ARTICLE

Noûs

Perspectivism

Jeremy Goodman[1] | **Harvey Lederman[2],***

[1]University of Southern California

[2]Princeton University

Correspondence
Harvey Lederman, Princeton University.
Email: harvey.lederman@princeton.edu

*The authors contributed equally to this paper.

Abstract

Consider the sentence "Lois knows that Superman flies, but she doesn't know that Clark flies". In this paper we defend a Millian contextualist semantics for propositional attitude ascriptions, according to which ordinary uses of this sentence are true but involve a mid-sentence shift in context. Absent any constraints on the relevant parameters of context sensitivity, such a semantics would be untenable: it would undermine the good standing of systematic theorizing about the propositional attitudes, trivializing many of the central questions of epistemology, the philosophy of mind, and the philosophy of action. In response to this problem, we prove a series of tenability results. We show that, given certain constraints on the parameters of context sensitivity, there is a broad class of principles of propositional attitude psychology whose good standing follows from corresponding claims about people's mental representations. But these constraints also have some surprising consequences: they are jointly incompatible with coarse-grained theories of propositions, and they are in tension with a natural picture of how speakers and hearers coordinate on the interpretation of attitude ascriptions. In light of these consequences we explore different ways in which the contextualist picture might be developed, and argue that our preferred way compares favorably with Fregeanism and neo-Russellianism.

1 | INTRODUCTION

Lois doesn't realize that the superhero she knows by the name "Superman" is the same person as the reporter she knows by the name "Clark". It's natural to describe her state of mind as follows:

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

But it is puzzling how this sentence could be true. After all, Superman *is* Clark. Shouldn't it follow that, if Lois knows that Superman flies, she thereby knows that Clark flies?

Not according to Fregeans. According to them, the semantic contribution of "Superman" in 1 is not its ordinary referent, Superman, but rather a *sense* of him. The semantic contribution of "Clark" is a different sense of him, and this difference allows the sentence to come out true.

Neo-Russellians, by contrast, think that 1 is false, despite being a natural thing to say. According to them, different names of the same person make the same semantic contribution to the sentences in which they occur. So 1 must be false, since its first conjunct affirms precisely what its second conjunct denies.

We will be exploring a third, contextualist option. On this view, the proposition that Superman flies is the proposition that Clark flies, but Lois knows this proposition in some ways and not in others. Moreover, ordinary uses of 1 involve equivocation: in such uses, the kind of knowing ascribed by "Lois knows that Superman flies" is different from the kind of knowing denied by "she doesn't know that Clark flies". Fregeans are right that ordinary uses of sentences like 1 are true and neo-Russellians are right about the semantics of names. Contextualism promises the best of both worlds.[1]

After presenting a version of this contextualist proposal in section 2, we explain in section 5 what we take to be the central challenge facing it: to make sense of the good standing of the kind of systematic theorizing about the propositional attitudes that is pervasive in epistemology, the philosophy of mind, and the philosophy of action. (Sections 3 and 4 present a framework for theorizing about context-sensitivity and about the good standing of principles of propositional attitude psychology; they may be skipped by the impatient reader.) In response to this challenge, in sections 5-8 we prove a series of results showing that, given certain simple and natural assumptions, the good standing of a broad class of principles of propositional attitude psychology can be secured by the truth of corresponding generalizations about people's mental representations. In section 9 we draw out some surprising consequences of these assumptions and explore the resulting tradeoffs, which suggest three alternative ways of developing the contextualist framework. In section 10 we argue for one of these alternatives on the grounds that it fits with an attractive picture of how speakers and hearers resolve the context-sensitivity of attitude ascriptions. In sections 11 and 12 we argue that the contextualist proposal also compares favorably with neo-Russellianism and with Fregeanism.

2 | MENTAL REPRESENTATIONS

We start with two sets of background assumptions: the first is a toy model of our cognitive psychology; the second connects the model to the context-sensitivity of propositional attitude ascriptions.

Here is the toy model. We have a language of thought, *mentalese*. It is an interpreted language (i.e., its expressions have meanings), its sentences are strings of symbols, and its lexicon includes symbols corresponding to truth-functional connectives. So, for example, if s and t are mentalese sentences that respectively mean that p and that q, and $\hat{\wedge}$ is the mentalese symbol corresponding to conjunction, then $\ulcorner s \hat{\wedge} t \urcorner$ is a mentalese sentence meaning that p and q. For every propositional attitude verb X, there is

a corresponding relation between mentalese sentences s and people a, which we will pronounce "s is in a's X-box". For example, in the case of Lois, a mentalese correlate s_0 of "Superman flies" is in her "know"-box, but her mentalese correlate c_0 of "Clark flies" is not in her "know"-box; indeed, its negation $\ulcorner \hat{\neg} c_0 \urcorner$ is in her "believe"-box.[2] In general, the idea is that cases of 'identity confusion' like Lois's are ones where a person is insensitive to the co-reference of two expressions in their mental lexicon.[3] This picture can be accepted by Fregeans, neo-Russellians, and contextualists alike. (Note that in adopting the jargon of "mentalese" and "boxes" we are not committing ourselves to any of the further doctrines of Fodor's (1975) influential 'language of thought hypothesis' beyond those appealed to above.)

Let's now turn to the proposed context-sensitivity of propositional attitude ascriptions. We will assume that, for every propositional attitude verb X, there is a corresponding three-place relation X^* between people, propositions, and third entities of some sort – which we will call *perspectives* – such that, on a given occasion of use, the sentence $\ulcorner A \ X \text{s that } \varphi \urcorner$ expresses the proposition that X^* relates a, p and π, where a is the referent of A, p is the proposition expressed by φ, and π is the perspective supplied by the context of use.[4] Call this abstract proposal *perspectivism*. In this paper we will explore a version of perspectivism according to which perspectives are sets of mentalese sentences, and X^* is the relation that holds between a person a, proposition p, and set of mentalese sentences π just in case some s in π both means p and is in a's X-box.[5]

As advertised, perspectivism allows ordinary uses of 1 to come out true, provided that such uses involve a mid-sentence context-shift. In particular, such uses will be true so long as the perspective supplied by context for its first conjunct contains s_0 (or any other mentalese sentence that both means that Superman flies and is in Lois's "know"-box) and the perspective supplied by context for its second conjunct contains neither c_0 nor any other mentalese sentence that both means that Clark flies and is in Lois's "know"-box.

3 | CONTEXT SENSITIVITY

In this section we first review a standard way of giving a compositional semantic theory for a language with context-sensitive expressions, and then turn to the question of what determines which perspective is supplied by context for a given use of an attitude ascription.

We will be working within the following broadly Kaplanian framework. *Semantic values* (what Kaplan (1989) called *contents*) are assigned, not to expressions, but to particular uses of them. The semantic value of a use of a complex expression is determined compositionally from the semantic values of the corresponding uses of its immediate constituents. So to determine the semantic value of a use of a complex expression, it suffices to determine the semantic values of the corresponding uses of its elementary constituents. These are determined by context as follows. Every elementary expression has a *conventional meaning* (what Kaplan called its *character*), and every use of an elementary expression is *in* a particular context. The conventional meaning of an elementary expression is identified with a function from contexts to semantic values, and the semantic value of a given use of an elementary expression is the result of applying its conventional meaning to the context that it is in. We will assume that the semantic values of uses of declarative sentences are propositions, which we say those uses *express*.[6] A use of a sentence (including as a subsentence of a larger sentence) is *true(/false)* just in case it expresses a true(/false) proposition, and a person who uses a sentence unembedded (i.e., not as a subsentence of a larger sentence) thereby asserts whatever proposition is expressed by that use.

Let *Millianism* be the view that the semantic value of a use of a name is the relevant individual whose name it is. Given uncontroversial background assumptions, Millianism implies that a true use

of 1 must involve a mid-sentence context-shift – i.e., it implies that at least two elementary expressions of 1 are used in different contexts.[7] In addition to perspectivism, we will assume Millianism throughout this paper. Perspectivism and Millianism are a natural pair, since perspectivism accounts for the truth of ordinary uses of sentences like 1 without requiring any difference in the semantic values of uses of different names of the same person. (Conversely, anyone antecedently committed to Millianism could appeal to the truth of ordinary readings of 1 to motivate perspectivism.[8])

The existence of mid-sentence context-shifts in English should not be controversial. For example, the sentence

2. Nour is sitting now, and Nour is not sitting now.

can be true when used as Nour is standing up. The truth of such a use requires a mid-sentence context-shift for the same reason that a true use of 1 requires a mid-sentence context-shift if Millianism is true.

In characterizing perspectivism, and in explaining our preferred treatment of 1, we appealed to the notion of a perspective being 'supplied by' context. We will now consider two strategies for making this notion more precise: *verbalism* and *hidden indexicalism*.[9] (The rest of this paragraph can be skipped without loss of continuity.) Verbalists think that the context-sensitivity of propositional attitude ascriptions that perspectivism posits is due to attitude verbs' conventional meanings delivering different semantic values for different contexts. Hidden-indexicalists, by contrast, think that attitude verbs' semantic values are context-invariant, but that attitude verbs take an additional phonologically null third argument (in addition to a grammatical subject and a complement clause), and it is these 'hidden indexical' arguments whose different semantic values in different contexts are responsible for the relevant context-sensitivity of attitude ascriptions. Verbalists think there is a unique function f from contexts in which attitude verbs are used to perspectives such that the semantic value of a use of an attitude verb X in a context c is the two-place relation that holds between a person and a proposition just in case they are related by X^* to $f(c)$. Hidden indexicalists, by contrast, think the semantic value of any use of an attitude verb X is the three-place relation X^*, regardless of the context it is in; but they think every such use is accompanied by a hidden indexical, and that there is a unique function f from contexts hidden indexicals are in to perspectives such that a hidden indexical's semantic value in a context c is $f(c)$. We will remain neutral on the question of whether perspectivists should be verbalists or hidden-indexicalists, since for our purposes all we need is the notion of a perspective supplied by context, and both views allow us to make sense of that notion in terms of the aforementioned functions from contexts to perspectives.

What determines which perspective is supplied by context for a given use of an attitude ascription?[10] This is a hard question. Consider the following case, modified from Schiffer (1979):

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

The accomplice's uses of "she knew Shorty limped" and "she didn't know Shorty limped" must be in different contexts (assuming he speaks truly). Perspectivists will think that these contexts supply different perspectives, for the same reason that they think that the contexts of the two attitude ascriptions in typical uses of 1 supply different perspectives. What determines which perspectives these are? It cannot be merely the speaker's choice of words in the complement clause of the relevant attitude ascriptions, since these are the same for both uses. Rather, it seems that it has to do with a difference in which of Thelma's ways of thinking about Shorty are salient to the conversational participants when those sentences are uttered. This shift in salience is achieved, in turn, by a combination of conversational cues and pragmatic factors. (We will not speculate about the details of how this works, since the phenomenon in question is one that any theory of attitude ascriptions and identity confusion needs to explain, not just perspectivist ones.) In section 10 we will sketch some ways in which the perspective that parametrizes a given attitude ascription might be partially determined by the conversationally salient ways the subject of the ascription has of thinking about the objects of her confusion.

