

Transparency and the Context-Sensitivity of Attitude Reports

Cian Dorr

Forthcoming in *Thinking and Talking about Nothing*, ed. Manuel Garcia-Carpintero and Genoveva Martí (Oxford University Press, 2012)

1 Transparency

Kripke convinced most of the philosophical community that proper names are *rigid*. Kripke glosses rigidity in terms of the highly theoretical concept of “designation relative to a possible world”. But the basic idea behind rigidity is not at all technical. It is simply that the following schema is valid:

Modal Transparency Schema: For all x , if $x = N$, then necessarily $\phi(x)$ iff necessarily $\phi(N)$.

Here ‘ N ’ is to be replaced with some ordinary proper name; ‘ $\phi(N)$ ’, with some sentence involving this name, in which it is embedded only under truth-functional operators; and ‘ $\phi(x)$ ’, with the result of replacing some occurrences of ‘ N ’ in that sentence with occurrences of ‘ x ’.

The main thing to be said in favour of the Modal Transparency Schema is that it just seems obviously valid. Consider the following dialogue:

A: There is one person who is necessarily a philosopher

B: Really? Who?

A: Aristotle.

B: So it is necessary that Aristotle is a philosopher?

A: No.

B: What? Didn't you just tell me that the person you mentioned earlier, who is necessarily a philosopher, was Aristotle?

A: Yes. But I don't think it's necessary that Aristotle is a philosopher. I only think that there is some x who is identical to Aristotle, such that it's necessary that x is a philosopher.

B: I don't get it. Who is this person x ?

A: Aristotle, of course.

B: ???

We should sympathise with B's reaction here. Our ordinary competence with modal operators simply doesn't equip us to process the kind of position A is trying to stake out. Of course, with sufficient exposure to A's discourse, B could eventually overcome the disposition to react in a way that takes the Modal Transparency Schema for granted. If we want to engage with logical heretics, we must train ourselves not to rely implicitly on inferences which the heretics reject. But the distinctions that we would need to draw in order to follow A's view are profoundly alien: ordinary usage just doesn't seem to keep track in the relevant way of the scope relations between names and modal operators.

That is just about all that there is to the case for the rigidity of ordinary proper names. Kripke's achievement was not to uncover some clever new argument, but simply to make the appearances vivid, while clearing away some theoretical obstacles which had previously made people think that names *couldn't* be rigid (see Stalnaker 1997). It is a testament to the deep-rootedness of our dispositions to reason in accordance with the Modal Transparency Schema that this was enough to convince most philosophers of the rigidity claim.

Now consider the following close cousin of the Modal Transparency Schema:

Attitude Transparency Schema: For all x and y , if $x = N$, then y ψ s that $\phi(x)$
iff y ψ s that $\phi(N)$.

Here ' N ', ' $\phi(N)$ ' and ' $\phi(x)$ ' are to be substituted for as before, while ' ψ ' can be replaced with any propositional attitude verb, such as 'believe' or 'desire' or 'know' or 'say'.¹

There is a famous argument against the Attitude Transparency Schema, which we will engage with in the next section. For now, try to set it aside, so as to take note of the strong *prima facie* appearance that the Attitude Transparency Schema is valid. In fact, the intuitive case for the validity of the Attitude Transparency Schema seems roughly as strong as the (almost universally accepted) intuitive case for the validity of the Modal Transparency Schema. Attempts to soberly spell out views which conflict with the Attitude Transparency Schema seem apt to induce befuddlement:

A: There is a famous world leader whom Joe believes to have a black belt in judo.

B: Really? Who?

A: Vladimir Putin.

B: And why does Joe believe that Putin has a black belt in judo?

A: No—Joe *doesn't* believe that Putin has a black belt in judo.

B: What? I thought you just said he did!

A: No, all I said was is that there is a famous world leader, identical to Putin, whom Joe believes to have a black belt in judo.

B: I'm lost! Who is this famous world leader believed by Joe to have a black belt in judo, if not Putin himself?

A: No, it *is* Putin, that's what I'm telling you.

¹I have to admit that the present chapter, rather incongruously given the title of the volume in which it appears, has nothing to say about empty names. I have been careful to formulate the Attitude Transparency Schema in such a way that it is compatible with a broad range of views about attitude reports involving empty names. Let me however record my sympathy for the view that quantifiers have multiple admissible interpretations, some of which validate the inference from $\lceil y \psi \text{ that } \phi(N) \rceil$ to $\lceil \exists x(y \psi \text{ that } \phi(x)) \rceil$. See Hofweber 2000 and Dorr 2005.

B: So Joe does believe that Putin has a black belt in judo?

A: No. All that's true is that Putin is someone Joe believes to have a black belt in judo.

B: ???

Although views that make the kinds of distinctions A is relying on have been widely accepted amongst philosophers, it is hard not to sympathise with B's reaction here. Our ordinary competence just does not equip us to keep careful track of distinctions between sentences such as the following:

Joe believes that Putin has a black belt

Joe believes Putin to have a black belt

Putin is believed by Joe to have a black belt

Putin is such that Joe believes he has a black belt

Putin is someone Joe believes to have a black belt

Of course we can learn to be sensitive to some of the differences here, as many philosophers have. But in view of the parallels between the case of modal operators and that of attitude verbs, it is natural to wonder whether this sensitivity is as misplaced in the latter case as it was in the case of the former.

2 The standard argument against Attitude Transparency

The big difference between the Attitude Transparency Schema and the Modal Transparency Schema lies not in their intuitive obviousness when considered in isolation, but in the existence of a powerful and influential argument against the claim that the Attitude Transparency Schema is valid, to which we must now turn. Here is one version of the argument, which has the advantage of vividness and the disadvantage that it requires us to pretend that the Superman stories are true. It takes the form of a *reductio*:

- (1) For all x and y , if $x = \text{Superman}$, then y believes that x flies iff y believes that Superman flies.
- (2) For all x and y , if $x = \text{Clark Kent}$, then y believes that x flies iff y believes that Clark Kent flies.
- (3) So, for all x and y , if $x = \text{Superman}$ and $x = \text{Clark Kent}$, then y believes that x flies iff y believes that Superman flies, and y believes that x flies iff y believes that Clark Kent flies.
- (4) So, if $\text{Superman} = \text{Clark Kent}$, then for all y , y believes that Superman flies iff y believes that Clark Kent flies.

(1) and (2) are instances of the Attitude Transparency Schema. If it is a logically valid schema, then they are logically valid sentences. (3) is an obvious logical consequence of (1) and (2). (4) is an obvious logical consequence of (3). Logical consequences of validities are themselves valid. So if the Attitude Transparency Schema is valid, so is (4). And this is a problem, because (4) seems obviously false, while (5) seems obviously true:

- (5) Even though Superman is Clark Kent, lots of people believe that Superman flies without believing that Clark Kent flies.

We are thus faced with a paradox: a seemingly valid argument from seemingly true (indeed, seemingly valid) premises to a seemingly false conclusion. The most widespread reaction to this paradox has been simply to deny that the Attitude Transparency Schema is valid.

There is, however, a highly visible minority which has held on to the Attitude Transparency Schema, and endorsed the validity of (4). Members of this tradition have generally said one of two things about ordinary speakers' dispositions to assent to (5) and reject (4).

- (i) *Pragmatic Millianism* (Salmon 1986, Soames 1987): When ordinary speakers utter a sentence like (5), they are not speaking literally. They are not really

endorsing, or committing themselves to, the false proposition semantically expressed by (5). Rather, they are using (5) as a tool for conveying some other, true proposition, which they believe. It is because (5) would typically be used in this way, to communicate a truth, that it strikes us as true: our untutored reactions are not a reliable guide to the distinction between sentences which semantically express true propositions and those which are merely standardly used to convey true propositions. Similarly, if anyone were to utter (4) outside of a philosophical context, they would probably be using it to communicate not only the true proposition that it semantically expresses, but some further false propositions which are inconsistent with the true propositions that would be communicated by an ordinary utterance of (5). This fact about (4) explains our disposition to classify it as false.

- (ii) *Error-Theoretic Millianism* (Braun 1988, Braun 2002, Saul 2007): Ordinary speakers are systematically mistaken about the facts. For example, they believe the false proposition that even though Superman is Clark Kent, some people believe that Superman flies without believing that Clark Kent flies. They express their belief in this false proposition by uttering (5), which semantically expresses it; such utterances are perfectly literal. These false beliefs are due to some systematic defect in our untutored abilities to engage in elementary reasoning. By doing philosophy, we can strive to overcome these false beliefs about belief, and attain an enlightened view.

To my mind, both of these options are quite unattractive. Views like (ii) which attribute rampant error to non-philosophers are hard to take seriously. (One is uncomfortably reminded of Hume's attribution to 'the vulgar' of a theory of perception which, according to him, could be refuted merely by an easy experiment involving pressing on one's eyeball.) Option (i), meanwhile, threatens to make the notion of semantic content into a mere theoretical construct of no explanatory significance, related only in some tenuous and indirect way to the facts about what we actually use sentences to do.

Are these the only options for defenders of the Attitude Transparency Schema? It might seem that they must be. For suppose the Attitude Transparency Schema is valid. Then (4) is logically valid. So the proposition it semantically expresses must be true. Similarly the proposition (5) semantically expresses must be false. Either this false proposition is widely believed, or it is not. If it is widely believed, we have option (ii). If it is not, then ordinary people who utter (5) are using it to do something other than express belief in its semantic content. But ordinary utterances of (5) are not lies or deliberate attempts to mislead: the utterer is not endorsing or committing herself to a proposition that she does not believe. Hence these utterances must be nonliteral, as in option (i). Speaking literally involves endorsing/committing oneself to the semantic content of the sentence one utters.

There is a problem with this argument: it ignores the phenomenon of context sensitivity. When a sentence is context-sensitive, there is no such thing as *the* proposition it semantically expresses. It is semantically associated with more than one proposition, in such a way that different speakers who utter the sentence may commit themselves to different propositions without any of them speaking non-literally. My aim in this chapter is to appeal to context-sensitivity to resolve our paradox. Section 3 will briefly argue that sentences containing attitude verbs are in fact systematically context-sensitive. Section 4 will consider how we should understand notions like 'logically valid sentence' when context-sensitivity is in play, and section 5 will apply this discussion to the Attitude Transparency Schema. The final three sections will discuss questions about how the relevant form of context-sensitivity should be accounted for syntactically and semantically.

3 Attitude reports are context-sensitive

How does one ever argue for the presence of context-sensitivity? Many familiar arguments can be seen as conforming to the following argument-schema.

(P1) $\{\phi_1, \dots, \phi_n\}$ is a logically inconsistent set of declarative sentences.

(P2) It is possible for there to be utterances u_1, \dots, u_n of ϕ_1, \dots, ϕ_n , respectively, such that

- a. Each of u_1, \dots, u_n is *literal* (meant literally).
- b. Each of u_1, \dots, u_n is *sincere* (free of pretence and deception).
- c. None of the utterers of u_1, \dots, u_n has any relevant false beliefs.

(C) At least one of ϕ_1, \dots, ϕ_n is context-sensitive.

I will call arguments of this form *No-Error Arguments*.

Are No-Error Arguments valid? To defend the claim that C follows from P1 and P2, we can appeal to the following non-schematic principles:

Inconsistency: When a set of context-insensitive declarative sentences is logically inconsistent, it is not possible that all the propositions they semantically express are true.

