

Verbalism

Jeremy Goodman and Harvey Lederman*

December 25, 2017

Abstract

Superman is Clark, and Lois knows that Superman flies, but she doesn't know that Clark flies. This looks like a counterexample to Leibniz's law. According to *verbalism*, this appearance is misleading because the above sentence involves equivocation on "know". The proposition that Superman flies is the proposition that Clark flies, and Lois knows this proposition in some ways but not in others. The fact that rational people can be confused or unsure about which things are which corresponds to a distinctive kind of context-sensitivity in propositional attitude verbs. We show how to reconcile this ubiquitous contextualism with a strong logic of propositional attitudes, and also how it allows us to reconcile coarse-grained theories of propositions with failures of logical omniscience.

1 Verbalism

We begin with a story:¹

THELMA

Thelma, a German who doesn't speak English, is traveling in New York when Shorty steals her purse. She doesn't get a good look at him, but she sees him limping away. The next day the police round up some suspects and call Shorty in for a lineup. Knowing that Thelma saw him get away, Shorty wisely shows up early. Thelma is looking for someone who limps, but doesn't know how to ask the police to make the suspects walk around. So the lineup fails – she can't pick anyone out. Later, as Shorty celebrates his ill-gotten gains, Shorty's friend is telling the story in the bar. "You won't believe how smart this guy is. Since the lady saw him getting away, she knew Shorty limped, so Shorty got there early. His trick

*The authors contributed equally to this paper. Thanks to Cian Dorr, Ben Holguin, Mike Rieppel and Daniel Rothschild for comments on versions of this paper. Thanks also to audiences at CMU, NYU, Oxford, Pittsburgh, UCL, UCSD, and the Formal Epistemology Workshop as well as participants in seminars at NYU, Pittsburgh, Princeton and USC for comments and questions.

¹The story is adapted from [Schiffer \(1979\)](#); cf. [Dorr \(2014\)](#).

worked: she didn't see him walk in, so she didn't know Shorty limped, and he got off scot-free!"

Thelma is subject to what we will call *identity confusion*: roughly, she takes seriously the hypothesis that what is in fact one person is two people. In describing her, Shorty's friend uses the two clauses "she knew Shorty limped" and "she didn't know Shorty limped". These sentences are contradictory: the latter is the negation of the former. But both are true in his mouth, and both are used to describe the same state of affairs: Thelma's state of mind at the time of the lineup.

If a sentence and its negation are both true when used to describe the same state of affairs, then the sentence must be *context sensitive*, in the sense that it can be used to express different propositions on different occasions of use.² The above story illustrates a general phenomenon: sentences reporting propositional attitudes are context-sensitive in a dimension related to identity confusion. We take this observation as our starting point.³

Assuming that the semantics of attitude reports is compositional, and that their syntactic structure is invariant across the relevant occasions of use, the fact that they are context sensitive entails that some of their elementary constituents are context sensitive, in the sense that their semantic contribution is different on different occasions of use. So we can ask which elementary constituents of "she knew Shorty limped" are responsible for the aforementioned context sensitivity – and, more generally, which expressions' context-sensitivity is responsible for the context-sensitivity of attitude ascriptions related to identity confusion. Two answers suggest themselves:⁴

Clausalism: Something in the complement clause (e.g., "Shorty");

Verbalism: The attitude verb (e.g., "knew").

The following example helps to illustrate the distinction between clausalism and verbalism. Frege (1892) is associated with the view that there are two aspects of the conventional meaning of a name, a referent and a sense that determines that referent. When a name occurs in the complement clause of a propositional attitude ascription, its sense is what matters for determining the truth or falsity of the ascription. For Frege, "The Greeks believed that Hesperus rises in the evening" is true because the Greeks were belief-related to the Fregean thought obtained by composing the sense of "Hesperus" with the sense of "rises in the evening", and "The Greeks did not believe that Phosphorus rises in the evening" is true because they were not belief-related to the thought that results from composing the sense of "Phosphorus" with the sense

²This follows from the assumptions (i) that a sentence is true when used to describe a situation only if it expresses a proposition that is true in that situation and no proposition that is not true in that situation, (ii) that the negation of a sentence can be used to express only propositions that are negations of ones that can be expressed by the sentence of which it is the negation, and (iii) that a proposition and its negation are never both true in the same situation.

³Sosa (1970) is the *locus classicus* of this observation. Cf. Schiffer (1979) and Dorr (2014).

⁴This terminology is due to Dorr (2014). We'll consider a third view, *complementizerism*, below.

of “rises in the evening”. So sentences built in the same way from different co-referring names can nevertheless differ in truth-value.

Thelma is confused about Shorty in the same way that the ancient Greeks were confused about Venus. It is natural, then, for Fregeans to extend their analysis of sentences containing “Hesperus” and “Phosphorous” to the sentences “she knew Shorty limped” and “she didn’t know Shorty limped” from the above vignette. Since the latter sentences involve the same name of the object of Thelma’s confusion (“Shorty”), the Fregean should postulate that this name is context-sensitive: it is used with different senses on these two occasions.⁵ In “she knew Shorty limped”, the Fregean might say that the name “Shorty” expresses a Fregean sense similar to that expressed by the definite description “the limping purse-snatcher”, whereas in “she didn’t know Shorty limped”, the name expresses a Fregean sense similar to that expressed by the definite description “the shortest person in the lineup”. This Fregean picture is clausalist: it attributes the identity-confusion-related context-sensitivity of attitude reports to expressions in their complement clauses (in this case, to the proper name “Shorty”).⁶

Verbalists offer a different diagnosis. According to them, Shorty’s friend’s use of “knew” in “she knew Shorty limped” expresses a different psychological relation from the one expressed by his use of “know” in “she didn’t know Shorty limped”. Informally, there are different ways of knowing that Shorty limps. Thelma knows in one way that Shorty limps, but in another way she doesn’t know that Shorty limps. Shorty’s friend exploits this context-sensitivity of “know” to describe these different aspects of Thelma’s psychology.

Clausalism and verbalism are views about the context-sensitivity of attitude reports, but they are intimately related to questions in the philosophy of mind about the nature of beliefs, desires, intentions and so on. To locate the relevant context-sensitivity of attitude reports in their complement clauses is, roughly, to say that the contents of our thoughts are fine-grained in a certain distinctive way. By contrast, to locate this context-sensitivity in the attitude verbs allows one to avoid drawing such distinctions in mental content by instead proliferating ways of believing, desiring, etc. So while we will focus on clausalism and verbalism in the interest of tractability, much of the interest of these theses derives from the competing pictures of propositional attitude psychology to which they correspond.

This paper develops and defends a version of verbalism. Our attitude is exploratory. Verbalism is rarely discussed in the literature on identity confusion, and most extant discussions cast the view in an unfavorable light. We think the view is actually quite promising, and that its relative obscurity is undeserved. One of our main goals is to correct the mistaken impression that verbalism is incompatible with systematic theorizing about the propositional attitudes. To this end, we use a simple language-of-thought model

⁵They should say the same about Kripke’s (1979) famous Paderewski example.

⁶Although Frege’s theory of senses is the most familiar clausalism-friendly framework, it is not the only one. See [Caie et al. \(2017\)](#) for discussion of some clausalism-friendly alternatives.

of our cognitive psychology to show how verbalism can be reconciled with a strong logic of knowledge, belief, desire, intention, and other propositional attitudes. A perhaps surprising upshot of our discussion is that the distinctive predictions of verbalism are largely independent of whether propositions are individuated modally, or are structured in the manner of sentences, or have some intermediate level of granularity. In all of these ways our assessment of verbalism departs markedly from those of [Schiffer \(1979, 1992\)](#), [Crimmins and Perry \(1989\)](#), and [Crimmins \(1992\)](#), the main exponents of the view to date. We conclude by showing that verbalism, but not clausalism, reconciles apparent failures of logical omniscience with coarse-grained theories of propositions.

2 Leibniz's law

Our opening example THELMA involved a pair of contradictory attitude reports. Such examples are atypical in the literature on identity confusion. It is much more common to focus on conjunctions of attitude reports that can be turned into contradictions by substituting a proper name occurring in one complement clause for a coreferring proper name occurring in another complement clause. To take a familiar example, consider:

1. Superman is Clark, but although Lois knows that Superman flies, she doesn't know that Clark flies.

What should verbalists say about such sentences?

Notice first that the second and third conjuncts of **1** are context-sensitive in the same way that Shorty's friend's descriptions of Thelma are. Consider the second conjunct, "Lois knows that Superman flies". Suppose that before getting into costume Superman is chatting to Wonder Woman and Catwoman at a superhero conference. Lois walks up and asks him sarcastically what his superpower is. Amused, Wonder Woman whispers to Catwoman:

2. Lois doesn't know that Superman flies; she thinks he's with the press.

The first clause in **2** is the negation of the second conjunct of **1**. Yet both can be used to truly describe Lois's state of mind at the time. So "Lois knows Superman flies" must be context sensitive, just like "she knew Shorty limped". The same goes for the third conjunct of **1**, "she doesn't know that Clark flies". Suppose later at the conference two journalists who have learned about their colleague Clark's superhero persona see Lois fawning over him in costume. Both have been trying to impress Lois, and one laments to the other

3. Nothing we do can possibly compete: she knows that Clark *flies*!

The third conjunct of **1** is the negation of the second clause of **3**. Yet both can be used to truly describe Lois's state of mind at the time. So "she doesn't know that Clark flies" must be context sensitive too (and not merely owing to the context-sensitivity of "she").

Bearing this context-sensitivity in mind, we make the following natural hypothesis: the most salient reading of **1** is one on which the clause “Lois knows that Superman flies” expresses the same proposition as “she knows that Clark flies” does in the journalist’s utterance of **3** (or near enough) and the clause “she doesn’t know that Clark flies” expresses the same proposition as “Lois doesn’t know that Superman flies” does in Wonder Woman’s utterance of **2** (or near enough). Both clausalists and verbalists can accept this hypothesis. But they will explain it differently. Clausalists might say that, on the relevant reading of **1**, “Superman” expresses the same sense as the one expressed by “Clark” in the reporter’s mouth, while “Clark” expresses a different sense that is the one expressed by “Superman” in Wonder Woman’s mouth. By contrast, verbalists will say that the occurrence of “knows” in the second clause of **1** expresses the same psychological relation to the one expressed by “knows” in the reporter’s mouth, while the occurrence of “know” in the third clause of **1** expresses a different psychological relation which is the one expressed by “know” in Wonder Woman’s mouth. The respective senses of “Superman” and “Clark” (according to the clausalist) or relations expressed by “knows” and “know” (according to the verbalist) must be distinct in order to reconcile our hypothesis with the truth of **1**, on its most salient interpretation, and the truth of **2** and **3** as uttered respectively by Wonder Woman and the journalist. So just as the verbalist treatment of Shorty’s friend’s speech postulates a shift in the kind of knowing expressed by “know” and its cognates from one sentence to the next, a parallel verbalist treatment of **1** postulates a shift in the kind of knowing expressed by occurrences of “know” and its cognates within a single sentence.

The truth of **1**, at least on its most natural interpretation, might look like a counterexample to:

Leibniz’s Law: If $a = b$ and φ , then $\varphi[b/a]$.⁷

If verbalism is correct, this appearance is misleading. For to reject Leibniz’s law on this basis would then be like rejecting the law of non-contradiction on the basis of a true utterance of “Nour is sitting now and Nour is not sitting now” by someone talking as Nour stands up. The truth of this utterance is no counterexample to the law of non-contradiction because it involves interpreting the two occurrences of “now” differently. Similarly, according to verbalists, the truth of **1** on its most natural interpretation involves interpreting the two occurrences of (cognates of) “know” differently, and so is no counterexample to Leibniz’s law.⁸

A bit of terminology will be helpful. Say that a schema is *valid* just in case every English sentence that is an instance of the schema is valid; an

⁷An instance of this schema is obtained by replacing φ with a declarative English sentence, a and b with proper names, and $\varphi[b/a]$ with some sentence obtained from φ by replacing some occurrences of the name substituted for a with the name substituted for b .

⁸The last two paragraphs set us apart from Richard (1990), who thinks the context-sensitivity of attitude verbs accounts for the true readings in examples such as THELMA, which feature the same name, but does not invoke context-sensitivity to account for examples such as **1**, and rejects Leibniz’s law as a result.

English sentence is *valid* just in case every use of it is true on every uniform interpretation of its constituent expressions consistent with the conventions of English; an interpretation of some expression-occurrences is *uniform* just in case any two occurrences of the same expression are interpreted in the same way. We should stress that our use of the word “valid” to denote this status is stipulative: it is not offered as an analysis of some antecedently understood notion of validity. We are interested in this status because it is helpful in comparing clausalism and verbalism. A distinctive feature of verbalism is that it allows us to reconcile the true reading of 1 with the validity (in our sense) of Leibniz’s law. Indeed, in what follows we will treat the validity of Leibniz’s law as a further commitment of verbalism. By contrast, clausalism is incompatible with the validity of Leibniz’s law. For the sentence

4. If Superman is Clark and Lois knows that Superman flies, then Lois knows that Clark flies.

is false on its most natural interpretation. If Leibniz’s law were valid, then this interpretation would have to be non-uniform. And verbalists should think that it *is* non-uniform, since they should posit the same mid-sentence variation in the interpretation of “knows” in 4 that they are already committed to in the case of 1. But clausalists should *deny* that the relevant interpretation of 4 is non-uniform: they think the relevant difference in interpretation concerns “Superman” and “Clark”, and these are different words, so a difference in their interpretation is not a failure to interpret 4 uniformly.