The Thelma/Shorty example helps to illustrate two other important facts about attitude ascriptions and identity confusion. The first is that examples like 1 are quite special. They are convenient because we associate the different names in the two complement clauses with different ways the subject has of thinking about the object of her confusion. But typical descriptions of people's identity confusion are more like Shorty's accomplice's description of Thelma. Typically, when we give such descriptions, we don't have multiple names for the object of the person's confusion, let alone multiple names that we robustly associate with the confused person's different ways of thinking about that object (in contrast to how we robustly associate "Superman" and "Clark" with different ways Lois has of thinking about Superman).

Second, the Thelma/Shorty case shows that Fregeans should agree with perspectivists that identity confusion is associated with a distinctive dimension of context-sensitivity in attitude ascriptions. They should say that the name "Shorty" denotes different senses in its two uses in the accomplice's discourse. Indeed, Fregeans should agree with perspectivists that we often describe cases of identity confusion using sentences that cannot be true without a *mid-sentence* shift in context. For example, the sentence

3. Peter knows that Paderewski (the pianist) is musically talented, but doesn't know that Paderewski (the politician) is musically talented.

is a natural way to describe the famous case from Kripke (1979). Fregeans should say that the name "Paderewski" is context-sensitive: its two uses have different senses of Paderewski as their semantic values, and so must be in different contexts.[11]

4 | UNEQUIVOCAL TRUTH

The following schema seems like a good principle of propositional attitude psychology:[12]

4. If A knows that φ, then A believes that φ.

But what does the 'goodness' of such a principle amount to? Being a schema, any good status it has it inherits from its instances. These instances are context-sensitive English sentences. So we cannot say that the schema is good only insofar as all of its instances are true: truth is a property not of sentences but of their uses. Nor can we say that the schema is good only insofar as every use of any of its instances must be true, since that condition is far too demanding. For consider:

The Law of Non-Contradiction: It is not that case that φ and not-φ.

The negation of 2 is an instance of this schema. It has untrue uses, since 2 has true uses. So not every use of any instance of the Law of Non-Contradiction is true. But this fact clearly does not undermine the good standing of that schema.

Let the semantic value of an expression *e relative to* a context *c* be the result of composing the values for *c* of the conventional meanings of *e*'s elementary constituents. Let a sentence be true(/false) relative to *c* if the proposition it expresses relative to *c* is true(/false). Call a sentence *unequivocally true* if it is true relative to every context; call a schema *good* if all of its instances are unequivocally true. The Law of Non-Contradiction is in this sense good, despite having instances with uses that are not true.

Now consider the schema:

Substitution: If $A = B$ and φ, then $\varphi[B/A]$.[13]

A use of an identity statement involving a pair of names is true only if both names are used to name the same individual. Given Millianism, this individual is also the semantic value of both name-uses. It follows that an identity statement involving a pair of names is true relative to a context in which both names are used only if both names have the same semantic value relative to that context. And this is true not only of contexts in which each name happens to be used: in general, if Millianism is true, then an identity statement involving a pair of names is true relative to a context only if both names have the same semantic value relative to that context. So, if Millianism is true, then Substitution is good.[14]

Goodness is not the only interesting status we might wish to claim for these schemas. They are arguably valid too. But 'validity' is a contested notion, and in the case of principles of propositional attitude psychology like 4, it is controversial whether they are valid, even if they are good. Some will claim, for example, that their validity is precluded by the 'non-logicality' of propositional attitude verbs. We will focus on goodness in what follows in part to sidestep these issues. The challenges for perspectivism we will be considering concern the mere goodness of the schemas of interest, and so *a fortiori* are challenges for more demanding statuses one might consider, like being valid, or (if this comes to something different) having only instances that express *necessary* truths relative to every context (being 'modally good' for short). Moreover, the results we will prove below concerning goodness can be mechanically generalized to establish parallel results about modal goodness; see footnote 17.[15]

5 | PROPOSITIONAL ATTITUDE PSYCHOLOGY

Why care about the good standing of systematic theorizing about the propositional attitudes? Because such principles figure centrally in epistemology, action theory, and the philosophy of mind. Is believing a proposition compatible with believing its negation? Does knowing a proposition imply knowing that you know it? Do people try to do what they believe to be necessary to realize their most desired ends, at least when they believe it to be something they can do? If a person intends to do something must they think that they will do it? Is to believe that something is probable just to have high subjective confidence in it? Indeed, our initial example 4 – does knowledge imply belief? – has been a central question of epistemology at least since Plato. All of these principles are controversial, and perhaps hold at best conditional on certain idealizing assumptions. But, as we will see, perspectivists face a challenge to make sense of the good standing of such theorizing even given quite powerful idealizing assumptions. It is to this challenge that we will now turn.[16]

We start by asking: is 4 good? Consider the corresponding generalization about mentalese sentences:

5. For all people a and mentalese sentences s, if s is in a's "know"-box, then s is in a's "believe"-box.

If perspectivism is correct and 5 is true, then 4 is good. This example is an instance of the following more general result:

Proposition 1. ⌜*If* A *X*s *that* φ, *then* A *Y*s *that* φ⌝ *is good if, for all people a and mentalese sentences s, if s is in a's X-box, then s is in a's Y-box.*

Proof. In what follows let $\Pi(c)$ denote the perspective supplied by c and $[\![e]\!]^c$ be the semantic value of e relative to c. Suppose ⌜A Xs that φ⌝ is true relative to c. So some mentalese sentence s in $\Pi(c)$ both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s X-box. So it is also in their Y-box. So ⌜A Ys that φ⌝ is true relative to c.[17] □

This result raises the question: under what conditions should perspectivists accept the goodness of a schematic principle of propositional attitude psychology provided they accept the corresponding generalization about people's mental representations? This question will be our focus in what follows.[18] We will prove a number of generalizations of Proposition 1, and present counterexamples to certain further generalizations of it. We will then consider how these results bear on whether perspectivists can meet the challenge of vindicating the good standing of propositional attitude psychology.

Perspectivists might hope to generalize Proposition 1 to cover more complex entailment relations between propositional attitudes. Consider the schemas:

6. If A suspects that φ, then A both thinks that φ and is unsure whether φ.

7. If A perceives that φ, then A either sees that φ or can hear that φ or . . . .

and the corresponding mentalese generalizations

8. For all people a and mentalese sentences s, if s is in a's "suspect"-box, then s is both in a's "think"-box and in a's "be unsure"-box.

9. For all people a and mentalese sentences s, if s is in a's "perceive"-box, then s is either in a's "see"-box or in a's "can hear"-box or. . . .

Should perspectivists who think that 8 and 9 are true think that 6 and 7 are good? Proposition 1 is silent on this question, since it does not cover conditionals whose consequents are conjunctions or disjunctions. But it can be generalized to cover such principles. Call σ a *positive combination* of $\{\varphi_1, \ldots, \varphi_n\}$ just in case it is either a member of that set or can be constructed from its members using only conjunction and disjunction.

Proposition 2. ⌜*If* A *X*s *that* φ, *then* σ⌝ *is good if the corresponding generalization about mentalese sentences is true, where* σ *is positive combination of* $\{$⌜A Y_1s *that* φ⌝$, \ldots,$ ⌜A Y_ns *that* φ⌝$\}$.

A precise definition of the notion of correspondence and a proof of this result are in a footnote.[19]

6 | INCOMPATIBLE ATTITUDES

But Proposition 2 does not cover all entailment relations between propositional attitudes that we might be interested in. Consider:

10. If A is unsure whether φ, then A is not sure that φ,

and the corresponding generalization

11. For all people a and mentalese sentences s, if s is in a's "be unsure"-box, then s is not in a's "be sure"-box.

Should perspectivists who think that 11 is true think that 10 is good? Proposition 2 is silent on this question, since it does not cover conditionals with negated consequents.

In fact, perspectivism does *not* imply that 10 is good if 11 is true. Suppose Lois has a mentalese correlate c_0 of the English sentence "Clark flies" in her "be unsure"-box. 11 implies that c_0 is not in her "be sure"-box. Suppose further that a mentalese correlate s_0 of the English sentence "Superman flies" *is* in her "be sure"-box. By Substitution, the proposition that Superman flies is the proposition that Clark flies, so s_0 and c_0 mean the same thing.[20] "If Lois is unsure whether Clark flies, then Lois is not sure that Clark flies" is therefore false relative to any context that supplies a perspective that contains both c_0 and s_0. Perspectivism is compatible with there being such perspectives. So 11 and perspectivism are compatible with this instance of 10 being false relative to some contexts. So they do not imply that 10 is good.

The above argument turned on the possibility of a context supplying a perspective containing distinct mentalese sentences that express the same proposition. Call a perspective *injective* just in case any distinct mentalese sentences it contains express distinct propositions. Although perspectivism is compatible with the existence of non-injective perspectives, nothing we have said so far implies the existence of any such perspectives. Let *injectivity* be the hypothesis that all perspectives are injective. Perspectivism and injectivity *do* imply that 10 is good if 11 is true, and this is again an instance of a more general result:

Proposition 3. *Given injectivity, ⌜If A Xs that φ, then A does not Y that φ⌝ is good if, for all people a and mentalese sentences s, if s is in a's X-box, then s is not in a's Y-box.*

Proof. Suppose ⌜A Xs that φ⌝ is true relative to c. So some mentalese sentence s in $\Pi(c)$ both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s X-box. So it is not in their Y-box. By injectivity, s is the only sentence in $\Pi(c)$ that means $[\![\varphi]\!]^c$. So there is no s' in $\Pi(c)$ that both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s Y-box. So ⌜A does not Y that φ⌝ is true relative to c. □

Not all entailment relations between propositional attitudes we might be interested in are covered by Propositions 2 and 3. Consider:

12. If A is unsure whether φ, then (if A desires that φ, A hopes that φ),

and the corresponding mentalese generalization

13. For all people a and mentalese sentences s, if s is in a's "be unsure"-box, then if s is in a's "desire"-box, s is in a's "hope"-box.

Should perspectivists who think that 13 is true think that 12 is good? Both Propositions 2 and 3 are silent on this question, since neither covers conditionals whose consequents are themselves conditionals. But those results can be generalized to cover such principles. Call σ a *Boolean combination* of $\{\varphi_1, \ldots, \varphi_n\}$ just in case it is either one of $\varphi_1, \ldots, \varphi_n$ or can be constructed from them using only negation, conjunction, disjunction and material implication.

Proposition 4. *Given injectivity, ⌜If A Xs that φ, then σ⌝ is good if the corresponding generalization about mentalese sentences is true, where σ is a Boolean combination of $\{⌜A\ Y_1 s\ that\ \varphi⌝, \ldots, ⌜A\ Y_n s\ that\ \varphi⌝\}$.*

The proof of this result is in a footnote.[21]

7 | COMPLEX COMPLEMENTS

Our results so far only cover principles about the pattern of attitudes a person has towards a given proposition. They do not cover principles like:

14. If A doubts that φ, then A thinks that $\neg\varphi$;
15. If A believes that φ and A believes that ψ, then A believes that $\varphi \wedge \psi$;
16. If A believes that φ and A believes that ψ, then A believes that $\varphi \vee \psi$;
17. If A believes that φ and A believes that ψ, then A believes that $\varphi \rightarrow \psi$.