Literality: Necessarily, when an utterance of a context-insensitive declarative sentence is literal, its utterer asserts the proposition semantically expressed by that sentence.

Sincerity: Necessarily, when one asserts a proposition by means of an utterance which is sincere, one believes that proposition (and this belief is relevant to the utterance).²

²Other ways of defending the inference from P1 and P2 to C would rely on variants of *Literality* and *Sincerity* in which 'asserts' is replaced with 'intends to communicate', or 'is committed to', or 'represents him- or herself as believing'. I will stick with 'asserts', but the work I am using it to do could be done just as well by any of these other expressions.

Arguments for context-sensitivity often bypass the need for *Sincerity*, by using a modified version of P2 which omits clause (b), and replaces clause (c) with 'None of the speakers of u_1, \dots, u_n asserts anything false' (or '... intends to communicate anything false', or 'is committed to anything false', or 'represents him- or herself as believing anything false'...). However, it seems to me that our pre-theoretic judgments about whether people have false *beliefs* are firmer, and less easy to shape to fit our favourite theories, than our beliefs about whether people are *asserting* anything false, so that the argument-scheme using P2 is likely to be more dialectically effective.

These principles look hard to deny. As regards *Inconsistency*, whatever it might mean to call a set of context-sensitive sentences ‘logically inconsistent’—we will consider this in section 4—the situation with context-insensitive sentences is relatively clear. Context-insensitive sentences can be counted “true” (in a derivative sense) when the propositions they semantically express are true; and a set of sentences all of which are true cannot be logically inconsistent. *Literality* and *Sincerity*, meanwhile, encapsulate central connections between the technical expression “semantically express” and other expressions we understand. There are possible views that either give up on one of these principles altogether, or allow it to be true only when ‘literal’ or ‘sincere’ is understood in some special technical sense, concerning which we are not entitled to rely on any pre-theoretic intuitions. The Pragmatic Millians we met in section 2 are pushed towards some such view. According to them, there is a wide class of utterances, such as ordinary utterances of (5), in which speakers utter context-insensitive sentences which semantically express propositions they do not believe, despite the fact that the utterances in question seem literal and sincere by ordinary standards. Without doing philosophy of language, no-one would think to classify these utterances with paradigms of non-literal speech like metaphor, exaggeration and sarcasm, or with paradigms of insincere speech like lies and deliberate attempts to mislead. Any view that takes this attitude to *Literality* or *Sincerity* faces the objection that it makes the notion of semantic content into an explanatorily idle theoretical construct.

(I am using ‘context-sensitive sentence’ broadly, in such a way that any declarative sentences which does not semantically express a unique proposition, *simpliciter*, counts as context-sensitive. I am not assuming any particular answer to the question how we should characterise the semantic functioning of these sentences. In particular, I am not assuming that context-sensitive sentences should be characterised using a notion of semantic value that is relativised to entities called ‘contexts’. Section 7 below will express some qualms about this dominant approach, and mention some alternatives.)

The literature contains many arguments for the context-sensitivity of attitude reports that can be reconstructed as No-Error Arguments.³ In general, premise P2 will be supported as possibility claims often are in philosophy, by telling a little story. Here, for example, is a version of a story due to Schiffer (1979). A purse-snatcher, Shorty, snatches Thelma's purse; she sees him limping away. Later, she is at a lineup in which the purse-snatcher is present, but fails to pick him out. Talking to one of his associates before the lineup, Shorty utters

(6) She knows that I limp.

He follows up this utterance by saying: 'So, I will make sure to show up early to the lineup, so she doesn't see me walking in.' Whispering to another associate during the lineup, Shorty utters

(7) She doesn't know that I limp.

He goes on: 'I made sure of that by showing up early to the lineup, so that she wouldn't be able to figure out that I am the one who stole her purse.' Each of these utterances is literal and sincere. And Shorty has no relevant false beliefs: nothing except for the desire to avoid making the story needlessly far-fetched prevents us from stipulating that he has no false beliefs at all. Since this story describes a possible state of affairs, and since {'She knows that I limp', 'She doesn't know that I limp'} is an inconsistent set of sentences, we can conclude by the argument-schema above that at least one of those two sentences is context-sensitive. And since we have independent reasons to think that the negation operation that takes (6) to (7) is not itself a source of context-sensitivity, we can conclude that *both* (6) and (7) are context-sensitive.

Hold on—it's already obvious that (6) and (7) are context-sensitive, since they both contain the first person pronoun 'I'. So the conclusion of our argument is less interesting than we wanted it to be. One way to address this issue would be to change

³Examples include Sosa 1970, Lewis 1979, Schiffer 1979, Crimmins and Perry 1989, Richard 1990, Richard 1993, Crimmins 1992, Oppy 1992, Recanati 1993, Bach 1997, Soames 2005, and Dorr 2011.

the examples so as to eliminate sources of context-sensitivity other than the one we are interested in:

(6') At noon on January 21st 1976, Thelma knows that Shorty sometimes limps during January 1976.

(7') At noon on January 21st 1976, Thelma doesn't know that Shorty sometimes limps during January 1976.

But intuitions about sentences that only a philosopher would ever utter will be a weak basis for argument; and besides, there is nothing we can do that will uncontroversially eliminate irrelevant sources of context-sensitivity. (For example, if we accept the appearance that different people can bear the same name, so that (6') and (7') can be used literally to make claims about different people named 'Thelma' and 'Shorty', we will be pushed to regard them as context-sensitive for that irrelevant reason.) Another strategy is to refine our argument-schema so that it can be used not only against theories which deny that certain sentences are context-sensitive at all, but also against theories which admit that the sentences are context-sensitive, but provide some too-short list of the relevant kinds or sources of context-sensitivity. For example, we can consider a theory according to which the context-sensitivity of (6) and (7) is exhausted by the facts that (i) the pronoun 'she' can be used to talk about different people, (ii) the pronoun 'I' can be used to talk about different people, (iii) the verb 'limp' can be used to talk about different periods of time, and (iv) the verb 'know' can be used to talk about different periods of time. According to this theory, there is no one pair of propositions such that one is committed to the truth of the first by any literal utterance of (6), and committed to the truth of the second by any literal utterance of (7). Rather, there is a set *A* of propositions associated with (6), and a set *B* of propositions associated with (7), such that any literal utterance of (6) commits the speaker to the truth of some member of *A*, and any literal utterance of (7) commits the speaker to the truth of some member of *B*. Each of *A* and *B* contains propositions that differ only in the four ways corresponding to the four posited "sources" of context-sensitivity. And whenever a member of *A* is

consistent with a member of B , these propositions differ in at least one of the four ways: they are about different people or different periods of time. We can argue against this theory using the same general kind of argument that we gave against the theory that posits no context-sensitivity at all. To do so, we must simply ensure that the possible state of affairs that we exhibit at step P2 involves literal utterances u_1 and u_2 of (6) and (7), in which the propositions asserted do not differ in any of the four specified ways: they are about the same people and the same periods of time.

If my main purpose were to argue that attitude reports are context-sensitive in the relevant way, I would need to cover a wider array of examples, and I would need to say more to defend the premise that the speakers in my stories really are speaking literally and sincerely (*pace* the pragmatic Millians), and really do lack relevant false beliefs (*pace* the error-theoretic Millians). Also, I would need to respond to a range of strategies for arguing *against* the presence of context-sensitivity that have recently enjoyed considerable influence, for example those given in Cappelen and Lepore 2005. But I will not do any of these things here. Rather, I am going from now on to take this context-sensitivity as established, so as to investigate how it bears on the status of the Attitude Transparency Schema.

4 Uniformity

When a sentence is context-sensitive, its role in our linguistic practice cannot usefully be encapsulated by associating it with a unique proposition as “semantic value”. The semantically interesting relation between context-sensitive sentences and propositions is a one-many relation. There are many different propositions which count as “admissible interpretations” of a context-sensitive declarative sentence; those who utter such a sentence literally always assert at least one of its admissible interpretations, but not always the same one.

We can sometimes make a useful distinction within the class of admissible interpretations of a given context-sensitive sentence, between those which are *uniform* interpretations of the sentence, and those which are not. Consider for example the

sentence

(8) Every superhero is strong, but he is not strong.

(8) is context-sensitive: there is no single proposition asserted by all those who utter it literally. For example, one could use (8) to assert, concerning Batman, something roughly to the following effect: whereas every superhero is stronger than typical human beings, he is not stronger than typical human beings. This is an example of a *uniform utterance* of (8). A hallmark of such utterances is that they invite the conclusion that the person in question is not a superhero. There are other uses of (8) which do not invite such a conclusion. For example, one could use (8) to assert, concerning Batman, that while every superhero is stronger than typical human beings, he is not stronger than typical *superheroes*. This is an example of a non-uniform interpretation of (8). Non-uniform interpretations of (8) are easier to access when the fact that the person being discussed is a superhero is uncontroversial common knowledge; it also helps if the speaker emphasises the second occurrence of 'strong'.

This contrast is one we can readily understand even before we have introduced any technical machinery for characterising what the distinction between uniform and non-uniform interpretations of a sentence consists in. In the remainder of this section I want to make some fairly simple obvious remarks about the phenomenon, staying at an informal level.

(i) *Non-uniform interpretation does not require repeated words.*

There is more to uniform interpretation than interpreting occurrences of the same word in the same way. A useful distinction between uniform and non-uniform interpretations can be drawn even for sentences with no repeated words, like (9):

(9) Every superhero is strong, but he is weak.

It seems natural to say that the uniform interpretations of (9) are like those of (8) in entailing that the person in question is not a superhero. This is not entailed by non-uniform interpretations of (9): for example, one could use (9) to assert, concerning

Batman, that whereas every superhero is stronger than typical human beings, he is weaker than typical superheroes.

(ii) *Non-uniform interpretations are sometimes hard to access, sometimes easy.*

In examples like (8) and (9), speakers have to work pretty hard, e.g. by using gestures or variations in tone of voice, in order to get across a proposition that is not a uniform interpretation of the sentence we are uttering. In such cases, hearers would typically reach the intended, non-uniform intention only after first considering what would have to be the case for a uniform interpretation of the sentence to be true, and concluding that none of these interpretations is likely to be intended. However, there are other cases where we can expect non-uniform interpretations to come immediately to our hearers' attention. Indeed, for some sentences, *uniform* interpretations are hard to access even when they would be perfectly sensible things to assert. For example, we naturally expect multiple occurrences of a demonstrative in a sentence to be used to refer to different objects:

(10) Everyone who works for that guy has to make coffee for that guy.

Something similar is true for proper names: out of the blue, (11) seems like a way of talking about two people called 'John':

(11) Everyone who works for John has to make coffee for John .

By analogy with (8), I want to say that the uniform interpretations of (10) and (11) are propositions which entail that there is someone such that everyone who works for him has to make coffee for him. But if you wanted to communicate one of those propositions, (10) and (11) would not be good choices: you will do better to use a pronoun instead of repeating the demonstrative or proper name:

(12) Everyone who works for that guy [John] has to make coffee for him.

Names and demonstratives are not the only cases in which hearers can easily access non-uniform readings. Here is an example (from Soames 1986) of non-uniformity in

the interpretation of quantifier-domains:⁴

(13) Everyone is asleep and is being monitored by a research assistant.

