truth conditions, not topic.

saying that the content of ψ is part of that of φ ; then $(\varphi \supseteq \psi) \wedge (\psi \supseteq \varphi)$ expresses same-saying. We use $x, y, z (x_1, x_2, \dots)$ for topics; w, w_1, w_2, \dots for possible worlds; P, Q, R, \dots for propositions. The semantics will recursively assign a thick proposition $[\varphi]$ to well-formed sentences φ of \mathcal{L} . We use $\lfloor \varphi \rfloor = \mathbb{W}_{[\varphi]}$ for the truth conditions of φ , $\lceil \varphi \rceil = \mathbb{T}_{[\varphi]}$ for its topic. (This slick notation is due to Peter Hawke, who co-authored that chapter.)

A *frame* for \mathcal{L} is a triple $\mathfrak{F} = \langle W, \mathcal{T}, \oplus \rangle$. W is a non-empty set of possible worlds; \mathcal{T} is a non-empty set of topics; and $\oplus : \mathcal{T} \times \mathcal{T} \rightarrow \mathcal{T}$ is topic fusion: an idempotent ($x \oplus x = x$), commutative ($x \oplus y = y \oplus x$), associative ($(x \oplus y) \oplus z = x \oplus (y \oplus z)$) operation making topics part of larger topics. So $\langle \mathcal{T}, \oplus \rangle$ is a join semilattice. For simplicity, fusion is unrestricted: $\forall xy \in \mathcal{T} \exists z \in \mathcal{T} (z = x \oplus y)$. One can then define topic parthood as $x \leq y := x \oplus y = y$, a partial ordering.

A *model* $\mathfrak{M} = \langle W, \mathcal{T}, \oplus, c, \iota \rangle$ adds to a frame two interpretation functions. To each atom p , the first assigns a truth set, $c(p) \subseteq W = \lfloor p \rfloor$, giving the Conditions under which p is true; the second, a Topic $\iota(p) = \lceil p \rceil \in \mathcal{T}$.⁷ The two are extended to the truth-functional composites as follows:

- $\lfloor \neg\varphi \rfloor = W \setminus \lfloor \varphi \rfloor$
- $\lfloor \varphi \wedge \psi \rfloor = \lfloor \varphi \rfloor \cap \lfloor \psi \rfloor$
- $\lfloor \varphi \vee \psi \rfloor = \lfloor \varphi \rfloor \cup \lfloor \psi \rfloor$
- $\lceil \neg\varphi \rceil = \lceil \varphi \rceil$
- $\lceil \varphi \wedge \psi \rceil = \lceil \varphi \vee \psi \rceil = \lceil \varphi \rceil \oplus \lceil \psi \rceil$

The left-hand recursion gives the usual Boolean algebra of thin propositions corresponding to the truth-functional vocabulary. The right-hand recursion secures the topic-transparency of that vocabulary. $\supseteq, \sqsubseteq, \approx$ are global operators:

- $\lfloor \varphi \supseteq \psi \rfloor = W$, if $\lfloor \varphi \rfloor \subseteq \lfloor \psi \rfloor$ and $\lceil \psi \rceil \leq \lceil \varphi \rceil$. Else: $\lfloor \varphi \supseteq \psi \rfloor = \emptyset$.
- $\lfloor \sqsubseteq \varphi \rfloor = W$, if $\lfloor \varphi \rfloor = W$. Else: $\lfloor \sqsubseteq \varphi \rfloor = \emptyset$.
- $\lfloor \varphi \approx \psi \rfloor = W$, if $\lceil \psi \rceil = \lceil \varphi \rceil$. Else: $\lfloor \varphi \approx \psi \rfloor = \emptyset$.