And perspectivism does not imply that such schemas are good, even if the corresponding mentalese generalizations are true. Consider the generalization corresponding to 14:

18. For all people a and mentalese sentences s, if s is in a's "doubt"-box, then $\ulcorner \hat{\neg} s \urcorner$ is in a's "think"-box.

Suppose Lois has c_0 in her "doubt"-box. Perspectivism is compatible with $\{c_0\}$ being the perspective supplied by some context. "If Lois doubts that Clark flies, then Lois thinks that Clark does not fly" will be false relative to any such context. The fact that 18 implies that Lois has $\ulcorner \hat{\neg} c_0 \urcorner$ in her "think"-box is irrelevant, since $\ulcorner \hat{\neg} c_0 \urcorner$ is not a member of the perspective supplied by this context. The case is similar for 15-17: If perspectives can fail to contain the conjunctions, disjunctions, or material implications of pairs of sentences they contain, then perspectivism does not imply that these schemas are good even on the assumption that anyone with both s and t in their "believe"-box has $\ulcorner s \hat{\wedge} t \urcorner$, $\ulcorner s \hat{\vee} t \urcorner$ and $\ulcorner s \hat{\rightarrow} t \urcorner$ in their "believe"-box too.

Call a perspective *closed* just in case it contains the negations, conjunctions, disjunctions, and material implications of any mentalese sentences it contains. Although perspectivism is compatible with the existence of non-closed perspectives, nothing we have said so far implies the existence of any such perspectives. Let *closure* be the hypothesis that all perspectives are closed. Perspectivism and closure *do* imply that 14-17 are good if the corresponding mentalese generalizations are true, by the following result:

Proposition 5. *Given closure, $\ulcorner$If A X_1s that φ_1, ..., and A X_ms that φ_m, then A Ys that $\psi\urcorner$ is good if the corresponding mentalese generalization is true, where $\varphi_1, \ldots, \varphi_m$ are pairwise-distinct and ψ is a Boolean combination of $\{\varphi_1, \ldots, \varphi_m\}$.*

The proof is in a footnote.[22] (Parallel reasoning establishes that, given closure, the result of adding "and A is ideally rational" to the antecedent of any such schema will be good if the generalization corresponding to the original schema is true of all ideally rational people; the same goes for all of our other results.)

This result contrasts sharply with the pessimistic assessment of Crimmins and Perry (1989), the most prominent defenders of perspectivism. They claim that the rampant context-sensitivity posited by perspectivism means that "there can be no simple logic of belief-sentences" (710):

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)

On the contrary, Proposition 5 shows that the goodness of schemas like 15-17 (and hence an "interesting logic of belief sentences") is not merely compatible with perspectivism, but is in fact implied (given closure) by the kind of generalizations about mental representations (i.e., "logic of beliefs, notions, and ideas") which Crimmins and Perry claim are "available".[23]

8 | TRANSPARENCY

But Crimmins and Perry's pessimism is not unfounded. This is because injectivity and closure together imply that no proposition expressed by any member of any perspective is identical to the result of applying any combination of Boolean operations to it. Suppose the mentalese sentence s means p and is a member of perspective π. By closure, every Boolean combination of $\{s\}$ is also a member of π. By injectivity, any pair of distinct such sentences – such as s and $\ulcorner \dot{\neg}\dot{\neg}s \urcorner$, or $\ulcorner s \dot{\wedge} (s \dot{\wedge} s) \urcorner$ and $\ulcorner (s \dot{\wedge} s) \dot{\wedge} s \urcorner$ – express distinct propositions. So p must be a distinct proposition from its double-negation, its conjunction with its self-conjunction must be distinct from its self-conjunction's conjunction with it, and so on. Many perspectivists will want to deny that propositions are this fine-grained. They must either reject injectivity or closure. One might then conclude that such perspectivists cannot accept the goodness of all the schemas covered by Propositions 4 and 5 whose corresponding mentalese generalizations they accept.

But this reaction is too quick. Say that *a grasps s* if, for some X, s is in a's X-box. Now consider the following weakening of injectivity:

Transparency: For all perspectives π, propositions p, mentalese sentences s and s', people a, and attitude verbs X: if s and s' are both in π and both mean p, s is in a's X-box and a grasps s', then s' is in a's X-box.

We can strengthen Proposition 3 as follows (and likewise for Proposition 4):

Proposition 6. *Given transparency, $\ulcorner$If A Xs that φ, then A does not Y that $\varphi \urcorner$ is good if, for all people a and mentalese sentences s, if s is in a's X-box, then s is not in a's Y-box.*

Proof. Suppose $\ulcorner A$ Xs that $\varphi \urcorner$ is true relative to c. So some mentalese sentence s in $\Pi(c)$ both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s X-box. So they grasp s, and it is not in their Y-box. By transparency, no sentence s' in $\Pi(c)$ both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s Y-box. So $\ulcorner A$ does not Y that $\varphi \urcorner$ is true relative to c. □

There are principled reasons to accept transparency that do not extend to injectivity. In particular, transparency can be motivated by combining assumptions about perspectives with assumptions about mentalese, as follows. For any relation $\approx$ between mentalese sentences, transparency is a consequence of:

P$\approx$: For all perspectives π, propositions p, and mentalese sentences s and s': if s and s' are both in π and both mean p, then $s \approx s'$.

M$\approx$: For all mentalese sentences s and s', people a, and attitude verbs X: if $s \approx s'$, s is in a's X-box and a grasps s', then s' is in a's X-box.

If we take $\approx$ to be the identity relation, then M$\approx$ becomes a logical truth and P$\approx$ collapses to injectivity. But this is not the only interesting interpretation of these two principles. For example, suppose we instead interpret $\approx$ as provable equivalence in the propositional calculus. Unlike injectivity, P$\approx$ (so interpreted) does not, given closure, entail the controversial principles about propositional granularity mentioned above. Indeed, it is compatible with the identity of any pair of propositions expressed by sentences that are equivalent in the propositional calculus. As for M$\approx$ (so interpreted), it is a substantive but natural idealizing assumption. And even if it is denied, so long as there is an interesting class of people (say, those who are ideally rational) and propositional attitude verbs (perhaps "believe", "know", and the like) for which M$\approx$ is true when $\approx$ is interpreted as logical equivalence, then given P$\approx$ we can establish a qualified version of transparency restricted to such agents and such attitudes, which in turn will allow us to prove a correspondingly qualified version of Proposition 4 concerning such attitudes of such agents.[24] (Other potentially interesting interpretations of $\approx$ include some sort of *a priori* equivalence, sameness of cognitive significance, etc.)

We are now in a position to unify our earlier results, by means of the following principle:

Grasping: People grasp all Boolean combinations of mentalese sentences they grasp.

Proposition 7. *Given closure and either (i) injectivity, or (ii) transparency and grasping, ⌜If **A** X_1s that φ_1, ..., and **A** X_ms that φ_m, then σ⌝ is good if the corresponding mentalese generalization is true, where $\varphi_1, \ldots, \varphi_m$ are pairwise distinct and, for some Boolean combinations ψ_j of $\{\varphi_1, \ldots, \varphi_m\}$, σ is a Boolean combination of $\{⌜\textbf{A} Y_1s \text{ that } \psi_1⌝, \ldots, ⌜\textbf{A} Y_ns \text{ that } \psi_n⌝\}$.*[25]

To illustrate the role of grasping in allowing us to combine the proofs of Propositions 5 and 6, consider the following schema covered by Proposition 7:

19. If **A** knows that φ, then **A** does not know that $\overbrace{\neg \ldots \neg}^{1000 \text{ times}} \varphi$.

Suppose propositions are identical to their double negations (which as discussed above is consistent with closure and transparency). Then 19 is clearly not good. We now argue that, given closure, transparency, and grasping, the corresponding mentalese generalization is not true. Consider some a, s and π such that s is in a's "know"-box and is a member of π. By closure, π contains the kilo-negation of s, which (we are supposing) means the same as s. So by transparency, the kilo-negation must be in a's "know"-box too if a grasps it. And by grasping, a does grasp it, since they grasp s. So it is not true that no one has in their "know"-box the kilo-negation of anything they have in their "know"-box.

Those who reject grasping as psychologically unrealistic may still accept a weaker generalization in the vicinity. The idea is that some propositional attitude verbs express *explicit* attitudes, and although what people grasp isn't closed under arbitrary Boolean combinations, they at least implicitly grasp all not-too-complicated Boolean combinations of things they explicitly grasp. Say that *a explicitly grasps s* if, for some explicit attitude X, s is in a's X-box; and let σ be a *Boolean k-combination* of $\{\varphi_1, \ldots, \varphi_n\}$ just in case it is either one of $\varphi_1, \ldots, \varphi_n$ or can be constructed from them using negation, conjunction, disjunction and material implication at most k times. The above idea can be formalized as:

k-**Grasping:** People grasp all Boolean k-combinations of mentalese sentences they explicitly grasp.

This principle allows us to prove the following variant of Proposition 7:

Proposition 8. *Given closure, transparency and k-grasping, ⌜If **A** X_1s that φ_1, ..., and **A** X_ms that φ_m, then σ⌝ is good if the corresponding mentalese generalization is true, where $X_1, \ldots, X_n$ express explicit*

attitudes, $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_m$ are pairwise distinct, and, for some Boolean k-combinations ψ_j of $\{\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_m\}$, σ is a Boolean combination of $\{\ulcorner \boldsymbol{A} \ Y_1 s \text{ that } \psi_1 \urcorner, \dots, \ulcorner \boldsymbol{A} \ Y_n s \text{ that } \psi_n \urcorner\}$.

Consider two perspectivists Nicol and Raneem who both think that closure is true, that "know" expresses an explicit attitude, that k-grasping is true for small k, and that no one grasps the kilo-negation of anything they grasp. Nicol accepts injectivity, and so thinks that no proposition is identical to its double negation. Raneem thinks that all propositions are identical to their double-negations, and so only accepts transparency. Although they both accept the mentalese generalization corresponding to 19, only Nicol will think the schema is good. This is exactly how we would expect competing views of propositional granularity to manifest themselves in disagreements about propositional attitude psychology.

Propositions 7 and 8 cover a wide range of schemas. But they don't cover every principle of propositional attitude psychology we might be interested in. Consider:

20. If $\boldsymbol{A}$ believes that $\boldsymbol{\varphi} \wedge \boldsymbol{\psi}$, then $\boldsymbol{A}$ believes that $\boldsymbol{\varphi}$.