The natural interpretations of (13) are non-uniform. Uniform interpretations, which require the research assistants to be doing their monitoring while asleep, are not especially salient. In the case of (13), of course, the far-fetchedness of the uniform interpretations gives an extra boost to the non-uniform ones; but even if we eliminate this by replacing 'asleep' with 'awake', it is still natural to assume that the domain relevant for 'a research assistant' is distinct from the one relevant for 'everyone'.⁵

(iii) *Semantically equivalent sentences can differ in the ease with which they suggest non-uniform interpretation.*

I take it that there is some good sense in which (14a) and (14b) are semantically equivalent:

- (14) a. Michael is stronger than Jennifer.
b. Jennifer is weaker than Michael.

In whatever sense (14a) and (14b) are equivalent, so are (15a) and (15b):

- (15) a. Michael is stronger than Mark, and Mark is stronger than Jennifer,
but it's not true that Michael is stronger than Jennifer.

⁴I should mention that Soames himself denied that quantifier-domain restriction involves semantic context-sensitivity. On Soames's favoured view, (13) context-insensitively expresses a proposition that is false unless everyone in the universe is asleep, but typical utterers of (13) do not assert this proposition. (Given *Literality*, this entails that a typical utterance of (13) would not be literal.) Arguing in favour of a contextualist treatment of domain restriction is beyond the scope of the present chapter. Some prominent contextualists about domain restriction have given examples similar to (13): e.g. 'Everyone waved to everyone' (Stanley and Szabó 2000).

⁵We can find something similar with possessives. Suppose that Alfred is watching horse racing on television with his friend Bert, who owns a slow old horse. Alfred has taken a fancy to a certain horse which he thinks likely to win the race. In this setting, it is very easy for Alfred to say 'My horse is nothing like your horse', as a way to convey the proposition that the horse *he has taken a fancy to* is nothing like the horse *Bert owns*. Here there is little pressure to find a single relation which could provide the interpretation for both of the possessives.

- b. Michael is stronger than Mark, and Mark is stronger than Jennifer,
but it's not true that Jennifer is weaker than Michael.

While (15a) and (15b) both sound pretty bad, it is somewhat easier to imagine a conversational setting in which (15b) might be uttered by a reasonable and co-operative speaker. I take it that this to be explained by saying that non-uniform interpretations of (15b) are somewhat easier to access than the corresponding non-uniform interpretations of (15a). The choice to use a different word rather than repeating 'stronger' three times helps, a little bit, to cue the reader in to the kind of non-uniform interpretation that is called for. If two different potential interpretations of 'stronger'/'weaker' are salient, hearers of (15b) will at least be able to guess that they are supposed to use one for both occurrences of 'stronger' and the other for the one occurrence of 'weaker', whereas hearers of (15a) will still be at a loss as regards what is being claimed to be the case, unless the speaker provides further cues. It would be a bad idea if someone took this difference between (15a) and (15b) to motivate rejecting the claimed semantic equivalence of (14a) and (14b). This sort of contrast is better explained by a pragmatic theory about the ways in which choices among semantically equivalent wordings can make a difference to the accessibility of different interpretations.

Choices in wording can affect the ease with which we access non-uniform interpretations even when the words between which we are choosing are not themselves relevantly context-sensitive:

- (16) a. Batman is strong, but Batman is not strong.
b. Bruce Wayne is strong, but Batman is not strong.

Given that Batman is Bruce Wayne, neither (16a) nor (16b) is true on any uniform interpretations. In both cases, true non-uniform interpretations are possible. But they are easier to access for (16b): the associations of the two names naturally suggest an interpretation where the reference classes relevant to the two occurrences of 'strong' are, respectively, the class of human beings and the class of superheroes. It would be a very bad idea to take this contrast between (16a) and (16b) to motivate some exotic

semantic account of 'Bruce Wayne' and 'Batman'.

(iv) *It is not always obvious which words are being interpreted differently when one speaks non-uniformly.*

Often, when a sentence is shown to be context-sensitive by means of a No-Error Argument, it will remain an open question which words in the sentence are responsible for the relevant dimension of context-sensitivity. Consider (17), for example:

(17) There are exactly seven rooms in Frank's house.

We can argue for the context-sensitivity of (17) using a run-of-the mill No-Error Argument, in which the relevant set of sentences is {(17), (18)}:

(18) There are exactly six rooms in Frank's house.

The conclusion of this argument is that at least one of these two sentences is context-sensitive; we can then appeal to considerations of symmetry to argue that if either is context-sensitive, both are. To support for the crucial premise P2 of the No-Error Argument, we can describe a particular possible case, such as the following. Frank owns a house with three bedrooms, a bathroom, a kitchen, a living room, and an attached garage, built later than the main house, that he uses to store his books. Jane is asked how many rooms there are in Frank's house; she counts them off on her fingers, including the garage, and without further reflection utters (17). Meanwhile, Mary, who also knows Frank's house similarly well, is having a similar conversation but instead utters (18), after having decided that it would be misleading to include the garage. Both speak sincerely and literally; neither has any relevant false beliefs.

This context-sensitivity in (17) could be explained in different ways. Perhaps it is due to the word 'room', which can express several different properties, some of which are instantiated by Frank's garage and others of which are not. Perhaps it is due to the word 'in', which can express several different relations, some of which Frank's house bears to its garage and some of which it does not. Perhaps it is due to the word 'house', which can express several different properties, some of which are had by a

relatively large object that includes the garage, and others of which are only had by a smaller object that does not include the garage (see Lewis 1993). Or perhaps several of these words are relevantly context-sensitive. Deciding between these options requires systematic philosophical investigation; our competence as speakers does not equip us to know which option is correct in the absence of such investigation.

When the source of a sentence's context-sensitivity is unclear, there can be cases where we can tell that a certain utterance is non-uniform without being able to pinpoint any particular pair of expressions as the ones that are being treated non-uniformly. Consider for example

(19) There are exactly seven rooms in Frank's house, and there are exactly six rooms in Jim's house.

It is easy to think of circumstances in which it would be plausible that (19) is false on all uniform interpretations, even though each of its conjuncts admits true interpretations. For example, Jim's garage might be somewhat less flimsy than Frank's, connected to the main house by a somewhat tighter seam than Frank's, etc. In such circumstances, anyone who uttered (19) literally would either be asserting something false or speaking non-uniformly. One can, with difficulty, imagine a conversational setting in which such an utterance would be reasonable for a sincere speaker without false beliefs. (It becomes much easier to do this if one imagines a longer discourse, with many other intervening claims which help to make the repetition less salient.) But until we do the philosophical work required to decide which words to blame for the relevant context-sensitivity, we will not know which words are being treated non-uniformly in these utterances of (19). We do not need to do this work to be justified in classifying these utterances as non-uniform, or in believing that the given conditions suffice for (19) to be false on all uniform interpretations.

(v) *Multi-sentence discourses can also be interpreted uniformly or non-uniformly.*

Something will have gone very wrong if our semantic theory forces us to tell a different kind of story about (8), on the one hand, and the two-sentence discourse (8'), on the

other:

(8) Every superhero is strong, but he is not strong.

(8') Every superhero is strong. He is not strong.

We thus need a way of talking about uniformity of interpretation across as well as within sentences. One option is to speak of arbitrary *sequences* of sentences as having “admissible interpretations”, which are sequences of propositions; every sequence of sentences has some uniform interpretations, but often, not all the admissible interpretations are uniform. For $\langle p_1, \dots, p_n \rangle$ to be a uniform interpretation of $\langle \phi_1, \dots, \phi_n \rangle$, it is necessary but not sufficient that p_1 be a uniform interpretation of ϕ_1 , ... and p_n a uniform interpretation of ϕ_n . Sometimes, we literally utter a sequence of sentences without asserting all the members of any uniform interpretation of that sequence; in these cases we speak non-uniformly.

Provided we can make sense of the notion of a uniform interpretation of a sequence of sentences, we should also be able to make sense of a contrast between uniform and non-uniform interpretations for *sets* of sentences, even infinite sets of sentences. One option is to define uniformity for such sets in terms of uniformity for sequences: a function f from sentences to propositions is a uniform interpretation iff whenever $\langle \phi_1, \dots, \phi_n \rangle$ is a sequence of members of the domain of f , $\langle f(\phi_1), \dots, f(\phi_n) \rangle$ is a uniform interpretation of $\langle \phi_1, \dots, \phi_n \rangle$. In this way we can make sense of the notion of a uniform interpretation of the entire language. However we approach this notion, it seems obvious that every uniform interpretation of a subset of the sentences of a language can be extended to a uniform interpretation of the entire language.

(vi) *Non-uniformity is not a kind of non-literality.*

Cases where we utter a sentence without committing ourselves to any of its uniform interpretations do not seem interestingly similar to paradigmatic cases of non-literal discourse, like metaphor, sarcasm and exaggeration: *prima facie*, it does not look promising to group them together as examples of non-literal speech. If our only examples of non-uniform discourse involved sentences like (8), where there is a strong default

presumption in favour of uniform interpretation, such an assimilation might be tempting. In these cases, the process whereby hearers reconstruct the intended non-uniform interpretation would often involve first considering what would be required for the sentence to be true on a uniform interpretation, and then using Gricean reasoning to rule out the possibility that such an interpretation is intended by the speaker. This is interestingly similar to the process whereby, in stereotypical cases of metaphor, sarcasm and exaggeration, hearers end up at an intended non-literal interpretation by first figuring what would be required for the sentence to be true on a literal interpretation, and then using Gricean reasoning to establish that such interpretations are not intended. However, once we broaden our focus to include the many cases discussed under (ii) above, where non-uniform interpretation comes naturally and without any sense of a presumption that needs to be overcome, the analogy looks much weaker. And given that we want to be able to speak of non-uniformity across sentences as well as within sentences, it seems hopeless to think of the phenomenon as a kind of non-literality: as the focus of a conversation shifts, we will naturally and unselfconsciously shift our use of context-sensitive vocabulary in a way that has nothing interesting in common with metaphor, sarcasm or exaggeration.

(vii) *Validity and inconsistency for sentences has to do with uniformity*

If we can get a grip on the concept of uniform interpretation, it is natural to put it to work in making sense of notions like validity and inconsistency, understood as features of sentences as opposed to propositions. Suppose that we already understand what it means for a proposition to be “valid” or “inconsistent”—e.g. we might characterise a valid proposition as one that is necessarily true, and an inconsistent one as one that is necessarily false, in some appropriate sense of ‘necessarily’.⁶ Then we can say that a *sentence* is “valid” (“inconsistent”) iff all its uniform interpretations are valid (inconsistent) propositions.

This is not the only possible way to extend the notions of validity and inconsistency

⁶It is not uncontroversial that validity should be connected in this way to truth. For an argument that it should not, see Field 2008, Sect. 9.2.

from propositions to sentences. We could also define a valid [inconsistent] sentence as one *all* of whose admissible interpretations, whether uniform or non-uniform, are valid [inconsistent]. The problem with this is that it makes the notions of validity and inconsistency of sentences uselessly strong, when context-sensitive sentences are in question. Even sentences of the form ‘ P and it is not the case that P ’ will not, in general, be inconsistent on such a definition, since such sentences will normally have consistent non-uniform interpretations. If validity and inconsistency at the level of sentences are supposed to be notions we can understand and have justified beliefs about prior to doing much theorising, we had better not understand them in this way.⁷

If, as suggested above, we can make the concept of uniform interpretation for sequences and sets of sentences as well as single sentences, the suggested definitions of validity and inconsistency for single sentences can be extended in a natural way to definitions of inconsistency as a property of sets of sentences, and logical consequence as a relation between sets of sentences and sentences (assuming that we already understand these notions as applied to sets of propositions). A set of sentence is “inconsistent” iff its image under every uniform interpretation is an inconsistent set of propositions. And when Γ is a set of sentences and ϕ is a sentence, ϕ is a consequence of Γ iff for every uniform interpretation f whose domain includes $\Gamma \cup \{\phi\}$, $f(\phi)$ is a consequence of $f(\Gamma)$.⁸

⁷But see section 6 below for a way of individuating sentences that restores the usefulness of this account of validity and inconsistency.

