

Scott Soames
10/13/2022 Draft
Cognitive Propositions:
What's structure got to do with it?

Today I will present a thought experiment on the metaphysics of propositions. Having previously rejected propositions as sets of truth-supporting circumstances, or structured Russellian propositions, I recently turned to propositions as representational cognitive act or operation types. I did so for four main reasons. First, it allows us to avoid the Platonic conception of *entertaining a proposition*, which I never understood. Second, it explains the ability of primitive agents to assume, believe, or know things, even if they can't cognize propositions and predicate properties of them. Third, it recognizes that just as propositions impose conditions on the world that must be satisfied if the proposition is true, so they impose conditions on the mind that bears cognitive relations to them. Fourth the fine-grained identity conditions of acts allow us to make progress on Frege's puzzle.

Recently I revisited Wittgenstein's tractarian doctrine that an elementary proposition is "*a propositional sign in its projective relation to the world.*" For him, propositional signs are often sentences, which he took to be facts in which symbols stand in syntactic relations. We use these facts to picture, or represent, non-linguistic facts, in which objects and properties stand in relations. The next step, which I took, but he didn't, is to take *uses of atomic sentences* to be elementary propositions. Since uses are acts, I took these propositions to be repeatable cognitive act or operation types, performances of which are events in which sentences are used in accord with their governing conventions. This led me to wonder whether the tractarian claim that *truth-functionally equivalent propositions* are identical could be maintained, while rejecting the claim that necessarily equivalent propositions are too. Exploring this requires thinking about the metaphysics properties and propositions. What follows is my thought experiment.

Imagine a language, **L1**, the vocabulary of which consists of names, general terms used to form predicates and operators on predicates. L1 is used to talk about a tiny geometrical universe containing squares, circles and triangles each of which is entirely red, green, or

yellow. Names are Millian and simple general terms stand for *being red/green/yellow* plus *being triangular/circular/square*. Every object has just one of each trio; ‘and’, ‘or’ and ‘not’ denote functions from properties to properties, giving us colors and shapes definable from simple ones. I ignore null properties like *being red and green*.

All sentences consist of a name and a predicate, but some predicates are complex. ‘Not’ maps properties onto their complements; ‘and’ maps pairs of properties onto intersective properties e.g. *being red and triangular*; ‘or’ maps pairs onto “additive” properties; e.g. *being green or yellow*. *Being not red* is true of all and only circles, squares, or triangles that are either green or yellow. *Being green or yellow* is a color, called ‘grelow’ in the metalanguage; it is true of all green plus all yellow shapes and nothing else. *Being grellow* isn’t composed of *being green* and *being yellow* plus any metaphysical building block corresponding to ‘or’. It’s metaphysically simple. *Not being red = being green or yellow*. Similarly for *not being square* and *being circular or triangular*. *Being red and triangular*, and *being green or circular* are also simple.

Now consider *being red*, *being red or red*, *being red and red*. Each is just being red, as is, *not being not being red*. Since *not being red is just being grellow*, and *not being not being red is just not being grellow* it follows that *not being grellow is being red*. What about *not being red* and *not being square*. The former is *being yellow or green (grelow)*; the latter, *not being square*, is *being circular or triangular (cirtangular)*. *Not being red and not being square* is *being grellow and cirtangular*. Also *not being red or not being square* is *being grellow or cirtangular*. Finally *not being (grelow or cirtangular) = being red and square*. **For any general terms P and Q, *being (not P) and (not Q) = not being (P or Q)* and *being (P or Q) = not being (not P and not Q)*.** In short, *truth-functionally equivalent predicates* in L1 express the same properties.

For each name and predicate, there are infinitely many sentences that predicate the same property of the same thing. If we ignore cognitive complexity in reaching representationally identical ends, we have different ways of expressing same proposition --though we could, if

we wished, build cognitive complexity into identity conditions of some propositions. Later I will do that, but for now I put it aside.

The result is tractarian. In the Tractatus, propositions are pictures of worldly facts. Their chief interest is in what they represent. Hence, propositions that represent the same things as being the same ways are identical. That must change when we consider attitudes of agents. But even in L1 we shouldn't take representationally identical sentences to be synonymous, provided that we tie *knowing the meaning of S* to *understanding S*, which requires understanding its vocabulary and its grammatical constructions. Although the propositions expressed by truth functionally equivalent sentences of L1 are identical, that proposition isn't the meaning of those sentences.