Some Fregeans might deny that the relevant interpretation of 4 is uniform, on the grounds that two occurrences of a name are interpreted in the same way only if they have the same referent (*Bedeutung*), and that while the first occurrence of “Superman” refers to Superman, the second occurrence of “Superman” refers not to Superman but to a Fregean sense of Superman. Lacking *any* uniform interpretations, 4 would then be trivially valid (in our sense). Notice that this proposal does not ‘save’ Leibniz’s law in any interesting way, since it also classifies as valid (in our sense) such profundities as:

5. Superman is an egg and everyone knows that Superman is a pig.

Moreover, there is clearly an important notion of a sentence being true on any *non-equivocating* interpretation – a notion that does not apply to 5 but does apply to

6. If Lois knows that Superman flies, then Superman flies.

In a Fregean setting, this kind of non-equivocation is not captured by sameness of *Bedeutung*. Instead, Fregeans should say that a sentence φ (containing no iterated attitude verbs) is interpreted uniformly just in case, for every expression e , there is a sense s such that all occurrences of e in φ in the scope of an attitude verb have s as their *Bedeutung* and all other occurrences of e in

φ have the referent determined by s as their *Bedeutung*.⁹ This Frege-friendly way of thinking about uniform interpretation does not prevent 4, 5, and 6 from having uniform interpretations. The most salient such interpretations will be ones on which 6 is true, 5 is false, and (crucially) 4 is also false.¹⁰ Contrary to popular belief, the Fregean distinction between sense and reference does not reconcile Leibniz’s law with ‘opacity’ in any interesting sense.¹¹

Of course, there are other notions in the vicinity of our notion of validity that one might be interested in. Here is one salient example.¹² It is common to theorize about context-sensitivity by postulating certain things called Contexts, which we might model as sequences consisting of a person (the speaker of the context), a time (the time of the context), a set of possibilities (the ‘modal base’ of the context), a function from variables to entities of the appropriate kind (the ‘variable assignment’ of the context), etc. For every use of a syntactically simple expression, there is a unique Context that it is ‘in’. The conventional meaning of such an expression is a function on Contexts, such that the interpretation of an occurrence of a syntactically simple expression is the result of applying its conventional meaning to the Context the occurrence is in. The proposition expressed by an utterance of a sentence is then determined compositionally from the interpretations of its elementary constituents. Say that a sentence is *valid*^C just in case an occurrence of the sentence is true whenever all of its elementary constituents are in the same context.

Provided the notion of validity^C is in good order, it is a weaker notion than validity in our sense. This is because distinct elementary expressions might

⁹It is a vexed question how Fregeans should think about iterated attitude ascriptions. Some Fregeans think that interpreting such ascriptions requires invoking a hierarchy of senses; for example, that the *Bedeutung* of an occurrence of an expression in the scope of two attitude verbs is the customary sense of its customary sense. Other Fregeans think that such an expression’s *Bedeutung* is simply its customary sense. Fregeans of the first sort should generalize the notion of non-equivocation to sentences involving iterated attitudes as follows: for every expression e , there is a sequence of entities x_0, x_1, \dots such that (i) x_{n+1} is a sense of x_n , for all n , and (ii) the *Bedeutung* of an occurrence of e in the scope of exactly n attitude verbs is x_n . Fregeans of the second sort should think that the notion of non-equivocation in the main text is the right notion of non-equivocation for sentences involving iterated attitudes too.

¹⁰Although the Fregean could reconcile the existence of uniform interpretations of 4 with the identification of sameness of interpretation with sameness of reference (*Bedeutung*) by postulating the existence of bizarre interpretations on which either all occurrences of “Superman” and “Clark” refer to the superhero or both occurrences of “Superman” refer to a sense, and likewise for “Clark”, doing so still fails to carve out an interesting status for Leibniz’s law, since the status it carves out not only fails to apply to 6, but *does* apply to the absurd “If Lois knows that Superman flies, then Superman does *not* fly” (since senses don’t fly).

¹¹Call a sentence *Frege favorable* if any two occurrences of the same expression in the sentence are embedded under the same number of attitude verbs; call a schema *valid*^F if every Frege favorable instance of it is valid. Fregeans have no reason to deny that Leibniz’s law is valid^F. But again, they shouldn’t delude themselves into thinking that Leibniz’s law thereby has the status we usually care about in writing down schemas when doing logic, since clearly “If A knows that φ , then $\neg\varphi$ ” lacks that status despite being valid^F. To illustrate the point another way: When we write down “ $\varphi \rightarrow \varphi$ ” in axiomatizing the propositional calculus, we clearly are not leaving it open that “(Superman sings or I know Superman flies) \rightarrow (Superman sings or I know Superman flies)” might have a false reading that involves neither equivocation nor a deviant interpretation of “ \rightarrow ”; but the mere validity^F of that schema is compatible with the existence of such a reading.

¹²We will discuss different weaker notions of validity in sections 4 and 7.

depend for their interpretation on the same parameters of the Context they occur in. (For example, “If he is tall, then he is not short” could turn out to be valid^C despite not being valid in our sense, if “tall” and “short” correspond to elementary syntactic constituents of the sentence and both expressions depend for their interpretation on the value of some ‘height-threshold’ parameter of Contexts.) This observation raises the question of whether clausalism is compatible with the validity^C of Leibniz’s law.

The answer depends on how clausalism is implemented in the framework of Contexts. One option would be to have every Context supply a function from expressions of the language to Fregean senses, on the model of an assignment function. 4 will then be a counterexample to the validity^C of Leibniz’s law for the same reason that it is a counterexample to the validity of Leibniz’s law, since “Superman” and “Clark” can co-refer in a Context that supplies a sense-function that maps them to different senses in a way that makes the sentence false. A different option would be to have Contexts supply a function that maps any *thing* (or at least any thing that could be the referent of a syntactically simple expression) to a Fregean sense of that thing – the sense of an occurrence of a name would then be determined by applying the function supplied by the Context the name occurs in to the referent of the name in that Context. On this way of doing things, 4 will *not* be a counterexample to the validity^C of Leibniz’s law. For according to this clausalist, any false interpretation of 4 must be one on which the initial occurrences of “Superman” and “Clark” co-refer while their later occurrences have different senses, and so cannot both occur in the same Context as their earlier occurrences did (since then they would co-refer, and so given the present proposal would have the same sense). The false interpretation of 4 would have to involve a mid-sentence Context-shift.¹³

We will operate with our notion of validity instead of validity^C in part because it allows us to sidestep the question of how best to implement clausalism in the framework of Contexts. That question is a distraction from the issues we are concerned with in this paper.¹⁴ It is the validity, not the validity^C, of Leibniz’s law that represents a distinctive commitment of verbalism in comparison to clausalism.

The validity of Leibniz’s law also provides a helpful point of comparison with other views in the literature. The most prominent defenders of Leibniz’s law in the face of sentences like 1 are not verbalists but so-called ‘neo-Russellians’ (e.g., Salmon (1986), Soames (1987), Braun (1998), Saul (2010)).

¹³Aloni (2005) gives the former treatment of proper names but the latter treatment of free variables, thereby making Leibniz’s law invalid^C but its universal closure valid^C. This sort of failure of universal instantiation is also defended by Bacon and Russell (2017) and is criticized by Caie et al. (2017).

¹⁴It also strikes us as a bad question to focus on given the extent to which “Context” is a term of art. A more interesting question for clausalists is whether, in whatever sense it is a linguistic convention that related context-sensitive words like “tall” and “short” receive coordinated interpretations, there is likewise a convention that co-referring names receive the same interpretation. Orthodox Fregeans will deny that there is any such convention, but see Dorr (2014) for a version of clausalism that is heterodox in this respect.

Unlike verbalists, they think that, on all relevant interpretations, the sentence **1** is simply *false* while **4** is simply true. Verbalists avoid this implausible result by embracing contextualism. Note: although someone could use the validity of Leibniz's law and the truth of **1** on its most natural interpretation as premises in an argument for verbalism, that is *not* what we are doing. On the contrary, we think pretheoretical judgments about the validity (in our sense) of Leibniz's law are of little probative value given the stipulative way we introduced the operative notion of validity. And we don't think any very theoretical argument is needed to establish the context-sensitivity of attitude reports in a dimension related to identity confusion; pretheoretical judgments about the truth of utterances like Shorty's friend's suffice, together with the basic assumptions about truth and falsity from footnote 2.

The views in the literature closest to verbalism are those defended by [Crimmins and Perry \(1989\)](#) and considered (though ultimately rejected) by [Schiffer \(1979, 1992\)](#) under the heading of the 'hidden-indexical' theory. This substantive similarity is somewhat obscured by presentational differences related to Leibniz's law. Schiffer does not mention the compatibility of the hidden-indexical theory and the validity of Leibniz's law, despite it being one of the most distinctive features of the view. More strikingly, Crimmins and Perry describe (and indeed *motivate*) their view as a way of rejecting Leibniz's law and endorsing what they call "opacity" ([Crimmins and Perry \(1989\)](#), e.g. 686-7)). However, the sense in which their view involves 'rejecting' Leibniz's law is one in which it also involves rejecting the law of non-contradiction – it is not a sensible way of thinking about what it means to accept a schema in a language with context-sensitive vocabulary.¹⁵

3 Non-uniformity

Consider the sentences:

7. Lois believes that Superman flies if and only if she believes that Clark flies.
8. Lois believes that either Superman flies or Superman does not fly if and only if she believes that either Clark flies or Superman does not fly.
9. Lois believes that Superman flies and Clark doesn't if and only if she believes that Superman flies and Superman doesn't.
10. Lois believes that Superman is Superman if and only if she believes that Superman is Clark.
11. Lois believes that Superman isn't Clark if and only if she believes that Superman isn't Superman.

¹⁵See [Predelli \(2000\)](#).

Since verbalists are committed to the validity of Leibniz’s law, they are committed to the truth of all of these sentences on any uniform interpretation on which all occurrences of “Superman” and “Clark” co-refer. Unlike neo-Russellians, verbalists accept that these sentences are often false on their most salient interpretations. According to verbalists, this is because such sentences often invite interpretations on which “believe” is interpreted differently in its two occurrences.

A common initial objection to verbalism is that such context-sensitivity is implausible on its face. This objection leaves us cold. We don’t have plausibility judgments one way or another about the existence of such context-sensitivity, and the issue is a sufficiently theoretical one that we are skeptical about the probative value of any plausibility judgments one might have concerning it. Moreover, the objection can’t be motivated by a general aversion to context-sensitivity, since we are taking for granted a contextualist treatment of Shorty’s friend’s discourse about Thelma. Whatever the worry is, then, it must be something specific to *mid-sentence* differences in the interpretation of context-sensitive vocabulary.

The worry might be sharpened as follows: the framework of Contexts (discussed in the previous section) is the right one for theorizing about context-sensitivity, and all occurrences of expressions in a sentence are in the same Context. We ourselves have some reservations about the framework of Contexts, but we can set them aside here. For it should be uncontroversial (within the that framework) that different occurrences of the same expression in a sentence can be in different Contexts, if for no other reason than to handle temporally extended utterances containing multiple occurrences of “now”, as discussed in the last section.¹⁶

A different way of sharpening the worry appeals to the principle that interpretations of a sentence that involve mid-sentence context-shifts are difficult to access when the sentence is not pronounced in some special way. For example, it is hard to access non-uniform interpretations of “now” without either an exaggerated pause or focal stress on the second “now”; similarly, non-uniform interpretations of pronouns can be difficult to access without accompanying pointing gestures or related cues. Since nothing like this is needed to access the relevant true interpretation of 1, the objector claims that that interpretation must be uniform.

This objection is unconvincing. First, it is implausible that the principles govern the non-uniform interpretation of expressions in a single sentence are fundamentally different from the principles governing non-uniform interpre-

¹⁶Such cases are sometimes thought to be a problem for theorizing about context-sensitivity in terms of Contexts; see [Glick \(2017\)](#) for discussion. But this worry is confused, since nothing about that framework requires that all occurrences of expressions in the same sentence be in the same Context. The fact that, e.g., [Kaplan \(1989, §§XVIII-XIX\)](#) doesn’t allow such flexibility in the model theoretic semantics for his formal language is easily explained by his interest in characterizing the *logic* of a context-sensitive fragment of English, which he does in terms of sentences’ validity^C. As we saw earlier in connection to the law of non-contradiction, the interpretations of a sentence relevant to assessing its validity need not include all interpretations of the sentence that are possible given the conventions of English.

tations across short stretches of discourse. Since cues like focus, pointing, and the like were not needed for Shorty’s friend’s audience to access the non-uniform interpretation of his short discourse about Thelma, the verbalist has a principled reason to doubt that any such cues should be needed to access the true interpretation of **1**. Second, contrary to what the objector claims, there are in fact special cues present in **1**. The use of different names of the same person in the respective complement clauses is a salient hint to the reader that identity confusion is at issue and hence, according to the verbalist, to be on the lookout for non-uniform interpretations of attitude verbs. Third, we can adapt examples from the literature used to argue that mid-sentence Context-shifts often occur in the absence of overt cues to argue that non-uniform interpretations of sentences can be preferred in the absence of overt cues. We have in mind cases where two quantifier expressions in the same sentence are tacitly restricted in different ways. Soames (1986, p. 357) gives the example of “Everyone is sleeping and monitored by a lab assistant” – on its most natural interpretation, the quantifier “everyone” receives a different restricted interpretation than “a” does, so that the truth of the sentence doesn’t require that all lab assistants be asleep. This isn’t an example of non-uniform interpretation in our sense, since the quantifiers “everyone” and “a” are different words. But we can get such an example by considering “Everyone is monitoring everyone”, which in the absence of focal stress or the like can, with the relevant setup, invite an interpretation according to which every lab assistant is keeping tabs on each patient. Fourth, everyone should think that identity confusion begets within-sentence non-uniform interpretation. Consider

12. Thelma knows that Shorty (the purse-snatcher) limps, but doesn’t know that Shorty (the guy in the lineup) limps.

The parenthetical descriptions function the way that verbalists think the use of distinct names does in **1** – they clue the reader in to how to non-uniformly resolve the context-sensitivity of expressions in the sentence (either “know”, if verbalism is true, or “Shorty”, if clausalism is true).^{17,18}

Granting these precedents for non-uniform interpretation, the clausalist might still claim that it is an advantage of their view that it doesn’t require

¹⁷One might question whether **12** indeed involves non-uniform interpretation by proposing that $\lceil S \text{ knows that } A \text{ (the } F \text{) } Gs \rceil$ means the same as $\lceil S \text{ knows that } A \text{ is the } F \text{ and } Gs \rceil$. But this purported equivalence is mistaken. Suppose there is someone other than Superman who flies, but who isn’t a reporter, and suppose that Luthor doesn’t know anything about this person and indeed doesn’t know that there are any reporters. The situation is naturally described by “Luthor knows that Superman (the reporter) flies”, but there seems to be no corresponding true reading of “Luthor knows that Superman is the reporter and flies”.

