Hyperintensionality and Topicality: Remarks on Berto's *Topics of Thought*

Jens Christian Bjerring · Mattias Skipper
Penultimate draft, forthcoming in *Analysis*

These remarks are going to be about Francesco Berto's excellent new book *Topics of Thought: The Logic of Knowledge, Belief, Imagination* (Berto 2022). We begin with some context.

The standard possible worlds framework has often been criticized for being "too coarse-grained." The problem, the critics say, is that many purported applications of the framework require us to draw *hyperintensional* distinctions, i.e., distinctions between contents that have the same extension in every possible world. An example, to which we will return, concerns the individuation of semantic content. On the standard possible worlds semantics, two sentences say the same thing if they are truth-functionally equivalent, i.e., if they are true in precisely the same possible worlds. But many writers have urged that two sentences can say different things even if they are truth-functionally equivalent. If so, the standard possible worlds framework is too coarse-grained for the purposes of individuating semantic content.

Another example, to which we will also return, concerns the logic of propositional attitudes such as knowledge and belief. Take belief. On the standard possible worlds semantics of belief, an agent believes that p if, and only if, p is true in every possible world which, for all the agent believes, might be the actual one. From this it follows that the agent's set of beliefs is closed under logical entailment: the agent believes everything that follows logically from what she believes, including all logical necessities. But many writers have urged that agents need not believe everything that follows logically from what they believe. If so, the standard possible worlds framework is too coarse-grained for the purposes of developing a logic for belief (and likewise for knowledge).

In response to such challenges, semanticists and logicians have explored a variety of ways to draw hyperintensional distinctions. Two of the most widely discussed approaches invoke *impossible worlds* (Cresswell 1975; Hintikka 1975) or *structured propositions* (Soames 1987; King 1995). While there is much to say about each of these approaches and how they stack up against each other, our focus here will be on a different approach, which takes the *topic* or *subject-matter* of a sentence—what the sentence is *about*—to be the missing link in accounting for various hyperintensional phenomena. This "topic-based" approach to hyperintensionality has been explored in great depth in Yablo's seminal book *Aboutness* (2014), with predecessors in Goodman (1961) and Lewis (1988). And with *Topics of Thought*, Berto has managed to develop the topic-based approach in exciting new ways.

In what follows, we will first give a brief introduction to the semantic framework that Berto develops before raising what we see as some potential problems and limitations of the framework. We have divided our critical remarks into two parts. The first part focuses on issues having to do with principles of closure under implication. The second part focuses on issues having to do with hyperintensional phenomena arising from cognitive limitations. The critical nature of our remarks

should not be mistaken for an overall lack of sympathy towards Berto's project. We find ourselves in agreement with much of what Berto has to say about hyperintensionality. But in the hope of moving the debate a small step forward, we have found it useful to focus on points of contention rather than agreement, even if this means overemphasizing what divides us.

1. Berto's Semantic Framework

Berto opens his book by announcing its main goal: to explore a *new logic of thought*. More specifically, he asks:

> "[G]iven that one thinks (believes, knows, etc.) that φ, what other ψs does one think (believes, knows, etc.) by the logic of one's thought? Under which logical operations is one's thought closed?" (Berto 2022: 9.)

This question has a venerable history going back to Hintikka's pioneering work in *Knowledge and Belief* (Hintikka 1962), where he developed the first logic of knowledge and belief based on the possible worlds semantics. For present purposes, there is no need to rehearse the full technical details of Hintikka's logic, but it will be useful to have an informal statement of his semantics fresh in mind. Here is his semantics for *belief*:

> (**Hintikkan semantics**) Bp is true at a world w iff p is true at all possible worlds v such that wRv.

As usual, we write "Bp" to say that the agent believes p, and we write "wRv" to say that world v is *accessible* from world w, which in this context amounts to saying that v is not ruled out by anything the agent believes at w.

Berto's logic of thought can be seen as building a *topicality filter* on top of the Hintikkan semantics. To a first approximation, the function of the topicality filter is to prevent claims like "Bp" from being true if p is not "on the right topic." To illustrate, consider:

(1a) Apples are red.
(1b) Apples are red or Martians are hiding behind the International Space Station.