These only get truth sets, as we just care about the conditions under which the relevant sentences are true.⁸ So, relative to an interpretation in \mathfrak{M} , the (thick) content of a sentence φ , what it says, is $[\varphi] = \langle \lfloor \varphi \rfloor, \lceil \varphi \rceil \rangle$. Entailment is, completely standardly, truth preservation at all worlds of all models: $\varphi_1, \dots, \varphi_n \vDash \varphi$ if for every model \mathfrak{M} and $w \in W$, if $w \in \lfloor \varphi_i \rfloor$ for every i , then $w \in \lfloor \varphi \rfloor$. Validity is truth at all worlds of all models: $\vDash \varphi$ if for every model \mathfrak{M} and $w \in W$, $w \in \lfloor \varphi \rfloor$.

Here’s some (in)validities delivered by the semantics (the proofs are easy):

24. $\vDash (\varphi \wedge \psi) \supseteq \varphi$
25. $\not\vDash \varphi \supseteq (\varphi \vee \psi)$
26. $\vDash \varphi \sqsubseteq \varphi \wedge \varphi$
27. $\vDash (\varphi \wedge \psi) \sqsubseteq (\psi \wedge \varphi)$

⁷ Notice that distinct atoms can be about the same topic. How to account for topicality for the subsentential constituents of sentences, of course, is an issue that doesn’t pop up for \mathcal{L} . Interesting ideas can be found in Hawke (2016); Badura (2021); Plebani and Spolaore (2023).

⁸ And, currently I take no stance on how topicality may work for them. However, for some very promising ideas on the topics of non-extensional operators, in a two-component account of content like the one just sketched, see Ferguson (2023a, b, c).

28. $\not\models \varphi \supseteq (\varphi \wedge (\varphi \vee \psi))$
29. $\not\models \varphi \supseteq \neg(\neg\varphi \wedge \psi)$
30. $\models \varphi \approx \neg\neg\varphi$
31. $\models (\varphi \wedge \psi) \approx (\varphi \vee \psi)$
32. $\Box\varphi \wedge \Box\psi \not\models \varphi \approx \psi$
33. $\Box(\varphi \equiv \psi) \not\models \varphi \approx \psi$

Here's how these capture intuitions of same-saying. To begin with, *saying that* should transmit down to the parts of what is said, though not to the mere entailments (Yablo, 2014, Chap. 1). Aisha says: 'Midori is tall and thin'. Bethany says: 'Midori is tall'. What Bethany said has already been said by Aisha, who also said more: that Midori is thin as well. (Why Bethany said that after Aisha is an interesting issue better left to pragmatics.)

But Bethany hasn't thereby said that Midori is tall or a footballer. Sure, that easily follows from what she has said. Bethany may be rationally committed to that, supposing one is committed to all the logical consequences of what one says. She hasn't said *that*, however. (24.) and (25.) capture this by marking a difference between conjunction and disjunction: the content of a conjunction includes, and not just entails, that of its conjuncts; but the content of a disjunct does not perforce include that of the disjunction, in spite of entailing it. They also tell us why: the other disjunct can bring in extra topic. Following Yablo (2014) again, $\varphi \vee \psi$ can *say less about more* than φ : the disjunction can address a larger topic than that of one of its disjuncts, even while being less informative in that it rules out fewer worlds. $\varphi \wedge \psi$ can *say more about more* with respect to φ : it can both address a larger topic and be more informative, if it rules out more worlds. (24.) and (25.) are a widely recognized mark of a topic-sensitive semantics:

A paradigm of inclusion, I take it, is the relation that simple conjunctions bear to their conjuncts—the relation *Snow is white and expensive* bears, for example, to *Snow is white*. A paradigm of noninclusion is the relation disjuncts bear to disjunctions; *Snow is white* does not have *Snow is white or expensive* as a part. (Yablo, 2014, p. 11)

A guiding principle behind the understanding of partial content is that the content of A and B should each be part of the content of $A \wedge B$ but that the content of $A \vee B$ should not in general be part of the content of either A or B . (Fine, 2016, p. 200)

Aisha says: 'Scottish grass is green'. Bethany says: 'Scottish grass is green and it is green'. This sounds marked as redundant. That's because what Bethany said is but what Aisha just said. (26.) captures this: φ and $\varphi \wedge \varphi$ say the same thing. As Bethany is sticking to what has been said redundantly, we may want to step outside semantics and into pragmatics to make sense of her move: perhaps she wanted to stress the pervasiveness of green as the colour of Scottish grass.