This schema is not covered by Proposition 7. And even assuming closure and injectivity, the corresponding mentalese generalization does not imply that it is good. For example, suppose s means that Jahangir drinks milk and Jansher drinks water, Ramy has s in his "believe"-box, and the set of Boolean combinations of $\{s\}$ is an injective perspective. "If Ramy believes that Jahangir drinks milk and Jansher drinks water, then Ramy believes that Jahangir drinks milk" is then false relative to any context supplying this perspective, since no member of the perspective means that Jahangir drinks milk. So 20 is not good (irrespective of whether the corresponding mentalese generalization is true).

Rather than search for some yet further condition on perspectives to secure the inference from the relevant mentalese generalization to the goodness of 20, we think that perspectivists should take such examples to show that 20 is simply not good. In its place, they might point to the weaker schema:

21. If $\boldsymbol{A}$ understands the question whether $\boldsymbol{\varphi}$ and $\boldsymbol{A}$ understands the question whether $\boldsymbol{\psi}$, then (if $\boldsymbol{A}$ believes that $\boldsymbol{\varphi} \wedge \boldsymbol{\psi}$, then $\boldsymbol{A}$ believes that $\boldsymbol{\varphi}$).

This schema *is* covered by Proposition 7 and (given 1-grasping and that "understand the question" expresses an explicit attitude) Proposition 8.

9 | ASSESSING TRANSPARENCY AND CLOSURE

Propositions 7 and 8 show that, given transparency and closure, a wide range of schemas of propositional attitude psychology are good if the corresponding mentalese generalizations are true. But are transparency and closure true?

Consider the principles:

Articulation: Every mentalese sentence anyone grasps is in some perspective.

Decomposition: For every perspective π and mentalese sentences s and s', if s is in π and s' is a subsentence of s, then s' is in π.

Coarseness: For all p, $p \wedge \neg p$ is the same proposition as $(p \wedge \neg p) \wedge \neg(p \wedge \neg p)$.

Here are two limitative results:

Proposition 9. *Articulation, Decomposition and Transparency are not all true.*

Proof. Lois has $\ulcorner s_0 \hat{\wedge} \hat{\neg} c_0 \urcorner$ in her "believe"-box. By articulation, it is in some perspective π. By decomposition, s_0 and c_0 are both in π. And s_0 and c_0 both mean that Superman flies and Lois has s_0 in her "know"-box and grasps c_0. By transparency, she has c_0 in her "know"-box too. But she doesn't. □

Proposition 10. *Articulation, Coarseness, Transparency and Closure are not all true.*

Proof. Lois has $\ulcorner s_0 \hat{\wedge} \hat{\neg} c_0 \urcorner$ in her "believe"-box. By articulation, it is in some perspective π. By closure, $\ulcorner (s_0 \hat{\wedge} \hat{\neg} c_0) \hat{\wedge} \hat{\neg} (s_0 \hat{\wedge} \hat{\neg} c_0) \urcorner$ is in π. By coarseness, they mean the same thing. By transparency, if Lois grasps $\ulcorner (s_0 \hat{\wedge} \hat{\neg} c_0) \hat{\wedge} \hat{\neg} (s_0 \hat{\wedge} \hat{\neg} c_0) \urcorner$, then she has it in her "believe"-box. But (we may suppose) she does grasp this mentalese sentence and does not have it in her "believe"-box. □

How should perspectivists respond?[26]

9.1 | Rejecting decomposition and closure

One way that perspectivists might respond to the limitative results of Propositions 9 and 10 is to reject decomposition and closure in order to accept transparency and coarseness.

This response naturally goes along with Crimmins and Perry's pessimism about there being an "interesting logic of belief-sentences", since not only does it reject the assumption of Proposition 5, but it is hard to see how alternative assumptions could achieve the same effect without falling prey to an analogue of Proposition 10. This strikes us as a severe cost. We also think there are strong reasons to accept decomposition, which we will consider in section 10. That being said, we think the view is worth exploring.

Despite being incompatible with an "interesting logic of belief-sentences", the present proposal is compatible with there being an interesting "logic" of the entailment relations between different propositional attitudes, since closure is not an assumption of Proposition 6 (and it is not needed to strengthen that result to cover Boolean combinations of attitudes towards a common content, just as nothing beyond injectivity was needed to generalize Proposition 3 to Proposition 4). Such principles are a central topic in the study of propositional attitudes, especially when we consider how some attitudes might be 'analyzed' in terms of others.

There are also reasons independent of decomposition for those who accept principles like coarseness to reject closure. Consider an analogue of coarseness for tautologies: for all p, $p \rightarrow p$ is the same proposition as $(p \rightarrow p) \rightarrow (p \rightarrow p)$. Now consider:

22. Lois doesn't know that, if Superman flies, then Clark flies.

We would expect ordinary uses of such ascriptions to be true and also to be in contexts whose supplied perspective contains $\ulcorner s_0 \hat{\rightarrow} c_0 \urcorner$. But this is ruled out by transparency for any closed perspective, since Lois grasps $\ulcorner s_0 \hat{\rightarrow} c_0 \urcorner$, and any closed perspective containing it also contains $\ulcorner (s_0 \hat{\rightarrow} c_0) \hat{\rightarrow} (s_0 \hat{\rightarrow} c_0) \urcorner$, which means the same thing (by our granularity assumption) but unlike $\ulcorner s_0 \hat{\rightarrow} c_0 \urcorner$ is in Lois's "know"-box. A parallel argument applies to sentences like

23. Lois doesn't know that Superman is Clark.

given the assumption that, for all x, $x = x$ is the same proposition as $x = x \rightarrow x = x$.

Similarly, proponents of very coarse-grained theories of propositions might reject closure in order to solve one version of the so-called 'problem of logical omniscience'. The problem is to provide pretheoretically plausible truth conditions for ordinary uses of negated attitude ascriptions whose complement clauses correspond to mentalese sentences that, though not in the same 'boxes' as

their self-implications, express the same propositions as their self-implications if propositions are sufficiently coarse-grained. For example, the challenge arises with sentences like

24. In 1900, Frege didn't know that naïve set theory is inconsistent.

for those, like Stalnaker (1984), who think that mathematics is non-contingent and that there is only one necessary truth. Only by rejecting closure can they hold that this sentence is true relative to a perspective that contains a mentalese sentence expressing the proposition that naïve set theory is inconsistent and whose self-implication was in Frege's "know"-box in 1900.

9.2 | Rejecting decomposition and coarseness

A second way that perspectivists might respond to Propositions 9 and 10 is to reject decomposition and coarseness in order to accept transparency and closure.

We think there are strong considerations in favor of theories of propositional granularity that imply coarseness. That being said, avoiding the limitative result posed by Proposition 10 does not require thinking that propositions are anywhere near as fine-grained as they are according to the most familiar theories of structured propositions. It requires only that no proposition is the result of applying any non-trivial polyadic Boolean operation to itself (the operation in this case being conjunction composed with negation in its second argument). This prohibition is consistent, for example, with φ and $\ulcorner \neg \neg \varphi \urcorner$ always expressing the same proposition, and similarly for $\ulcorner \varphi \wedge \psi \urcorner$ and $\ulcorner \psi \wedge \varphi \urcorner$.[27]

Setting aside the attractions of coarseness and of decomposition (the latter of which will be discussed in section 10), there is a tension between transparency and closure considered on their own. Consider the sentence:

25. That Superman flies and that Clark doesn't fly are both things that Lois believes.

A use of this sentence is true just in case the perspective supplied by context contains two mentalese sentences in Lois's "believe"-box, one of which expresses the proposition that Superman flies and the other of which expresses the proposition that he doesn't fly. But transparency and closure imply that there are no such perspectives (at least given the extremely natural assumption that, for any mentalese sentence in Lois's "believe"-box that expresses the proposition that Superman flies, its negation is not in her "believe"-box but is something that she grasps). For example, no perspective can contain both s_0 and $\ulcorner \dot{\neg} c_0 \urcorner$, since it would have to contain $\ulcorner \dot{\neg} s_0 \urcorner$, by closure, which is grasped by Lois but, unlike $\ulcorner \dot{\neg} c_0 \urcorner$, is not in her "believe"-box, in violation of transparency (since both express the proposition that Superman doesn't fly). We arrive at the surprising conclusion that ordinary uses of 25 are false. Unlike 1, its truth cannot be salvaged by invoking a mid-sentence shift in the perspective supplied by context, since there is only one attitude-verb occurrence in the sentence. The view thereby convicts ordinary speakers of error concerning such sentences.[28]

9.3 | Rejecting transparency

A third way in which perspectivists might respond to Propositions 9 and 10 is to reject transparency in order to accept decomposition and closure.

In assessing this view, note first that the proofs of Propositions 2 and 5 can be combined to establish the following result:

Proposition 11. *Given closure, $\ulcorner$If A X_1s that φ_1, ... , and A X_ms that φ_m, then $\sigma\urcorner$ is good if the corresponding mentalese generalization is true, where $\varphi_1, \ldots, \varphi_m$ are pairwise distinct and, for some*

Boolean combinations ψ_j of $\{\varphi_1, \ldots, \varphi_m\}$, σ is a positive combination of $\{\ulcorner A \; Y_1 s \; that \; \psi_1 \urcorner, \ldots, \ulcorner A \; Y_n s \; that \; \psi_n \urcorner \}$.

The view thereby vindicates "positive psychology": schemas covered by this result will be good provided the corresponding mentalese generalizations are true.[29]

Perspectivists who reject transparency might still uphold the goodness of some schemas of propositional attitude psychology not covered by Proposition 11.[30] We will now consider two strategies for doing so.

One strategy is to hold that the naïve box-based semantics applies only to "unanalyzable" attitude verbs.[31] For example, one might think that "be unsure" is analyzable in terms of "understand the question" and "be sure", so that $[\![A$ is unsure whether $\varphi]\!]^c = [\![A$ understands the question whether φ and A is neither sure that φ nor sure that $\neg\varphi]\!]^c$ = the proposition that some s in $\Pi(c)$ that means $[\![\varphi]\!]^c$ is in $[\![A]\!]^c$'s "understands the question"-box and no s in $\Pi(c)$ that either means $[\![\varphi]\!]^c$ or means $[\![\neg\varphi]\!]^c$ is in $[\![A]\!]^c$'s "be sure"-box. This analysis guarantees that 10 is good, despite not being covered by Proposition 11.

This strategy can be applied quite widely. But it cannot be applied to all schemas covered by Proposition 6, such as:

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

Even if the mentalese generalization corresponding to this schema is true, decomposition implies that the schema is not good: some perspective contains $\ulcorner s_0 \hat{\wedge} \hat{\neg} c_0 \urcorner$, by articulation and the fact that Lois grasps it, and hence contains both s_0 and $\ulcorner \hat{\neg} c_0 \urcorner$, by decomposition, both of which are in Lois's "believe"-box; so "If Lois believes that Superman flies, then Lois does not believe that Superman does not fly" is not unequivocally true.