Accommodating Sentential Truth-Functional Operators

Can we extend our results to a language L2, with the same names and simple predicates as L1, but a different syntax? In L2 'and', 'or', and 'not' combine with *sentences* to form compound *sentences*. First, conjunction and disjunction. Given the sentences 'A is red' and 'B is square', expressing the propositions *that a is red* and *that b is square*, we first map those propositions onto the properties-- *being such that a is red* and *being such that b is square*, each of which is true of everything or nothing. Next, we conjoin (or disjoin) these properties, giving us *being such that a is red and/or b is square*. Predicating them of a and b, gives us the propositions *that a is red and/ or b is square*.

One of these represents a as being red and b as being square, and nothing more; the other represents a as being red or b as being square, and nothing more. In moving from 1-place to what we might call *0-place* properties, we haven't changed the metaphysics of properties. For any properties P, Q, their conjunctions/disjunctions are metaphysically simple properties that are true of objects that the original properties are both true of or that at least one is true of. Similar results follow for negation.

Next consider the conjunction/disjunction of the proposition *that a is red* with itself. It predicates, of a, *being such that a is red and/or being such that a is red* which is just *being such that a is red*. Predicating this property of a and predicating *being red* of a are cognitively

different but representationally identical acts. When we have two ways of reaching a single representational end, we can abstract a higher-order act *representing a as red*, which can be performed either way. Let it be the ur-proposition *that a is red*--the act type that is performed no matter whether one performs Act 1, predicating *being red* of *a*, or one performs Act 2, by also generating *being such that a is red* and predicating it of *a*, or one performs Act 3 by generating *being such that a is red*, and *being such that a is red*, conjoining/disjoining the two, and predicating the result of *a*. Each is a way of performing the higher-order act type A4 of *representing a as red*, just as traveling to work is a higher-order act that is always performed by performing a more specific act.

Similarly for double negation. I start by predicating *being red* of *a*, from which I form *being such that a is red*. Negating this property, I get *not being such that a is red*; predicating it of *a* gives me the proposition *that a is not such that a is red*. Repeating the process, I predicate *not being such that a is not such that a is red* of *a*, which is just *being such that a is red*. Predicating this of *a* gives us the proposition *that a is such that a is red*, making the entire cognitive act another instance of ur-proposition *that a is red*.

That is how we replicate results of L1 in L2. Next consider sentences 1-4.

1. $\sim(a \text{ is red or } (b \text{ is circular or } b \text{ is triangular}))$
2. $\sim(a \text{ is red}) \text{ and } \sim(b \text{ is circular or } b \text{ is triangular})$
3. $\sim(a \text{ is red}) \text{ and } (\sim b \text{ is circular and } \sim b \text{ is triangular})$
4. $(a \text{ is green or } a \text{ is yellow}) \text{ and } (b \text{ is square})$.

3a λ . $\lambda x [\sim x \text{ is red}] a \quad \rightarrow \quad (a \text{ is green or } a \text{ is yellow})$

3b λ . $(\sim b \text{ is circular and } \sim b \text{ is triangular}) \quad \rightarrow \quad \lambda x [\sim(x \text{ is circular and } \sim x \text{ is triangular})]b$

The moves from 1-3 just play with negation, conjunction, and disjunction. They are fine. If 3a λ and 3b λ were both fine they would get us to 4. But whether the application of 3b λ to 3 preserves propositional identity is questionable, which requires us to look at propositional attitudes

IV: Cognitive Propositions and Propositional Attitudes

Consider 5, which is used to represent X as asserting/believing one or more propositions expressed by a reporter's use of a sentence S in which a simple singular term designates an

individual *a*, of which a property *P* is predicated. The semantic content of *S* represents *a* as having *P*, without constraining how *a* or *P* are cognized.

5. *X* asserted/believed that *S*

Since propositions are cognitive acts that involve *sub acts* of identifying objects and properties, we **always** entertain bare propositions by entertaining cognitively enriched versions of them. The enrichments are ways of bringing objects and properties to mind. These include identifying objects and properties perceptually, linguistically (using specific words or phrases) and identifying them via *de se* and/or *present-tense* cognitions. Abstracting from these, we reach bare propositions the entertainment of which places *no restrictions on how the properties or objects are brought to mind*. By entertaining the enriched proposition, we entertain the bare proposition too. This allows us to use fine-grained identity conditions of act-types in dealing with instances of Frege's puzzle.

Consider uses of (6a,b), which express representationally identical but cognitively distinct propositions, making different truth values of (7a,b) possible.