Aisha says: 'Midori is an accountant and a football player'. Bethany says: 'Midori is a football player and an accountant'. Perhaps Bethany wanted to stress that playing football is what really matters for Midori. We resort to pragmatics again, because what

Bethany has said just is what Aisha has said. (27.) captures this: $\varphi \wedge \psi$ and $\psi \wedge \varphi$ say the same thing. (That happens, to be sure, when ‘and’ encodes order-insensitive, truth-functional conjunction. Sometimes ‘and’ can encode some kind of—e.g., temporal—ordering: ‘John went to the hospital and got ill’ can then say something different from ‘John got ill and went to the hospital’.)

Aisha says: ‘Midori is happy’. She hasn’t thereby said that Midori is happy and either Midori is happy or extremally disconnectedness is no hereditary property of topological spaces. (28.) captures this. Aisha says: ‘The car is out of fuel’. She hasn’t thereby said it’s not the case that the car has fuel but the gauge is stuck. (29.) captures this. (You may already have guessed where these take us once knowledge or belief ascriptions step in. We’ll get there in the next section.)

(30.)–(33.), taken together, give the core of the hyperintensional semantics. (30.) and (31.) express the transparency of the truth-functional logical vocabulary: that grass is green and that grass isn’t green are about the exact same topic. Of course, they say opposite things about that same topic, hence they are different contents. That Carlos is short and handsome and that Carlos is handsome and short are about the exact same topic, too, say, Carlos’ height and looks. (32.) and (33.) guarantee the possibility of topic-diverging necessities and co-necessities: that $2 + 2 = 4$ and that $3 + 3 = 6$ are both necessary, but about different things. That vixens are female foxes and that clopen sets are both open and closed—ditto. That grass is green and that it’s true that grass is green are about different things too, albeit co-intensional.

This is the barest sketch of a hyperintensional semantics. The view is developed in more detail (and confronted with some problems) in Berto (2022).

6 ... Goes a long way

Topic-sensitivity can do a lot of work in epistemic logic. In Berto (2022), operators are added to a language essentially like \mathcal{L} above, expressing conditional belief ($B^\varphi\psi$: one believes that ψ conditional on φ), belief revision ($[\varphi]\psi$: after revising one’s beliefs by φ , ψ is the case), knowability relative to information ($K^\varphi\psi$: ψ would be knowable for one given information φ), suppositional thinking ($I^\varphi\psi$: supposing φ , one imagines that ψ). They are all modals (variously restricted quantifiers over possible worlds, or constructions thereof) whose truth conditions are given in terms of topic-sensitive contents; hence they are labeled as Topic-Sensitive Intentional Modals (TSIMs).

Here’s a first application: we can ground a distinction, often made in the literature, between two kinds of logical closure principles for attitudes. E.g., Holliday (2012) takes \wedge -Elimination within knowledge operators (in the form: $K(\varphi \wedge \psi) \supset K\psi$) as a *pure* (contrast *deductive*) epistemic closure principle (see also Yablo (2017)). A deductive closure principle from $\varphi_1, \dots, \varphi_n$ to ψ has it that if an agent comes to believe ψ starting from $\varphi_1, \dots, \varphi_n$, by competent deduction, and all the while knowing each of $\varphi_1, \dots, \varphi_n$, then the agent knows ψ . This can always go wrong for realistic agents: the deduction may be too complex for our Joe Bloggs. But \wedge -Elim, qua pure closure principle, is such that ‘an agent cannot know $\varphi \wedge \psi$ without knowing ψ —regardless of whether the agent came to believe ψ by “competent deduction” from $\varphi \wedge \psi$ ’ (Holliday,

2012, p. 15). Pure closure is (as one referee appropriately asked to mention) mindful of the inferential skills or behaviour of the relevant agent.