A second strategy is to retreat from schemas like 26 to weaker ones obtained from them by adding a special proviso.[32] Suppose there is a sentence Ω that is true relative to all and only contexts that supply perspectives that satisfy transparency.[33] Then for every schema σ covered by Propositions 7 and 8, corresponding results hold for $\ulcorner$If Ω, then $\sigma\urcorner$ without the assumption of transparency. This fact suggests one way in which perspectivists can vindicate the legitimacy of remarks like "Setting aside identity confusion, …" that philosophers are wont to make. Note that this way of setting aside the issue of identity confusion for the purposes of doing propositional attitude psychology is very different from restricting the scope of such theorizing to idealized agents who are not subject to any such confusion.[34] That restriction is extremely draconian; the resulting theory would not apply to any of us.[35] The present proposal involves no such restriction: in contexts supplying transparent perspectives, instances of Ω-qualified schemas can express non-trivial claims about, for example, Lois's attitudes towards Superman (despite her confusion about his identity).

10 | ELEMENTARY GENERATION

In section 3 we promised a sketch of how the perspective that parametrizes a given attitude ascription might be partially determined by the conversationally salient ways the subject of the ascription has of thinking about the object of her confusion. We will now give that sketch. Although it leaves much unsettled, it is quite predictive. In particular, it implies decomposition. The naturalness of this picture thereby provides strong reason to prefer the view discussed in section 9.3 (rejecting transparency) to the views discussed in section 9.1 (rejecting decomposition and closure) and in section 9.2 (rejecting decomposition and coarseness).

David Braun has influentially objected to views like ours on the grounds that "ordinary speakers [do not] have […] sophisticated thoughts and intentions about mental representations when they utter belief sentences" (Braun, 1998, pp. 560-1). We agree that ordinary speakers don't think about sets of sentences of mentalese "as such" when making attitude ascriptions. But ordinary speakers often do have in mind particular ways that the subject of their attitude ascriptions has of thinking about the object of their confusion, as Shorty's accomplice's description of Thelma in section 3 demonstrates. These ways of thinking about objects correspond to elementary expressions in the subjects' of those ascriptions mental lexicons.[36] Moreover, every set of elementary expressions of mentalese determines a set of mentalese sentences: namely, the set of sentences that can be built from the elementary expressions in that set. These considerations suggest the following natural principle:

Elementary generation: For every perspective π there is a set of elementary expressions of mentalese E such that π is the set of all mentalese sentences all of whose elementary constituents are in E.

More specifically, they suggest the following picture of how speakers and hearers succeed in resolving the context-sensitivity of attitude ascriptions. In a given conversational context, certain ways the person being talked about has of thinking about the object they are confused about are saliently *relevant* and others are saliently *to be ignored*. For example, for an ordinary utterance of 1, the context in which the second conjunct will be uttered is one in which the "Clark"-associated way Lois has of thinking about Superman is saliently relevant and the "Superman"-associated way she has of thinking about him is saliently to be ignored. Moreover, these facts determine that the perspective supplied by context is generated from a set of elementary mentalese expressions that includes all mentalese expressions corresponding to ways of thinking that are saliently relevant and no mentalese expressions corresponding to ways of thinking that are saliently to be ignored.

This story is by no means a full account of which perspectives are supplied by context. For example, it says nothing about which mentalese expressions that do not correspond to conversationally salient ways of thinking are in the set that generates the perspective supplied by context.[37] The story is also incomplete because it does not explain in general how speakers and hearers coordinate on saliently relevant/irrelevant ways the person under discussion has of thinking about the objects of their confusion (although, as we noted earlier, the Thelma case shows that this is something that any theory of identity confusion should explain, and so is not a special challenge for perspectivists). However, in the special case of names like "Superman" and "Clark" which we conventionally associate with particular ways Lois has of thinking about the object of her confusion, it seems clear that these ways of thinking are saliently relevant when we use the corresponding names to describe her state of mind. So an ordinary use of an attitude ascription about Lois in which "Superman" appears in the complement clause will be in a context that supplies a perspective generated from a set of elementary expressions that contains the corresponding name in Lois's mental lexicon, and likewise for "Clark". In the next section we will draw out some implications of this fact.

As advertised, this picture requires rejecting transparency, since elementary generation implies decomposition, and decomposition is incompatible with transparency by Proposition 9.[38] The picture also allows those who accept Braun's dictum that ordinary attitude ascriptions should not "express propositions that are partly about believers' mental representations" to be perspectivists, by allowing them to identify perspectives with sets of ways of thinking about entities rather than with sets of sentences of mentalese, and modifying the truth conditions for attitude ascriptions accordingly.[39]

11 | ERROR

We think perspectivists should embrace this picture of how facts about conversationally salient ways of thinking constrain which perspectives are supplied by context. But the picture also has some surprising implications. It implies that ordinary utterances of 22 ("Lois doesn't know that, if Superman flies, then Clark flies") are in contexts that supply perspectives containing $\ulcorner s_0 \dot\to c_0 \urcorner$, Lois's mentalese counterpart of that ascriptions' complement clause. This is not surprising by itself. But given elementary generation, it follows that $\ulcorner s_0 \dot\to s_0 \urcorner$ is also a member of this perspective. And this mentalese sentence *is* in Lois's "know"-box. It follows that ordinary utterances of 22 are false. A parallel argument applies to ordinary utterances of 23, which have often been taken to be the starting point in the literature on identity confusion.[40]

It is certainly a cost of the present picture that it convicts ordinary speakers of error when they judge that uses of sentences like 22 and 23 are true. But we don't think it is a fatal cost. Identity confusion is extremely puzzling, and although in any particular case we would prefer not to attribute error to ordinary speakers, it does not strike us as particularly implausible that ordinary speakers sometimes make mistakes as a result, perhaps quite systematically.[41] We will now consider and respond to two more specific objections to the combination of perspectivism and the error hypothesis under consideration.

Objection 1: Anyone willing to grant that ordinary speakers make mistakes about sentences like 22 and 23 might as well be a neo-Russellian. Reply: One version of this objection presupposes that postulating context-sensitivity in attitude ascriptions is a substantial theoretical cost, which is not worth paying for anything less than full vindication of ordinary truth-value judgments. We are unmoved by the objection so stated, since we don't consider contextualism to be particularly costly in itself. A more concrete worry is needed. Three such worries are that the kind of context-sensitivity we are postulating (i) rests on an unappealing mentalese foundation, (ii) requires too much sophistication on the part of ordinary speakers, or (iii) is too unconstrained to be predictive. Regarding (i), neo-Russellians use a similar framework to give truth conditions for attitude ascriptions, and we suspect that alternative neo-Russellian frameworks would support parallel perspectivist alternatives. Regarding (ii), neo-Russellians face a more acute challenge than perspectivists do to explain ordinary conversational patterns involving attitude ascriptions, given the much more widespread error they postulate in speakers' truth-value judgments about attitude ascriptions.[42] Regarding (iii), we deny the charge: the versions of perspectivism we have been exploring make strong and sometimes surprising predictions both about the truth-values of ordinary attitude ascriptions and about the good standing of principles of propositional attitude psychology.

Our preferred version of perspectivism not only convicts ordinary speakers of error in far fewer cases than neo-Russellianism does: it also fails to convict them of error in canonical cases like ordinary uses of sentences like 1. According to perspectivists, it is by exploiting the context-sensitivity of such sentences that ordinary speakers succeed in conveying the contours of people's identity confusion using ordinary English. If neo-Russellianism is true, it is far less clear how we routinely succeed in doing so without explicitly talking about people's underlying mental representations.

Objection 2: Although a limited amount of error on the part of ordinary speakers is tolerable, error about sentences like 23 is not tolerable because such ascriptions are central to the phenomenon under discussion. Reply: Sentences like 23 are much less central to the phenomenon under discussion than their prominence in the literature on identity confusion suggests. As discussed at the end of section 3, most real cases of identity confusion do not involve a pair of proper names that are canonically associated with the respect in which the person in question is confused about the thing they are confused about.[43] Real cases of identity confusion are rarely naturally described using negated knowledge

ascriptions with identity statements involving two occurrences of proper names as their complement clauses. Moreover, we agree with Saul (1997) that, given the pathological pattern of ordinary truth-value judgments for sentences involving pairs of names like "Superman" and "Clark" but not involving any attitude verbs, we should be especially cautious in drawing sweeping conclusions on the basis of examples that seem to essentially rely on such pairs of names. The present picture only convicts ordinary speakers of error in cases of negated attitude ascriptions whose complement clauses contain co-referential expressions different uses of which are associated by the speaker with distinct ways that the subject of the ascription has of thinking about those expressions' common referent. Such ascriptions do occur in the wild, but they are unusual.

12 | CONCLUSION

Like Fregeanism, perspectivism provides a framework in which the English attitude ascriptions that we ordinarily use to characterize the states of mind of identity-confused people can succeed in conveying true and relevant features of those people's underlying psychology. Like neo-Russellianism, it is naturally combined with a Millian treatment of proper names and an unqualified endorsement of Substitution.[44] But is perspectivism compatible with the good standing of systematic theorizing about the propositional attitudes? After sharpening this question in a way that is sensitive to the context-sensitivity the theory posits, we established a number of tenability results in this direction, showing how principles about perspectives like closure and injectivity, principles about people's cognitive psychology like grasping and k-grasping, and principles relating both like transparency, can secure the good standing of principles of propositional attitude psychology conditional on the truth of corresponding generalizations about people's mental representations. We then explored tradeoffs between these principles on the one hand, and vindicating ordinary truth-value judgments about attitude ascriptions on the other. After considering views that prioritize truth-value judgments over general principles (section 9.1), and vice versa (section 9.2), we considered (section 9.3) and advocated for (section 10) a version of perspectivism that makes some surprising predictions on both fronts, but none of them so extreme as to threaten the communicative utility of attitude ascriptions or the intellectual respectibility of propositional attitude psychology.

The results of this paper are also relevant to Fregeans and neo-Russellians. The usual way of developing neo-Russellianism is a limiting case of our way of developing perspectivism: in effect, it is the view that perspectivism is true but there is only one perspective, the set of all mentalese sentences.[45] This perspectivist recasting of neo-Russellianism implies closure, so neo-Russellians can avail themselves of Proposition 11 to vindicate the goodness of "positive psychology", as discussed in section 9.3. But since the set of all mentalese sentences is not transparent, neo-Russellians cannot appeal to Propositions 7 and 8 to establish the goodness of other principles of propositional attitude psychology; nor can they use the strategy mentioned in section 9.3 of vindicating the non-trivial goodness of such principles when qualified by the suitably context-sensitive understanding of the proviso "If we set aside identity confusion, then …". Neo-Russellians are also committed to the falsity of all uses of 22 and 23.[46]

Finally, let us return to Fregeans. They have no trouble vindicating the goodness of the sort of schematic principles we have been discussing.[47] But it is a vexed question how they should think about the universal generalizations of those schemas, since it is a vexed question how they should handle quantification into the scope of attitude ascriptions. The orthodox treatment is that of Kaplan (1968), who suggests that sentences that involve quantification into the scope of attitude ascriptions' complement clauses have truth conditions involving existential quantification over what he calls "vivid

names" – which, for ease of comparison with our version of perspectivism, we may identify with mentalese names. For example, consider the following universal generalization of an instance of 26:

27. For all x, if Lois believes that x flies, then Lois does not believe that x does not fly.

According to Kaplan, this sentence is true just in case, for all x, if Lois has $\ulcorner Fn \urcorner$ in her "believe"-box for some mentalese name n of x, then Lois does not have $\ulcorner \dot{\neg} Fn \urcorner$ in her "believe"-box for any mentalese name n of x (where F is her mentalese analogue of "flies"). And this generalization is false, since Superman is a counterexample to it.[48] This result is not surprising: Kaplan's truth conditions for quantified attitude ascriptions introduce existential quantification over mental representations in the same place that perspectivist truth conditions for non-quantified attitude ascriptions do, so the challenge for perspectivists concerning principles like 26 arises for Kaplanians when we consider universal generalizations of such principles' instances like 27. If Fregeans wish to vindicate more generalizations than perspectivists can, they will have to adopt a non-Kaplanian treatment of quantifying in.[49]

There is no parallel issue for pespectivists. They can treat quantification into the scope of attitude ascriptions in the same way that they treat quantification into any other context, since they are Millians. So when perspectivists take a schema to be good, they can straightforwardly take universal generalizations of it to be good too.[50] Since, as we have shown, perspectivism is compatible with the good standing of a wide range of schematic principles, it is likewise compatible with the good standing of the universal generalizations of those schemas.