¹⁸The example also shows that, in cases where focal stress helps us access the intended interpretation of a context-sensitive sentence, the location of focal stress is not a foolproof guide to which expressions are the source of that context-sensitivity. This is because accessing the relevant interpretation of **12** is facilitated by focusing “purse-snatcher” and “guy in the lineup”, but neither verbalists nor clausalists should think that accessing the intended interpretation of the sentence is a matter of resolving the context-sensitivity in these expressions. By the same token, the fact that accessing the intended interpretation of **1** is facilitated by focusing “Superman” and “Clark” but not by focusing “know” is not much of a consideration against verbalism.

the phenomenon to be as widespread as the verbalist does, since it doesn't require it to be present in the true reading of 1. But this claim is unpersuasive for a number of reasons. First, it is unclear why the relative prevalence of non-uniform interpretation is something we should care about, once the existence of the phenomenon is granted. Second, in most real cases of identity confusion (if not in philosophers' hypothetical examples) the object of confusion does not have two names saliently associated with the way that the relevant people are confused about that object. So the case at hand is in fact exceptional in allowing clausalists to convey the contours of Lois's confusion without non-uniform interpretation by using different names of the same person. Third, there are also ways in which verbalists but not clausalists can convey the contours of confusion without non-uniform interpretation, by using different attitude verbs to describe the confused person's state of mind. For example, just as the distinctness of "Superman" and "Clark" allows clausalists, but not verbalists, to claim that the salient true interpretation of 1 is uniform, so too the distinctness of the verbs "remembers" and "know" allows verbalists, but not clausalists, to claim that the salient true interpretation of

13. Thelma still remembers that he limps, but she doesn't now know that he limps.

is uniform.

A more theoretical objection to the sort of non-uniform interpretations posited by verbalists invokes the principle:

Ellipsis Principle: In cases of ellipsis involving context-sensitive expressions, the deleted occurrence of the expression receives the same interpretation as its most recent occurrence.

The putative problem concerns variants of 1 where the second occurrence of the attitude verb is deleted:

14. Lois knows that Superman flies but not that Clark flies, even though Superman is Clark.

Like 1, this sentence seems to have a true interpretation. Given verbalism, the existence of such an interpretation requires non-uniformly interpreting "knows". But since the second occurrence of "knows" in 14 is deleted, the Ellipsis Principle rules out such an interpretation of that sentence. So verbalists must either deny that 14 has the true interpretation that it seems to have, or reject the Ellipsis Principle.

Denying that 14 has a true interpretation recalls the non-contextualist neo-Russellian views mentioned earlier, since both imply that ordinary speakers' judgments of attitude ascriptions' truth and falsity are in error. However, it convicts ordinary speakers of error in a much narrower range of cases, and the errors it does predict can arguably be explained by a hypothesized tendency to equate sentences like 14 with (the most natural interpretation of) their ellipsis-free counterparts. So pairing this circumscribed attribution of error with the

verbalist's rejection of neo-Russellianism is not an unprincipled combination of commitments.¹⁹

Still, we think verbalists would do better to reject the Ellipsis Principle. Apparent counterexamples abound. Consider

15. In Japan I'm tall, but in Norway I'm not.

This sentence invites an interpretation where the deleted occurrence of "tall" expresses a more demanding height requirement than the voiced occurrence. One could resist this conclusion by holding that the property expressed by "tall" here is not the property of exceeding any particular height, but rather the property of being tall for where you are. Following [Dorr \(2014, p. 63-4\)](#), we think such proposals cannot be generalized to cover all relevantly similar examples in a plausible way. Consider

16. Pliny the Elder was so-called for genealogical reasons, and so was Pliny the Younger.

The voiced and deleted occurrences of "so-called" seem to express respectively the properties of being called "the Elder" and of being called "the Younger". The analogue of the rejoinder about "tall" in 15 would be to claim that both occurrences of "so-called" express the property of being called the (salient) name that you are known by. But if this were an available interpretation of "so-called", then we should be able to access a corresponding true interpretation of

17. # Everyone is so-called because of how people use words.

But this sentence doesn't seem to have a true interpretation. Also, notice that the proposal cannot handle cases involving two names of the same person, such as

18. ? Superman is so-called because he has superpowers, and Clark because his parents named him that.

This might seem like a good prediction, since it would be infelicitous for us to utter this sentence. However, this infelicity seems to be because we know that Superman is Clark. In the mouth of a citizen of Metropolis who is subject to the relevant identity confusion, 18 can be used felicitously and truly, and this cannot be explained by the proposal that both occurrences of "so-called" express the property of being called the (salient) name you are known by.²⁰

¹⁹A worry: one might think that, if this hypothesized tendency is robust in the manner of linguistic conventions, then such sentences would thereby literally mean what the proposal says they are mistakenly thought to mean. We needn't take a stand on this issue here.

²⁰It is puzzling why 18, unlike 1, should be infelicitous in the mouths of people not subject to the relevant identity confusion. Note that, whatever is going on, it doesn't depend on the ellipsis, since the same pattern is exhibited in "Superman is so-called because of his powers, and Clark is so-called because his parents named him that".

So the precedent of “so-called” provides a strong basis for verbalists to reject the Ellipsis Principle in the case of 14.²¹

Yet another option would be to reconcile the Ellipsis Principle with a true reading of 14 by rejecting both verbalism and clausalism in favor of a third view we call *complementizerism*. According to this view, the complementizer “that” is the source of the identity-confusion-related context-sensitivity of sentences like “Lois knows that Superman flies”, the complementizer “whether” is the source of such context-sensitivity in “Thelma is unsure whether Shorty limps”, etc. The key observation is that 14 does not have a true reading if the second occurrence of “that” is deleted. The Ellipsis Principle could therefore be used to argue for complementizerism over verbalism. While we think it is an interesting view, we are going to set complementizerism aside for the remainder of this paper because it is more-or-less indistinguishable from verbalism as far as the issues we will be focussing on are concerned.²²

4 Attitudinal validity

Consider the principles:

19. If A knows that φ , then A believes that φ .
20. If A fears that φ , then A hopes that it is not the case that φ .
21. If A doubts that φ , then A thinks that it is not the case that φ .
22. If A intends to φ , then A expects to φ .
23. If A wants above all else to φ , then A will try to φ .

These are the sorts of schemas philosophers of mind, language, and action write down in formulating their theories. Some of these theories are mistaken, but the practice of writing down such principles doesn’t seem to be inherently misguided. In particular, the practice is presumably not undermined by the kind of context-sensitivity of attitude ascriptions we have been investigating.

²¹In rejecting the Ellipsis Principle, we aren’t denying that some restricted version of the principle is true. Consider the sentence “Laura showed a drawing to every teacher, but Lena didn’t” from Heim and Kratzer (1998, p. 250). This sentence cannot receive a true interpretation if Laura and Lena both showed pictures to the same teachers, and the reason why seems to be that the deleted occurrence of the quantifier “every” must be restricted in the same way as the voiced occurrence.

²²One might think that complementizerism faces a parallel problem with the Ellipsis Principle from sentences like “Lois believes that Superman flies but doesn’t believe Clark flies”. We think complementizerists should respond by distinguishing two sorts of ellipsis. In the case of omitting “that” after “believe”, the deleted complementizer is not in any sense anaphoric on some previous occurrence thereof, since such deletion is acceptable discourse initially. By contrast, most cases of deletion are anaphoric, in that their felicity depends on an earlier occurrence of the deleted expression. The complementizerist can then plausibly reply that the Ellipsis Principle is only well-motivated in the case of anaphoric ellipsis.

But if attitude verbs are context-sensitive in the way that verbalists claim, then these principles are not valid – i.e., they are not true on every uniform interpretation. The reason is that, for example, “know” and “believe” are different words, and uniform interpretation only requires that each word be interpreted in the same way in all of its occurrences. So the natural true interpretation of “Lois knows that Superman flies, but doesn’t believe that Clark flies” counts as uniform. Since every instance of Leibniz’s law is true on every uniform interpretation, it follows that:

24. Lois knows that Superman flies but doesn’t believe that Superman flies.

is true on this uniform interpretation. So we have a counterexample to the validity of 19.

Moral: if verbalism is true, then in putting forward principles like 19-23 philosophers must not be committing themselves to their being valid in our sense, since that would make such theorizing inherently misguided, which it is not. So verbalists need some other way of understanding what they are doing. Here is one idea. Just as there is some interesting sense in which occurrences of the context-sensitive modal adverbs “necessarily” and “possibly” can be interpreted in the same way, the verbalist should recognize a sense in which occurrences of different attitude verbs can likewise be interpreted in the same way. Just as “Lois knows that Superman flies but doesn’t know that Clark flies” is true only on non-uniform interpretations, the verbalist should say that “Lois knows that Superman flies but doesn’t believe that Clark flies” is true only when “know” and “believe” fail to receive appropriately coordinated interpretations.

Let an *attitudinally uniform* interpretation be one that interprets all propositional attitude verbs ‘in the same way’, in the sense just gestured at. Call a sentence *attitudinally valid* just in case it is true on every attitudinally uniform interpretation. Since not every uniform interpretation is attitudinally uniform, being attitudinally valid does not imply being valid. For example, while verbalists must deny that 19 is valid, they have no reason to deny that it is attitudinally valid.

We can make the notion of propositional attitude verbs being interpreted ‘in the same way’ more precise. We’ve already encountered the idea that one can know a proposition in one way but fail to know it in another. The same goes for believing, intending, and other propositional attitudes. In general, the idea (which is not original to us by any means) is that underlying the two-place relation expressed by a propositional attitude verb X in context there is a three-place relation α_X between a person, a proposition, and a third entity, which we’ll call a *perspective*. For example, the meaning of ‘ A knows that φ ’ on a given interpretation is $K(x, p, \pi)$, where $K = \alpha_{\text{“know”}}$ (the underlying three-place relation), x is the referent of A on that interpretation, p is the proposition expressed by φ on that interpretation, and π is a perspective. For any perspective π and interpretation i of an occurrence o of an attitude verb X , say that π parametrizes o on i just in case $i(o)$ is the relation of α_X -ing relative to π (i.e., $i(o) = \lambda xp.\alpha_X(x, p, \pi)$). An interpretation i is

then attitudinally uniform only if, for some π , π parametrizes o on i for all propositional-attitude-verb-occurrences o on which i is defined.²³

(We say “only if” rather than “if and only if” because the notion of attitudinal uniformity applies not only to propositional attitude verbs, but also to adjectives (“known”, “certain”, “likely”, “probable”, ...), adverbs (“must”, “might”, “certainly”, “probably”, ...), and attitude verbs taking interrogative complements (“is unsure”, “wonders”, “knows [who/where/whether...]”, ...) that are clearly context-sensitive in the same identity-confusion-related way.^{24,25} Verbalists should hold that the interpretation of such expressions

²³In the framework of Contexts, we might say that the Context supplies a value of this perspective parameter. Unlike that framework, however, the present one does not assume that relative to an interpretation every attitude-verb-occurrence can be associated with a unique perspective, since it does not assume that parametrization is a functional relation. Having said that, everything we say in what follows is compatible with the functionality of parametrization: i.e., that for every interpreted attitude-verb occurrence, we can speak of *the* corresponding perspective. The reader who prefers to think in terms of Contexts may thus understand the following discussion as concerned with the validity^C of these schemas.

²⁴Perhaps this category of expressions also includes psychological verbs that take nominal direct objects (“fears”, “likes”, “looks for”, “knows about”, etc.) and even nouns (e.g. “suspect”). However, whatever the status of verbs like “fears”, verbs like “knows” seem not to be context-sensitive in the relevant way when they take nominal complements, as in “Lois knows Superman” – this fact poses a challenge for any account of acquaintance knowledge in terms of propositional knowledge. The question of which expressions exhibit such context sensitivity is the analogue, for verbalists, of the question, for clausalists, of which expressions generate identity-confusion-related counterexamples to the validity of Leibniz’s law. Note too that complementizerism has trouble accounting for apparent identity-confusion-related context-sensitivity in sentences with no psychological expressions taking complement clauses as arguments.

²⁵A subtle question for verbalists is how to treat speech-reporting verbs like “says that”. On the one hand, there is some pull to thinking that the following schemas are attitudinally valid:

- (i) If A says that φ and A is not insincere, then A believes that φ .
- (ii) If A says that φ in making an utterance u , then the proposition that φ is the unique proposition expressed by u .

But now we face a puzzle. Suppose Lois not insincerely utters “Superman flies”. As we have previously argued, there is a true interpretation of “Lois does not believe that Superman flies”. Assuming that, on this interpretation, “Lois’s utterance expresses the proposition that Superman flies” is still true (which is naturally denied by clausalists, but not by verbalists), we must either reject the attitudinal validity of one of (i) and (ii) or accept the surprising conclusion that there is a true interpretation of “There is nothing that Lois said when she uttered ‘Superman flies’.”

One option for verbalists is to accept this surprising conclusion, which may be easier to swallow once we notice (for the reasons discussed in section 7) that it is compatible with “If A understands the proposition that φ and knows that σ means that φ , then A is insincere if they utter σ without believing that φ ” being a good principle of propositional attitude psychology.

A second option is to deny that (i) is attitudinally valid. The idea would be to assimilate Lois’s situation to the following case. Suppose a French speaker Jacques memorizes and utters the English sentence “Frenchmen are gullible” without knowing what it means or believing the proposition it expresses. It is difficult to access a true reading of “Grégory didn’t say that Frenchmen are gullible”. But Grégory wasn’t being insincere. So we seem to have a counterexample to the attitudinal validity of (i). Verbalists might say that, as in the case of Grégory, “Lois doesn’t know what she is saying” is true on the relevant interpretation, although, unlike in the case of Grégory, this sentence considered on its own has a more natural interpretation on which it is false, corresponding to a different interpretation of “know”.