Since (1a) logically entails (1b), Hintikka's semantics implies that anyone who believes the former proposition must also believe the latter. Yet there is a clear intuitive sense in which (1a) and (1b) concern rather different topics. True, (1b) is in part about red apples, but it also concerns aliens and man-crafted objects in space. Seeing that the two propositions are about different topics in this intuitive sense, Berto wants to claim that one can believe (1a) without believing (1b).

By contrast, consider:

(2a) Apples are red and Martians are hiding behind the International Space Station.
(2b) Apples are red.

This time Hintikka's semantics gets the right result, but for the wrong reason. It is true, says Berto, that anyone who believes (2a) must also believe (2b). But the reason for this is not merely that (2a) logically entails (2b), but also that (2a) is "on-topic" with respect to (2b) in the sense that (2a) is

about everything that (2b) is about. In general, while the topic of p is always included in the topic of $p \lor q$, the topic of $p \lor q$ is not always included in the topic of p. This is why, according to Berto, belief is closed under conjunction elimination but not closed under disjunction introduction, although both closure principles are truth-preserving.

What goes for belief goes for other propositional attitudes, as well. For example, on Berto's view, whether an agent *knows* or *imagines* that p depends, in part, on whether p is on the right topic. And whether p is "on the right topic" is again relative to a proposition under consideration: p is on-topic relative to q if, and only if, p is about everything that q is about.

At a more formal level, Berto suggests that we model propositional attitudes by means of what he calls "Topic-Sensitive Intentional Modals" (TSIMs): two-place operators of the form "$X^p q$" with the generic reading "Given p, one Xs that q," where X is some propositional attitude such as belief or knowledge. The basic semantics for $X^p q$ is similar to Hintikka's semantics, but with an added topicality constraint:

> (**Bertonian semantics**) $X^p q$ is true at a world w iff the following two conditions are satisfied:
> (1) q is true at all worlds v such that $wR_p v$, and
> (2) q's topic is wholly included in p's topic.

As witnessed by condition (1), Berto's semantics still relies on quantification over possible worlds, restricted by an accessibility relation. Furthermore, it still allows us to impose different properties on the accessibility relation (e.g., reflexivity, symmetry, and transitivity) depending on the application of interest. But the first new invention is that the accessibility relation, R_p, is indexed to a sentence, p, which means that there are as many accessibility relations as there are sentences in the formal language. Intuitively, we can think of the index p as determining which class of worlds the agent "attends to" given input p. (What this means, exactly, depends on the application of interest, as we will see later.)

An important consequence of indexing the accessibility relation to sentences is that $X^p q$ does not in general entail $X^{p \& r} q$, because q need not be true at all $R_{q \& r}$-accessible worlds, even if q is true at all R_q-accessible worlds. In other words, TSIMs can behave *non-monotonically*, just as they can violate various closure principles.

We find the second departure from Hintikka's semantics in condition (2), which says that $X^p q$ is true only if q is "on-topic" with respect to p in the sense that q is about everything that p is about. This constraint ensures, among other things, that TSIMs can be closed under conjunction elimination without being closed under disjunction introduction (once again, because p's topic is included in that of $p \& q$, whereas $p \lor q$'s topic need not be included in that of p).

More generally, when it comes to invalidating various inferential patterns, much of the interesting work is done by the topicality constraint. So an important question is when one topic is included in another. While Berto embraces Yablo's initial characterization of aboutness as "the relation that meaningful items bear to whatever it is that they are *on* or *of* or that they *address* or *concern*" (Yablo 2014: 1), he advises us not to focus (for present purposes) on the metaphysical nature of topics, but rather to look at what structural constraints they should obey (Berto 2022: 35). We

shall not here summarize everything Berto has to say about these structural constraints on topics, but a few central points are worth highlighting.

First of all, we have already seen that Berto—along with Yablo (2014), and others—thinks of topics as having *mereological structure*: we can meaningfully talk about whether one topic is included in another, whether two topics are disjoint or partially overlapping, what topic results from merging one topic with another, and so on. For example, we might ask whether the fusion of the topics of "Apples are edible" and "Pears are edible" is included in the topic of "All fruits are edible."