ACKNOWLEDGEMENTS

ENDNOTES

[1] Such contextualist views have been defended by Crimmins and Perry (1989) and Crimmins (1992) and explored by Schiffer (1979); Richard (1990) holds a similar contextualist view about Kripke's "Paderewski" example, which is similar to the Thelma/Shorty example described in section 3.

[2] In what follows, all talk of mentalese correlates of public language sentences and expressions is relative to a particular person at a particular time. Everything that follows is therefore compatible with a hyper-individualist conception of mental representations according to which, if a has s in their X-box at t and b has s in their Y-box at t', then $a = b$ and $t = t'$.

As should be clear from our talk of Lois's "know"-box, we are not thinking that which mentalese sentences are in which boxes is something 'internal' that could be revealed, say, by ultra-advanced brain-scans. To the extent that "box"-talk carries with it this unintended internalist connotation, the reader is invited to substitute "s is an X-vehicle for a", or some other idiom that lacks such connotations.

[3] More precisely: a person a is subject to identity confusion when, for some attitude verb X, mentalese sentence s, and co-referring expressions e_1 and e_2 of a's mental lexicon, a has s but not $s[e_2/e_1]$ in their X-box, where $s[e_2/e_1]$ results from s by replacing some occurrence of e_1 in s with an occurrence of e_2 (and the occurrence is in a position that licenses the intersubstitution of co-referential terms *salva veritate*). Lois counts as identity confused because s_0 results from c_0 by such a replacement of occurrences of co-referential expressions of her mental lexicon – namely, mentalese correlates of "Superman" and "Clark".

[4] Our assumption that uses of English attitude ascriptions express the holding of a three-place relation between a person, proposition, and perspective will be rejected by those like King (2007) who think that sentences express propositions

whose structure reflects those sentences' syntactic structure, since assuming English has a binary-branching syntax, *no* English sentence ever expresses the holding of any polyadic relation between any entities. Such theorists should reformulate perspectivism so that it fits their preferred view about English syntax. We ignore such subtleties in what follows since we will not be considering embeddings of attitude ascriptions in non-extensional contexts that might be sensitive to structural differences among truth-conditionally equivalent propositions. Thanks to an anonymous referee for raising this issue.

[5] This version of perspectivism closely corresponds both to the "notion constraint" proposal from Crimmins and Perry (1989) and to the appeal to "types of modes of presentation" in Schiffer (1979, p. 32); cf. Schiffer (1992). A simpler proposal would be to identify perspectives with particular mentalese sentences, and let X^* be the relation that holds between a person a, proposition p, and mentalese sentence s just in case s both means p and is in a's X-box. But this simple proposal makes the wrong predictions about quantified attitude ascriptions; for example, it predicts the truth of "There is at most one person who Alice knows runs", since for any mentalese sentence s, there is at most one person such that s means that they run.

A natural generalization of the present framework would be to identify perspectives not with sets of mentalese sentences but with relations between people and mentalese sentences. X^* would then be the relation that holds between a person a, proposition p, and perspective π just in case a is related by π to some s that both means p and is in a's X-box. In this way, which mental representations are relevant for the truth or falsity of a given attitude ascription could differ from person to person, from time to time, and from world to world. We ignore this proposal in what follows not because we have any aversion to it, but because the sort of generality it affords is orthogonal to our discussion, which concerns attitude ascriptions involving a single person at a given time.

Another natural generalization of the present framework would be to allow perspectives to contain sub-sentential mentalese expressions, so that psychological verbs taking non-clausal complements like "thinks about" and "looks for" could be given a parallel contextualist treatment. We will discuss a proposal like this in section 10. Note that in what follows we will treat attitude ascriptions whose complement clauses are polar questions, as in example 10, in the same way that we treat attitude ascriptions whose complements are 'that'-clauses.

[6] There are important reasons to question this assumption, but they are orthogonal to the issues we will be concerned with here; see Ninan (2010); Rabern (2012); Yli-Vakkuri (2013); Yalcin (2014).

[7] The first assumption is that the relevant individual named "Superman" and the relevant individual named "Clark" are identical. The second assumption is that a negated sentence expresses the negation of the proposition expressed by the unnegated subsentence, and similarly for other logical connectives. (We will harmlessly pretend that, e.g., "Saurav swims" is a subsentence of "Saurav doesn't swim"; reframing our discussion in a more syntactically faithful way is orthogonal to the main concerns of this paper.) The third assumption is that the truth value of "she doesn't know that Clark flies" in ordinary uses of 1 doesn't depend on any difference there may be between the semantic value of the use of "she" and the semantic value of the earlier use of "Lois".

[8] Sociologically, however, most Millians reject contextualism in favor of neo-Russellianism; e.g. Salmon (1986), Soames (1987b), Braun (1998), Saul (2010).

[9] The labels are taken respectively from Dorr (2014) and Schiffer (1992). In the influential terminology of Perry (1986), those who think that perspectives are constituents of the propositions expressed by attitude ascriptions if perspectivism is true will think that they are 'unarticulated' constituents of those propositions if verbalism is true.

[10] By the context in which an entire attitude ascription is used, we mean the context in which the relevant perspective-sensitive elementary constituent of the ascription is used (be that the main verb, if verbalism is true, or the accompanying hidden indexical, if hidden indexicalism is true).

[11] One might question whether this sentence really does involve a mid-sentence context-shift by proposing that ⌜A knows that B (the F) Gs⌝ means the same as ⌜A knows that B is the F and Gs⌝. 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 might be described by saying "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".

Dorr (2014) makes a persuasive and much more detailed case that everyone, whatever their views about the semantics of names, should think that identity-confusion breeds mid-sentence context-sensitivity, and moreover that views like perspectivism don't violate any well-attested generalizations about such context-sensitivity in English.

[12] We use boldface for schematic letters to distinguish them from variables ranging over expressions.

[13] Instances of this schema are obtained by replacing $\boldsymbol{\varphi}$ with a declarative English sentence, $\boldsymbol{A}$ and $\boldsymbol{B}$ with proper names, and $\boldsymbol{\varphi}[\boldsymbol{B}/\boldsymbol{A}]$ with some sentence obtained from $\boldsymbol{\varphi}$ by replacing some occurrences of the name substituted for $\boldsymbol{A}$ with the name substituted for $\boldsymbol{B}$.

[14] Suppose $\ulcorner A = B \urcorner$ is true relative to c. Then, assuming Millianism, A and B have the same semantic value relative to c; so φ and $\varphi[B/A]$ have the same semantic value relative to c (by compositionality); so $\ulcorner$If $A = B$ and φ, then $\varphi[B/A]\urcorner$ is true relative to c.

This argument does not require the full strength of Millianism; it requires only the weaker assumption that two names have the same semantic value whenever they are used *in the same syntactic position* to name the same individual. This point is relevant for those who, following Montague (1973), think that when "Superman" occurs as the direct object of "seek" as in "Lois seeks Superman" its semantic value is not Superman but rather the generalized quantifier *being a property of Superman*: the view is not Millian, but like Millianism it implies that Substitution is good, since the semantic value of a use of a name is a function of the individual named and the name's syntactic position in the sentence in which it is used.

[15] There are other positive statuses for English sentences that we might have considered instead. Following Kaplan (1989), assume that propositions can be evaluated as true or false relative to *indices* and that every context uniquely determines an index. Say that a sentence is true *in* a context just in case the proposition it expresses relative to that context is true relative to the index of that context; say that a sentence is *inevitably true* just in case it is true in all contexts.

Unequivocal truth is neither sufficient nor necessary for inevitable truth: any non-context-sensitive sentence expressing a truth that is false relative to the index of some context is unequivocally true but not inevitably true, and (given Kaplan's treatment of "actually") the sentence $\ulcorner \varphi$ if and only if actually $\varphi \urcorner$ is inevitably true but not unequivocally true for any φ such that, for some c, it is not the case that φ is true relative to c if and only if φ is true in c. More generally, theorems of Kaplan's logic of demonstratives will be inevitably true but may not be unequivocally true. We think that inevitable truth is one interesting notion of validity for English, with the caveat that it is a rather liberal one (since, in the case of non-context-sensitive sentences, inevitable truth coincides with being true relative to the index of every context).

But the difference between unequivocal and inevitable truth is orthogonal to our main concerns in this paper, for the following reason. The potential counterexamples to the goodness of schemas of propositional attitude psychology we discuss below are (or could be replaced with) sentences that express the same proposition relative to any pair of contexts that supply the same perspective. We will assume that there is an 'accurate' index α such that all and only the true propositions are true relative to it, and that, if π is the perspective supplied by c, then it is also the perspective supplied by some context c' whose index is α. It follows that if any of the sentences of interest is not unequivocally true, then it is not inevitably true either: there is a context c relative to which it expresses a proposition that isn't true, and hence a proposition that isn't true at α; but the proposition it expresses relative to a context depends only on the perspective of the context; so the sentence is false in any context with index α that supplies the same perspective as c; and there is such a context; so the sentence is not inevitably true. In the other direction, a sentence $\ulcorner \Box \varphi \urcorner$ is inevitably true if it is unequivocally true, where $\Box$ is a modal operator that expresses a notion of necessity equivalent to being true relative to every index: suppose $\ulcorner \Box \varphi \urcorner$ is true relative to every context; then φ expresses relative to any given context c a proposition that is true relative to every index, and hence a proposition that, relative to every index (and hence relative to the index of c), is true relative to every index (and so is necessarily true; we are here assuming that, if a proposition is true relative to every index, then it is true relative to every index that it is true relative to every index); so $\ulcorner \Box \varphi \urcorner$ is inevitably true. So our results about goodness, when extended to parallel results about modal goodness as advertised above, imply parallel results about inevitable truth.