A third option is to deny that (ii) is attitudinally valid. The idea would be to assimilate Lois’s situation to the following case. Suppose Ramy sincerely utters “The Nile floods frequently and flows north”. It is difficult to access a true reading of “Ramy didn’t say that the Nile flows

is similarly parametric on perspectives, and that an interpretation is attitudinally uniform *if and only if* some perspective parametrizes all interpreted occurrences of such expressions.)

In the next section we will present a toy model of perspectives that is compatible with 19-23 being attitudinal valid. This is not to say that verbalists should see attitudinal validity as the aim of ordinary theorizing about the propositional attitudes. Indeed, in section 6 we will argue that attitudinal validity is still too demanding a notion for that purpose. Nevertheless, attitudinal validity is the natural starting point for verbalists to try to make sense of what we are doing when we put forward systematic principles about how different propositional attitudes are related.

5 Mentalese

What are perspectives? And what does it take for an agent to believe/know/etc. a proposition from a perspective? (I.e., what is the mapping α from attitude verbs to underlying three-place relations?) The answer depends on the nature of people's underlying cognitive psychology. For concreteness and tractability, we will operate under the idealizing assumption that people have a 'language of thought' – call it *mentalese*. They stand in various relations to sentences of this language, and it is by standing in relations of this kind that they count as having the propositional attitudes that they do. More precisely, for every attitude verb X , there is a relation β_X between people and mentalese sentences which we might pronounce 'being in one's X -box'. In the interest of brevity, say that one *accepts* a mentalese sentence just in case it is in one's "believe"-box. Verbalism is not hostage to this particular mentalese framework, but it is a helpful toy model for testing the theory's tenability and developing a more predictive theory of identity-confusion-related context-sensitivity. More sharply: if verbalism cannot be made to work using the rich structure allowed by this toy model, it is unlikely to work given more realistic assumptions. So even if the mentalese framework is unrealistic, it is a good starting point for testing the workability of verbalism.

Sometimes philosophers and cognitive psychologists (most prominently, Jerry Fodor²⁶) seem to think that once one adopts this mentalese picture, one doesn't need to say anything more about identity confusion: we can simply observe that people sometimes have two mentalese names for the same thing. These theorists are mistaken. To describe the mental representations of people subject to identity confusion is not yet to give an account of the corresponding

north" or "Ramy didn't say that a river floods frequently and flows north". But since neither the proposition that the Nile flows north nor the proposition that a river floods frequently and flows north is identical to the proposition that the Nile floods frequently and flows north, we seem to have a counterexample to (ii). Verbalists might say that, just as Ramy said that a river floods frequently and flows north, "Lois said that someone flies, and hence said something", is true on the relevant interpretation, reconciling the purported attitudinal validity of (i) with there being no true interpretation of "There is nothing that Lois said when she uttered 'Superman flies'."

²⁶See, e.g. Fodor (2008, Ch. 3).

context-sensitivity of attitude ascriptions (or a theory of what people believe, want, etc.) any more than to describe them neurophysiologically is. Such descriptions may succeed in conveying what the person’s mental life is like, but as Kripke (1979, p. 259) emphasized, it is not hard to describe confused people’s states of mind in some idiom or other – vignettes like THELMA do just that. The framework of mentalese is at best a first step in theorizing about the identity-confusion-related context-sensitivity of attitude ascriptions.

Indeed, some philosophers who think that mentalese is the right framework for theorizing about the underlying psychology of people subject to identity confusion nevertheless object to our way of using that framework to theorize about the semantics of propositional attitude ascriptions. For example, Braun (1998, p. 560-1) writes:

According to these theories, belief reports often express propositions that are partly about believers’ mental representations. [...] Moreover, these theories say that speakers routinely think about other people’s mental representations, and intend to talk about those representations when they utter belief sentences. But I seriously doubt that ordinary speakers have such sophisticated thoughts and intentions about mental representations when they utter belief sentences.

We too doubt that ordinary speakers have sophisticated thoughts and intentions about mental representations. Nothing we have said so far implies otherwise. To characterize the context-sensitivity of a class of expressions in terms of entities of a certain sort is not to postulate that speakers have thoughts or intentions about such entities. The project of specifying sentences’ truth conditions is not the project of spelling out the inner monologues of competent users of those sentences. For example, those who think that the context-sensitivity of modal adverbs is parametrized by sets of possible worlds are not committed to the claim that ordinary speakers have sophisticated thoughts about possible worlds or sets thereof. What goes for possible worlds goes for mental representations.

The most flatfooted strategy for relating perspectives to mentalese sentences would be to identify them. On this picture, the three-place relation $\alpha_{\text{“believe”}}$ underlying “believe” would be the relation that holds between a person x , mentalese sentence π and proposition p just in case the person accepts the sentence (i.e., has it in their “believe”-box) and the sentence expresses the proposition. More generally, the proposal would be to identify the three-place relation α_X associated with a propositional attitude verb X with the relation $\lambda x \pi p. (\beta_X(x, \pi) \wedge \pi \text{ means that } p)$.²⁷

²⁷Note that this proposal requires not counting negations of attitude verb-phrases as attitude verb-phrases. For “ A is sure that φ or A is not sure that φ ” should be true when interpreted using any perspective π . But if the second disjunct requires for its truth that π (be in one’s “not sure”-box and) express the proposition expressed in context by φ , then the disjunction will be false on interpretations where π does not express that proposition. This point carries over to the proposal (to be discussed presently) that perspectives be identified with sets of mentalese sentences. It also

Unfortunately, this proposal is a non-starter, since (assuming any given mentalese sentence expresses at most one proposition²⁸) it entails the validity of the following clearly invalid schemas:

25. If A believes that φ and B believes that ψ , then φ iff ψ .
26. A believes at most one thing.

For if “ A believes that φ ” is true when “believes” is interpreted relative to a perspective (i.e. mentalese sentence) π , then the present proposal entails that π expresses the proposition expressed by φ ; likewise for “ B believes that ψ ” and ψ ; so if both belief ascriptions are true then φ and ψ express the same proposition, and so are materially equivalent. A parallel argument establishes the validity of 26.²⁹

A proponent of the present proposal might reply that verbalists are already in the business of accepting the surprising validity of certain instances of Leibniz’s law. Why should 25 and 26 be any different? There are three reasons. First, when considered as a schema, Leibniz’s law is extremely attractive; by contrast, 25 and 26 are not instances of any similarly compelling schema. Second, verbalists agree that the puzzling instances of Leibniz’s law are indeed false on their most natural interpretations because these interpretations are non-uniform. The same strategy is unavailable to the defender of 26, since that sentence contains no repeated expressions and so has only uniform interpretations. Third, and most importantly, our present goal is to reconcile the context-sensitivity of attitude ascriptions with the clearly theoretically fruitful practice of writing down general principles about propositional attitudes, like 19-23. Any proposal that classifies 25 and 26 with 19-23 has manifestly failed to make sense of that practice.³⁰

The natural response to this problem is to identify perspectives not with individual mentalese sentences, but with sets of mentalese sentences.³¹ For example, the three-place relation underlying “believe” will be the relation that holds between a person, a set of mentalese sentences, and a proposition just in case the person accepts some sentence in the set and that sentence

foreshadows the fact that, when considering identity confusion, we should expect differential behavior of attitude verbs like “unsure” and negated attitude verbs like “not sure” – a point we will discuss at length below.

²⁸This assumption is not beyond question. However, even if a single mentalese sentence can express multiple propositions, presumably there are pairs of propositions (e.g., that the earth is flat and that the sky is blue) both of which are believed by someone but which are not both expressed by any single mentalese sentence, which is enough for the present argument.

²⁹A version of the problem illustrated by 25 applies to the “notion provision” proposal for interpreting attitude verbs from Crimmins and Perry (1989, pg. 705). However, it does not apply to a different way of interpreting attitude verbs that they postulate, under the heading of “notion constraint”; see footnote 33.

³⁰Another problem for the single-sentence proposal is that it can’t handle non-propositional attitude verbs like “looks for”, “thinks about”, etc. See footnote 32 for a solution to this problem.

³¹A different solution to the problem, considered and persuasively criticized by Richard (1993, p. 147), assumes that perspectives are denoted by unpronounced variables in the logical forms of sentences containing attitude verbs, and postulates that in sentences like 26 these variables are somehow bound by the quantifier “at most one proposition”.

expresses the proposition in question. More generally, the proposal is that $\alpha_X = \lambda x \pi p. \exists s (s \in \pi \wedge \beta_X(x, s) \wedge s \text{ means } p)$ for every propositional attitude verb X .³²

This proposal fails to validate either 25 or 26.³³ It also has a much more general theoretical advantage over the previous single-sentence proposal: given natural (if idealized) assumptions about our underlying cognitive psychology, it is compatible with there being a non-trivial logic of propositional attitudes. This point is worth pausing over, since some in the literature are confused about it. For example, Crimmins and Perry (1989, 710) remark that on their view “there can be no simple logic of belief-sentences”. They write:

Whereas there is little possibility of an interesting logic of belief sentences, the logic of beliefs, notions, and ideas [i.e., mental representations] is available. Such issues as logical and analytic closure of belief, explicit versus implicit belief, and inferential issues in belief change really belong to the logic of beliefs rather than to the logic of belief sentences. We can explore the logic of the relations we have seen as underlying our ordinary talk about beliefs – but this logic will not be a logic of ordinary language. (711)

(Note that in this passage Crimmins and Perry use “beliefs, notions, and ideas” *not* to talk about what people believe – i.e., about the mental states ascribed using the word “believe” – but rather to talk about people’s underlying mental representations; their talk of “belief-sentences” concerns ordinary attitude ascriptions.)

This pessimism is unjustified. We will now explain one way that general principles governing which mentalese sentences an agent accepts (and, more generally, which mentalese sentences are in which attitude boxes) can support a non-trivial logic of belief and other propositional attitudes. The following schemas are quite attractive (at least when the antecedent is conjoined with “ A is ideally rational”, a proviso we will hereafter suppress for brevity):³⁴

- 27. If A believes that φ and A believes that ψ , then A believes that $\varphi \wedge \psi$.
- 28. If A believes that φ and A believes that ψ , then A believes that $\varphi \vee \psi$.
- 29. If A believes that φ and A believes that ψ , then A believes that $\varphi \rightarrow \psi$.

³² This proposal can be generalized to solve the problem raised in footnote 30 by identifying perspectives with sets of mentalese terms not all of which are sentences. For example, we might say that “Jack is looking for Jill” is true when “looking for” is interpreted relative to the perspective π just in case Jack’s “looking for”-box contains a mentalese singular term $t \in \pi$ that denotes Jill. We will ignore this complication in the main text.

³³ The proposal can be seen as a contextualist variant of Salmon (1986), and closely corresponds both to the “notion constraint” proposal from Crimmins and Perry (1989) and the appeal to “types of modes of presentation” in Schiffer (1992).

³⁴ Of course, even so qualified these principles are controversial, but mainly for reasons orthogonal to the issue of identity confusion. For example, some philosophers reject 27 for reasons having to do with the preface paradox; see Makinson (1965). Note that even these philosophers often accept the validity of parallel schemas with “is certain” in place of “believes”.

We can vindicate their attitudinal validity by combining assumptions about cognitive psychology with assumptions about the nature of perspectives. Concerning cognitive psychology: we assume that mentalese contains binary connectives $\hat{\wedge}$, $\hat{\vee}$, and $\hat{\rightarrow}$ expressing conjunction, disjunction, and material implication, and that whenever an ideally rational person accepts two mentalese sentences s and t , they also accept $\ulcorner s \hat{\wedge} t \urcorner$, $\ulcorner s \hat{\vee} t \urcorner$, and $\ulcorner s \hat{\rightarrow} t \urcorner$. Similarly, to vindicate the attitudinal validity of 20 and 21, we assume that mentalese contains a unary sentential connective $\hat{\neg}$ expressing negation and that, if s is in A 's "fear"-box, then $\ulcorner \hat{\neg}s \urcorner$ is in their "hope"-box, and that if s is in their "doubt"-box, then $\ulcorner \hat{\neg}s \urcorner$ is in their "think"-box.³⁵ Concerning perspectives: we assume that if a perspective contains s and t , then it also contains $\ulcorner s \hat{\wedge} t \urcorner$, $\ulcorner s \hat{\vee} t \urcorner$, $\ulcorner s \hat{\rightarrow} t \urcorner$, and $\ulcorner \hat{\neg}s \urcorner$. These assumptions entail the attitudinal validity of 20, 21, and 27-29.³⁶

Just as importantly, identifying perspectives with sets of mentalese sentences closed under the formation of conjunctions, disjunctions, material implications, and negations – *closed sets*, for short – is compatible with the pattern of context-sensitivity of attitude reports that motivated the introduction of perspectives in the first place. For example, let s be Lois's mentalese counterpart of the sentence "Superman flies" and c her mentalese counterpart of "Clark flies". Let π_s be the smallest closed set containing s and π_c be the smallest closed set containing c . Restricting our attention to interpretations on which "Lois" refers to Lois, "flies" means flies, and "he" refers to Superman, "Lois knows that he flies and doesn't believe that he doesn't fly" is true on any interpretation parametrized by π_s and false on any interpretation parametrized by π_c ; conversely, "Lois believes that he doesn't fly, and doesn't know that he flies" is true on any interpretation parametrized by π_c and false

³⁵We don't actually need to assume that mentalese contains such symbols (which might be denied, for example, by someone who thought that mentalese had the same syntactic structure as English). It suffices that there is a well-defined operation of conjoining two mentalese sentences to form a mentalese sentence such that if s means that p and t means that q then the result of conjoining s with t means that p and q (and likewise for other Boolean operations).

³⁶If we want to allow perspectives to contain non-sentential mentalese expressions, as mentioned in footnote 32, we should restrict the closure condition so that it applies only to members of perspectives that are sentences. Moreover, the sorts of schemas that motivate requiring that perspectives be closed in this sense might in this more general setting motivate additional closure requirements, such as that perspectives be closed under the formation of well-formed terms. This requirement implies the closure condition in the main text provided that all perspectives contain $\hat{\wedge}$, $\hat{\vee}$, $\hat{\rightarrow}$ and $\hat{\neg}$.