Second, topicality is *hyperintensional*: necessarily co-extensional expressions may differ in topic. For example, while both "Water is H_2O" and "Socrates is human" are (arguably) *a posteriori* necessary and hence true in all possible worlds, they are not about the same topic: one is about the chemical composition of water, the other about Socrates and whether he is human. An exception, to which we shall return, concerns necessarily co-extensional expressions containing distinct terms that rigidly designate the same object. Compare "Cicero was a great orator" with "Tully was a great orator." On Berto's view, as we understand it, these sentences not only have the same truth-conditions, but also the same topic: both concern whether a certain person was a great orator; and the person in question is the same in both cases.

Third, Berto maintains—again in line with Yablo (2014)—that logical connectives such as negation and conjunction are *topic-transparent*: they do not add any topic of their own. For example, "Jane is a lawyer" and "Jane is *not* a lawyer" concern the same topic, just as the topic of the conjunction "Apples are red and pears are green" is identical to the merged topic of the conjuncts "Apples are red" and "Pears are green." Although logical connectives are topic-transparent in this sense, the topic of a sentence is not simply a function of the topics of its subsentential constituents. For example, although "2+2=4" and "4+2=2" have the same subsentential constituents, they nonetheless differ in topic: if we are discussing whether 2 and 2 add to 4, it would seem off-topic to start talking about whether 4 and 2 add to 2 (Berto 2022: 32). Or, to take an amusing example from Yablo (2014: 24): a headline saying "Man Bites Dog" concerns a rather more interesting—and hence different—topic than one saying "Dog Bites Man," although the two headlines have the same subsentential constituents.

With the basic semantics in place, and with a rough understanding of how topics operate, we can now illustrate two of several interesting applications of the semantic framework that Berto discusses in his book. Given the richness of each application, it is futile to go into all of the fine details, but we can do enough to give a sense of the kinds of issues that Berto tackles.

The first application, which is the topic of chapter 4 (co-authored with Peter Hawke[1]), centers around a notion of *knowability relative to information*. Here the "X" in $X^p q$ is replaced by a knowability operator, K, and "$K^p q$" gets read as: "Given total information p, one is in a position to know q." Berto and Hawke use the semantics of $K^p q$ to shed new light on a number of longstanding topics in epistemology, including Cartesian skepticism and the Kripke-Harman Dogmatism Paradox. Here we will focus on skepticism. Since closure under entailment fails in the basic semantics for TSIMs, it fails for $K^p q$ in particular: even if $K^p q$ is true and q logically entails r, there is no guarantee that $K^p r$ is true, because r's topic need not be contained in q's topic. Berto and Hawke argue that this sort of

---

[1] The chapter is based on Berto and Hawke (2021).

closure failure is friendly to familiar fallibilist attempts to fend off Cartesian skepticism. Roughly speaking, if closure under entailment fails, then we can no longer appeal to the conditional that having hands entails not being a brain-in-a-vat to establish knowledge that we are not brains-in-a-vat from our knowledge of having hands. But then neither can the skeptic appeal to the contrapositive of this conditional to infer that we do not know that we have hands from our incapacity to know that we are not brains-in-a-vat.

The second application we want to mention at this initial stage is developed in chapter 5 and centers around a notion of *imagination given a supposition*. Here the "X" in $X^p q$ is replaced by an imagination operator, I, and "$I^p q$" gets read as: "In an act of imagination starting with suppositional input p, one imagines that q." Importantly, Berto is not simply interested in daydreaming or unconstrained speculation, but rather in a regimented form of suppositional thought where "we keep reality more or less as we know or believe it to be, compatible with what is needed for our initial supposition to hold" (Berto 2022: 111). On this "reality-oriented" way of understanding imagination, it is easy to appreciate the suppositional component of $I^p q$: it is through the act of supposing that p (together with our background knowledge and beliefs) that we can keep what we imagine in check with reality.