⁸Arguably, this way of defining logical vocabulary at the sentential level will not work if we are trying to understand this vocabulary in a “narrow” way on which, for example, ‘There are groundhogs’ is not logically equivalent to ‘There are woodchucks’. For any two context-insensitive sensitive expressions which semantically express the very same proposition will be counted as logically equivalent by the above definition; and it is arguable that ‘There are groundhogs’ and ‘There are woodchucks’ are context-insensitive and semantically express the same proposition. If this is right, we need some other idea in order to recover the “narrow” logical notions. The most obvious thought is the Tarskian one, according to which narrow logical consequences requires truth-preservation across a wide family of interpretations, not all of which are admissible given the actual semantics of the relevant vocabulary. (Or we could define it as *validity* on each interpretation in this wide family, using some antecedently understood notion of validity at the level of propositions.) The Tarskian approach looks especially natural in the present setting, where *any* notion of validity and inconsistency at the sentential

5 The paradox resolved

These reflections suggest a flat-footed response to section 2's "paradox of transparency". Recall that the paradox turned on the fact that there is a valid argument from some instances of the Attitude Transparency Schema to (4):

- (4) If Superman is Clark Kent, then for all y , y believes that Superman flies iff y believes that Clark Kent flies.

If the instances of that schema are valid, then (4) itself must be valid. And this looks objectionable, since (4) seems false.

What I want to suggest is that (4) *is* valid, so that all of its uniform interpretations are true. However, (4) also admits non-uniform interpretations, many of which are false. Moreover, some false non-uniform interpretations of (4) are more salient, easier to access, than any of its uniform interpretations. The predicates 'believes that Superman flies' and 'believes that Clark Kent flies' have exactly the same range of admissible interpretations. But the use of the name 'Clark Kent' tends to favour some of these interpretations—properties which are instantiated only by those who have penetrated Superman's secret identity—whereas the use of the name 'Superman' tends to favour other interpretations—properties instantiated by ordinary folk who see Superman flying around in his Superman suit.

For the same reason, (5) is acceptable:

- (5) Even though Superman is Clark Kent, lots of people believe that Superman flies without believing that Clark Kent flies.

Being logically inconsistent, (5) cannot be true on any uniform interpretation. But this is beside the point, since utterers of (5) would be very unlikely to intend a uniform interpretation, and such interpretations are unlikely to occur to anyone hearing (5). The most easily accessible interpretations of (5) are non-uniform, and true.

On this account, the sentences in (20) are all semantically on a par:

level involves quantification over interpretations.

- (20) a. Lois believes that Superman flies but does not believe that Clark Kent flies.
- b. Lois believes Superman to fly but does not believe Clark Kent to fly.
- c. Superman is believed by Lois to fly, but Clark Kent is not believed by Lois to fly.
- d. Superman is such that Lois believes he flies, but Clark Kent is not such that Lois believes he flies.

All these sentences admit true non-uniform interpretations, while none admits true uniform interpretations (given that ‘Superman is Clark Kent’ is true). Whatever differences there may be in our reactions to the sentences are best explained as arising from differences in the relative salience and naturalness of the uniform interpretations. Because the names in (20c) and (20d) occur in subject position, the fact that these sentences are of the right form to provide the premise for a valid argument by Leibniz’s Law, whose conclusion is ‘Superman is not Clark Kent’ is more vivid and eye-catching than it is in the case of (20b) or (20a). Because of this, the danger that addressees will favour a uniform interpretation of (20c) or (20d) is greater, making speakers less willing to assert these sentences. It would be a mistake, I think, to take these differences of degree as a basis for a semantic account that treats some of the sentences in (20) very differently from the others.

The idea of non-uniformity in the interpretation of *multiple* context-sensitive sentences is one of the most familiar and popular tools in the philosopher’s paradox-solving toolbox. A number of philosophers have had the idea of using this tool to resolve certain puzzles connected with substitutivity.⁹ But we have been less comfortable appealing to non-uniform interpretation for *single* sentences. The usual view seems

⁹One example is Lewis’s (1986, sect. 4.5) invocation of context-sensitivity in ‘counterpart’ to rebut modal Leibniz’s Law arguments for the distinctness of people and their bodies, bronze statues and their constituent lumps of bronze, etc. As regards attitude reports, Richard (1990, pp. 190–96) endorses a “context-switching” strategy for resolving the famous phone booth puzzle introduced in Richard 1983.

to be that non-uniformity within a sentence is a rather unusual phenomenon, with a distinctive phenomenology—that of not uttering a sentence ‘in the same breath’—which is obviously absent in typical utterances of (5). This view will seem natural if our paradigms of non-uniformity are utterances of sentences like (16a):

(16a) Batman is strong, but Batman is not strong.

A better paradigm, I think, would have been

(13) Everyone is asleep and being monitored by a research assistant

which shows that the fact that the natural interpretations of a speech are non-uniform ones can easily escape our notice.

My strategy for vindicating the validity of the Attitude Transparency Schema might seem like a cheat. I can imagine someone feeling that if non-uniformity is as widespread as I claim, and logically inconsistent sentences can have true non-uniform interpretations, logical statuses like inconsistency and validity become much less interesting than they are normally supposed to be. Perhaps this is right. But I can’t see any other meanings for ‘valid’ and ‘inconsistent’ that would render them any more interesting: I can think of no theoretically interesting status that would distinguish paradigmatic logical principles like the non-contradiction schema from the Attitude Transparency Schema. If recognising the prevalence of non-uniformity requires you to adjust your sense of the interest and dialectical importance of notions like validity and inconsistency, you had better go ahead and make that adjustment.

6 Fine-grained individuation of sentences

So far, I have been deliberately naïve in my use of ‘sentence’. I have assumed that I can uniquely pick out a particular sentence (-type) for future reference by writing it down with a number to its left. This ignores the possibility of *homographs*: distinct sentences which can be written down using very same string of symbols (letters, spaces, punctuation marks). But homographs are clearly possible, since sentences with different *sounds* can be represented by the same string of symbols:

(21) She took a bow.

More importantly, philosophical orthodoxy also recognises cases of simultaneous homography and homophony, in cases like these:

(22) The banks in this town are made of stone.

(23) Visiting relatives can be annoying.

(24) Some girl loves every boy.

There is something attractive about the suggestion that some or all of the cases which I described as involving a single sentence admitting both uniform and non-uniform interpretations really involve a multiplicity of homonymous (homophonic and homographic) sentences. For example, I treated

(10) Everyone who works for that guy has to make coffee for that guy

as a sentence that admits uniform interpretations (which entail that there is someone x such that everyone who works for x has to make coffee for x) as well as non-uniform interpretations (which do not). One could say, instead, that there are at least two different sentences associated with the string of symbols inscribed above after '(10)'. There is at least one co-ordinated sentence, whose admissible interpretations are the propositions I called 'uniform interpretations of (10)' (i.e. propositions which entail that there is someone x such that everyone who works for x has to make coffee for x). And there is also at least one unco-ordinated sentence, whose admissible interpretations do not all entail this.¹⁰ On this approach, the syntactic contrast between

¹⁰On one version of the view, the *only* admissible interpretations of the unco-ordinated sentence, or sentences, are what we were previously calling 'non-uniform interpretations of (10)': thus no proposition is both an admissible interpretation of a co-ordinated sentence and an admissible interpretation of an unco-ordinated homonymous sentence. On another version of the view, all the propositions we were previously calling 'admissible interpretations of (10)' are admissible interpretations of the unco-ordinated sentence(s); the set of admissible interpretations of the co-ordinated sentence(s) thus a proper subset of the set of admissible interpretations of the unco-ordinated sentence. The first version of the view would look extremely problematic if we identified propositions with sets of worlds, or with Russellian structures built

co-ordinated and unco-ordinated sentences takes over the theoretical work previously done by a semantic contrast between uniform and non-uniform interpretations of a single sentence.

The standard way of working out this idea is to embrace what I will call an “extreme homonymy theory”. According to an extreme homonymy theory of (10), this string of symbols can be used to write down *infinitely many* distinct sentences, in the sense of ‘sentence’ that matters to semantics.¹¹ A notation in which strings of symbols corresponded one-to-one with sentences would need to replace expressions like ‘that guy’ with infinite collections of different symbols. This is standardly achieved by attaching arbitrary numerical subscripts to such expressions. In such a notation, we could pick out particular members of the relevant infinite class of sentences using strings of symbols like the following:

- (25) a. Everyone who works for (that guy)₁₇ has to make coffee for (that
guy)₁₀₄₂
- b. Everyone who works for (that guy)₃ has to make coffee for (that
guy)₃

The apparatus of indices is normally used only for to pronouns and pronoun-like elements, and expressions capable of binding them. One could, however, extend the multiplication of indices to context-sensitive expressions of other sorts.¹² This introduces a few complications. First, we might want to use different families of indices for different families of expressions, so as to make it impossible in some cases for two expressions to share an index. For example, we want to allow an occurrence of

out of objects and properties: think of an utterance of the unco-ordinated version of ‘this ship is the same length as that ship’, where the speaker is in fact unknowingly demonstrating the same ship twice over (cf. Fiengo and May 1994, ch. 1). However, in a finer-grained theory of propositions—e.g. a “Fregean” theory, such as the one I will consider in section 8, or the “relationalist” theory of Fine 2007—the prospects for such a view might be better.

¹¹‘Sentences’ in this sense might be identified with the LFs of syntactic theory, or with sequences of syntactic representations one of whose elements is an LF.

¹²If there are any. Stanley (2007) argues that the only syntactically simple context-sensitive elements in natural languages are pronouns and unpronounced pronoun-like elements.

'strong' to share an index with an occurrence of 'weak', but there is no clear motivation for allowing an occurrence of 'strong' to share an index with an occurrence of 'that guy'. Second, we might need to allow multiple indices on the same word, when we can discern different 'dimensions' of context-sensitivity. For an example of this, consider

(26) Michael is strong, and everyone stronger than Michael is tall.