Note that perspectives could correspond to sets of terms closed under the operation of term-formation without corresponding to sets of *atomic expressions* closed under the operation of term-formation. For example, the former but not the latter proposal is compatible with there being a perspective that contains " $a = b$ " but does not contain " $a = a$ ". This is something we might want in order to get a true reading of "It is temporarily true that Lois does not know that Superman is Clark" (at a time when Lois has met him in his superhero persona and in his reporter persona but not yet figured out that these are personas of the same person) from the fact that it is temporarily true that the mentalese correlate of "Superman = Clark" is not in Lois's "know"-box. This would not be possible if any perspective containing that sentence had to also contain "Superman = Superman" (which is already in Lois's "know"-box). However, as we will discuss in section 9, this strategy for achieving the desired true reading may be untenable if propositions are too coarse-grained.

on any interpretation parametrized by π_s .

We will take as a working hypothesis that perspectives are closed sets of mentalese sentences. (We revisit this assumption in section 9.) Note that we are *not* making the converse assumption that all closed sets of mentalese sentences are perspectives, in the sense of parametrizing some interpretation of attitude verbs consistent with the conventions of English. We do not accept this further claim, for reasons we will discuss in section 7.

6 The limits of attitudinal validity

Schemas like 19-23 and 27-29 are instances of a more abstract schema³⁷

30. If A X_1 s that φ_1, \dots , and A X_n s that φ_n , then A Y s that $\mathcal{B}(\varphi_1; \dots; \varphi_n)$,

instances of which are obtained by replacing X_1, \dots, X_n, Y with attitude verbs and $\mathcal{B}(\varphi_1; \dots; \varphi_n)$ with a schematic sentence built using only negation, conjunction, disjunction, material implication, and the pairwise-distinct schematic letters $\varphi_1, \dots, \varphi_n$. We can now prove a general result: given our assumptions about perspectives, an instance of 30 is attitudinally valid whenever the corresponding principle about mentalese sentences

31. Anyone who has s_1 in their X_1 -box, \dots , and s_n in their X_n -box has $\ulcorner \hat{\mathcal{B}}(s_1; \dots; s_n) \urcorner$ in their Y -box,

is true, where $\hat{\mathcal{B}}(s_1; \dots; s_n)$ results from $\mathcal{B}(\varphi_1; \dots; \varphi_n)$ by substituting “ $\hat{\neg}$ ”, “ $\hat{\wedge}$ ”, “ $\hat{\vee}$ ”, and “ $\hat{\supset}$ ” for their English correlates and the variables s_1, \dots, s_n for the schematic letters $\varphi_1, \dots, \varphi_n$.³⁸ (*Proof*: consider a schema Σ that is an instance of 30, and consider a sentence σ that is an instance of Σ . Suppose the antecedent of σ is true when interpreted using the perspective π . Then π contains some mentalese sentences s_i respectively meaning the same as the respective complement clauses of σ and in the respective X_i -boxes of the subject of σ . It also contains all Boolean combinations of these s_i , since perspectives are closed. So if Σ^* is true, then the appropriate Boolean combination of the s_i will be in the attitude box needed to make the consequent of σ true when interpreted using π .)

The situation is thus quite different from the one imagined by Crimmins and Perry. According to them, although systematic theorizing is possible when thinking about mental representations, verbalism renders parallel generalizations formulated using propositional attitude verbs untenable. But we have just shown that, to the contrary, for a large class of schematic principles

³⁷This isn’t quite true, since the attitude verbs in 22 and 23 take infinitival complements, but this subtlety is orthogonal to the issues we will be discussing. The same point applies to attitude verbs like “unsure” and “wonders” that take interrogative complements.

³⁸This isn’t quite right when φ is a schematic predicate in an infinitival complement, as in 22 and 23. The mentalese generalization corresponding to 22, for example, ought to be that anyone with $\ulcorner F\iota \urcorner$ in their “intend”-box has $\ulcorner F\iota \urcorner$ in their “expect”-box, where ι is a mentalese first-person pronoun. This subtlety does not affect the proof in the main text.

formulated using propositional attitude verbs, their attitudinal validity is *entailed* by the truth of corresponding generalizations about mentalese sentences (given our identification of perspectives with closed sets of mentalese sentences). In this way, generalizations about mentalese actually vindicate large swaths of propositional attitude psychology.

However, this way of vindicating propositional attitude psychology also has severe limitations, since many attractive principles formulated using propositional attitude verbs are not instances of the abstract schema 30. Consider, for example:

32. A does not believe that $\varphi \wedge \neg\varphi$.
33. A doesn't prefer that φ than that φ .
34. If A believes that $\neg\varphi$, then A doesn't believe that φ .
35. If A is unsure whether φ , then A is not sure that φ .
36. If A assigns probability x to $\varphi \wedge \psi$ and assigns probability y to ψ , then A assigns probability $\frac{x}{y}$ to φ given ψ .
37. If A wants above all else to φ and A thinks that they can φ , then A will try to φ .

These schemas have the same flavor as 19-23 and 27-29: they too are reasonably attractive generalizations, at least concerning people who are in some intuitive sense perfectly rational. But their attitudinal validity is not entailed by the truth of the corresponding principles about mentalese sentences. In the case of 32-35, this turns out to be because the attitude verbs occur under negation. In the case of 36-37 it is (most importantly) because the same schematic letter φ is repeated in the complement of multiple attitude ascriptions in their antecedents. In this section we will argue that these schemas are *not* attitudinally valid, even if the corresponding mentalese generalizations are true. We will consider them in turn, drawing out some general morals as we go.

Regarding 32: Every uniform interpretation of "Lois believes that Superman flies and Clark does not fly" is attitudinally uniform because there is only one occurrence of an attitude verb in it. This sentence is true on some interpretation on which "Superman" and "Clark" co-refer. Since that interpretation is uniform, it follows from the validity of Leibniz's law that on this interpretation, "Lois believes that Superman flies and Superman does not fly" is true. Since it is true on this interpretation, its negation "Lois does not believe that Superman flies and Superman does not fly" is false on this interpretation. So not every instance of 32 is true on every attitudinally uniform interpretation.

Regarding 33: Every uniform interpretation of "Lois prefers that she be around Superman than that she be around Clark" is attitudinally uniform, again because there is only one occurrence of an attitude verb in it. This sentence is true on some interpretation on which "Superman" and "Clark" co-refer. It follows from the validity of Leibniz's law that on that interpretation,

“Lois prefers that she be around Superman than that she be around Superman” is true. Since this sentence is true on this interpretation, its negation “Lois does not prefer that she be around Superman than that she be around Superman” is false on that interpretation. So not every instance of 33 is true on every attitudinally uniform interpretation.

(There is some room to resist this argument. Suppose, for example, one thought that to prefer that φ than that ψ just is for there to be an amount of money that you’d be willing to pay to make it the case that φ but would not be willing to pay to make it the case that ψ . Since “Lois is willing to pay x to be around Superman but would not be willing to pay x to be around Superman” is false on any uniform interpretation, we get the surprising conclusion that there is no true uniform interpretation of “Lois prefers that she be around Superman than that she be around Clark” on which “Superman” and “Clark” co-refer. This consequence strikes us as ample reason for verbalists to reject the above analysis of “prefers” and accept the above argument against the attitudinal validity of 33.)

Regarding 34: It is natural to think that “Lois believes that Superman flies and believes that Superman does not fly” is a counterexample to the attitudinal uniformity of 34 when it receives an interpretation uniform with the true uniform interpretation of “Lois believes that Superman flies and Superman does not fly” used to argue above against the attitudinal validity of 32. Nevertheless, it is worth considering how this natural thought might be established from general principles.

The most obvious such argument would be to appeal to the attitudinal validity of the schema:

38. If A believes that $\varphi \wedge \psi$, then A believes that φ and A believes that ψ .

The attitudinal invalidity of 34 would then follow from the already established attitudinal invalidity of 32. However, it is not clear that 38 is attitudinally valid. For if we accept a coarse-grained theory of propositions according to which necessarily material equivalent propositions are identical, the principle implies that anyone believes every proposition modally weaker than any proposition they believe – even ones they lack the cognitive resources to express. This implication is incompatible with the truth conditions we have been assuming for belief ascriptions in terms of what is meant by mentalese sentences in one’s “believe”-box, since in general there will be propositions not expressed by any sentence of mentalese that are modally weaker than some proposition expressed by some mentalese sentence in one’s “believe”-box.³⁹

However, such coarse-grained theories of propositions support a different argument for the attitudinal invalidity of 34. For such theories imply the attitudinal validity of:

³⁹Friends of coarse-grained theories of propositions could accommodate the attitudinal validity of 38 within a broadly mentalese framework by replacing “means that” with “means something that necessarily implies that” in the truth conditions for “believe”. However, since parallel modifications are not tenable for other attitude verbs like “doubts”, this proposal breaks the systematic connection between propositional attitude verbs and ‘boxes’ for mentalese sentences.

39. If A believes that $(\varphi \wedge \neg\varphi) \rightarrow (\varphi \wedge \neg\varphi)$, then A believes that $\neg(\varphi \wedge \neg\varphi)$.

Assuming the attitudinal validity of 29, the attitudinal invalidity of 34 then follows from the attitudinal invalidity of 32.⁴⁰

A different argument for the attitudinal invalidity of 34 starts with the fact that both ‘Lois believes that Superman flies’ and ‘Lois believes that Superman does not fly’ have true uniform interpretations, and appeals to the principle that, if two positive attitude ascriptions both have true attitudinally uniform interpretations, then their conjunction has a true attitudinally uniform interpretation. This principle is a consequence of the following natural hypothesis about which closed sets of mentalese sentences count as perspectives:

Perspectival Plenitude: For all perspectives π_1 and π_2 and $s \in \pi_1$ and $t \in \pi_2$, there is a perspective π_3 such that $s \in \pi_3$ and $t \in \pi_3$.⁴¹

We therefore conclude that, whatever the status of 38 and 39, verbalists should think that 34 is not attitudinally valid.

Regarding 35: The situation is similar to that concerning 34. Suppose that Lois is unsure whether the reporter she knows lives a double life as the local superhero. Then “Lois is unsure whether Clark flies” has a true uniform interpretation, and hence so does “Lois is unsure whether Superman flies”, by the validity of Leibniz’s law. Since “Lois is sure that Superman flies” also has a true uniform interpretation, Perspectival Plenitude entails that the conjunction “Lois is unsure whether Superman flies and Lois is sure that Superman flies” has a true attitudinally uniform interpretation, in which case “If Lois is unsure whether Superman flies, then Lois is not sure that Superman flies” has a false attitudinally uniform interpretation, so 35 is not attitudinally valid.

A different argument for the same conclusion appeals to the attitudinal validity of the schema:

40. If A is unsure whether φ , then A is sure that $\varphi \rightarrow \varphi$,

together with a coarse-grained theory of propositions according to which $\varphi \rightarrow \varphi$ and $(\varphi \rightarrow \varphi) \rightarrow (\varphi \rightarrow \varphi)$ are the very same claim. Consider a true uniform interpretation of “Lois is unsure whether Clark flies if Superman flies”.

⁴⁰*Proof:* If A believes that $\varphi \wedge \neg\varphi$, then A believes that $(\varphi \wedge \neg\varphi) \rightarrow (\varphi \wedge \neg\varphi)$ (by 29), and so A believes that $\neg(\varphi \wedge \neg\varphi)$ (by 39).

⁴¹There are ways of rejecting Perspectival Plenitude that allow schemas like 34 to come out attitudinally valid. For example, one could hold that π is a perspective only if, for some mentalese sentence φ , π is the smallest closed set containing φ . But this proposal faces the same problem as the earlier proposal to identify perspectives with individual mentalese sentences: it makes the notion of attitudinal validity so indiscriminating as to be uninteresting. For just as the single-sentence proposal implied the attitudinal validity of 25 and 26, the proposal under consideration implies the attitudinal validity of the equally unattractive schema:

If A believes that φ and A believes that ψ , then either $\Box\varphi$, $\Box\neg\varphi$, $\Box\psi$, $\Box\neg\psi$, $\Box(\varphi \leftrightarrow \psi)$, or $\Box(\varphi \leftrightarrow \neg\psi)$,

where \Box abbreviates “it is metaphysically necessary that”. We see no principled proposal about perspectives that both avoids this sort of problem and is compatible with the attitudinal validity of schemas like 34.

Leibniz’s law, 40 and our granularity assumption imply that “Lois is sure that Superman flies if Superman flies” is also true on this interpretation, contradicting the attitudinal validity of 35. (Compare the earlier argument from 29 and principles of coarseness of grain against the attitudinal validity of 34.)

Regarding 36: For simplicity, imagine we have a continuum of attitude verbs Pr_x , one for each real number $x \in [0,1]$, each with a corresponding Pr_x -box. We say that A assigns probability x to mentalese sentence s just in case s is in their Pr_x -box. We assume that ideally rational people are coherent in their assignments of probabilities to mentalese sentences, in the following sense: they assign each mentalese sentence exactly one probability; they assign probability 1 to every theorem of propositional logic; they assign $\lceil s \vee t \rceil$ the sum of the probabilities they assign to s and to t whenever $\lceil \neg(s \wedge t) \rceil$ is a theorem of propositional logic; and they assign at least as high a probability to t as they do to s whenever $\lceil s \rightarrow t \rceil$ is a theorem of propositional logic. We say that A assigns probability x to s given t just in case $x = \frac{y}{z}$, where y is the probability A assigns to $\lceil s \wedge t \rceil$ and z is the probability A assigns to t .⁴² This allows us to interpret both conditional and unconditional subjective probability ascriptions relative to perspectives just as we do for other propositional attitude ascriptions.

We may suppose “Lois assigns probability .5 to Superman = Clark” has a true uniform interpretation. By the validity of Leibniz’s law, so does “Lois assigns probability .5 to Superman = Superman”. Now “Lois assigns probability 1 to Superman = Superman \wedge Superman = Superman” also has a true uniform interpretation. Assuming Perspectival Plenitude, it follows that the conjunction “Lois assigns probability 1 to Superman = Superman \wedge Superman = Superman and assigns probability .5 to Superman = Superman” has a true attitudinally uniform interpretation. Since “Lois assigns probability 2 to Superman = Superman” is not true on any interpretation, we have a counterexample to the attitudinal validity of 36.