This is so, TSIM theory explains, because when P is exactly about x , and one thinks (believes, knows, supposes, etc.) that P , one must be—taking a couple of metaphors from (Yablo, 2014, p. 39)—‘attentive to everything within x ’. But one can be ‘oblivious to matters lying outside of x ’ although there are propositions Q entailed by P which are about those (subject) matters. That is, the topic-sensitivity of attitudes, inherited from that of the propositions making for their contents, delivers some natural logical closure and non-closure properties. *Thinking that* (believing that, knowing that, supposing that, etc.) should transmit down to the parts of what one thinks (Yablo calls this ‘immanent closure’) though not to the mere entailments.

On the immanent closure side, e.g., when Aisha believes (plainly, or conditional on something else) that Midori is tall and thin, she thereby automatically believes that Midori is tall ($B^\varphi(\psi \wedge \chi) \models B^\varphi\psi$ is one validity in the TSIM semantics). That works also for supposedly anarchic mental activities like imagining, which is subject to voluntary control in ways belief is not: try and *imagine* that Midori is tall and thin without imagining that she is tall. That would be a bit like imagining that Midori is tall without imagining that she is tall, wouldn’t it? And so we find Williamson endorsing \wedge -Elim as a pure or immanent closure principle, first for knowledge:

... Knowledge of a conjunction is *already* knowledge of its conjuncts. ... There is no obstacle here to the idea that knowing a conjunction *constitutes* knowing its conjuncts, just as, in mathematics, we may *count* a proof of a conjunction as a proof of its conjuncts, so that if $p \wedge q$ is proved then p is proved, not just provable. (Williamson, 2000, pp. 282–283)

He then generalizes and conjectures that \wedge -Elim may hold for all positive attitudes (Ibid.): in believing a conjunction, one believes the conjuncts; in conceiving a conjunction, one conceives the conjuncts, etc.

This is so, TSIM theory explains, because what φ is about is (a proper) part of what $\varphi \wedge \psi$ is about. And by thinking about the whole, one has already thought about the parts: there’s nothing more for one to do, such that if one failed to do it one would be thinking that $\varphi \wedge \psi$ without thinking that φ . Thinking that φ is understood here as having a contentful mental state, endowed with intentionality and directed towards what it’s about. This draws a wedge with a merely syntactic conception of thought. If thinking that $\varphi \wedge \psi$ was having a sentence (say, of mentalese) tokened in the head, one may need to do something to move from thinking that $\varphi \wedge \psi$ to thinking that φ . And if one failed to apply \wedge -Elim to one’s mentalese sentence, one would have ‘ $\varphi \wedge \psi$ ’ in the head without having ‘ φ ’ there.

But immanent closure is weaker than full closure under entailment, and the TSIMs generally are not fully closed. Immanent closure is well-suited to capture pure closure as characterized above. One may think that φ without thinking that ψ although the former entails the latter, because one is not thinking about what ψ is about. That may happen for different reasons. One may think that φ (plainly, or given or supposing or conditionally on something else) without thinking that $\varphi \vee \psi$ although the former easily entails the latter ($X^\chi\varphi \not\models X^\chi(\varphi \vee \psi)$ is one invalidity in the semantics for a

number of TSIMs X), because one lacks some concept needed to grasp what ψ is about; and one cannot have attitudes such as knowing, believing, or even supposing, towards contents one cannot grasp.⁹ Williamson again:

\wedge -elimination has a special status. It may be brought out by a comparison with the equally canonical \vee -introduction inference to the disjunction $p \vee q$ from the disjunct p or from the disjunct q . Although the validity of \vee -introduction is closely tied to the meaning of \vee , a perfect logician who knows p may lack the empirical concepts to grasp (understand) the other disjunct q . Since knowing a proposition involves grasping it, and grasping a complex proposition involves grasping its constituents, such a logician is in no position to grasp $p \vee q$, and therefore does not know $p \vee q$. In contrast, those who know a conjunction grasp its conjunct, for they grasp the conjunction. (Williamson, 2000, pp. 282–283)