Speaking uniformly using (26) requires both a certain co-ordination between 'strong' and 'tall' (one could not use 'strong' to mean *strong for a wrestler* while using 'tall' to mean *tall for an adult human being*) and a certain co-ordination between 'strong' and 'stronger' (one could not use 'strong' to mean *strong-willed* while using *stronger* to mean *better at weightlifting*). (26) thus presents at least two different possibilities for non-uniformity. Using indices, we would need to associate at least two indices with 'strong', one of which could also occur on 'stronger', and another of which could also occur on 'tall'.¹³

The extreme homonymy theory has several advantages. First, it allows us to keep the notion of a uniform interpretation out of our semantic ideology. Second, it allows for notions of validity and inconsistency for sentences that are more firmly tied to usage: we no longer have to allow that someone could utter a logically inconsistent sentence while speaking literally without asserting anything false. Third, it makes life safe for a straightforward form of semantics in which each of the admissible interpretations of a sentence is derived from some assignment of semantic values to the expressions of the language which obeys simple compositional rules. However, the approach also has serious drawbacks. First, it just seems absurd to think that natural languages contain these infinite families of homonymous sentences—this seems to gloss over an important contrast between 'bank'/'bank' and run-of-the-mill context-sensitivity. Second, it is unclear how our pre-theoretic opinions about (linguistic) validity and

¹³Alternatively, we might want to posit several silent constituents associated with the overt word 'strong', to be the bearers of the different indices. But on this approach, there is a danger that we will find ourselves positing such a large quantity of silent material that questions will arise over whether we are really talking about the same thing that *syntacticians* are talking about when they posit "empty categories".

inconsistency are supposed to be brought to bear in investigating the logical properties of sentences as conceived of by the extreme homonymy theorist, given that the pre-theoretic judgments do not seem to be directly about such sentences. Third, if we adopt an extreme homonymy theory it will be hard to see why we should think that there is such a thing as context-sensitivity in natural language at all. For how are we to decide when, if ever, speakers on different occasions uttering the same member of an infinite class of homonyms? Since we are committed to saying that distinct expressions are being used whenever we encounter differences in interpretation within a single sentence, the obvious way to answer this question is to extend this policy more broadly, so that any two utterances whose intended interpretations are different are held to be utterances of distinct expressions. This would result in a picture where we can, after all, assign each declarative sentence to a single proposition which it semantically expresses, and which is asserted by all those who utter it literally: the multiplicity has been shifted from the relation between sentences and propositions to the relation between sound- and shape-types and sentences. Running together the phenomena of syntactic disambiguation and semantic interpretation in this way seems bizarre. But it is hard to think of any *other* principled answer to the question under what circumstances different speakers count as using the same member of an infinite class of homonyms. The only alternative I can think of is to allow such questions about identity of expressions across conversations to be massively vague. But note that on such a view it will still never *definitely* be the case that two different speakers are using the same expression while intending different interpretations. It is hard to believe that a perspicuous account of the phenomena could require the introduction of a technical use of 'uttering the same sentence' which involves so much *more* vagueness than the ordinary one.

A less extreme, and to my mind more attractive, option is to adopt I will call a *moderate* homonymy theory. A moderate homonymy theory would posit just *two* homonymous sentences associated with (10): one co-ordinated sentence and one unco-ordinated one. The admissible interpretations of the co-ordinated sentence are the propositions

we were previously calling ‘uniform interpretations of (10)’; the admissible interpretations of the unco-ordinated sentence are, or include, what we were previously calling ‘non-uniform interpretations of (10)’.

For a moderate homonymy theorist, an ideally perspicuous notation which established a one-to-one mapping between sentences and patterns of symbols might look something like the ‘telegraphic’ notation for first-order logic mentioned by Quine (1940), in which syntactic co-ordination is indicated by drawing lines connecting the occurrences which are required to be interpreted uniformly. In this system, the string of symbols (10) would unambiguously inscribe the unco-ordinated sentence, while the corresponding co-ordinated sentence would look something like this:

(10') $\overbrace{\text{Everyone who works for that guy has to make coffee for that guy}}$

Alternatively, we could use numerical indices to indicate the co-ordination relations within a sentence, while treating permutations of the numerical values of the indices as making no more difference to the identity of the sentence represented than the choice of a colour or typeface. (It is arguable that this account of the representational function of numerical indices better fits their use in linguistic theory than an extreme homonymy theory.)

Like the extreme homonymy theory, the moderate homonymy theory can be extended beyond the case of pronouns, so that any sound-pattern for which the contrast between uniform and non-uniform interpretation arises is resolved into at least two distinct expressions.¹⁴ Assuming we want a parallel treatment of the phenomena of uniformity and non-uniformity whether we are dealing with single sentences or larger discourses, the moderate homonymy theory for sentences will need to be extended to discourses. This will have the somewhat surprising result that discourses are more finely individuated than sequences of sentences: given a sequence of particular sentences, there are often several different discourses in which those sentences occur in

¹⁴To deal with the multiple sentences associated with (26) above, the telegraphic notation would need to be extended so as to allow different sorts of telegraph lines, and multiple telegraph lines coming out of the same expression.

that order, which differ as regards the co-ordination relations they impose between expressions in different sentences.

The moderate homonymy theory avoids all of the worries I raised about the extreme homonymy theory, while preserving two of its central advantages (no need for ‘uniform interpretation’ as a primitive in semantics; no need to speak of inconsistent sentences that can be used literally without committing the speaker to any false proposition). It does, however, lack one putative advantage of the extreme homonymy theory, in that it is not compatible with certain interpretations of the slogan that semantics is compositional (‘the meaning of a sentence is determined by the meanings of its constituent expressions and the way they are put together’). The problem is that the usual ways of making this slogan precise (see Szabó 2008) are tailored to fit the assumption that the syntactic structure of an expression—the way it is put together—can be fully captured by a tree diagram which shows how it is formed out of other expressions, by the application of “modes of syntactic composition”. It is just not clear how to describe the structures employed by moderate homonymy theories in the language of “modes of composition”.¹⁵ My own sense is that the most interesting interpretations of the compositionality slogan should be compatible with moderate homonymy theories and their non-standard account of ‘how expressions are put together’. But a proper exploration of this issue is beyond the scope of the present chapter.

¹⁵For a moderate homonymy theorist, there are two ways of combining the sentences ‘Obama is tall’ and ‘Romney is tall’ with the word ‘and’ to yield a bigger sentence that can be written ‘Obama is tall and Romney is tall’. So if we want to talk about “modes of syntactic composition”, we need to say that there at least two such modes in play in this case. Suppose that we say there are *exactly* two such modes. Then we will face awkward questions. Consider the mode of composition that yields the co-ordinated version of ‘Obama is tall and Romney is tall’. Can this also be applied to the unco-ordinated sentence ‘Gingrich is tall and Romney is tall’, and the sentence ‘Obama is tall’? If it can, presumably the resulting sentence can be written as ‘Gingrich is tall and Romney is tall and Obama is tall’; but it which of the occurrences of ‘tall’ in this sentence should we take to be co-ordinated? And it is not obvious how we could do any better by saying that there are more than two modes of semantic combination, some of which yield the very same sentence when applied to ‘Obama is tall’ and ‘Romney is tall’.

If I were to adopt an extreme or moderate homonymy theory, I would need to modify the official formulation of section 5's solution to the transparency paradox. For example, instead of saying that (5) is logically inconsistent, I would say merely that its *co-ordinated* disambiguations are logically inconsistent.

- (5) Even though Superman is Clark Kent, lots of people believe that Superman flies without believing that Clark Kent flies.

The unco-ordinated disambiguations, which are the ones which would typically be uttered, are logically consistent. Similarly, instead of saying that (1) is logically valid, I would say that its co-ordinated disambiguations are logically valid:

- (1) For all x and y , if $x = \text{Superman}$, then y believes that x flies iff y believes that Superman flies.

(1) also has unco-ordinated disambiguations which are not logically valid, although out of the blue we are not likely to take an utterance of (1) to be an utterance of one of these unco-ordinated sentences. I would be happy with these revised formulations. What I really care about is on all these views, the Attitude Transparency Schema still turns out to have the same positive status as (e.g.) the Non-Contradiction Schema. In an extreme homonymy theory, the point is straightforward: a genuine instance of a schema containing multiple occurrences of the same schematic letter must of course substitute the very same expression (represented by the same choice of indices) for each occurrence of that letter; for this reason the unco-ordinated disambiguations of (1) are not instances of the Attitude Transparency Schema. In a moderate homonymy theory, we need to go beyond the standard logic-textbook notion of 'instance of a schema', which makes no provision for the syntactic contrast between co-ordinated and unco-ordinated sentences. If we build it into our conception of schemas that a genuine instance of a schema must not only substitute the same expression for each occurrence of a schematic letter, but must also involve co-ordination relations connecting each occurrence of a context-sensitive expression in one substituent with its counterpart in every other substituent, then we can again say that all the genuine instances of the

Attitude Transparency Schema are valid. This seems like the most sensible way to talk about schemas in the setting of a moderate homonymy theory, since if we do not require such co-ordination relations, even the Non-Contradiction Schema will have instances which are not valid. But the central point is that our investigation of the individuation of sentences has not suggested any interpretation of ‘valid schema’ on which the Non-Contradiction Schema counts as valid and the Attitude Transparency Schema does not.

7 Kaplanian semantics for context-sensitive sentences

When a sentence ϕ is context-sensitive, there is no way to interpret ‘semantically expresses’ as standing for a binary relation R such that (a) ϕ bears R to exactly one proposition, and (b) inquiry into the facts which sentences bear R to which propositions serves the purposes of semantic theorising. What, then, should we do if we want a theory which does serve these purposes?

Among those who are sympathetic with the basic idea that semantics should characterise sentences in terms of their relations to propositions, the dominant approach is the one developed by Kaplan (1989). This approach involves giving a central theoretical role to a class of entities which I will call ‘Contexts’. (I use the capital letter to flag that the word is being used in a special technical sense, since much mischief has been done by slipping back and forth between the technical use of ‘context’ and various more informal uses.) Contexts feature as relata of two central theoretical relations: *sentence ϕ semantically expresses proposition p relative to Context c* , and *utterance u is in Context c* . These relations are subject the following basic principles:

- (K1) Each declarative sentence of a language semantically expresses at most one proposition relative to each Context.
- (K2) Each utterance of a sentence is in exactly one Context.

Because of (K1), each sentence ϕ is associated with a (perhaps partial) function, its *character*, which maps each Context c to the proposition semantically expressed by

ϕ relative to c . Because of (K2), the three-place relation of semantic expression for declarative sentences induces a two-place relation of semantic expression for utterances of declarative sentences: u “semantically expresses” p iff there is some declarative sentence ϕ such that u is an utterance of ϕ , and ϕ semantically expresses p relative to the Context u is in. Given that each utterance is an utterance of at most one sentence, it follows from (K1) and (K2) that each utterance semantically expresses at most one proposition.

This notion of semantic expression for utterances provides one place where the theoretical semantic vocabulary gets linked up with vocabulary we already understand. The basic idea of the connection is captured by (K3):

(K3) Necessarily, when an utterance of a declarative sentence is literal, the utterer asserts the proposition the utterance semantically expresses.¹⁶

Note that for the special case of a context-insensitive sentence, which semantically expresses the same proposition relative to every Context, (K3) reduces to principle *Literality* from section 3.

Another important connection with already-understood vocabulary involves logical notions like validity and inconsistency. It is worth distinguishing two rather different ideas about how this might work. On one approach, we begin with some notion of consequence for propositions, perhaps cashed out in terms of necessary truth preservation. Then we can characterise logical consequence for sentences as follows:

(K4a) Sentence ϕ is a logical consequence of set of sentences Γ iff for each Context c relative to which ϕ and the members of Γ express propositions, the proposition expressed by ϕ relative to c is a consequence of the set of propositions expressed by members of Γ relative to c .¹⁷

¹⁶One might prefer to replace ‘asserts’ here with something like ‘intends to communicate’ or ‘is committed to’: see note 2.