Berto goes on to argue that the semantics for $I^p q$ can be used to capture various features that research in cognitive science and philosophy of mind associates with a reality-oriented type of imagination. For instance, such imagination is plausibly not closed under entailment: even if one imagines that "Hitler gets killed" on the supposition that "Stauffenberg had put the bomb on the other side of the table," one need not imagine on that same supposition that "Either Hitler gets killed or there is life on Kepler-442b." Nor is imagination of the relevant reality-oriented type monotonic: even if one imagines that "Birdy flies" on the supposition that "Birdy has wings," one need not imagine that "Birdy can fly" on the supposition that "Birdy has wings and is a penguin."

After arguing that knowledge relative to information ($K^p q$) and imagination given a supposition ($I^p q$) can be modeled using the semantics for TSIMs, Berto investigates several other applications of the semantics to issues having to do with hyperintensional belief revision (chapter 6), framing effects (chapter 7, co-authored with Aybüke Özgun), and the connection between probabilities and conditionals (chapter 8, also co-authored with Aybüke Özgun). But rather than considering these further applications in detail, let us instead raise what we see as some potential problems and limitations of the semantic framework.

2. Closure Under Implication

The first set of issues we want to raise concern the following closure principles (where "$\rightarrow$" is the material conditional, and "$\Rightarrow$" is the strict conditional):

(**Closure Under Xed Implication**) $\{X^p q, X^p (q \rightarrow r)\} \vDash X^p r$.

(**Closure Under Strict Implication**) $\{X^p q, q \Rightarrow r\} \nvDash X^p r$.

Regardless of how *X* is interpreted, the first of these principles holds in the basic Bertonian semantics while the second does not. Berto takes these facts as a strength of his semantics. We are more doubtful, as we shall argue now.

2.1. Closure Under Known Implication

When "*X*" is interpreted as "*K*", we get the following instance of Closure Under *X*ed Implication:

(**Closure Under Known Implication**) $\{K^p q, K^p(q \rightarrow r)\} \vDash K^p r$.

On the face of it, this principle might seem plausible enough. But one of Berto and Hawke's main reasons for denying Closure Under *Strict* Implication is equally a reason to deny Closure Under *Known* Implication. Recall that one of Berto and Hawke's main motivations for denying Closure Under Strict Implication stems from anti-skeptical considerations. If Closure under Strict Implication fails to hold, you might know that you have hands without knowing that you are not a brain-in-a-vat (BIV), although the former entails the latter. To many, this is a desirable result. After all, the thought goes, while it seems clear that we can know ordinary facts such as the fact that I have hands, it also seems plausible that we cannot know that we are not BIVs precisely because the skeptical BIV hypothesis is constructed in a way that seemingly prevents us from gaining any evidence of its falsity.

There are various ways to develop this anti-skeptical thought. Nozick (1981), for instance, argued that *sensitivity* is a necessary condition on knowledge, where one's belief that *p* is sensitive if, and only if, one would not have believed *p*, if *p* had been false. On this picture, I can know that I have hands, because my belief in this proposition is sensitive (if I had not had hands, I would not have believed that I did), whereas I cannot know that I am not a BIV, because my belief in this proposition is not sensitive (if I had been a BIV, I would still have believed that I was not). In a similar vein, Dretske (1970) argued that *knowing that*—unlike *entails that*—is only a "semi-penetrating" operator: one knows *some*, but not *all* of the logical consequences of what one knows. In particular, one's knowledge that one has hands does not penetrate to the fact one is not a BIV.

Regardless of the exact theoretical underpinnings, suppose we agree with this broadly anti-skeptical motivation for denying Closure under Strict Implication. Given this, we also seem to have compelling reason to deny Closure Under Known Implication. Why? Well, just as I can know the simple empirical fact that I have hands, I can presumably know the simple analytical fact that if I have hands, then I am not a BIV.[2] Yet, even if I *do* know both of these facts, I might still not be in a