This argument appeals to Perspectival Plenitude, but the same conclusion can be established assuming only that “Lois assigns probability .5 to Superman = Clark given Superman = Superman” has a true interpretation (compare the case of “Lois prefers being around Superman to being around Clark”) and that “Lois assigns probability 1 to Superman = Superman \wedge Superman = Superman and assigns probability .5 to Superman = Clark” is also true relative to the perspective that parametrizes that interpretation. We can then appeal to the validity of Leibniz’s law to establish the existence of a true attitudinally uniform interpretation of “Lois assigns probability 1 to Superman = Superman \wedge Superman = Superman and assigns probability .5 to Superman = Superman”, and then argue as before.

Regarding 37: We again appeal to Perspectival Plenitude, but it is worth going through the details since they illustrate a more general phenomenon. Suppose Lois is confronted by a bespectacled reporter and wants more than anything else to thank the local superhero who recently saved her. Let π_1 be

⁴²For a broadly similar proposal, see Braun (2016).

the perspective raised to salience by the sentence “Lois wants above all else to thank Superman for saving her”, and let π_2 be the perspective raised to salience by the sentence “Lois believes that she can thank Clark”. Since these sentences are true on the corresponding interpretations, π_1 must contain a sentence s in Lois’s “want”-box that means that she thanks Superman for saving her, and π_2 must contain a sentence t in her “believe”-box that means that she can thank Clark (i.e., Superman). By Perspectival Plenitude, there must then be a perspective π_3 that contains both s and t . Since there is no mentalese sentence in Lois’s “try”-box that means that she thanks Clark, the sentence “If Lois believes that she can thank Clark and she wants above all else to thank Superman for saving her, then she will try to thank Clark” is false on this interpretation. So by the validity of Leibniz’s law, “If Lois believes that she can thank Clark and she wants above all else to thank Clark for saving her, then she will try to thank Clark” is a counterexample to the attitudinal validity of 37 when interpreted using π_3 .⁴³

There is a general moral here. Recall our earlier principle 23. It is logically weaker than 37. So it must be attitudinally invalid too. But its attitudinal invalidity is unsurprising: rational people don’t try to do things that they don’t think they can do, irrespective of their desires. The surprise is that, assuming Perspectival Plenitude, identity confusion means that we can’t make 23 attitudinally valid by modifying it in the manner of 37. In general, we aren’t going to be able to repair oversimple principles of propositional attitude psychology by making them conditional on having other propositional attitudes towards related contents, assuming the standard of repair is attitudinal validity. (Formally, observe that adding such provisos results in principles that are not instances of the general pattern 30, because schematic letters will be repeated in the complement clauses in their antecedents.) But such repairs are part and parcel of our ordinary practice of theorizing using propositional attitude verbs. Since they do not generally yield attitudinally valid principles, we cannot vindicate the good standing of such theorizing by understanding its goal as the production of attitudinally valid principles.

7 Confusion-free validity

Call a set of mentalese sentences *transparent* just in case two members of the set express the same proposition only if they are logically equivalent. The arguments just given against the attitudinal validity of 32-37 all appeal to a sentence having an interpretation parametrized by a perspective that is not transparent. This observation leads us to make the following proposal. Call a

⁴³Braun (2000) makes a similar point in the context of neo-Russellianism, which for present purposes we can think of as the non-contextualist limiting case of verbalism according to which the only perspective is the set of all mentalese sentences. More generally: principles of propositional attitude psychology that are false according to neo-Russellians will be attitudinally invalid according to those verbalists who think (naturally enough) that the set of all mentalese sentences is a perspective.

sentence *confusion-free valid* just in case it is true on every attitudinally uniform interpretation that is parametrized by a transparent perspective. We propose that verbalists take confusion-free validity, not attitudinal validity, to be the aim of ordinary theorizing about the propositional attitudes.⁴⁴

Moving from attitudinal validity to confusion-free validity allows us to improve on our earlier result relating the validity of schemas containing propositional attitude verbs to generalizations about sentences of mentalese. Consider the following schema:

41. If A X_1 s that φ_1, \dots , and A X_n s that φ_n , then $\mathcal{B}_0(A$ Y_1 s that $\mathcal{B}_1(\varphi_1; \dots; \varphi_n); \dots; A$ Y_m s that $\mathcal{B}_m(\varphi_1; \dots; \varphi_n))$,

where, as before, $\lceil \mathcal{B}_i(\varphi_1, \dots, \varphi_n) \rceil$ schematically stands for any Boolean combination of the schematic sentences letters $\varphi_1, \dots, \varphi_n$. 41 generalizes 30 by allowing in its consequent arbitrary Boolean combinations of attitudes towards Boolean combinations of complement clauses of conjuncts in the antecedent. For example, although 34 and 35 are not instances of 30, they are instances of 41. And 33, though not an instance of 41, is logically equivalent to the following pair of schemas:

- 33a. If A understands the proposition that φ , then A does not prefer that φ than that φ .
- 33b. If A prefers that φ than that φ , then A understands the proposition that φ .

each of which are instances of 41 (or, at least, of the obvious generalization of 41 that makes allowance for polyadic attitude-verb-phrases). Just as we proved that an instance of 30 is attitudinally valid if the corresponding instance of 31 is true, we can prove that an instance of 41 is confusion free-valid if the corresponding generalization about sentences of mentalese is true, given the assumption that the contents of the relevant attitude boxes are closed under logical equivalence, in the sense that

⁴⁴To say that a sentence is confusion-free valid is not to say that it is attitudinally valid *modulo* cases of identity confusion, whatever that might mean. The sentence “No one is identity-confused about anything” is presumably attitudinally valid *modulo* cases of identity-confusion if anything is. But it is not confusion-free valid, and it is not the sort of claim that ordinary theorizing about the propositional attitudes aims to uncover, since it has only false interpretations.

Indeed, perhaps being a rational person implies being identity confused, in the sense of either having a sentence meaning that $x = x$ in one’s “unsure”-box or a sentence meaning $x \neq x$ in one’s “sure”-box, for some x . Plausibly, any rational person A will be unsure whether p , for some p , and know about at least two individuals, a and b . Let c be the unique thing that is identical to a if p and identical to b if not- p . Arguably, the following sentence has a true interpretation:

There is something identical to c (a if p , and b otherwise) such that A is unsure whether that thing is identical to c .

If so, the present framework implies that there is some s in A ’s “unsure”-box that means that $c = c$. Another example, appealing to a more idealized conception of rationality: perhaps the continuum hypothesis is true but any rational person should be unsure about it, in which case they will have the mentalese analogue of the true identity “ $\aleph_1 \neq \beth_1$ ” in their “unsure”-box.

42. Any rational person with s_1 in their X_1 -box, \dots , s_n in their X_n -box and t in their Y -box has $\mathcal{B}(s_1; \dots; s_n)$ in their Y -box if it is logically equivalent to t .⁴⁵

(The proof is in a footnote.⁴⁶) So provided the mentalese generalizations corresponding to 33a, 34 and 35 are true, this result implies that those schemas are confusion-free valid, despite not being attitudinally valid.

Not all of 32-37 are instances of 41 (or, as in the case of 33, consequences of plausible instances thereof). Consider 32. It not an instance of 41, and, moreover, its confusion-free validity is not an uncontroversial consequence of the corresponding mentalese generalization (i.e., that no rational agent accepts a mentalese sentence of the form $\lceil s \hat{\wedge} \hat{\neg} s \rceil$). Let α be a syntactically simple mentalese sentence-letter meaning that $p \wedge \neg p$. It is controversial whether any transparent perspective contains α . Those who think that $p \wedge \neg p$ and $(p \wedge \neg p) \wedge \neg(p \wedge \neg p)$ are the same proposition will hold that no closed set containing α is transparent, since any such set will contain $\lceil \alpha \hat{\wedge} \hat{\neg} \alpha \rceil$, which is not logically equivalent to α but, by their lights, does express the same proposition as α . But those who deny that propositions are so coarse-grained may well think, e.g., that the smallest closed set containing α is a transparent perspective. If they are right, and some rational agent accepts α , then 32 is not confusion-free valid.⁴⁷

Verbalists who reject the confusion-free validity of 32 needn't reject all

⁴⁵Here is why we need this assumption. For all we have said there is an attitude *+belief* such that s is in a perfectly rational person's "+believe"-box just in case s is in their "believe"-box and $s \neq \lceil \hat{\wedge} t \rceil$ for any t . Then the schema

If A +believes that φ , then A does not +believe that $\neg\neg\varphi$

corresponds to a true mentalese generalization. But whether it is confusion-free valid will be controversial; for example, its confusion-free validity will be denied by those who think that propositions are identical to their double negations.

⁴⁶*Proof:* Let Σ be an instance of 41, σ be an instance of Σ whose antecedent is true relative to an attitudinally uniform interpretation parametrized by the transparent perspective π . We will argue that the consequent of σ is true on this interpretation if the generalization Σ^* about sentences of mentalese corresponding to Σ is true. This has already been established in the case where the consequent of σ is a Y -ascription not embedded under any Boolean connectives, since then Σ is an instance of 30. Now suppose the consequent of σ is a negated Y -ascription. We know that some mentalese sentences s_i are in the relevant attitude boxes, since the antecedent of σ is true. So the appropriate Boolean combination of them β is not in the person's Y -box, since Σ^* is true. So no sentence logically equivalent to β is in their Y -box, by 42. So no $t \in \pi$ with the same meaning as β in their Y -box, since π is transparent. So the consequent of σ is true when interpreted relative to π . Now consider the case where the consequent of σ is a disjunction of either unembedded or negated attitude ascriptions. We know that $s_i \in \pi$ witness the truth of the disjunctive consequent of Σ^* , and so witness the truth of one its disjuncts, in which case the corresponding disjunct of the consequent of σ is true on the present interpretation, and hence so is σ . This covers all cases, since any Boolean combination of sentences is equivalent to its conjunctive normal form, and a conditional with a conjunctive consequent is confusion-free valid if the conditionals with the same antecedent but with the individual conjuncts as consequents are each confusion-free valid.

⁴⁷This example also shows that unrecognized co-reference between syntactically simple (mentalese) expressions is not the only way in which a rational agent might come to accept a mentalese sentence that expresses a contradiction: for all we've said it can also happen (absent any unrecognized co-reference) when the contradictoriness of the proposition expressed by a mentalese sentence is not reflected in that sentence's syntactic structure.

versions of the idea that rational agents don't believe contradictions. For they can accept the following weaker schema:

43. If A understands the proposition that φ , then A does not believe that $\varphi \wedge \neg\varphi$.

Unlike 32, 43 is an instance of 41, so its confusion-free validity is implied by the corresponding principle about mentalese (i.e., that no rational agent who has a mentalese sentence in their "understands the proposition that"-box accepts the conjunction of that sentence with its negation). The possibility of the kind of scenario described in the previous paragraph does not threaten the confusion-free validity of 43. This is because no transparent perspective containing α can contain any sentence s meaning that p . (Any such perspective would be closed, and so would contain $\lceil s \wedge \hat{\alpha}s \rceil$, which expresses the same proposition as the logically inequivalent α , making the perspective non-transparent.) So 43 will be vacuously true when interpreted relative to any transparent perspective containing α .

Next consider 36. It is also not an instance of 41, since the complement clause in its consequent is not a Boolean combination of complement clauses in its antecedent.⁴⁸ And unlike 32, its failure to be confusion-free valid does not depend on tendentious principles about propositional granularity. Consider a perspective π containing mentalese sentences s meaning $p \wedge q$ and t meaning q but not containing any mentalese sentence meaning p . That such a perspective be transparent is compatible with coarse-grained theories of propositions.⁴⁹ Then any rational person who assigns some probability to s and to t presents an immediate counterexample to the confusion-free validity of 36. However, as in the case of 32, we can retreat to an appropriately "understanding"-qualified principle

44. If A understands the proposition that φ and understands the proposition that ψ , then, if A assigns probability x to $\varphi \wedge \psi$ and assigns probability y to ψ , A assigns probability $\frac{x}{y}$ to φ given ψ ,

which is an instance of 41.⁵⁰

Finally, consider 37. It too is not an instance of 41, since the non-Boolean operator "can" is embedded under the attitude verb "think". But we don't

⁴⁸Another reason why 36 is not an instance of 41 is that one of the complement clauses in its antecedent is logically complex. Such complexity can prevent the truth of a mentalese generalization from implying the confusion-free validity of the corresponding schema. For example, in Goodman and Lederman we argue that

If A is sure that $\varphi \wedge \psi$ and is unsure whether $\varphi \wedge \chi$, then A is unsure whether $\psi \wedge \chi$ is not confusion-free valid even on the assumption that anyone with $\lceil s \wedge t \rceil$ in their "sure"-box and $\lceil s \wedge u \rceil$ in their "unsure"-box has $\lceil t \wedge u \rceil$ in their "unsure"-box.

⁴⁹For example, suppose propositions are individuated modally, that p and q are contingent, that q necessitates p but not vice versa, and that π is the smallest closed set containing s .