Aisha may believe that Midori is tall without believing that either Midori is tall or extremally disconnectedness is no hereditary property of topological spaces, because Aisha has no idea what topological spaces are. Topology is an alien topic to Aisha. Thus, the TSIMs are good for modeling agents with certain conceptual limitations. As Williamson remarks, these can affect a ‘perfect logician’. So they must be of a different kind from limitations due to the boundaries of one’s deductive capacities. Aisha’s conceptual incompetence has little to do with her bounded inferential resources: the disjunction is just one basic inferential step away.

You may think that Aisha is rationally committed to believing the disjunction. You may also think that, by believing that Midori is happy, Aisha is committed to believing that Midori is happy or Midori is happy and extremally disconnectedness is no hereditary property of topological spaces: one is rationally committed to being on top of *all* concepts. ‘Conceptually omniscient’ agents—agents who can think about anything, that is, who can entertain any propositional content—represent a normative ideal.

Or, you may have a Harmanian view of rationality: sometimes one should not think that ψ although one thinks that φ and the former entails the latter, *and* one is even perfectly on top of the concepts needed to think about what ψ is about. Chased by a predator, you run towards a small stream of water. Will you make it if you try and jump to the other side? Before blindly trying, you quickly simulate jumping in your mind; you bring in your knowledge or beliefs on the width of the stream, your physical abilities, etc., and you come to believe that you will make it if you jump. Are you committed to coming to believe that either you will make it or grass is green, or there’s life on Kepler-442b, or... Well, sure, you can entertain all such contents. But Harmanian normativity has it that you’d better focus on jumping now, before the predator is on you. Thus the TSIMs are good for capturing agents whose mental states are sensitive to relevance, in that they keep their suppositional and belief management procedures *on-topic*.

⁹ For a recent work on how to tackle the problem of conjunction elimination and disjunction introduction in attitude logic, that also addresses how to accommodate (hyperintensional counterparts of) these, see Jespersen (2023).

There's a wave of push-backs against Harman in formal epistemology (Christensen, 2004; Smithies, 2015; Titelbaum, 2015). It is sometimes stressed that the theoretical costs of coming up with formal epistemologies modeling non-logically-omniscient agents outweigh the benefits of sticking to standard Bayesianism or normal modal-epistemic logic. I don't think this debate should prevent one from attempting to formally capture ideas concerning Harmanian agents. Epistemic logicians won't wait for normative theorists to come to an agreement on the principles of rationality before they start building their models.

Aisha knows that the car is out of fuel on the basis of having checked the gauge, which reads 'empty'. Is she thereby positioned to know that it's false that the car has fuel but the gauge is stuck? Bethany may ask Aisha: the car is old; could it be the gauge is stuck? Aisha may then retract her knowledge claim. This connects to the debate on epistemic closure, 'one of the most significant disputes in epistemology over the last forty years' (Kvanvig, 2006, p. 256). One may list Williamson (2000), Hawthorne (2004), Roush (2010), Kripke (2011) among the yea-sayers; Dretske (1970), Nozick (1981), Lawlor (2005), Holliday (2015), Hawke (2016), Alspector-Kelly (2019) among the nay-sayers. I think the jury is out on this.

But look at the putative counterexamples to closure: you know it's a zebra on the basis of your sensory perception; you're not thereby positioned to know it's no cleverly disguised mule (Dretske, 1970). You know the table is red on the same kind of basis; you're not thereby positioned to know it's no white table under a deceiving red light (Cohen, 2002). You know where your car is parked on the basis of your memory of having left it there a few minutes ago. You're not thereby positioned to know it's false that it has been stolen and it's not there (Vogel, 1990). And so on. These are all of the form: one can know on a certain basis that φ without being positioned to know, on that same basis, that $\neg(\neg\varphi \wedge \psi)$. So they are all rooted in the addition of subject matter the information positioning you to know that φ is supposedly insensitive to, or incapable to provide you with evidence for, or, plainly, not *about*.