---

[2] Berto and Hawke might perhaps try to resist this claim. Elsewhere, in their discussion of the Dogmatism Paradox, they claim that a proposition can be knowable *a priori* without being knowable relative to a body of empirical information. In support of this claim, they observe that, although it is knowable *a priori* that 2 + 2 = 4, it "would be odd to conclude that 2 + 2 = 4 can be known on the basis of the news that Beth's grandmother is ill" (Berto 2022: 102). Odd indeed—but not, in our view, a compelling reason to deny that 2 + 2 = 4 is knowable on the basis of the news that Beth's grandmother is ill. Rather, it seems natural to explain the oddness in pragmatic terms: the reason why it seems odd to say "I know that 2 + 2 = 4 on the basis of such-and-such empirical information" is that it would seem to imply that you would not have been in a position to know that 2 + 2 = 4 without this empirical information. Given that a natural pragmatic explanation is available, we see little reason to deny that if a proposition is knowable *a priori*, it will also be knowable relative to a body of empirical information. In particular, we see little reason to deny that if you are in a position to know that you have hands based on your total body of evidence, then you are also in a position to know the analytical fact that if you have hands, then you are not a BIV, based on that total body of evidence.

position to know that I am not a BIV; or so the anti-skeptic would tell us. So we should not expect friends of the anti-skeptical strategy to be friendly towards Closure Under Known Implication.

The upshot, then, is that Berto and Hawke seem to be caught in an uncomfortable halfway position with respect to the anti-skeptical reasons for denying various closure principles for knowledge. True, one can think of reasons to retain Closure Under Known Implication, just as one can think of reasons to reject Closure Under Strict Implication. But in the context of avoiding skepticism, the two of them typically stand or fall together.[3]

2.2. Closure Under Strict Implication

But let us set aside Closure under Known Implication and focus on Closure under Strict Implication. While Berto's semantics does not validate the latter principle in general, it validates a restricted version of it:

(**Closure Under Strict "On-Topic" Implication**) $\{X^p q, q \Rightarrow r\} \vDash X^p r$, whenever the topic of r is part of the topic of q.[4]

There are, however, reasons to worry about this principle.

As mentioned, Berto seems to think that topics are not always transparent to agents who have beliefs that involve these topics (cf. §2.2.4). Consider again:

(3a) Tully was a great orator.
(3b) Cicero was a great orator.

Not only do these sentences have the same extension in all possible worlds; they also concern the same topic, namely whether the person designated by "Tully" and "Cicero" was a great orator. Yet, as Frege famously pointed out, "Tully" and "Cicero" might carry very different cognitive associations for an agent. Indeed, one might be completely in the dark about the fact that "Tully" and "Cicero" refer to one and the same person. If so, one will not be in a position to tell that (3a) and (3b) are on the same topic.

Suppose we grant that topics can be opaque in this way. Consider then a case where q strictly entails r and where q and r are on the same topic, but where these facts are not transparent to an agent, relative to some total amount of information p that the agent possesses. According to Closure Under "On-Topic" Strict Implication, it now follows that $X^p q$ must hold whenever $X^p r$ does. In particular, when "X" is interpreted as "K," it follows that $K^p q$ must hold whenever $K^p r$ does. But this seems implausible. After all, the agent has no idea that q and r are on the same topic. For all the agent knows, q and r might be about completely different subject-matters. Given this, it is natural to think that the agent might be in no position to know r relative to p, even if they know q relative to p.

---

[3] We note here that Hawke (2016) offers a topic-sensitive account of unconditional knowability (as opposed to knowability relative to information) which invalidates both Closure Under Strict Implication and Closure Under Known Implication on similar anti-skeptical grounds.

[4] Here is a proof sketch: suppose $X^p q$ is true at w. Then q is true at each v such that $wR_p v$. Since q entails r, we know that r is true at all possible worlds at which q is true, which means that r must likewise be true at each v such that $wR_p v$. Under the assumption that r's topic is part of q's topic, the Bertonian semantics will thus tell us that $X^p r$ is true at w.

There are different ways one might try to hold on to Closure Under "On-Topic" Strict Implication in light of this worry. One option is to say that what an agent knows can be opaque in much the same way that topics can be opaque. On this view, someone who knows that "Cicero was a great orator" relative to information p will indeed know that "Tully was a great orator" relative to p. Yet this fact about the agent's knowledge need not be transparent to the agent himself, because he might not know that Cicero is Tully.