¹⁷We could, additionally, require that if ϕ is a logical consequence of Γ , ϕ expresses a proposition relative to every Context relative to which all the members of Γ express propositions. But given the kinds of considerations that are most often cited in support of the view that some sentences fail to express any propositions relative to certain

An alternative, favoured by Kaplan, instead defines logical consequence in terms of the notion of a sentence being ‘true at’ a Context. Kaplan defines this by way of an intermediate class of entities called ‘circumstances of evaluation’, each of which includes ‘a possible state or history of the world, a time, and perhaps other features as well’ (Kaplan 1989, p. 502). Each proposition can be evaluated as true or false with respect to any circumstance. Each Context determines a unique circumstance. This new ideology enables the following alternative account of logical consequence:

(K4b) ϕ is a logical consequence of Γ iff for each Context c relative to which ϕ and all the members of Γ express propositions, if all the propositions expressed relative to c by members of Γ are true with respect to the circumstance determined by c , then the proposition expressed relative to c by ϕ is true with respect to the circumstance of c .

Given the particular stipulations Kaplan makes about what counts as a Context, the result of (K4b) is to make ‘I am here’ but not ‘I am sentient’ or ‘I am uttering a sentence’ count as logically valid. However, it is unclear to me what the principled basis for this distinction is supposed to be.

Note that the framework as I have described it is compatible with a wide range of accounts of the nature of Contexts. On one view (associated with Lewis 1980), a Context is something like an ordered triple of a possible world, an interval of time in that world, and a person. This approach allows a relatively straightforward theory of the ‘in’ relation: at w' , utterance u is in $\langle w, t, i \rangle$ iff $w = w'$ and t is the interval occupied by u at w' and i is the utterer of u at w' .¹⁸ On a second approach, a Context is an ordered

Contexts, this would yield bizarre results. For example, standard accounts of the context-sensitivity of ‘tall’ suggest that when u is an utterance of a sentence that does not contain ‘tall’, and whose utterer is not thinking at all about the word ‘tall’, u is in a Context relative to which sentences involving ‘tall’ do not express propositions. On such a view, there are contexts relative to which ‘Snow is white’ expresses a proposition and ‘Snow is white or Obama is tall’ expresses no proposition. But it would be a bad idea to conclude, on this basis, that the argument from ‘Snow is white’ to ‘Snow is white or Obama is tall’—an instance of disjunction introduction—is not valid.

¹⁸The “short n -tuple” approach runs into trouble with cases where two utterances by the same agent have the same temporal boundaries: one might speak the sentence ‘he

n -tuple for some large n , and a lot more goes in to the characterisation of the ‘in’ relation. For example, a certain co-ordinate may be occupied by a place which plays the role of the centre in the interpretation of the verbs ‘come’ and ‘go’, and the fact about the value of this co-ordinate for the Context a given utterance is ‘in’ may depend in some complex way on such factors as the intentions of the utterer, the facts about what is salient to the utterer and the audience, the previous course of the conversation, the facts taken for granted by the relevant people, etc. On a third approach (associated with Stalnaker 1978), a Context is a proposition or set of propositions, and an utterance is in a Context iff the propositions in the Context are exactly those that are mutually accepted by the utterer and the (intended?) addressees.

How, in a Kaplanian framework, should we make sense of the contrast between uniform and non-uniform interpretation? There are several options. One option, is to characterise non-uniform utterances in terms of *Context-switching*. This is straightforward when we are dealing with non-uniformity *across* sentences: a non-uniform discourse can be characterised as one whose constituent sentence-utterances are in distinct Contexts, relative to which at least one of the sentences uttered expresses different propositions. However, it is not clear how this account could be applied to non-uniformity *within* sentences. The most obvious approach would be to modify (K2) so as to allow different temporal phases in the utterance of a single sentence to be “in” different Contexts, and extend (K1) so as to allow us to talk of the proposition semantically expressed by a sentence relative to a temporal sequence of Contexts, or something like that.¹⁹ But this does not seem very satisfactory. First, the idea that the central

won’t be happy’ on the phone, referring to one person, while simultaneously typing it into an email, referring to someone else. Further difficulties arise when applying the approach to written language. The process that results in a sentence being printed in a book may span hours or years, with long periods of hiatus, and periods during which only small edits are being made; in such cases it is very unclear where we should locate the temporal boundaries of the relevant “utterance”. Similar issues arise in connection with recorded speeches, such as answering machine messages: see Sidelle 1991, Predelli 2005, Cohen forthcoming.

¹⁹Considering written communication suggests that we should not be too literal-minded in thinking about Context-switching as a temporal matter. The process whose end result is a written sentence need not have different temporal phases corresponding to the structure of that sentence. For example, a kidnapper could build up a ransom

theoretical notion of semantics should be that of a sentence expressing a proposition relative to a temporal sequence of Contexts does not look very appealing: the relation seems to mix up considerations that belong at the level of types with considerations that belong at the level of tokens. Second, the approach forces us to answer awkward questions about *when* the relevant Context-shifts occur. For example, we will need to associate a typical utterance of (13) with at least two Contexts, such that the occurrence of ‘everyone’ is in one of them and the occurrence of the indefinite article is in another:

(13) Everyone is asleep and being monitored by a research assistant.

Are there just two Contexts, or more than two? If we say that there are only two, we face the awkward question at what point in the utterance of (13) the change of Context occurs—a question for which it is hard to provide a non-arbitrary answer.²⁰ If we say there are more than two—e.g. if we say that every word-occurrence in (13) is in a different Context—it looks like we are multiplying Contexts to such an extent that it will almost never happen that an utterance is in a single Context. And that would be odd. If it is standard for each word of an utterance to be in a distinct Context, we no longer have a satisfying story about uniform utterances, such as those produced by people who are explicitly trying to construct valid arguments.

A different approach is to hold on to (K2), while adopting one of the fine-grained approaches to the individuation of sentences discussed in section 6. On such a view, non-uniform utterances of sentences can just be utterances of unco-ordinated (or less-than-maximally co-ordinated) sentences. However, the objections to the extreme homonymy theory raised in section 6—most importantly, the fact that the only principled answer to the question when distinct utterances are of the same sentence entails that the relevant expressions are not context-sensitive at all are—are not in any way mitigated by first pasting on all the ‘a’s, then all the ‘b’s, etc. The idea of a *temporal* sequence gets no grip in such cases.

²⁰We could sidestep the question by claiming that it is *indefinite* when the change of Context in the relevant utterance occurs. But as I say below in connection with this sort of suggestion, the more we find ourselves having to do this, the more tempting will be the thought that we would do better to find a mode of theorising in which the relevant questions do not arise.

igated by the Kaplanian machinery. And it is hard to see any good way to apply the machinery to sentences as conceived of by moderate homonymy theories. For example, let ϕ be the unco-ordinated disambiguation of 'Obama is tall and Obama is not tall'. Let u be a typical utterance of ϕ , and let c be the Context u is in. Suppose that ϕ semantically expresses, relative to c , the proposition that Obama is tall for a politician and not tall for an amateur basketball-player. What proposition (if any) does 'Obama is tall' express relative to c ? There is no obvious basis for a principled answer, unless we are prepared to say that it expresses no proposition at all. But saying that risks making it uncomfortably rare for any sentence other than the sentence an utterance is *of* to semantically express a proposition relative to the Context that utterance is in. (Is there *any* Context relative to which ϕ and 'Obama is tall' both express propositions? Could any utterance be *in* such a Context?) Further, logical consequence as analysed by (K4a)/(K4b) is a relation between a set of sentences and a sentence, whereas in a moderate homonymy theory, it is much more natural to think of logical validity as a property of "arguments" or "discourses", entities given by a structure of co-ordination relations over some sentences.

Hybrid approaches are possible, which treat some cases of non-uniformity in terms of Context-shifting and other cases in terms of homonymy. Indeed, the Kaplanian framework seems to force us to adopt some sort of hybrid approach. We are clearly going to need something like the apparatus of indices to deal with the context-sensitivity of pronouns: Context-switching cannot helpfully explain the contrast between the bound and deictic readings of 'Every farmer wishes that his corn grew faster'. And the whole apparatus of Contexts would have little point if it did not sometimes happen that the same sentence is uttered in different Contexts relative to which it expresses different propositions. But hybrid approaches do not provide any obvious reply to the worries I raised about Context-shifting and homonymy. In fact, the need to posit two very different mechanisms in accounting for the apparently unitary phenomenon of non-uniform interpretation seems itself to be a flaw in the Kaplanian approach to semantics. The difficulty in finding a natural division of labour between the two mech-

anisms brings out how much nicer it would be to have a way of formulating semantic theories that did not impose any such dichotomy.

The hegemony of the Kaplanian framework helps to explain philosophers' unwillingness to rely on the idea that certain sentences are naturally interpreted non-uniformly as a tool for resolving paradoxes like our paradox of transparency. The fact that it is hard to give a natural, general description of non-uniform utterances of single sentences within the framework encourages the assumption that such utterances must be special, marginal phenomena which philosophers would do best to ignore. (One would probably want to make an exception for demonstratives and pronouns, since non-uniform uses of them are so commonplace; but this will result in an exaggerated sense of the difference between these expressions and other context-sensitive vocabulary.) But non-uniformity within sentences is a pervasive phenomenon. If it is hard to characterise it using the Kaplanian framework, that is a problem for the framework, and not a problem for views like mine which characterise certain utterances as non-uniform even though their non-uniformity is not obvious.

Let me briefly mention two other worries about the Kaplanian framework which might also help motivate the search for an alternative way of characterising the semantic roles of context-sensitive sentences.

(i) *Characters are not conventional enough.*

The use of 'I' and 'today' is (in most settings) is constrained by some rather simple rules, which can perspicuously be expressed by assigning these expressions their usual Kaplanian characters. But these "automatic indexicals" are quite atypical of the general run of context-sensitive expressions. Consider, by contrast, the expression 'What's-his-face', which seems to be capable of being used to refer to just about anything, or at least any male person. What is the Kaplanian character of (27)?

(27) What's-his-face has been elected Mayor of London.

Suppose for concreteness that our Contexts are ⟨world, time, individual⟩ triples. Then our account of what it takes for a proposition to be semantically expressed by (27)

relative to $\langle w, t, i \rangle$ will have to depend heavily on facts about the mental states (e.g. communicative intentions) of i during t at w , or on the speech acts being performed by i during t at w , or on some other facts that are closely correlated with these. For if we look at the range of propositions asserted by utterances of (27) across possible worlds, the only clear common factor we see is that they all are propositions to the effect that x has been elected Mayor of London, where x is the person concerning whom the speaker is attempting to communicate that they have been elected Mayor of London. Thus, the only way to respect (K3) is to allow such mental or speech-act relations to figure in the character of (27). And in this respect, most context-sensitive sentences, e.g. 'Obama is tall', are like (27). Facts about the external circumstances often *constrain* what one can reasonably try to get across by using a particular sentence, but they almost always allow the speaker considerable latitude.²¹

This looks odd if one thinks of semantics as centrally concerned with characterising the *conventional* element in language use. The fact that one can only use 'I' to refer to oneself is a distinctive fact about this particular word, something which students of English simply have to learn. By contrast, the fact that one can only use 'What's-his-face' to refer to *the thing one intends to be communicating something about by uttering 'What's-his-face' on this occasion* looks more like a necessary truth about what it is for a speaker to refer to something. The question what exact form this constraint takes doesn't look like the sort of question we should have to answer in characterising the semantic profile of this particular English expression.^{22,23}

²¹For example, 'now' normally has to be used to denote a period which overlaps the utterance, but speakers have free rein as regards how long this period is.