6a. He is in danger (said by me gazing at my unrecognized reflection in a window.

b. I am in danger. (considered in the same situation)

7a. I believe that he is in danger. (True, as said by me)

b. I believe that I am in danger. (False as said by me)

Next consider (8), in which 'logicism' is a Millian name for the proposition *L*, *that arithmetic is reducible to logic*, which is designated by the *that*-clause in (8a). Since the name and the clause make the same contribution to the representational contents of (8a,b), they may be used to express representationally identical, but cognitively different truths.

8a. Russell tried to prove (the proposition) *that arithmetic is reducible to logic*.

b. Russell tried to prove *logicism*.

Entertaining, asserting or believing those propositions requires predicating *trying to prove* of Russell and *L*. Uses of (8a) also require identifying *L* by entertaining it. Since uses of (8b) don't require any specific cognitive enrichments, the truth of (9a) is sufficient for the truth of (9b), but not conversely

9a. Mary believes that Russell tried to prove that arithmetic is reducible to logic.

b. Mary believes that Russell tried to prove logicism.

Next consider (10 a-c).

10a. Logicism is (the proposition) that arithmetic is reducible to logic.

b. Logicism is logicism.

c. That arithmetic is reducible to logic is that arithmetic is reducible to logic.

Anyone who accepts the semantic content of (10a), must identify the second argument, L, of the identity relation by entertaining L; for (10c) the requirement applies to both arguments; for (10b) it applies to neither. Anyone who believes/knows proposition semantically expressed by (10c) believes/knows the propositions semantically expressed by (10a) and (10b). Since (10c) is knowable apriori, the propositions expressed by all three are too, because there is *a way of having it in mind* that allows one to see its truth, without requiring empirical evidence.

This brings us to ‘Hesperus’/‘Phosphorus’. Like other names, their semantic contents are their bearers. But unlike most names, there is more to *understanding* them. One who uses them is expected know that uses of ‘Hesperus’ designate something visible in the evening while uses of ‘Phosphorus’ designate something visible in the morning. To mix this up is to *misunderstand* the names. Next consider (11).

11a. Hesperus is a planet.

b. Phosphorus is a planet.

c. x is a planet. (with Venus as value of ‘x’)

Let proposition P be the semantic content of all three. Let P_H and P_P be cognitive enrichments of P that require identification of the planet via the name ‘Hesperus’ and ‘Phosphorus’ respectively. Next consider A’s use of sentence (12), addressing B, each presupposing they *understand* the names.

12. Hesperus is Phosphorus

A asserts not only the bare proposition P that is the semantic content of (12), *but also the proposition entertainable only by identifying Venus via the two names*. Presupposing that A *understands* the names, B reasons that A knows A will be taken to be committed to the claim that *Hesperus, which is visible in the evening is Phosphorus which is visible in the morning*. Thus, B concludes that A *asserted* this enriched proposition (plus the bare, unenriched proposition).

This proposition is contingent even though A's assertive utterance of (13) is true.

13. Necessarily Hesperus is Phosphorus.

The extra representational content carried by A's utterance of (12) arises from the presupposition that A and B understand the names. This presupposition is also present when A utters (12), but it doesn't lead to the same enrichment. *Understanding the names requires knowing that most agents who use them take, and expect others to take, one to stand for something seen in the evening and the other to stand for something seen in the morning.* Presupposing that A and B understand the names, both add descriptive content to A's utterance of (12). Since taking the names to refer to things *actually* seen at certain times tells one *nothing* about when they are seen at other *possible* world-states, they *don't* descriptively enrich *under the modal operator* in (13).

Transparent vs. Non-Transparent Representational Identity

Now back explaining why we haven't shown that truth-functionally equivalent propositions expressed in L2 are identical. Our earlier argument used lambda abstraction to turn a conjoined sentence involving predications of different 1-place predicates of a single object into a sentence involving a single predication of a complex predicate. This allowed us to use the metaphysical assumptions of L1 to guarantee representational identity. But in doing so we collapsed a crucial cognitive difference.