Closure nay-sayers have it that whatever justifies your belief in the former fails to transmit to the latter. They will also think that there's no other route for you to get to know the latter. They may be wrong on this. But TSIM theory can capture, via our invalidity (29.) from the previous section, how knowing something on a certain basis doesn't mean being positioned to know something else which is logically entailed, on *that same* basis. In particular, in the TSIM knowability-relative-to-information setting, $K^x\varphi \not\vdash K^x\neg(\neg\varphi \wedge \psi)$ (see Berto and Hawke (2021) and Chap. 4 of Berto (2022) for some details).

7 Conclusion

Hyperintensional distinctions, motivated by purely semantic considerations of aboutness, can be fruitfully put to use in formal and mainstream epistemology. A hyperintensional account of content may be safe from the Objection from Overfitting, insofar as it's based on robust judgments of same-saying. These may be difficult to account for at a cognitive level alien to content, without introducing devices which,

once developed in some detail, will look as semantics under another name and/or just as complex as the supposedly more complex hyperintensional rivals.

Acknowledgements Ideas on the Objection from Overfitting that made it into this paper were presented to the Logic & Metaphysics Group in Arché at St Andrews and during my José Gaos Chair lectures at the UNAM in Mexico City. I am grateful to both audiences, as well as to three anonymous referees of this Journal, for helpful remarks and objections. This research is published within the project ‘The Logic of Conceivability’, funded by the European Research Council (ERC CoG), Grant No. 681404.

Declarations

Conflict of interest The author declares that there are no financial or non-financial conflicts of interest directly or indirectly related to the work submitted.

Ethical approval The work includes no research involving human participants and/or animals.

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/>.

References

- Alspector-Kelly, M. (2019). *Against knowledge closure*. Cambridge University Press.
- Anderson, A. R., & Belnap, N. (1975). *Entailment: The logic of relevance and necessity I*. Princeton University Press.
- Anderson, A. R., Belnap, N., & Dunn, J. M. (1992). *Entailment: The logic of relevance and necessity II*. Princeton University Press.
- Badura, C. (2021). More aboutness in imagination. *Journal of Philosophical Logic*, 50, 523–47.
- Barwise, J., & Perry, J. (1983). *Situations and attitudes*. CSLI.
- Berto, F. (2022). *Topics of thought. The logic of knowledge, belief: Imagination*. Oxford University Press.
- Berto, F., & Hawke, P. (2021). Knowability relative to information. *Mind*, 130, 1–33.
- Berto, F., & Hornischer, L. (2023). Cognitive synonymy: A dead parrot? *Philosophical Studies*, 180, 2727–52.
- Berto, F., & Jago, M. (2019). *Impossible worlds*. Oxford University Press.
- Branquinho, J. (1990). Are Salmon’s ‘guises’ disguised Fregean senses? *Analysis*, 50, 19–24.
- Carnap, R. (1947). *Meaning and necessity*. Chicago University Press.
- Chierchia, G., & McConnell-Ginet, S. (1990). *Meaning and grammar: An introduction to semantics*. MIT Press.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. Oxford University Press.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–29.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–23.
- Dunn, J. M., & Restall, G. (2002). Relevance logic. In D. Gabbay & F. Guenther (Eds.), *Handbook of philosophical logic* (2nd ed., Vol. 6, pp. 1–136). Kluwer Academic.
- Duží, M., Jespersen, B., & Materna, P. (2010). *Procedural semantics for hyperintensional logic*. Springer.
- Duží, M., Jespersen, B., & Kostelec, M. (2023). *Transparent intensional logic: Selected recent essays*. College.
- Edgington, D. (1995). On conditionals. *Mind*, 104, 235–329.