To judge the merits of this response, we would need to look at various other interpretations of "X" as well. In chapter 6, for example, Berto discusses a notion of conditional belief, where "$B^p q$" gets read as: "The agent believes q conditional on information p." In this context, we would have to say that anyone who believes that "Cicero was a great orator" conditional on information p must also believe that "Tully was a great orator" conditional on p. We find this difficult to accept, especially given that one of Berto's main motivations for rejecting Closure Under Strict Implication is that belief ascriptions do not in general seem to allow for substitution of necessarily co-extensional expressions *salva veritate*. Whether a response along these lines can be developed in a satisfying way remains to be seen.

Another option is to adopt a "subjective" or "perspectival" notion of topicality, which is sensitive to an agent's limited information about the world. On such a view, "Cicero was a great orator" and "Tully was a great orator" might be on the same topic for someone who knows that Cicero is Tully, while being on different topics for someone who does not have this knowledge. By relativizing topics to an agent's information, we can hold on to Closure Under "On-Topic" Strict Implication while respecting the intuition that one might know that "Cicero was a great orator" without knowing that "Tully was a great orator." In places, Berto seems to want to resist this way of thinking about topics (see, again, §2.2.4), but in light of the worries raised here, a subjective notion of topics might deserve further consideration.

3. Cognitive Limitations as a Source of Hyperintensionality

As Berto helpfully points out in the opening chapter of his book, hyperintensionality is closely connected to *failures of logical omniscience*. And, as he also points out, logical non-omniscience can have many different sources. Sometimes we might fail to believe a necessary truth because we lack a relevant piece of empirical information (say, that Cicero is Tully). Other times we might fail to believe a logical consequence of what we believe because we lack a relevant concept (say, the concept of a Martian hiding behind the International Space Station). And yet other times we might fall short of logical omniscience, because our cognitive abilities are limited. On the latter, Berto writes: "We sometimes know some necessary truths because they were easy to prove, whereas we don't know others because […] they involve sophisticated reasoning we are unable to carry out" (2022: 13).

As we have seen, Berto wants to capture different sources of logical non-omniscience—and the corresponding hyperintensional phenomena to which they give rise—through topicality constraints. But, as we shall argue now, there is a class of hyperintensional phenomena, ones grounded in cognitive limitations, which his semantic framework is ill-suited to handling. This has important—and, in our view, adverse—ramifications for what his framework allows us to say about the individuation of semantic content and the logic of belief.

3.1. The Individuation of Semantic Content

Chapter 2 (co-authored with Peter Hawke) defends the idea that the semantic content of a sentence—the proposition it expresses—depends on two factors: its truth-conditions, and its topic. This idea is expressed in the following thesis, which Berto and Hawke (2022: 30) name after Yablo:

> (**Yablo's Thesis**) The semantic content of p includes the semantic content of q iff (1) p entails q and (2) q's topic is included in p's topic.

Derivatively, on this account, two sentences have identical semantic contents if, and only if, they have identical truth-conditions and identical topics. Berto and Hawke do not go as far as to fully endorse Yablo's Thesis, since only the left-to-right direction is needed for many of their purposes. But if their framework is to provide us with a theory of semantic content, Yablo's Thesis (together with a theory of topic inclusion) is presumably it.

The worry we want to raise for such a theory of semantic content is the following. Remember that, on Berto's view, the standard connectives of propositional logic are topic-transparent: they do not add any topic of their own. Consider then an arbitrary proposition, p, and let E_p be a logical equivalent of p composed of just p and the standard logical connectives. So, for example, E_p might be $\sim\sim p$, $E_{p\&q}$ might be $\sim(p \rightarrow \sim q)$, $E_{p\rightarrow q}$ might be $\sim(p\&\sim q)$, and so on. By construction, E_p has the same truth-conditions as p. And since the connectives are topic-transparent, E_p also has the same topic as p. Thus, Yablo's Thesis will tell us that p and E_p have the same content. This is so regardless of the complexity of E_p. Indeed, even if it is far beyond the cognitive reach of any ordinary agent to establish that p and E_p are equivalent, Yablo's Thesis will tell us that p and E_p have the same content. But for someone who is entirely in the dark about the fact that p and E_p are equivalent, these sentences might have very different cognitive significance: one might feel very confident that p is true while being unsure about whether E_p is true; one might find it informative and even surprising to learn that p and E_p are equivalent; and so on. Insofar as a theory of semantic content should be sensitive to such differences in cognitive significance, Yablo's Thesis, conjoined with the idea that the logical connectives are topic-transparent, is in trouble.