Look again at the sentences in (1-3), uses or which are *different ways of entertaining the same abstract proposition, identified by its representational content alone.* Next, we must show (4) has the same representational content as (3), using $3a\lambda$ and $3b\lambda$

$3a\lambda. \lambda x [\sim x \text{ is red}] a \quad \rightarrow \quad (a \text{ is green or } a \text{ is yellow})$
 $3b\lambda. (\sim b \text{ is circular and } \sim b \text{ is triangular}) \quad \rightarrow \quad \lambda x [\sim(x \text{ is circular and } \sim x \text{ is triangular})]b$

($3a\lambda$) plus our account of properties tells us that for **a** not to be such that **a** is red, is for **a** not to be red, which is for **a** to be *green or yellow*. So far, so good. ($3b\lambda$) tells us that for **b** not to be such that **b** is circular or **b** is triangular is for **b** to be square; so predicating *not*

being circular or triangular of b represents b as square. This tempts us to think that propositions (3,4) are identical. But there is a well-known objection to this reasoning.¹

Imagine ancients who, when they looked at Venus in the evening, called it ‘Hesperus’, while calling it ‘Phosphorus’ in the morning -- not realizing they had named the same thing twice. Today John reports them by assertively uttering (14a).

14a. The ancients believed/asserted that ‘Hesperus’ referred to Hesperus and ‘Phosphorus’ referred to Phosphorus.

Assuming Millianism about names we get (14b).

14b. The ancients believed/asserted that ‘Hesperus’ referred to Venus and ‘Phosphorus’ referred to Venus.

In deriving this we aren’t saying that John didn’t **also** assert and believe other propositions too, including, perhaps one requiring Venus to be cognized once via the name ‘Hesperus’ and once via ‘Phosphorus’. Perhaps he did. But since that also counts as asserting/believing (14b), the move to (14b) stands.² Lambda abstraction gives us (14c,d) which are false.

14c. The ancients believed/asserted that λx [‘Hesperus’ referred to x and ‘Phosphorus’ referred to x] Venus.

14d. The ancients (believed/asserted) that ‘Hesperus’ and ‘Phosphorus’ are coreferential.

The problem isn’t with lambda abstraction *per se*, but with applications of it that collapse multiple occurrences of a name A in a clause, transforming it from one that predicates an n-place property of an n-tuple (containing multiple occurrences of the referent of A) into the predication of a related, m-place property of a related m-tuple, where m is less than n, and the referent of A occurs less often in the m-tuple than in the n-tuple. The case in which the original clause has just two occurrences of the name is illustrated by the contrast between (15 a,b,c). on one hand, and (15d) on the other.

15a ‘Hesperus’ referred to Venus and ‘Phosphorus’ referred to Venus

15b λx (‘Hesperus’ referred to Venus and ‘Phosphorus’ referred to x) Venus.

15c λxy (‘Hesperus’ referred to x and ‘Phosphorus’ referred to y) Venus, Venus

¹ The objection is presented in Soames ([1987] 2009). A prominent objection is raised in Edelberg (1994), and refuted in Soames ([2008] 2009).

² Soames (2015) pp. 18-19, 22-25, 73-79, 157-58.

If the complement clause (15a) of (14b) is replaced by (15 b,c), truth is preserved and (14c,d) are not derived. Only the move to (15d) as the content of the complement clause is questionable.

15d. λx ('Hesperus' referred to x and 'Phosphorus' referred to x) Venus.

The move is tempting because it is hard to deny that propositions (15a-d) are *representationally identical*. But they do differ in informativeness. Whereas entertaining and accepting (15d) requires **recognizing** Venus as the *common referent* of the names, entertaining and accepting (15a-c) may provide no way of recognizing this, as when one cognizes Venus via different Millian names or visual perceptions.³

The difference between propositions (15a-c) and (15d) is cognitive; (15a-c) are representationally identical to (15d), but not *transparently* so. Any way of cognizing and accepting (15d) provides one with a way of cognizing and accepting (15a-c), but not conversely. Propositions P, Q are transparently representationally identical *iff* agents who entertain and accept one are always in position to derive the other, as with (15a-c). If *recognition of recurrence* is included in the *cognitive contents of propositions* expressed by uses of sentences 15a-c, then the *representational identity* of those enriched propositions with proposition 15d *is transparent*. Otherwise not.

Disjunctive examples like 16 lead to the same result.

- 16a. 'Hesperus' refers to Venus **or** 'Phosphorus' refers to Venus
b. λx ('Hesperus' refers to Venus **or** 'Phosphorus' refers to x) Venus.
c. λxy ('Hesperus' refers to x **or** 'Phosphorus' refers to y) Venus, Venus
d. λx ('Hesperus' refers to x **or** 'Phosphorus' refers to x) Venus.

Next consider sentences in (17), containing names for Venus, the understanding of which *doesn't* require knowing when their referents thought to be visible.