⁵⁰See Goodman and Lederman for a general theory of where such "understanding"-provisos are needed and how to formulate them when the principles in question are more logically complex than 32 and 36.

think that this fact jeopardizes the principle's confusion-free validity. For just as reflection on 27 led us to postulate that perspectives are closed under the operation of forming sentences using the mentalese analogue $\hat{\wedge}$ of "and", reflection on 37 should likewise lead us to postulate that perspectives are closed under the operation of forming sentences using a mentalese analogue $\hat{\diamond}$ of "I can make it the case that". 37 will then be confusion-free valid so long as the corresponding mentalese generalization is true. This fact is an instance of a more general result. When considering attractive schemas in which constants c_1, \dots, c_n with respective mentalese correlates $\hat{c}_1, \dots, \hat{c}_n$ occur under attitude verbs, the operative notion of a perspective will be one that includes only sets of mentalese sentences closed under the formation of sentences using these mentalese correlates. Let $\mathcal{C}(\varphi_1; \dots; \varphi_m)$ schematically stand for any schema built from $\varphi_1, \dots, \varphi_m, c_1, \dots, c_n$, and Boolean connectives. An instance of

45. If A X_1 s that φ_1, \dots , and A X_n s that φ_n , then $\mathcal{B}(A$ Y_1 s that $\mathcal{C}_1(\varphi_1; \dots; \varphi_n); \dots; A$ Y_m s that $\mathcal{C}_m(\varphi_1; \dots; \varphi_n)$)

will then be confusion-free valid so long as the corresponding mentalese-generalization is true, for the same reason that instances of 41 are. Since 37 is logically equivalent to

46. If A wants above all else to φ , then (if A thinks that they can φ , then A will try to φ),

which is an instance of 45, its confusion-free validity is implied by the corresponding mentalese generalization (i.e., that anyone with s in their "want above all else"-box and $\ulcorner \hat{\diamond}s \urcorner$ in their "think"-box will have s in their "try to make it the case that"-box). The shift from attitudinal validity to confusion-free validity also resolves the more general issue mentioned in connection to 37 at the end of section 6: we now *can* repair oversimple principles of propositional attitude psychology by adding provisos saying that one also has certain propositional attitudes towards related contents, since any instance of

47. If A X_1 s that φ_1, \dots , and A X_n s that φ_n , then $\mathcal{C}_0(A$ Y_1 s that $\mathcal{C}_1(\varphi_1; \dots; \varphi_n); \dots; A$ Y_m s that $\mathcal{C}_m(\varphi_1; \dots; \varphi_n)$) if A Y_{m+1} s that $\mathcal{C}_{m+1}(\varphi_1; \dots; \varphi_n)$

is also an instance of 41.

We need to ensure that imposing these further closure conditions on perspectives is compatible with there being enough transparent perspectives to make intuitively unacceptable schemas not confusion-free valid. To illustrate the issue, suppose (i) that there is only one necessarily true proposition, (ii) that $\ulcorner s \urcorner \rightarrow \hat{\diamond}s \urcorner$ expresses it, for any mentalese sentence s , and (iii) that mentalese sentences are logically equivalent only if they are provably equivalent in first-order logic. Since $\ulcorner s \urcorner \rightarrow \hat{\diamond}s \urcorner$ is not first-order-equivalent to any tautology, the requirement that perspectives be closed under the formation of sentences using $\hat{\diamond}$ then implies that the only transparent perspective is the empty set, in which case all negated attitude ascriptions come out confusion-free valid!

The upshot is that verbalists who accept coarse-grained theories of propositions must deploy a correspondingly expansive notion of logical equivalence in characterizing confusion-free validity. This is to be expected, since in general treating more expressions (e.g., modal operators) as logical constants goes along with treating more sentences as logically equivalent.⁵¹

8 Introspection⁵²

An important class of principles of propositional attitude psychology that we haven't yet considered are those concerning people's attitudes about their own attitudes. For example, consider

48. If A believes that φ , then A believes that they believe that φ .

The question of how the confusion-free validity of such principles relates to generalizations about mentalese is complicated by the context-sensitivity of the embedded attitude verbs. We will now explain how to apply the strategy from the previous section to schemas like 48. The proposal will fit the usual pattern of combining a constraint on perspectives with a generalization about mentalese to secure the confusion-free validity of such principles.

Our proposal relies on a way of thinking about mentalese that can be illustrated by an analogy with public language. On one way of thinking about words, you can get a pretty good estimate of the number of words you have typed in a word processor by using its word-count function. But it is contentious whether this way of thinking about words is the one relevant for semantic theorizing. For example, in cases of vague language, some philosophers have wanted to say that a single utterance or inscription constitutes many sentences, each of which is used to assert a different proposition. On this view, in uttering "Harry is bald" one will thereby utter many sentences, \lceil Harry is bald \rceil_1 , \lceil Harry is bald \rceil_2 , etc., where bald $_1$, bald $_2$, etc. are distinct words, in the sense relevant to compositional semantics, each expressing slightly different properties. In this sense of 'word'(/'sentence'), one will have uttered not three(/one) but infinitely many of them. A similar situation arguably also arises in cases where someone mistakes two things for one. Erotium fails to realize that the twin Menaechmi she knows are different people. Were she to say "I think that's Menaechmus up ahead" (when in fact it's a cleverly disguised pig), she would arguably have uttered two sentences, containing different names respectively referring to the different twins, and thereby asserted the two propositions those sentences respectively express. This idea, that multiple representations can be constituted by a common vehicle, is perhaps even more plausible when we shift from public language to

⁵¹There are limits to what such logical ecumenicalism can achieve in very coarse-grained settings. For example, if there is only one logical truth, then the confusion-free validity of 35 and 40 implies the confusion-free validity of $\lceil A$ is not unsure whether $\varphi \rceil$ for any logically valid φ .

⁵²This section can be skipped without loss of continuity.

mentalese, since the fact that each of the relevant propositions is believed is a more secure datum than that each of them would be simultaneously asserted.

Our proposal appeals to a similar vehicle/representation distinction among mentalese correlates of propositional attitude verbs and sentences containing them. Where we previously spoke simply of mentalese sentences and their constituent expressions, we will now speak of *r*-sentences and their constituent *r*-expressions. It is these *r*-expressions that have compositional semantic values, which determine the meanings of the *r*-sentences containing them, and it is *r*-sentences that are members of perspectives. These representations are constituted by corresponding vehicles: an *r*-sentence is constituted by a *v*-sentence, its component *r*-expressions are respectively constituted by component *v*-expressions, and an *r*-sentence is in an agent's *X*-box just in case it is constituted by a *v*-sentence in their *X*-box. In the cases of expressions not containing mentalese analogues of attitude verbs, we will harmlessly identify *r*-expressions and *v*-expressions.

Every perspective π has a corresponding *v*-perspective: the set of *v*-sentences that constitute some *r*-sentence in π . It might seem natural to characterize *v*-perspectives by first characterizing the constitution relation between elementary *v*-expressions and *r*-expressions, using this relation to characterize the constitution relation between *r*-sentences and *v*-sentences, and finally saying which sets of *r*-sentences are perspectives, thereby characterizing which sets of *v*-sentences are *v*-perspectives. But in fact, our proposal is holistic in a way that makes such an approach unworkable. We will instead proceed in the opposite direction. We begin by characterizing *v*-perspectives directly, and then use this notion to characterize the *r*-sentence analogues of English attitude ascriptions and how they are constituted by *v*-sentences, at no point appealing to a relation of constitution between elementary expressions.

We assume that every attitude verb *X* has a corresponding *v*-expression \hat{X} in mentalese. We identify *v*-perspectives with sets of *v*-sentences that are closed under the formation of *v*-sentences using mentalese analogues of Boolean connectives, other relevant constants (e.g., $\hat{\diamond}$, discussed above), and prefixes $\ulcorner \iota \hat{X} \text{ that} \urcorner$, where ι is a mentalese first-person pronoun. We identify *r*-expressions corresponding to attitude verbs with ordered pairs $\langle \hat{X}, \rho \rangle$, where *X* is an attitude verb and ρ is a *v*-perspective. For any *v*-perspective ρ , let f_ρ be a function that maps each *v*-sentence φ to the *r*-sentences obtained by replacing all occurrences of \hat{X} in φ with $\langle \hat{X}, \rho \rangle$, and let $\rho^\dagger = \{f_\rho(\varphi) : \varphi \in \rho\}$. Note that ρ^\dagger is a set of *r*-sentences closed under the formation of *r*-sentences using mentalese analogues of Boolean connectives, other relevant constants, and prefixes $\ulcorner \iota \langle \hat{X}, \rho \rangle \text{ that} \urcorner$; it also contains no *r*-sentence containing the *r*-expression $\langle \hat{X}, \rho' \rangle$ for any $\rho' \neq \rho$. We can now characterize perspectives as follows: π is a *perspective* just in case $\pi = \rho^\dagger$ for some *v*-perspective ρ . We stipulate that $\langle \hat{X}, \rho \rangle$ expresses the relation of *X*-ing relative to ρ^\dagger . A *v*-sentence φ constitutes an *r*-sentence ψ just in case $\psi = f_\rho(\varphi)$ for some *v*-perspective ρ containing φ .⁵³

⁵³Constitution is defined in this way in part to ensure that agents with the *v*-sentence

We can now strengthen our earlier result relating the confusion-free validity of schemas to corresponding mentalese generalizations. Consider:

49. If A X_1 s that φ_1, \dots , and A X_n s that φ_n , then $\mathcal{A}(\varphi_1; \dots; \varphi_n)$,

where $\mathcal{A}(\varphi_1; \dots; \varphi_n)$ is to be replaced by any schematic sentence built from the schematic sentence letters $\varphi_1; \dots; \varphi_n$, Boolean connectives, attitude verbs, and the schematic letter A .⁵⁴ We can prove that an instance of this schema is confusion-free valid whenever the corresponding mentalese generalization about v -sentences is true and the relevant attitude boxes are closed under logical equivalence in the sense of 42.⁵⁵

$\ulcorner \iota$ believe $s \urcorner$ in their “believe”-box will not thereby have the r -sentence $\ulcorner \iota$ (believe, ρ) $f_\rho(s) \urcorner$ in their “believe”-box when $s \notin \rho$, since this r -sentence may be false even if s is in their “believe”-box.

Although the notion plays no essential role in our proposal, we can say that an occurrence o of a v -expression in the v -sentence φ constitutes an occurrence of the r -expression α just in case φ constitutes an r -sentence ψ in which α is the r -expression an occurrence of which in ψ corresponds o . Notice that occurrences o and o' of the same v -expression in different sentences may differ in whether they constitute an occurrence of a given r -expression. For example, if $\varphi \in \rho$ but $\varphi' \notin \rho$, then occurrences of \hat{X} in φ will constitute occurrences of (\hat{X}, ρ) , but occurrences of \hat{X} in φ' will not constitute occurrences of (\hat{X}, ρ) .

We are *not* assuming that if Φ expresses X -ing relative to π , then $\ulcorner A \Phi$ that $\varphi \urcorner$ is constituted by $\ulcorner A \hat{X}$ that $\varphi \urcorner$ – i.e., we are not assuming that (\hat{X}, ρ) is the *only* r -expression that expresses X -ing relative to ρ^\dagger . This is because, underlying your knowledge that Lois believes relative to π_s that Superman flies and that she doesn’t believe relative to π_c that he flies, there are a pair of r -sentences $\ulcorner L B_s$ that $\varphi \urcorner$ and $\ulcorner \neg(L B_c$ that $\varphi) \urcorner$ in your “know”-box, where B_s expresses believing relative to π_s and B_c expresses believing relative to π_c ; but if $B_s = \langle \text{“believe”, } \rho_s \rangle$ and $B_c = \langle \text{“believe”, } \rho_c \rangle$ (where $\pi_s = \rho_s^\dagger$ and $\pi_c = \rho_c^\dagger$), then this would require having both the v -sentence $\ulcorner L$ “believe” that $\varphi \urcorner$ and its negation in your “know”-box; but no v -sentence and its negation can simultaneously be in anyone’s “know”-box. This is not a problem, because the present proposal only concerns schemas about first-personal metacognition; it is not offered as a general story about the mentalese underpinnings of iterated attitude ascriptions, and in particular does not underpin non-confused people’s attitudes about identity-confused people’s attitudes.

⁵⁴We can also allow other constants to occur in complement clauses of attitude ascriptions in $\mathcal{A}(\varphi_1; \dots; \varphi_n)$, for the reasons discussed in connection to 45 above. Again, we need to ensure that imposing these further closure conditions on perspectives is compatible with there being enough transparent perspectives to invalidate intuitively unacceptable schemas. In particular, verbalists who accept a coarse-grained view of propositions according to which $\ulcorner \iota$ (“know”, ρ) $\varphi \urcorner \rightarrow \iota$ (“believe”, ρ) $\varphi \urcorner$ expresses the same proposition as its self-implication had better characterize confusion-free validity using a notion of ‘logical equivalence’ that relates that r -sentences and its self-implication. To achieve this, they might say that two r -sentences φ and ψ are logically equivalent just in case, for some v -sentence χ and attitudinally valid schema σ , χ constitutes $\ulcorner \varphi \leftrightarrow \psi \urcorner$ and χ is an instance of σ' , where σ' is the result of replacing every attitude verb X in σ with \hat{X} .

⁵⁵It is perhaps worth noting that Leibniz’s law makes trouble for negative introspection principles, according to which, when a rational person doesn’t have certain propositional attitudes, they will know(/be sure) that they don’t have those attitudes. Suppose Nicol is unsure whether Jahangir is identical to Jansher; in fact, they are distinct, and Nicol is not subject to any relevant identity confusion. She is sure that Jahangir drinks two glasses of milk a day (and knows this about herself) but she is not sure whether Jansher drinks two glasses of milk a day. This is a counterexample to negative introspection about being sure: Nicol should not be sure that she is not sure that Jansher drinks two glasses of milk a day, since she takes seriously the hypothesis that Jansher is Jahangir, in which case, by Leibniz’s law, her being sure that Jahangir drinks two glasses of milk a day implies that she is sure that Jansher drinks two glasses of milk a day.

9 Granularity

If verbalism is true, then propositional attitude verbs are importantly different from other families of context-sensitive expressions. For most families of context-sensitive expressions, holding fixed their parameter of context sensitivity is enough to make sense of our ordinary practice of using those expressions in systematic theorizing. By contrast, as we saw in section 6, our ordinary theorizing using propositional attitude verbs would be misguided if its aim were to uncover attitudinally valid principles. But such theorizing is not misguided, and its good standing is perfectly compatible with verbalism, since verbalists can reasonably take the aim of such theorizing to be the discovery of confusion-free valid principles. We have shown that, for a broad class of principles of propositional attitude psychology, their confusion-free validity is implied by the truth of corresponding generalizations about mental representations. We have done so without making any substantive assumptions about propositional granularity. Moreover, we have argued that confusion-free validity, though a weaker notion than attitudinal validity, is still a highly discriminating one, provided the notion of logical equivalence used to define it is understood in a reasonable way. And while questions of propositional granularity do bear on the good standing of principles like 32 and 38, which are not covered by the aforementioned theorems, this dependence holds for verbalists and clausalists alike.