Of course, it is a matter of debate whether a theory of semantic content *should* be sensitive to differences in cognitive significance. Those who deny this might respond that differences in cognitive significance, although genuine, do not track differences in semantic content, but are to be explained in a different way. We will not take a stance on this issue here, although we have elsewhere proposed a semantic framework specifically designed to capture a notion of semantic content which is sensitive to cognitive significance (Bjerring & Skipper 2020). Instead, let us be cautious and say that there seems to be a genuine phenomenon here, which it would be good to understand, regardless of whether it belongs to semantics proper. If Berto's framework is unable to capture this phenomenon, it is something to bear in mind.

3.2. The Logic of Belief

When it comes to the logic of belief, very similar concerns arise. Consider again the equivalent propositions p and E_p from above, and suppose that $B^q(p)$ is true at some world w. Then, by Berto's

semantics, $B^q(E_p)$ must also be true at w, because p and E_p have identical truth-conditions and identical topics. But this seems implausible for a cognitively limited agent who is completely in the dark about the fact that p and E_p are equivalent. As we have seen, Berto himself lists cognitive limitations as one of the main reasons why someone might fail to be logically omniscient.[5]

What might be said in response to this worry? One option, which Berto mentions in passing (see §2.5.1), is to take on board Stalnaker's (1987) metalinguistic approach to the problem of logical omniscience. On this view, the reason why agents appear to be logically non-omniscient is not that they are in fact logically non-omniscient, but rather that they are ignorant about what certain linguistic expressions mean. So, for example, if I believe each of the axioms of ZF set theory but apparently fail to believe Cantor's theorem, this is only because I am, at some level, ignorant about what Cantor's theorem actually expresses.

An initial worry one might have about this response is that one of Berto's main motivations for accommodating (rather than explaining away) various hyperintensional phenomena is that belief ascriptions do not in general seem to allow for substitution of necessarily co-extensional expressions *salva veritate*. But even if we set this worry aside, there is reason to doubt that the metalinguistic approach can ultimately help Berto in the way he imagines. Consider again the logically equivalent p and E_p, and let M_p be the metalinguistic proposition that "$p \leftrightarrow E_p$" expresses a truth. By appealing to the metalinguistic approach, Berto's idea is to say that although the agent in fact knows E_p, we can explain why it appears as if he does not by reference to the agent's lack of knowledge of M_p.

Granting this approach for now, let us ask what it would mean to fail to know M_p. Since M_p is composed of just p and the standard logical connectives, an agent would fail to know M_p in virtue of failing to know the metalinguistic propositions expressed by either p or one of the logical connectives. But, as far as we can tell, there is no good reason to think that agents should generally lack this metalinguistic knowledge. In fact, it seems plausible that any agent with basic training in logic or semantics will know the relevant metalinguistic facts about the connectives. Yet, due to cognitive limitations, such an agent might be unable to realize that p and E_p are equivalent.

Stalnaker himself was well aware that appealing to metalinguistic ignorance on its own is not enough to handle the problem of logical omniscience (for details, see Stalnaker 1987: ch. 5). We also need a notion of *fragmentation*. In general, we can think of a fragmented agent as someone who has multiple belief or knowledge systems that encode distinct bodies of information. These "fragments" are then differentially activated, depending on the circumstances in which the agent finds herself. To use one of Stalnaker's own examples: "Ralph's ability to find his way, unerringly, from home to work shows that he has the information that a certain route will take him from home to work and that he has access to that information to guide certain of his actions. But his inability to give directions to someone, or to draw a map of the route, shows that he lacks access to that same information for certain other purposes" (Stalnaker 2021: 190).