- 17a. *H is the only planet regularly seen in the evening*
b. *Ph is the only planet regularly seen in the evening*
c. *H is the only planet regularly seen in the evening or Ph is the only planet regularly seen in the evening*

³ These are cases in which the agent doesn't recognize the recurrence of the entity Venus in the agent's act of entertaining the proposition (or, sometimes, of recurrence of the same name in a sentence used to express it). See chapters 6-8 of Soames (2015) for discussion.

The representational contents of uses of (17a,b) are the same. Now consider two uses of sentence (17c); one involving *recognition of recurrence of Venus* in the proposition asserted/believed, and one that doesn't. Here we distinguish 3 act types. **A1** predicates being the only planet regularly seen in the evening of Venus, with no restrictions on how Venus is identified. **A2** involves performing A1 twice, disjoining the relevant 0-place property with itself, *recognizing recurrence* of Venus (as part of cognitive content), and predicating that property of Venus. **A3** is like **A2** but without recognition of recurrence.

How are these act types related? Consider the act **A1** of predicating *being the only planet regularly seen in the evening* of Venus with no cognitive restrictions on how property and object are identified. Its representational content represents Venus as having the property. It is an instance of *proposition Alpha*, which is the abstract act type with this representational content, without requiring recognition of recurrence (if such there be). I don't say that A1 is proposition Alpha; I say A1 is an instance of Alpha, by which I mean that to perform A1 is to perform Alpha, but not conversely.

Next consider the act of disjoining Alpha with itself. Since it is a proposition, it is natural to think we can do this. But there's a wrinkle. Since Alpha is a higher-order propositional act type that is performed by performing any of its lower-level instances, the *disjunction of Alpha with itself* can be performed in many ways – e.g., disjoining A1 with itself, which is just A3, by disjoining A3 with itself, and by disjoining A1 with A3, each of which is representationally redundant. Thus, it seems the disjunction of Alpha with itself is just Alpha. In this case, at this abstract level, the disjunctive operation disappears. Next consider a different, but representationally identical higher-level act **Beta**, which requires recognizing all recurrences, if any, of Venus. A1 and A2 are instances of Beta, while A3 is not. Thus, Alpha and Beta are representationally identical but cognitively distinct in a way that renders representational identity of Alpha and Beta non-transparent.

What about belief? It's a short step from affirmatively performing **A2** (which is an instance of Beta) to performing A1, which is an instance of *Alpha*. So believing *Beta* is tantamount to believing *Alpha*. This contrasts with affirmatively performing **A3** (which is an

instance of Alpha alone). When one believes *Alpha* in this way, there may be no direct route to believing *Beta*. Next, a delicate point. Although A2 and A3 are instances of different propositions, each may seem to be a candidate for being *the disjunction of the bare proposition of which A1 is an instance*. What's going on? Here is a point to remember. **If**

- (i) propositional identity is simply sameness of representational content plus sameness of cognitive content (involving requirements on how items are cognitively identified), and
- (ii) propositions are what we believe assert and know, **then**
- (iii) it seems that one can believe, assert, and know the bare proposition expressed by sentence (17a) -- *that H is the only planet regularly seen in the evening* -- by believing, asserting, or knowing the bare proposition expressed by sentence (17c) -- ***that either H is the only planet regularly seen in the evening or Ph is the only planet regularly seen in the evening*** – *without recognizing recurrence of Venus and without accepting any sentence or representation one uses merely to predicate the relevant property of Venus*.

Can we accept this result? If so, *what is the disjunction, with itself, of proposition Alpha, of which the performance of A1, and the performance of A3 are instances?* Presumably, the disjunction of *Alph* with itself, is the one the representational content of which simply represents Venus as being the only planet regularly seen in the morning, with no requirements on how that content is cognitively determined. This describes the bare proposition associated with uses of (17a, b, c). But then the assertive content of uses of these sentences should include a common proposition, and the assertive content of positive *belief or assertion ascriptions* involving those sentences *as complement clauses* should all attribute belief or assertion of this common proposition.

Still, one may wonder. Earlier I mentioned *the disjunction, with itself, of the bare proposition of which the performance of A1 is an instance*. What, if anything, does it mean to call a proposition *disjunctive*? The bare proposition (without extra cognitive content) expressed by uses of *sentence* (17c) is also expressed by single clause sentences (17a,b). What, then, is the disjunctive operator doing? If, in (17c), the words of the second disjunct don't add anything to the proposition expressed, why wouldn't one of the *sentential* disjuncts be used on its own by the speaker who chooses a disjunctive sentence? The answer

of course, is that the speaker may not know that the *sentential disjuncts* represent the same thing as being the same way.