[16] An interesting question for perspectivists is the extent to which their contextualism should extend to speech-reporting verbs like "says" in order to maintain the goodness of schematic principles relating speech and thought. Consider the two schemas:

 (i) If $\boldsymbol{A}$ says that $\boldsymbol{\varphi}$ and $\boldsymbol{A}$ is not insincere, then $\boldsymbol{A}$ believes that $\boldsymbol{\varphi}$.

 (ii) If $\boldsymbol{A}$ says that $\boldsymbol{\varphi}$ in making an utterance $\boldsymbol{u}$, then the proposition that $\boldsymbol{\varphi}$ is the unique proposition expressed by $\boldsymbol{u}$.

There is some temptation to think that (i) and (ii) are both good. But this generates surprising consequences for perspectivists. Suppose Lois not insincerely utters "Superman flies". Perspectivists think that there are contexts relative to which "Lois doesn't believe that Clark flies" is true, and hence (by Substitution) relative to which "Lois doesn't

believe that Superman flies" is true. Assuming "Lois's utterance expresses the proposition that Superman flies" is true relative to such contexts, the goodness of (i) and (ii) leads to the surprising conclusion that "There is nothing that Lois said when she uttered 'Superman flies'" is also true relative to such contexts.

Maybe perspectivists should accept this conclusion, the idea being (roughly) that saying that p isn't merely a matter of uttering a sentence meaning that p, but additionally requires knowing that the sentence means that p, and hence the context-sensitivity of such knowledge ascriptions (which patterns with the context-sensitivity of more basic knowledge ascriptions like those in 1) will lead to corresponding context-sensitivity of speech reports.

Alternatively, perspectivists might deny that (i) is good, and assimilate Lois's situation to the following case. Suppose a French speaker Grégory memorizes and utters the English sentence "Frenchmen are gullible" without knowing what it means or having any inclination to accept the synonymous French sentence. Arguably, there are contexts relative to which "Grégory said that Frenchmen are gullible and he was not insincere, but he does not believe that Frenchmen are gullible – rather, he doesn't know what he is saying" truly describes this situation. Perspectivists might say that, similarly, "Lois doesn't know what she is saying" is true relative to the relevant context (although, unlike "Grégory doesn't know what he is saying", a typical utterance of this sentence would be in a context relative to which it was false).

Alternatively still, perspectivists might deny that (ii) is good, and assimilate Lois's situation to the following case. Suppose Ramy sincerely utters "The Nile floods frequently and flows north". Arguably, there are contexts relative to which "In making this utterance Ramy said that the Nile floods frequently and also said that it flows North – which are obviously distinct propositions" truly describes this situation. If correct, this seems like a general phenomenon, whereby we can truly report people to have said things that are related to but logically weaker than the propositions expressed by the sentences they utter. Perspectivists might then say that, similarly, "Lois said that someone flies, and hence said something" is true relative to the relevant context. They can thereby hold that (i) is good while denying that "There is nothing that Lois said when she uttered 'Superman flies'" is true relative to any context. (The general strategy holds, roughly, that if, owing only to a's identity confusion about x, $\ulcorner A$ says that $\varphi \urcorner$ is true relative to some contexts but not to others, then there is a property f such that (a) φ expresses a proposition truth-conditionally-equivalent to x being f, (b) $\ulcorner A$ says that something is $F \urcorner$ is true relative to every context relative to which F expresses f, and (c) $\ulcorner A$ says something $\urcorner$ is true relative to every context.)

A fourth (and perhaps our favored) option holds that "says" is polysemous and different of the above three responses are correct for different of its disambiguations.

[17] As advertised, parallel reasoning shows that $\ulcorner$Necessarily, if A Xs that φ, then A Ys that $\varphi \urcorner$ is good if necessarily, for all people a and mentalese sentences s, if s is in a's X-box, then s is in a's Y-box. The same goes for our subsequent results.

[18] Braun (2000) poses a related question for neo-Russellianism.

[19] Let σ be a positive combination of $Z = \{\ulcorner A$ Y_js that $\varphi \urcorner : 1 \leq j \leq n\}$. The mentalese generalization *corresponding to* $\ulcorner$If A Xs that φ, then $\sigma \urcorner$ is the claim that, for any person a and mentalese sentence s, if s is in a's X-box, then $v(\sigma) = 1$ for all classical valuations v such that $v(\ulcorner A$ Y_js that $\varphi \urcorner) = 1$ iff s is in a's Y_j-box.

To prove Proposition 2: Suppose $\sigma = \ulcorner A$ Ys that $\varphi \urcorner$. Then the mentalese generalization corresponding to $\ulcorner$If A Xs that φ, then $\sigma \urcorner$ is the claim that, for any person a and mentalese sentence s, if s is in a's X-box, then $v(\ulcorner A$ Ys that $\varphi \urcorner) = 1$ for all classical valuations v such that $v(\ulcorner A$ Ys that $\varphi \urcorner) = 1$ iff s is in a's Y-box. This claim is equivalent to the claim that, for any person a and mentalese sentence s, if s is in a's X-box, then s is in a's Y-box. So by Proposition 1, if it is true then the corresponding schema is good.

Next, suppose σ is a disjunction of members of Z. Then the mentalese generalization corresponding to $\ulcorner$If A Xs that φ, then $\sigma \urcorner$ is true just in case for every person a and mentalese sentence s, if s is in a's X-box, then there is a Y such that $\ulcorner A$ Ys that $\varphi \urcorner$ is a disjunct of σ and s is in a's Y-box. Suppose that $\ulcorner A$ Xs that $\varphi \urcorner$ is true relative to c. So some s in $\Pi(c)$ both means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s X-box. So, for some Y such that $\ulcorner A$ Ys that $\varphi \urcorner$ is a disjunct of σ, s is also in their Y-box. So $\ulcorner A$ Ys that $\varphi \urcorner$ is true relative to c. So the relevant instance of σ is true relative to c, since one of its disjuncts is.

This covers all cases, since any positive combination of sentences is logically equivalent to its conjunctive normal form, and a conditional with a conjunctive consequent is unequivocally true if the conditionals with the same antecedent but with the individual conjuncts as consequents are each unequivocally true. □

[20] Indeed, this is a direct consequence of Millianism and compositionality. Note that the goodness of Substitution does not imply that ordinary uses of "The proposition that Superman flies is the proposition that Clark flies" are true: a

perspectivist could give a context-sensitive semantics for "the proposition that" modeled on their treatment of attitude verbs, and hold that such uses involve a mid-sentence shift in context.

[21] Let σ be a Boolean combination of $Z = \{\ulcorner A\ Y_j\text{s that } \boldsymbol{\varphi}\urcorner : 1 \leq j \leq n\}$. The mentalese generalization *corresponding to* $\ulcorner$If $\boldsymbol{A}\ X$s that $\boldsymbol{\varphi}$, then $\sigma\urcorner$ is defined as in footnote 19. The proof is then a straightforward combination of the proofs of Propositions 2 and 3.

[22] Let ψ be a Boolean combination of $\{\boldsymbol{\varphi}_1, \ldots, \boldsymbol{\varphi}_m\}$ for pairwise-distinct $\boldsymbol{\varphi}_1, \ldots, \boldsymbol{\varphi}_m$. Let $\hat{\psi}$ be the result of replacing $\neg, \wedge, \vee$, and $\rightarrow$ in ψ with their mentalese counterparts $\hat{\neg}, \hat{\wedge}, \hat{\vee}$, and $\hat{\rightarrow}$. The mentalese generalization corresponding to $\ulcorner$If $\boldsymbol{A}\ X_1$s that $\boldsymbol{\varphi}_1, \ldots,$ and $\boldsymbol{A}\ X_m$s that $\boldsymbol{\varphi}_m$, then $\boldsymbol{A}\ Y$s that $\psi\urcorner$ is the claim that any person with s_1 in their X_1-box, $\ldots,$ and s_m in their X_m-box has $\hat{\psi}[s_i/\boldsymbol{\varphi}_i]$ in their Y-box, where this is the result of replacing every occurrence of $\boldsymbol{\varphi}_i$ in $\hat{\psi}$ with s_i. (This is well-defined because $\boldsymbol{\varphi}_1, \ldots, \boldsymbol{\varphi}_m$ are pairwise-distinct.) Now suppose that $\ulcorner A\ X_1$s that $\varphi_1, \ldots,$ and $A\ X_m$s that $\varphi_m\urcorner$ is true relative to c. So for all i, $\Pi(c)$ contains a mentalese sentence s_i that both means $[\![\boldsymbol{\varphi}_i]\!]^c$ and is in $[\![A]\!]^c$'s X_i-box. So $\hat{\psi}[s_i/\boldsymbol{\varphi}_i]$ is in $[\![A]\!]^c$'s Y-box (by the mentalese generalization corresponding to our schema) and is also a member of $\Pi(c)$ (by closure). So $\ulcorner A\ Y$s that $\psi\urcorner$ is true relative to c. $\qquad\square$

[23] Goodman (2019) shows how Proposition 5 can be generalized to schemas with iterated attitude ascriptions in their consequents, such as "If $\boldsymbol{A}$ believes that $\boldsymbol{\varphi}$, then $\boldsymbol{A}$ believes that they believe that $\boldsymbol{\varphi}$". The challenge in accommodating such schemas is that the context-sensitivity of attitude ascriptions makes it non-obvious how to think about the mentalese generalizations corresponding to them.

[24] It is illustrative to consider how 10 might fail to be good even if 11 is true, if M$\approx$ is denied on an interpretation of $\approx$ as logical equivalence. Let π be the set of all Boolean combinations of $\{s\}$ for some mentalese sentence s expressing a contingent proposition. Suppose we accept the coarse-grained view that necessarily equivalent propositions are identical. Then π is not injective. But π is no counterexample to P$\approx$. Now suppose M$\approx$ fails for Marwan and "be sure": his "be sure"-box contains the self-implication $\ulcorner s\hat{\rightarrow}s\urcorner$ but fails to contain some complicated mentalese tautology $t \in \pi$. Suppose moreover that t is in Marwan's "be unsure"-box. Then 10 is not good, since "If Marwan is unsure whether grass is green if grass is green, then Marwan is not sure that grass is green if grass is green" is false relative to any context that supplies π. Yet for all we have said the corresponding mentalese generalization 11 is true.

[25] Grasping allows us to extend the argument in the proof of Proposition 6 to the case where the complement clause of the negated ascription in the consequent is a Boolean combination of complement clauses of ascriptions in the antecedent. The rest of the proof is a straightforward combination of the proofs of our earlier results.

[26] We won't consider views that reject articulation, since we think that, unless it is true, the connection between ordinary language and 'box'-ology is too tenuous to justify using the word "'believe'" to characterize the psychological relation between Lois and $\ulcorner s_0\hat{\wedge}\hat{\neg}c_0\urcorner$.

[27] Dorr (2016) defends such a view, and shows that it is not vulnerable to the paradoxes that afflict more naïve structured theories of propositions.