- Ferguson, T. (2023). Subject-matter and intensional operators i: Conditional-agnostic analytic implication. *Philosophical Studies*, 180, 1849–1879.
- Ferguson, T. (2023). Subject-matter and intensional operators ii: Applications to the theory of topic-sensitive intentional modals. *Journal of Philosophical Logic*, 52, 1673–1701.
- Ferguson, T. (2023). Subject-matter and intensional operators iii: State-sensitive subject matter and topic-sufficiency. *Review of Symbolic Logic*. <https://doi.org/10.1017/S1755020323000230>
- Fine, K. (2012). Counterfactuals without possible worlds. *Journal of Philosophy*, 109, 221–46.
- Fine, K. (2016). Angelic content. *Journal of Philosophical Logic*, 45, 199–226.
- Fine, K. (2017). A theory of truthmaker content I: Conjunction, disjunction and negation. *Journal of Philosophical Logic*, 46, 625–74.
- Fine, K. (2020). Yablo on subject matter. *Philosophical Studies*, 177, 129–71.
- Fine, K., & Jago, M. (2018). Logic for exact entailment. *Review of Symbolic Logic*, 12, 536–56.
- Forbes, G. (1987). Review of Frege's Puzzle, by Nathan Salmon. *Philosophical Review*, 96, 455–8.
- Gärdenfors, P. (1986). Belief revisions and the Ramsey test for conditionals. *Philosophical Review*, 95, 81–93.
- Gemes, K. (1994). A new theory of content I: Basic content. *Journal of Philosophical Logic*, 23, 595–620.
- Gemes, K. (1997). A new theory of content II: Model theory and some alternatives. *Journal of Philosophical Logic*, 26, 449–76.
- Goodman, N. (1961). About. *Mind*, 70, 1–24.
- Grice, H. P. (1989). *Studies in the way of words*. Harvard University Press.
- Hajek, D. (1989). Probabilities of conditionals. *Journal of Philosophical Logic*, 18, 423–8.
- Hawke, P. (2016). Questions, topics and restricted closure. *Philosophical Studies*, 73, 2759–84.
- Hawke, P. (2018). Theories of aboutness. *Australasian Journal of Philosophy*, 96, 697–723.
- Hawke, P., & Özgün, A. (2023). Truthmaker semantics for epistemic logic. Relevance, and non-classical logic. In F. Faroldi & F. Van de Putte (Eds.), *Outstanding contributions to logic: Kit fine on truthmakers* (pp. 295–335). Springer.
- Hawke, P., Hornischer, L., & Berto, F. (2024). Truth, topicality, and transparency: one-component versus two component semantics. *Linguistics and Philosophy*. Online first
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.
- Hintikka, J. (1962). *Knowledge and belief. An introduction to the logic of the two notions*. Cornell University Press.
- Holliday, W. (2012). *Knowing what follows: Epistemic closure and epistemic logic*. PhD Dissertation, Stanford.
- Holliday, W. (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.
- Humberstone, L. (2008). Parts and partitions. *Theoria*, 66, 41–82.
- Jago, M. (2014). *The impossible. An essay on hyperintensionality*. Oxford University Press.
- Jespersen, B. (2023). Is act theory a propositional logic without logic? *The Journal of Philosophy*, 120, 484–510.
- King, J. (1996). Structured propositions and sentence structure. *Journal of Philosophical Logic*, 25, 495–521.
- King, J. (2007). *The nature and structure of content*. Oxford University Press.
- King, J., Speaks, J., & Soames, S. (2014). *New thinking about propositions*. Oxford University Press.
- Kocurek, A. (2021). Counterpossibles. *Philosophy Compass*, 16, 12787.
- Krämer, S. (2022). Mighty belief revision. *Journal of Philosophical Logic*, 51, 1175–1213.
- Kripke, S. (1979). A puzzle about belief. In A. Margalit (Ed.), *Meaning and use* (pp. 239–83). Reidel.
- Kripke, S. (2011). Nozick on knowledge. *Philosophical troubles: Collected papers*. (Vol. 1). Oxford University Press.
- Kvanvig, J. (2006). *The knowability paradox*. Oxford University Press.
- Lawlor, K. (2005). Living without closure. *Grazer Philosophische Studien*, 697, 25–49.
- Lewis, D. (1973). *Counterfactuals*. Blackwell.
- Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 95, 581–9.
- Lewis, D. (1988). Relevant implication. *Theoria*, 54, 161–74.
- Lewis, D. (1988). Statements partly about observation. *Philosophical Papers*, 17, 1–31.
- Moltmann, F. (2018). An object-based truthmaker semantics for modals. *Philosophical Issues*, 28, 255–88.
- Nolan, D. (1997). Impossible worlds: A modest approach. *Notre Dame Journal of Formal Logic*, 38, 535–72.