The speaker who assertively utters (17c) uses the disjunctive operator on the pair consisting of the 0-place property *being such that Venus is the only planet regularly seen in the evening*, taken twice over, without recognizing the recurrence. Since the disjunctive operator maps property pairs. PP onto property P, the representational content of the use of (17c) is the same as uses of (17a,b). It simply represents Venus as being the only planet regularly visible in the evening, *without placing cognitive requirements on how that content is reached*. One who understands and accepts any of the sentences in (17) counts as believing this proposition. Earlier I identified it with Alpha, which encompasses any and all subordinate acts the representational content of which is that Venus is the only planet regularly seen in the evening (irrespective of further cognitive restrictions). Thus, it seems, propositional identity sometimes abstracts from sharp differences in sentence structure.

What about the truth functional equivalent sentence (17c*)?

17c*. $\sim [\sim H \text{ is the only planet regularly seen in the evening and } \sim Ph \text{ is the only planet regularly seen in the evening}]$

If our metaphysics of properties and relations carries over from the discussion of L1, we may take the representational content of (17c*) to match that of (17c). Thus, we seem to be well on the road to recapitulating the results of L1 in L2.

Let's put aside any further steps needed to demonstrate that the representational contents of uses of truth-functionally equivalent sentences of L2 will always be identical. Think instead about ascriptions *A believes that S* and *A believes that S**, where S and S* are used to express different cognitive enrichments of the same purely representational proposition. Such ascriptions can be used to report A as believing different *cognitively enriched* propositions associated with the complement clauses while simultaneously reporting belief in the same purely representational proposition.

With this in mind, consider the sentences in (18).

- 18 a. A asserts/believes that *H is the only planet regularly seen in the evening*
b. A asserts/believes that *Ph is the only planet regularly seen in the evening*

c. A asserts/believes that *H is the only planet regularly seen in the evening or Ph. is the only planet regularly seen in the evening.*

Suppose A understands sentence (17c) and knows it to be true because she knows that one of the two names designates the planet in question, while suspending judgment on (17 a,b) due to uncertainty whether ‘H’ and ‘Ph’ designate the same thing. Since the *representational content* of the complement clauses in (17a-c) is the bare proposition *that Venus is the only planet regularly seen in the evening*, I ought to be able to use (18c) to report either A’s belief in the cognitively enriched proposition (18c) that corresponds most closely to the meaning of its complement clause, sentence, or A’s belief in the higher order proposition that predicates, of Venus, *that it is the only planet regularly seen in the evening*, without imposing cognitive requirements on how this comes about. The latter belief can also be correctly reported using (18a,b), even though A doesn’t accept their complement clauses, (17a,b).

Can this be correct? Suppose A understands sentence (17c) and knows it to be true because she knows that one of the two names designates the planet in question, while suspending judgment on (17a,b). Since the *representational content* of the clauses in the belief reports (18a-c) is the bare abstract proposition *that Venus is the only planet regularly seen in the evening*, I should be able to use (18c) to report A’s belief in the higher-level proposition that simply represents Venus as the only planet regularly seen in the evening. That belief can, of course, also be correctly reported using (18a,b), even though A doesn’t accept their complement clauses (17 a,b). What should we say about this?

Suppose A is a contestant in a game in which players will be asked *What is the only planet regularly seen in the evening?* to which either answer *H*, *Ph* or *Venus* is correct. Suppose we care that A gets it right. Suppose we also know that A doesn’t know which of the names designate the planet in question. This may worry you for various reasons, perhaps because you are in the same position as A. If so, this may lead you to suspect that A will be inclined to say, “The planet is H or Ph,” while confessing that she can’t be more specific. I, on the other hand, know that ‘H’ and ‘Ph’ are coreferential, leading me to be confident that A’s prospective use of (17c) is enough. If all this is so, can I reassure you by truthfully

reporting that *A does know that Venus is the only planet regularly seen in the evening* and so will answer the question correctly?

I suspect I could, even though in other circumstances I might use the *negations* of (18a, b,) in reporting A's beliefs. Those negations may correctly be used report that A doesn't believe representationally identical propositions that require either the use of the name 'H', or the name 'Ph' to designate Venus. In such a case, my negative belief report will be true if the content of complement clause is understood as including such cognitive content, over and above representational content. Still, the fact that A doesn't believe those cognitively enriched propositions doesn't show that A doesn't believe a representationally identical proposition that places no cognitive requirements on the route to that content.