[28] A version of this problem can be generated without assuming transparency. The considerations that motivate closure apply to vocabulary other than Boolean connectives. For example, the schemas "If $\boldsymbol{A}$ believes that $\boldsymbol{\varphi}$, then $\boldsymbol{A}$ believes that the proposition that $\boldsymbol{\varphi}$ is true" and "If there is something that $\boldsymbol{A}$ believes is $\boldsymbol{F}$, then $\boldsymbol{A}$ believes that something is $\boldsymbol{F}$" have a similar flavor to schemas 15-17. So insofar as 15-17 are taken to motivate closure, these two schemas (or some minor weakening of them akin to the weakening 21 of 20) might likewise be taken to motivate the claim that perspectives are closed under truth introduction and existential generalization. If that is right, then any non-empty perspective contains a sentence s, and so contains $\ulcorner \hat{T}\hat{\#}s\urcorner$ (by truth introduction, where $\hat{T}$ and $\hat{\#}$ are the mentalese analogues of "is true" and "the proposition that") and hence contains $\ulcorner\hat{\exists}x\hat{T}x\urcorner$ (by existential generalization). So every non-empty perspective will contain $\ulcorner\hat{\exists}x\hat{T}x\urcorner$, which we may assume is in any rational person's "know"-box. But such a person may be subject to identity confusion with regard to the property of being true, just as an English speaker who doesn't know that "attorney" and "lawyer" are coextensive is identity-confused about being a lawyer. In particular, they may have a mentalese predicate V that expresses the property of truth without having $\ulcorner\hat{\exists}xVx\urcorner$ in their "know"-box. We might be inclined to convey this person's state of mind using the sentence "They don't know that something is veridical". But this ascription cannot express a true proposition relative to any non-empty perspective. Although one might quibble about the example, the point is general: the kind of considerations that motivate closure can be used to motivate other closure conditions on perspectives which may jointly imply that certain mentalese sentences which are in any sufficiently cognitively sophisticated person's "know"-box must be members of any non-trivial

perspective. We then risk falling into error when we attempt to use negated knowledge ascriptions to characterize such people's identity confusion concerning the properties and relations expressed by the elementary constituents of such mentalese sentences.

[29] This result should also be of interest for those who accept transparency but not grasping, since it shows that the restriction in Proposition 8 to Boolean k-combinations in complement clauses is unnecessary when the Boolean combination of attitude ascriptions is also a positive combination.

[30] All perspectivists will agree that there are some such schemas – e.g., "If A knows that φ, then φ" is not covered by our earlier results, but it is good provided every mentalese sentence in anyone's "know"-box is true.

[31] As a precedent for this proposal, note that all perspectivists must reject the box-based semantics for negated attitude verbs. That is, they cannot say that $[\![A$ does not X that $\varphi]\!]^c$ is true only if some $s \in \Pi(c)$ means $[\![\varphi]\!]^c$ and is in $[\![A]\!]^c$'s $\ulcorner$does not $X\urcorner$-box. This is because $\ulcorner$Either A X's that φ or A does not X that $\varphi\urcorner$ will then not be unequivocally true, despite being an instance of the law of excluded middle. (It will be false relative to any c such that no $s \in \Pi(c)$ means $[\![\varphi]\!]^c$.)

Williamson (2000, Ch. 1.3, Ch. 3) discusses different senses in which propositional attitudes might be analyzable in terms of others.

[32] Recall that even perspectivists who accept transparency should think that certain pre-theoretically interesting principles of propositional attitude psychology like 20 are in fact not good, for reasons having to do with identity confusion, and that we should instead focus our attention on appropriately qualified principles like 21.

[33] Does any sentence of pre-theoretical English have this property? One potential candidate might be "It is impossible that an ideally rational person believe the conjunction of two things they understand without also believing both conjuncts".

[34] This restriction should also be distinguished from the idea that principles of propositional attitude psychology are good only when restricted to some a special subject-matter about which identity confusion is impossible (e.g., one's 'sense data').

[35] Indeed, perhaps being ideally rational requires being identity confused: if the continuum hypothesis is true but any rational person should not be sure of it, rational people will be identity-confused, since they will have "$\hat{\aleph}_1 \hat{=} \hat{\aleph}_1$" but not "$\hat{\aleph}_1 \hat{=} \hat{\beth}_1$" in their "be sure"-box.

[36] Perhaps they sometimes correspond to complex expressions in their mental lexicons, but we will ignore this complication here, since a more general account would not issue different verdicts about Lois, whose identity confusion concerning Superman (we have been assuming) is a matter of her having distinct syntactically simple names of him in her mental lexicon.

[37] One might suppose that all such mentalese expressions should be included by default. We are reluctant to endorse this principle, however, since it would imply that we are never (or hardly ever) in contexts that supply transparent perspectives. Also, strictly speaking talk of *the* set of expressions that generate a perspective is illegitimate; for example, the empty set is generated by any set of elementary expressions insufficiently syntactically diverse to form sentences from.

[38] Elementary generation might also be thought to motivate closure, since it implies it given the assumption that every set of elementary expressions of mentalese that generates a perspective includes $\hat{\neg}$, $\hat{\wedge}$, $\hat{\vee}$, and $\hat{\rightarrow}$.

[39] Such perspectivists should hold that X^* is the relation that holds between a person a, proposition p, and set π of ways of thinking about entities just in case some mentalese sentence s means p, is in a's X-box, and contains only elementary expressions that (for a) correspond to some member of π.

[40] In the case of 22 the argument can be run with decomposition and closure in place of elementary generation. The same is not true of 23, since Lois's mentalese analogue of "Superman is Superman" is not a Boolean combination of subsentences of her mentalese analogue of "Superman is Clark".

[41] See Braun (1998) for a defense of a similar outlook.

[42] Salmon (1986) and Soames (1987b) are the classic neo-Russellian attempts to meet this challenge.

[43] See Schiffer (1979), Kripke (1979), and Dorr (2014).

[44] By contrast Fregeans are commmitted to rejecting Substitution; see Goodman and Lederman (forthcoming). Caie, Goodman, and Lederman (forthcoming) and Bacon and Russell (2019) discuss the challenges of developing an overall theory of identity and quantification without Substitution. These difficulties are our main reason for preferring perspectivism to Fregeanism.

[45] See Salmon (1986), Soames (1987b), and Braun (1998).

[46] The errors predicted by neo-Russellians in ordinary speakers' judgments about sentences like 23 are akin to the errors predicted by coarse-grained theories of propositions in ordinary speakers' judgments about attitude ascriptions like 24 whose complement clauses (according to such theories) non-obviously express logical truths. Given this kinship, it is not clear that neo-Russellians have a principled basis for objecting to coarse-grained theories on the grounds that they predict such errors, as, e.g. Soames (1987a) has influentially done.

[47] Most straightforwardly, they can identify senses with mental representations and have every attitude verb X express the relation of being in one's X-box.

[48] In a recent Kaplan-inspired discussion, Yalcin (2015) develops a semantics that, unlike Kaplan's, introduces existential quantification over "modes of presentation" where the variable is bound, rather than immediately before the attitude verb. As a result, Yalcin predicts that 27 is true just in case, for any x, there is a mentalese name n of x such that, if Lois has $\ulcorner Fn \urcorner$ in her "believe"-box, then Lois has $\ulcorner \hat{} Fn \urcorner$ in her "believe"-box. And this prediction is incorrect, since it implies that 27 will be true so long as every individual has some mentalese name which is not in Lois's mental lexicon, and that it must be false if any individual is such that there is no mentalese name of it.

[49] For work in this direction, see Aloni (2005), Dorr (2014), Bacon and Russell (2019) and Caie et al. (forthcoming).

[50] More precisely, they can hold that, if σ is good, then $\ulcorner \forall x_1 \dots \forall x_n \sigma[\varphi/\boldsymbol{\varphi}] \urcorner$ is good too, provided no variable free in φ becomes bound when it is substituted for an occurence of $\boldsymbol{\varphi}$ in σ, and each x_i is free in φ and not in σ.

REFERENCES

Aloni, M. (2005). Individual concepts in modal predicate logic. *Journal of Philosophical Logic*, *34*(1), 1–64.

Bacon, A., & Russell, J. S. (2019). The logic of opacity. *Philosophy and Phenomenological Research*, *99*(1), 81–114.

Braun, D. (1998). Understanding belief reports. *The Philosophical Review*, *107*(4), 555–595.

Braun, D. (2000). Russellianism and psychological generalizations. *Noûs*, *34*(2), 203–236.

Caie, M., Goodman, J., & Lederman, H. (forthcoming). Classical opacity. *Philosophy and Phenomenological Research*.

Crimmins, M. (1992). *Talk about Beliefs*. MIT Press.

Crimmins, M., & Perry, J. (1989). The prince and the phone booth: Reporting puzzling beliefs. *The Journal of Philosophy*, *86*(12), 685–711.

Dorr, C. (2014). Transparency and the context-sensitivity of attitude reports. In M. Garcia-Carpintero & G. Martí (Eds.), *Empty Representations: Reference and Non-existence* (pp. 25–66). Oxford University Press.

Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, *30*(1), 39–134.

Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.

Goodman, J. (2019). First-person perspectivism. In progress.

Goodman, J., & Lederman, H. (forthcoming). Sense, reference and substitution. *Philosophical Studies*, 1–6.

Kaplan, D. (1968). Quantifying in. *Synthese*, *19*, 178–214.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.

King, J. C. (2007). *The Nature and Structure of Content*. Oxford University Press.

Kripke, S. A. (1979). A puzzle about belief. In A. Margalit (Ed.), *Meaning and use* (pp. 239–283). Dordrecht: D. Reidel.

Montague, R. (1973). The proper treatment of quantification in ordinary english. In P. Suppes, J. Moravcsik, & J. Hintikka (Eds.), *Approaches to Natural Language* (pp. 221–242). Dordrecht.

Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, *33*(5), 355–380.

Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian society, supplementary volumes*, *60*, 137–152.

Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, *189*(1), 75–96.

Richard, M. (1990). *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*. Cambridge University Press.

Salmon, N. U. (1986). *Frege's Puzzle*. MIT Press.

Saul, J. M. (1997). Substitution and simple sentences. *Analysis*, *57*(2), 102–108.

Saul, J. M. (2010). *Simple Sentences, Substitution, and Intuitions*. Oxford University Press.

Schiffer, S. (1979). Naming and knowing. In A. F. Peter, E. U. Theodore, Howard Jr, & K. Wettstein (Eds.), *Midwest Studies in Philosophy II: Contemporary Perspectives in the Philosophy of Language* (pp. 28–41). University of Minnesota Press.

Schiffer, S. (1992). Belief ascription. *The Journal of Philosophy*, *89*(10), 499–521.

Soames, S. (1987a). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, *15*(1), 47–87.

Soames, S. (1987b). Substitutivity. In *On Being and Saying: Essays for Richard Cartwright* (pp. 99–132). Cambridge, MA: MIT Press.

Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.

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

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New Essays on the Foundations of Meaning* (pp. 17–54). Oxford University Press.

Yalcin, S. (2015). Quantifying in from a fregean perspective. *Philosophical Review*, *124*(2), 207–253.

Yli-Vakkuri, J. (2013). Propositions and compositionality. *Philosophical Perspectives*, *27*(1), 526–563.