Suppose that Berto followed Stalnaker's lead and added fragmentation to the mix. The resulting account would then say, roughly, that an agent believes p on the supposition q relative to a fragment if, and only if, p is true at all R_q-accessible worlds relative to that fragment. This would allow for the

---

[5] To lay our own cards on the table, we have elsewhere proposed a solution to the problem of logical omniscience which is sensitive to cognitive resources (Bjerring & Skipper 2019).

possibility that an agent might fail to know all the relevant metalinguistic facts relative to a *single* fragment. Instead, this metalinguistic knowledge might be distributed across *different* fragments: the metalinguistic information about "~" might be contained in one fragment, the metalinguistic information about "&" might be contained in another fragment, and so on. Insofar as the agent, for whatever reason, fails to put these pieces of metalinguistic information together into a single fragment, the agent might then fail to know the proposition M_p, as desired.

However, the appeal to fragmentation comes with its own set of challenges, and even in conjunction with metalinguistic ignorance, it is unclear whether the Stalnakerian strategy can ultimately be made to work.[6] So if Berto wants to saddle up with both metalinguistic ignorance and fragmentation in the pursuit of logical non-omniscience, he owes us an explanation of the details.

4.   Conclusion

In closing, let us reiterate that despite the issues raised we find *Topics of Thought* to be an immensely valuable book full of original ideas and insights. We have learned a great deal from engaging with it, and we warmly recommend it to anyone who has interests in hyperintensionality, topicality, and the logic of propositional attitudes such as knowledge, belief, and imagination.

References

Berto, F. (2022): *Topics of Thought: The Logic of Knowledge, Belief, Imagination*. Oxford: Oxford University Press.

Berto, F. & Hawke, P. (2021): "Knowability Relative to Information." *Mind* 130(517): 1–33.

Bjerring, J. C., & Skipper, M. (2019). A dynamic solution to the problem of logical omniscience. *Journal of Philosophical Logic 48*: 501–21.

Bjerring, J. C., & Skipper, M. (2020): "Hyperintensional Semantics: A Fregean Approach." *Synthese* 197: 3535–58.

Bjerring, J. C., & Tang, W. H. (2023): "Fragmentation, metalinguistic ignorance, and logical omniscience". *Philosophical Studies*: 1–23.

Borgoni, C., Kindermann, D., & Onofri, A. (eds.) (2021): *The Fragmented Mind*. Oxford: Oxford University Press.

Cresswell, M. J. (1975): "Hyperintensional Logic." *Studia Logica*, 34(1): 25–38.

Dreske, F. (1970): "Epistemic operators". *The journal of philosophy*, *67*(24), 1007–23.

Fine, K. (2016): "Angellic 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.

Goodman, N. (1961): "About". *Mind*, 70: 1–24.

Hawke, P. (2018): "Theories of Aboutness." *Australasian Journal of Philosophy*, 96(4): 697–723.

---

[6] For a recent critique of the Stalnakerian strategy, see Bjerring & Tang (2023). For a recent collection of essays on fragmentation, see Borgoni et al. (2021).

Hawke, P. (2016): "Questions, Topics, and Restricted Closure." *Philosophical Studies* 173(10): 2759–84.

Hintikka, J. (1962): *Knowledge and Belief.* Ithaca, NY: Cornell University Press.

Hintikka, J. (1975): "Impossible Possible Worlds Vindicated". *Journal of Philosophical Logic*, 4(4): 475–484.

King, J. C., (1995): "Structured Propositions and Complex Predicates". *Noûs*, 29(4): 516–535.

Lewis, D. (1988): "Statements partly about observation." *Philosophical Papers*, 17: 1–31.

Nozick, R. (1981). *Philosophical explanations.* Cambridge: Harvard University Press.

Soames, S. (1987): "Direct Reference, Propositional Attitudes, and Semantic Content." *Philosophical Topics*, 15(1): 47–87.

Stalnaker, R. (1987): *Inquiry* (a Bradford Book). Cambridge, MA, MIT Press.

Stalnaker, R. (2021): "Fragmentation and Singular Propositions". In Borgoni et al. (2021): 183–98.

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