²²Soames (2008; 2009, n. 3) makes a closely related point. For Soames, the point of semantic theory is to be something knowledge of which can play a certain role in 'a rational reconstruction of the ability of speaker-hearers to identify what is asserted and conveyed by an utterance' (2008, p. 322). This explanatory role cannot be played by a semantic theory in which the a certain expression is assigned a character whose value relative to the Context of an utterance depends on what is being asserted or conveyed by the relevant utterance: to apply one's knowledge of such a semantic theory, one would need already to have exercised the ability of which we are trying to give a rational reconstruction.

²³On the "long n -tuple" view of Contexts, one can avoid having to appeal to notions like *speaker's intentions* in the specification of characters, by including appropriate

(ii) *Too much arbitrariness in what utterances semantically express.*

As we have seen, it is a consequence of (K1) and (K2) that each utterance semantically expresses at most one proposition. A worry about this is that for many utterances, there may be no principled grounds for selecting one particular proposition as *the* proposition semantically expressed. If we have independent grounds for thinking that a particular utterance is literal, (K3) tells us that we had better choose the proposition semantically expressed from among those that the utterer *asserts* by means of that utterance.²⁴ But the claim that utterances of declarative sentences typically involve the assertion of just one proposition is quite controversial. Opponents (e.g. Cappelen and Lepore 2005, Soames 2009) point to the fact that in many cases, a single utterance can be described, truly, by means of a multifarious collection of sentences of the form ‘So-and-so said [/claimed/asserted/. . .] that *P*’.²⁵ Moreover, even if literal utterances normally involve the assertion of just one proposition, the arbitrariness worry arises with regard to the choice of a single proposition to be the proposition semantically expressed by a non-literal utterance. In some nonliteral utterances, we may be able to appeal to the speaker’s intentions to pick out a particular proposition as the one semantically expressed: for example, when I ironically utter ‘Fred is as smart as What’s-his-face’, I may intend my audience’s route to figuring out what I am trying to get across (that Fred

co-ordinates in the Contexts. For example, we might say that the character of ‘What’s-his-face’ is just the function that maps each Context *c* to the the 117th element of *c*; as part of our theory of what it takes for an utterance to be *in* a Context, we say that co-ordinate 117 of the Context an utterance is *in* is whoever the speaker is using the expression ‘What’s-his-face’ to refer to (if there is such a person; otherwise, we had better allow the Context to have a “gap” at position 117). But on this view, the *in* relation is not a language-independent matter; our account of what it takes for an utterance to be in a Context is just as specific to our theory of the particular language we are interested in as our account of what it takes for a sentence to express a proposition relative to a Context. So we still face the worry that our total semantic theory for a particular language will include a lot of details that look like they really belong in a general account of the nature of assertion and other speech acts.

²⁴Kaplan ties semantic expression to *saying*: one says that *p* iff one produces an utterances which semantically expresses the proposition that *p*. However, it is not clear to what extent ‘says’ as used by Kaplan is meant to be interpreted as a special philosophical term of art.

²⁵For another line of argument that it is common for many propositions to be asserted at once, see Dorr and Hawthorne MS.

isn't very smart) to involve, as an intermediate step, their entertaining the proposition that Fred is as smart as Richard Feynman. But in other cases, there may be no such specific intentions—for example, if I (metaphorically) utter 'Newton was a giant', it does not seem that my mental state needs to contain anything that would be helpful in answering the question just how tall Newton would have to have been for the proposition semantically expressed to be true.²⁶

Often in philosophy, the right rejoinder to objections based on arbitrariness is to invoke the ideology of vagueness. One could make this move in this case: maybe there are often several candidates to be the proposition semantically expressed by an utterance, none of which is *definitely* semantically expressed, although it is definite that exactly one of them is semantically expressed. Whether this is a satisfactory response to the worry depends on what kind of theoretical work we wanted the relativised notion of semantic expression to do. While even extremely vague theoretical terms can sometimes do useful explanatory work (e.g. in economics), it is hard to see the explanatory point of a notion of semantic content whose uniqueness is achieved simply by means of an extra dose of vagueness which corresponds to nothing in the phenomena to be explained.

To bring this discussion to a close, let me mention a couple of alternatives to the Kaplanian framework for characterising the semantics of context-sensitive sentences. One simple way to modify the framework is simply to give up (K1), allowing a sentence to express many propositions relative to a single Context. The thought is that the set of propositions semantically expressed by a sentence relative to a Context are those *available* to be asserted by a literal speech in that Context; in general, only some of them will in fact be asserted. Thus (K3) would be replaced with

(K3*) Necessarily, when an utterance of a declarative sentence is literal, the utterer asserts *at least one* of the propositions it semantically expresses.

We still need the apparatus of Contexts so that our semantic theory can represent the

²⁶Thanks to Emanuel Viebahn for discussions on this point.

distinctive constraints that govern the use of words like ‘I’ and ‘you’. But for many context-sensitive sentences, such as ‘Obama is tall’, we can take the set of propositions semantically expressed to be the same relative to any Context: there is no *conventional* constraint on what these sentences can be used to assert on a given occasion.

Unless we embrace the extreme or moderate homonymy theory, our semantic ideology will also need to make a division within the set of propositions semantically expressed by a sentence relative to a Context, between those that are uniform interpretations of the sentence and those that are not.

I am not aware of any defence in the literature of this “multiple-propositions” approach to semantic theory. However, a closely related proposal has been discussed under the label ‘semantic incompleteness’ (Bach 1994, Soames 2005). On this approach, the semantic values of certain sentences are taken to be ‘propositional radicals’ (Bach), or ‘propositional matrices’ (Soames), thought of as structures with “gaps” which can be filled in different ways to yield different propositions, its “completions”. Where a multiple-propositions theorist would say that a sentence semantically expresses, relative a certain Context *c*, all and only the propositions in a certain set *s*, an incompleteness theorist will say that the sentence’s semantic value relative to *c* is a propositional matrix whose completions are exactly the propositions in *s*. It is not obvious that anything substantive hangs on this difference.²⁷

Proponents of semantic incompleteness have not had much to say about how we should understand logical consequence and inconsistency for semantically incomplete sentences. As far as I can see, any reasonable account of such these notions will need to invoke some contrast between uniform and non-uniform completions of a propositional matrix. If we think of the gaps in a propositional matrix as coming in different types, we could define a “uniform completion” of a set of matrices to be a function that assigns each of the matrices to one of its completions in such a way

²⁷If there are sets of propositions which cannot be represented as the set of completions of some propositional matrix, the propositional matrices approach is less flexible than the multiple-propositions approach. To see whether this loss of flexibility is worrisome, one would need a more fully worked-out account of the ontology of propositional matrices.

that gaps of the same type are filled by the same entity; then logical consequence for matrices could be defined as propositional logical consequence under all uniform completions.²⁸

8 The source of the context-sensitivity in attitude reports

Adherence to the Kaplanian framework generally goes along with acceptance of a principle of compositionality, to the effect that the proposition expressed by a sentence relative to a Context is determined, in accordance with invariant rules, by the semantic values of its elementary constituents relative to that Context and the way they are put together.²⁹ However this idea is cashed out, it will surely have the following as a corollary:

Constituency: Every context-sensitive sentence must have at least one context-sensitive elementary constituent

In the recent literature, opposition to the Kaplanian hegemony has tended to coincide with rejection of *Constituency*. For example, incompleteness theorists who hold that the semantic value of ‘Tipper is ready’ is the propositional matrix *Tipper is ready for* __ often think that the semantic value of ‘Tipper’ is Tipper, and the semantic value of ‘is ready’ is the binary relation *being ready for*. The result is that the sentence has multiple admissible interpretations (the results of filling its gap in different ways) even

²⁸Characterising uniform interpretation in terms of gaps of the same type being filled in the same way may require us to posit quite a lot of structure in the propositions that are admissible interpretations of context-sensitive sentences. For example, since the possibility of uniform interpretation arises for occurrences of ‘strong’ and ‘weak’, we will need to posit gaps associated with these occurrences which can be filled by something like a ‘dimension of comparison’. The multiple-propositions approach is more flexible: it allows us, if we want, to treat ‘strong’ and ‘weak’ as simply contributing unstructured properties, which must be related in some interesting way in the propositions that count as uniform interpretations.

²⁹As Rabern (MS) and Yli-Vakkuri (MS) point out, compositionality for contents is *prima facie* inconsistent with the usual semantic account of quantification in terms of assignment functions. However, even if one treats quantifiers as “monsters” which are exceptions to compositionality for contents, the case for *Constituency* will remain in place—quantifiers can turn assignment-sensitive expressions into assignment-insensitive ones, but cannot do the reverse.

though none of its syntactic constituents has multiple admissible interpretations.³⁰ However, *Constituency* is certainly consistent with the non-Kaplanian approaches we are currently considering. An incompleteness theorist could assign “gappy” semantic values to expressions other than sentences, and endorse compositional rules which entail that every sentence expressing a gappy proposition relative to a Context has an elementary constituent whose semantic value relative to that Context is gappy—for example, ‘ready’ could express the “gappy property” *being ready for* __.³¹ And a multiple propositions theorist could, of course, claim that every sentence that expresses multiple propositions relative to a Context has an elementary constituent that has multiple semantic values relative to that Context.

I think we should be loath to give up on *Constituency*. Compositionality principles of the kind that entail *Constituency* are presupposed by a great deal of explanatorily successful work in formal semantics. And without some principle of this sort in the background, the project of assigning semantic values to expressions other than sentences looks worryingly underconstrained. Thus the burden of proof is on opponents of *Constituency*. And I do not think that the burden of proof has yet been met: putative counterexamples often rest on dubious assumptions about the syntax of the

³⁰If we are working with a picture of propositions as structured, this kind of failure of *Constituency* will have the result the admissible interpretations of the sentence have “unarticulated constituents” (Crimmins and Perry 1989)—constituents which do not correspond to anything in the syntactic structure of the sentence. However, the denial of *Constituency* does not have to go along with an embrace of unarticulated constituents. According to Recanati (2011, ch. 1), for example, the content of a complex expression relative to a Context is derived from the contents of its constituents relative that Context by first applying to them a “modulation” function determined by the Context, and then applying one of the usual structural rules. On this view, the content of a complex expression relative to a Context will have a structure which matches the syntactic structure of the expression, although its constituents need not be the contents relative to the Context of the expression’s constituents. This gives Recanati a way to implement one of his leading ideas: that among the many admissible interpretations of a context-sensitive expression, one (the one without modulation functions) can be singled out as a default which will be the semantic value of an utterance of that expression unless some “optional process” occurs. However, giving up *Constituency* is certainly not the only way to accept this picture.

³¹In the terminology of Bach (1994), this would amount to the claim that all incompleteness is *lexical* incompleteness.

relevant sentence, or about its range of admissible interpretations, or about the range of admissible interpretations of certain words in the sentence.³²

If *Constituency* is true, we have to face the question which words are responsible for the characteristic kind of context-sensitivity exhibited by attitude reports. Different answers to this question will lead to different explanations of the non-uniformity that, I have argued, characterises typical utterances of

- (5) Even though Superman is Clark Kent, lots of people believe that Superman flies without believing that Clark Kent flies.