Still, the question of how best to develop verbalism is not wholly separable from questions of propositional granularity. Suppose we accept a coarse-grained theory of propositions according to which, for any p , the proposition that $p \rightarrow p$ is the proposition that Superman is Superman. Now consider:

50. Lois does not believe that Superman is Clark.

This is exactly the sort of sentence we naturally reach for to characterize Lois's state of mind. What can we say about the perspective π that parametrizes its most salient interpretation? As an idealization, let us assume that Lois has every mentalese sentence of the form $\lceil s \rightarrow s \rceil$ in her "believe"-box. Then 50 is true on its most salient interpretation only if there is no $s \in \pi$ that Lois accepts and means that Superman is Clark, and so no $s \in \pi$ that Lois accepts and means that Superman is Superman (by Leibniz's law), and so no $s \in \pi$ that Lois accepts and means that $p \rightarrow p$, for any p (by our granularity assumption), and so no $s \in \pi$ of the form $\lceil s \rightarrow s \rceil$ that Lois accepts, and so no $s \in \pi$ of that form (by our idealizing assumption), in which case π is empty (since it is a closed set of mentalese sentences). It follows that "Necessarily, no one believes anything" is also true on this interpretation. But this seems like a terrible result. Any sense of "believe" in which it expresses an impossible attitude is not one worth denying of Lois. Moreover, even without the idealization about having every self-implication in one's "believe"-box, our granularity assumption together with the confusion-free validity of 29 implies the confusion-free validity of the unattractive schema "If $a = b$ and A does not believe that $a = b$, then A does not believe anything".

Proponents of such coarse-grained theories of propositions could simply deny that 50 is true on its most salient interpretation. The idea might be that ordinary speakers are tacitly committed to a false theory of structured propositions, which leads them to make false negative attitude ascriptions.⁵⁶ While this response is reminiscent of neo-Russellianism in the way that it convicts ordinary speakers of error, and in particular of the diagnosis by Braun (1998) of such purported errors in terms of an alleged tacit commitment to a mistakenly fine-grained theory of mental contents, it does not impute errors to ordinary speakers in cases like THELMA, and its combination with a verbalist treatment of such cases is not obviously unprincipled.

Alternatively, proponents of such coarse-grained views could relax our working hypothesis that all perspectives are closed. For example, if π were the singleton set containing the mentalese analogue of “Superman is Clark”, then 50 will be true on its most salient interpretation without “believe” expressing a trivial relation on that interpretation. Allowing for such interpretations means giving up the attitudinal validity of sentences like “If Luthor fears that Superman is Clark, then Luthor hopes that Superman is not Clark”. But it does not require giving up the confusion-free validity of such sentences, so long as closure is now added as an extra condition on the relevant perspectives in the definition of confusion-free validity. Having modified the definition of confusion-free validity in this way, allowing for the existence of non-closed perspectives does nothing to disrupt any of our earlier results about the connection between confusion-free validity and mentalese generalizations.⁵⁷

The above antinomy relied on the assumption that, for all p and q , the proposition that $p \rightarrow p$ is identical to the proposition that $q \rightarrow q$. But merely accepting a view of propositional granularity that rejects this principle doesn’t suffice to avoid all versions of the antinomy. For example, consider the view that there are many tautologies with different subject matters, but necessarily equivalent propositions with the same subject matter are identical and $p \rightarrow p$ has the same subject matter as p for any p .⁵⁸ This view entails that the proposition that Superman is Superman is identical to the proposition that $p \rightarrow p$ for any p with the same subject matter as it. The above mode of argument then only establishes that “believe” expresses relative to π a relation that is impossible to bear to any propositions with that subject matter. While this may be less objectionable than “believe” expressing an impossible relation, it still makes “believes that Superman is Clark” express an impossible property on the relevant interpretation, which is arguably a bad enough result.

Reflecting on this example, one might think that the only views of proposi-

⁵⁶That we are pre-theoretically committed to a false structured theory of propositions is one possible moral of the so-called Russell-Myhill paradox; see Russell (1903, Appendix B), Uzquiano (2015) and Goodman (2017).

⁵⁷A different option that avoids trivializing the relation expressed by “believe” on any true interpretation of 50 but doesn’t disrupt the attitudinal validity of principles like 20 would be to identify perspectives with relations between people and mentalese sentences such that, necessarily, for any person, the set of sentences to which they are so related is closed, modifying the truth conditions from section 5 accordingly (i.e., letting $\alpha_X = \lambda x \pi p. x$ bears π to $s \wedge \beta_X(x, s) \wedge s$ means p).

⁵⁸See Goodman (forthcoming) for more on such views.

tional granularity for which the above antinomy doesn't arise in some form or other are those according to which propositions are structured in a way that reflects the syntactic structure of the sentences we use to express them. But that would be a mistake. For example, consider the view of [Dorr \(2016\)](#), according to which, informally, no proposition is a non-trivial ingredient of itself and, in particular, no proposition is identical to its self-implication. This view reconciles the truth of [50](#) on its most natural interpretation with the natural hypothesis that this interpretation is parametrized by a closed set of mentalese sentences containing the mentalese analogue of "Superman = Clark", and hence with [50](#) expressing a contingent proposition with a sensible modal profile. But Dorr's view also implies that propositions are not structured in the aforementioned way, since it implies that, for any logical truth φ , $\ulcorner \varphi \rightarrow \varphi \urcorner$ expresses the same proposition as $\ulcorner \neg\varphi \rightarrow \varphi \urcorner$.

The above discussion has important implications for the so-called 'problem of logical omniscience'. It is easy to imagine cases naturally described using sentences of the form $\ulcorner A \text{ knows that } \varphi \text{ but } A \text{ does not know that } \psi \urcorner$ where φ and ψ are logically equivalent. For example, consider

51. Gödel knew that $ZF2$ but he did not know that $ZF2 \wedge FLT$,

where $ZF2$ stands for the conjunction of the axioms of second-order Zermelo-Fraenkel set theory and FLT stands for a formalization of Fermat's last theorem. The problem of logical omniscience, as we will understand it, is an objection to the view that logically equivalent sentences express identical propositions, on the grounds that, e.g., $ZF2$ and $\ulcorner ZF2 \wedge FLT \urcorner$ must express different propositions in order that [51](#) have the true reading it seems to have.

The natural rejoinder for proponents of coarse-grained views is to assimilate sentences like [51](#) to sentences like [1](#): they are true on their most natural interpretations, but these interpretations are non-uniform. For verbalists, this means that they involve non-uniformly interpreting "know". To implement this strategy, we need to establish the existence of a pair of perspectives the first of which contains the mentalese analogue s of $ZF2$ and the second of which contains the mentalese analogue t of $\ulcorner ZF2 \wedge FLT \urcorner$ and does not contain any sentence with the same meaning that was in Gödel's "know"-box. But in cases where s and t are not merely logically equivalent but are logical truths, the coarse-grained view implies that t expresses the same proposition as $\ulcorner t \rightarrow t \urcorner$. Since this sentence presumably was in Gödel's "know"-box, the existence of the second sort of perspective is then inconsistent with the assumption that perspectives are closed. And yet, as in the case of [50](#), two replies are available: denying that the relevant negated knowledge ascriptions with logically true complements are true on their most natural interpretations, or adopting a more liberal theory of perspectives on which they need not be closed sets. Either way, the problem of logical omniscience, understood as a challenge to make sense of the truth of sentences like [51](#) on their most natural interpretations, does not pose a new challenge for verbalist proponents of coarse-grained theories of propositions beyond that posed by [50](#).

Moreover, verbalists who hold that two sentences never express the same proposition when one is a non-trivial constituent of the other have no trouble making the relevant negated knowledge ascriptions have true interpretations parametrized by closed perspectives, just as they had no such problem with 50. As the consistency of Dorr’s view shows, this assumption about propositional granularity is compatible with propositions being significantly less fine-grained than they are according to traditional structured views. Contextualism is doing the heavy lifting here – the granularity assumption is simply greasing the wheels, allowing the technology of non-uniform interpretation to operate more smoothly. It is notable that the strongest advocates of structured propositions are usually motivated by the problem of logical omniscience (e.g. Soames (1987)) and yet they accept the validity of Leibniz’s law. Verbalism both fits naturally with their acceptance of Leibniz’s law and affords the resources to resist structured theories of propositions.

The situation is very different for clausalists. It is entirely unclear how the contextualist precedent of 1 might help them reconcile the truth of sentences like 51 on their most natural interpretations with propositions not being structured in the manner of the sentences we use to express them. Assuming a compositional semantics, the context-sensitivity of 1 is attributable to context-sensitivity in its elementary constituents. As discussed in section 1, clausalists should see this as a matter of the words “Clark” and “Superman” having semantic values associated with different ways of thinking about Superman. But the problem of logical omniscience can be raised using sentences of pure logic, and in such cases there doesn’t seem to be any relevant multiplicity in relevant ways of thinking about the corresponding logical operations. It seems bizarre that the truth of 51 on its most natural interpretation should have something to do with Gödel having more than one way of thinking about, e.g., negation, conjunction, or the number 0.⁵⁹ Verbalists can respond to the problem of logical omniscience by assimilating cognitive accomplishment in logic to the clearing up of identity confusion. Clausalists cannot.⁶⁰

⁵⁹The idea isn’t just strange, but also hard to model. The reason is that, assuming the validity of propositional identifications corresponding to logical equivalences, any shift in interpretation of the logical vocabulary must leave those equivalences fixed. This will require logical expressions to shift their interpretation in a coordinated way since, e.g., the identification of logically equivalent entities pins down conjunction uniquely in terms of its relation to negation and disjunction. But ordinary models of higher-order languages that allow us to talk about the interpretation of logical connectives and also validate the relevant principles of coarseness of grain do not have the feature that there is an interesting (in the sense of non-trivial and not glaringly unintended) range of families of interpretations we could assign to the logical vocabulary compatibly with the identification of propositions expressed by logically equivalent formulas.

⁶⁰It is worth noting that clausalists, in rejecting Leibniz’s law, have reason to give up standard principles of higher-order logic in a way that has the side-effect of insulating their view from paradoxes associated with structured propositions; for details, see Goodman (2017). The problem of logical omniscience therefore seems most acute for views that reject Leibniz’s law without relinquishing other standard principles of higher-order logic, such as the views developed in Caie et al. (2017).

References

- Maria Aloni. Individual concepts in modal predicate logic. *Journal of Philosophical Logic*, 34(1):1–64, 2005.
- Andrew Bacon and Jeffrey Sanford Russell. The logic of opacity. *Philosophy and Phenomenological Research*, 2017.
- David Braun. Understanding belief reports. *The Philosophical Review*, 107(4):555–595, 1998.
- David Braun. Russellianism and psychological generalizations. *Noûs*, 34(2):203–236, 2000.
- David Braun. The objects of belief and credence. *Mind*, 125(498):469–497, 2016.
- Michael Caie, Jeremy Goodman, and Harvey Lederman. Classical opacity. Manuscript, 2017.
- Mark Crimmins. *Talk about beliefs*. MIT Press, 1992.
- Mark Crimmins and John Perry. The prince and the phone booth: Reporting puzzling beliefs. *The Journal of Philosophy*, 86(12):685–711, 1989.
- Cian Dorr. Transparency and the context-sensitivity of attitude reports. In Manuel Garcia-Carpintero and Genoveva Martí, editors, *Empty Representations: Reference and Non-existence*, pages 25–66. Oxford University Press, 2014.
- Cian Dorr. To be F is to be G. *Philosophical Perspectives*, 2016.
- Jerry A Fodor. *LOT 2: The language of thought revisited*. Oxford University Press on Demand, 2008.
- Gottlob Frege. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 100:25–50, 1892.
- Ephraim Glick. Cross-contextual utterances. Unpublished Typescript, 2017.
- Jeremy Goodman. Reality is not structured. *Analysis*, 77(1):43–53, 2017.
- Jeremy Goodman. Agglomerative algebras. *Journal of Philosophical Logic*, forthcoming.
- Jeremy Goodman and Harvey Lederman. Understanding. in preparation.
- Irene Heim and Angelika Kratzer. *Semantics in generative grammar*. Blackwell, 1998.
- David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes from Kaplan*, pages 481–563. Oxford: Oxford University Press, 1989.
- Saul A Kripke. A puzzle about belief. In Avishai Margalit, editor, *Meaning and use*, pages 239–283. D. Reidel, Dordrecht, 1979.
- David C Makinson. The paradox of the preface. *Analysis*, 25(6):205–207, 1965.
- Stefano Predelli. Who’s afraid of substitutivity? *Noûs*, 34(3):455–467, 2000.

- Mark Richard. *Propositional attitudes: An essay on thoughts and how we ascribe them*. Cambridge University Press, 1990.
- Mark Richard. Attitudes in context. *Linguistics and Philosophy*, 16(2):123–148, 1993.
- Bertrand Russell. *The Principles of Mathematics*. Cambridge University Press, 1903.
- Nathan U Salmon. *Frege's puzzle*. MIT Press, 1986.
- Jennifer M Saul. *Simple sentences, substitution, and intuitions*. OUP Oxford, 2010.
- Stephen Schiffer. Naming and knowing. In A. French Peter, E. Uehling Theodore, Howard Jr, and K. Wettstein, editors, *Midwest Studies in Philosophy II: Contemporary Perspectives in the Philosophy of Language*, pages 28–41. University of Minnesota Press, 1979.
- Stephen Schiffer. Belief ascription. *The Journal of Philosophy*, 89(10):499–521, 1992.
- Scott Soames. Incomplete definite descriptions. *Notre Dame Journal of Formal Logic*, 27(3):349–375, 1986.
- Scott Soames. Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, 15(1):47–87, 1987.
- Ernest Sosa. Propositional attitudes de dicto and de re. *Journal of Philosophy*, 67(21): 883–896, 1970.
- Gabriel Uzquiano. A neglected resolution of russell's paradox of propositions. *The Review of Symbolic Logic*, 8(2):328–344, 2015.