- Nolan, D. (2014). Hyperintensional metaphysics. *Philosophical Studies*, 171, 149–70.
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
- Oppy, G. (1992). Semantics for propositional attitude ascriptions. *Philosophical Studies*, 67, 1–18.
- Perry, J. (1989). Possible worlds and subject matter. *The problem of the essential indexical and other essays* (pp. 145–60). CSLI.
- Plebani, M. (2020). Why aboutness matters: Meta-fictionalism as a case study. *Philosophia*. <https://doi.org/10.1007/s11406-020-00272-9>
- Plebani, M., & Spolaore, G. (2021). Subject matter: A modest proposal. *The Philosophical Quarterly*, 71, 605–22.
- Plebani, M., & Spolaore, G. (2023). Predicative subject matter. *Philosophical Studies*, 181, 247–65.
- Priest, G. (2016). *Towards non-being: The logic and metaphysics of intentionality* (2nd ed.). Oxford University Press.
- Ramsey, F. P. (1990). General propositions and causality. In D. H. Mellor (Ed.), *Philosophical papers* (pp. 145–63). Cambridge University Press.
- Ripley, D. (2012). Structures and circumstances: Two ways to fine-grain propositions. *Synthese*, 189, 97–118.
- Rothschild, D. (2013). Do indicative conditionals express propositions? *Nous*, 47, 49–68.
- Rothschild, D. (2021). Living in a material world: A critical notice of suppose and tell: The semantics and heuristics of conditionals by timothy williamson. *Mind*, 132, 2.
- Roush, S. (2010). Closure on skepticism. *Journal of Philosophy*, 107, 243–56.
- Salmon, N. (1986). *Frege's puzzle*. MIT Press.
- Schipper, A. (2018). Aboutness and negative truths: A modest strategy for truthmaker theorists. *Synthese*, 195, 3685–722.
- Schipper, A. (2020). Aboutness and ontology: A modest approach to truthmakers. *Philosophical Studies*, 177, 505–33.
- Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192, 2769–93.
- Soames, S. (1985). Lost innocence. *Linguistics and Philosophy*, 8, 59–71.
- Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, 15, 47–87.
- Soames, S. (2010). *What is meaning?* Princeton University Press.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory (American Philosophical Quarterly Monographs 2)* (pp. 98–112). Blackwell.
- Standefer, S. (2023). Hyperintensionality in relevant logics. Rationality and interaction - LORI 2023 In N. Alechina, A. Herzig, & F. Liang (Eds.), *Logic* (pp. 238–50). Springer.
- Titelbaum, M. (2015). Rationality's fixed point, volume 5. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 253–94). Oxford University Press.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–31.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary perspectives on skepticism* (pp. 13–29). Kluwer.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, T. (2020). *Suppose and Tell*. Oxford: The Semantics and Heuristics of Conditionals. Oxford University Press.
- Williamson, T. (2021). Degrees of freedom: Is good philosophy bad science? *Disputatio*, 13, 73–91.
- Wittgenstein, L. (1921–1922). *Tractatus logico-philosophicus*. Routledge & Kegan Paul.
- Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.
- Yablo, S. (2017). Open knowledge and changing the subject. *Philosophical Studies*, 174, 1047–71.
- Yablo, S. (2018). Reply to Fine on Aboutness. *Philosophical Studies*, 175, 1495–512.