There are essentially three options.³³

(i) *Verbalism*. The relevant form of context-sensitivity is due to context-sensitivity in *attitude verbs* like ‘believe’ and ‘say’. Thus, for example, in a typical non-uniform utterance of (5), the two occurrences of ‘believes’ are to be interpreted in different ways. Even in these non-uniform utterances, we thus have no need to posit any difference in the semantic contributions made by the embedded sentences ‘Superman flies’ and ‘Clark Kent flies’. We are thus free to adopt a simple account of the propositions expressed by such sentences, e.g. as Russellian complexes of objects and properties.³⁴

(ii) *Clausalism*. The relevant form of context-sensitivity is due not to the attitude verbs, but to their sentential complements. Thus, on the occasion of a typical non-

³²See Stanley 2007 for discussion of many of the examples.

³³Perhaps I should have considered a fourth option, on which the relevantly context-sensitive word in ‘Lois believes that Superman flies’ is the complementiser ‘that’. The problem I see with this view is that seems to predict that despite the fact that ‘Superman flies’ is not (relevantly) context-sensitive, there is a range of different entities that ‘the proposition that Superman flies’ can admissibly be interpreted as referring to. This seems to require us to make the implausible claim that ‘the proposition that Superman flies is not one of the admissible interpretations of “Superman flies”’ admits true, uniform interpretations.

³⁴Richard (1990) defends a contextualist account of attitude reports on which their context-sensitivity is due to context-sensitivity in attitude verbs. However he also posits a (context-invariant) difference between the semantic values of sentences involving different names: on his account, the proposition expressed by any sentence is a complex containing the very words out of which that sentence is composed. Richard is driven to this unattractive view of propositions by the need to explain how various felicitous attitude reports involving co-referential names can be true *within a single context* (i.e. on a uniform interpretation). From my perspective, this is not something we need to do.

uniform utterance of (5), the clauses ‘Superman flies’ and ‘Clark Kent flies’ must contribute two different propositions, despite the fact that they attribute the same property to the same person.³⁵ The clausalist thus need a fine-grained, broadly Fregean account of propositions, which allows that propositions can be necessarily equivalent, and structurally isomorphic, while differing in the conditions one needs to satisfy in order to bear propositional attitudes towards them. Given *Constituency*, the relevant form of context-sensitivity in ‘Superman flies’ must be based on context-sensitivity in one of *its* syntactically elementary constituents, presumably ‘Superman’. Thus the clausalist will need a fine-grained semantics for names, according to which the semantic contribution of a name includes something like a (not necessarily fully specific) “mode of presentation” as well as a referent.

(iii) *Hidden indexicalism* (Schiffer 1979, Ludlow 1995). The distinctive form of context-sensitivity exhibited by attitude reports is due neither to propositional attitude verbs nor to their clausal complements. Rather, it is due to some unpronounced constituents which are really present in the syntax of the relevant sentences, although they have no phonological or orthographic manifestation.

The task of deciding between these options raises many difficult issues which I hope to address in future work. Let me here mention just one consideration which might be thought to settle the question decisively in favour of the clausalist, although in fact, as we will see, its probative force is far from clear. This arises from the *prima facie* plausible claim that the possibility of non-uniformity in interpretation only arises for sentences that contain *at least two* appropriately related occurrences of relevantly

³⁵One might think that the clausalist could introduce a notion of *compositional semantic value* on which the compositional semantic value of a sentence is not a proposition, and characterise the relevant occurrences of ‘Superman flies’ and ‘Clark Kent flies’ as contributing the same proposition while having different compositional semantic values. But this proposal does not respect the equivalence of ‘believes that *P*’ and ‘believes the proposition that *P*’. Given that ‘Lois believes the proposition that Superman flies’ is context-sensitive in exactly the same way as ‘Lois believes that Superman flies’, the clausalist must think that different admissible interpretation assign different *referents* to ‘the proposition that Superman flies’. And it is plausible that each admissible referent for ‘the proposition that Superman flies’ is an admissible interpretation for ‘Superman flies’ (see n. 33).

context-sensitive expressions, which could either be interpreted in the same way or in different ways. Call this “the multiple occurrences generalisation”. The generalisation is supported by cases of ellipsis like (28) and (29):

- (28) a. My house is bigger than this but isn’t bigger than this.
b. My house is bigger than this but isn’t.
- (29) a. Batman is strong but he isn’t strong.
b. Batman is strong but he isn’t.

By taking the two occurrences of ‘this’ to refer to different items, we can easily interpret (28a) as expressing a truth. But these interpretations are simply not available for (28b), which simply sounds awful. Similarly, we can (with more difficulty) access non-contradictory interpretations of (29a), by associating the two occurrences of ‘strong’ with different properties; (29b) cannot be saved in this way. The operation of ellipsis, which takes (28a) to (28b) and (29a) to (29b), thus removes the possibility of non-uniform interpretation.³⁶

If the multiplicity generalisation is correct, it is a useful tool for pinpointing the source of context-sensitivity in a sentence. Start with a sentence that admits a non-uniform interpretation; modify it somehow (e.g. by ellipsis, or by replacing sentential conjunction with sub-sentential conjunction) to remove repeated constituents; the context-sensitive constituent is the one such that when it is not repeated, the non-uniform interpretation ceases to be available. What happens when we apply this methodology to attitude reports? Consider (30):

- (30) Thelma believes that I limp but doesn’t believe that I limp.

Once we have primed ourselves by going through section 3’s No-Error Argument for the context-sensitivity of ‘Thelma believes that I limp’, we can, with some difficulty, access a non-uniform interpretation for (30) on which it is true in the circumstances of

³⁶See Stanley 2003 for a deployment of the idea that non-uniform readings are disallowed in cases of ellipsis, to argue against a certain kind of philosophical deployment of context-sensitivity.

Schiffer's example. This interpretation seems still to be available when we modify (30) so that 'believes' is not repeated:

(31) Thelma believes that I limp but not that I limp.

(Imagine this uttered with emphasis on the second 'I'.) This seems to rule out verbalism, according to which 'believes' is the relevantly context-sensitive expression. It also, arguably, rules out hidden-indexicalism. On the usual version of that view, the unvoiced elements it posits are associated with attitude verbs, of which they provide additional arguments, so that (31) contains only half as many of them as (30) does. By contrast, when we change things around so that the attitude verb *is* repeated but 'I limp' is not, the relevant non-uniform reading seems to disappear:

(32) Thelma believes that I limp but doesn't believe it.

Even if we imagine (32) spoken with emphasis on the second 'believe', it does not seem to admit the interpretation we have been considering. Non-uniform interpretations of (32) are arguably available—e.g. that Thelma's level of confidence exceeds some low threshold but not some higher threshold. But the kind of non-uniform reading we are interested in, to the effect that Thelma attributes limping to Shorty under a certain kind of guise but not under a certain other kind of guise, is much less easily associated with (32) than with (30) or (31). Given the multiplicity generalisation, this seems to settle the debate in favour of clausalism.

The problem with this argument is that there seem to be exceptions to the multiplicity generalisation. Consider (33):

(33) This car is mine, and it isn't [mine].

If we pronounce both occurrences of 'mine', we can access non-uniform interpretations of (33)—e.g. that the car is in my possession but I am not its legal owner. And in this case, eliding the second 'mine' does not make these interpretations unavailable: neither sentence is very good, but the elision doesn't make things appreciably worse. Similarly with (34), eliding the second 'strong' does not seem to affect the availability of salient

non-uniform readings like ‘Bruce Wayne is strong for a human being, but Batman is not strong for a superhero’:

(34) Bruce Wayne is strong, but Batman is not [strong].

Apparent exceptions to the multiplicity generalisation also arise with demonstratives, as in (35):

(35) Giorgone was so-called because of his size, but Bellini was not [so-called because of his size].

Here the obvious account of the sentence without elision is that ‘so-called’ involves a demonstrative, whose two occurrences are naturally taken to refer, respectively, to the names ‘Giorgione’ and ‘Bellini’. But surprisingly, the relevant reading remains available when we elide the material in square brackets. We see something similar in (36) and (37):

(36) The marks on the wall were put there by my kids, but the stain on the carpet was not [put there by my kids].

(37) People who are anxious get to be that way because of their upbringing, but people who are depressed do not [get to be that way because of their upbringing].

One might try to reconcile the the acceptability of the elided version of (36) with the multiplicity generalisation by claiming that ‘put there by my kids’ can mean ‘put in the place where they are by my kids’. But note that (37) is perfectly compatible with there being people who are both anxious and depressed, so we cannot, analogously, account for its acceptability by glossing ‘get to be that way because of their upbringing’ as anything like ‘get to be the way they are because of their upbringing’.

I do not pretend to have an account of what is going on in these examples. Indeed I doubt that they are all to be explained in the same way. I mention them here just to illustrate the pitfalls faced by attempts to use the multiplicity generalisation as a tool for identifying the source of context-sensitivity. As a further reason for caution,

notice that in some cases, a non-uniform reading seems to be available even when *all* potentially relevant constituents are elided:

- (38) Superman is widely believed to fly, but Clark Kent is not [widely believed to fly].

However the consistent non-uniform interpretations of the unelided version of (38) work, they have to involve two different properties associated with the two occurrences of ‘widely believed to fly’. But eliding the second occurrence of this expression, which includes both the attitude verb and its clausal complement, makes no noticeable difference to the felicity of the speech.³⁷ ; but consistent non-uniform interpretation This looks like bad news for the multiplicity generalisation.³⁸

The question which constituents are responsible for the context-sensitivity of attitude-reports thus remains wide open. If *Constituency* were false, the question might have a false presupposition; but we have not seen any good reason to treat the context-sensitivity of attitude reports as a counterexample to *Constituency*. And my central thesis—that the Attitude Transparency Schema is valid—is compatible with all of the answers to the question. In particular, it is compatible with clausalism. Given the clausalists’ need for a fine-grained, “Fregean” view of propositions, one might have thought they would be committed to regarding the Attitude Transparency Schema as invalid. But this is a mistake. Validity only entails truth under uniform interpretations. And as we have seen, uniformity can require assigning the same, or appropriately related,

³⁷According to standard syntactic theory, the complement of ‘believe’ in (38) is a tenseless clause, ‘PRO to fly’.

³⁸As we have seen, clausalism seems to require a Fregean approach to semantics, according to which the true non-uniform readings of ‘Lots of people believe that Superman flies but not that Clark Kent flies’ involve assigning different semantic values to the occurrences of ‘Superman’ and ‘Clark Kent’, despite their having the same referent. But we cannot account for the true non-uniform interpretation of (38) by saying that ‘Superman’ and ‘Clark Kent’ make different semantic contributions while the two occurrences of ‘widely believed to fly’ make exactly the same semantic contribution. For no matter what exotic entities we might assign as *semantic values* to names and predicates, it remains overwhelmingly plausible that the *truth value* of the proposition expressed by a subject-predicate sentence on a particular interpretation depends only on the *referent* of the subject and the *property expressed* by the predicate.

semantic values not only to multiple occurrences of the same word, but to occurrences of different words, e.g. the occurrences of ‘Superman’ and ‘Clark Kent’ in (5).

Up to now, the question whether the Attitude Transparency Schema is valid has played a central role in the debate between Russellian and Fregean theories about the nature of propositions. It has been generally assumed that if the Attitude Transparency Schema is valid, the Russellian theory (or some even coarser-grained theory, such as the view that propositions are sets of worlds) should be endorsed, while the converse claim—that if the Russellian theory is true the Attitude Transparency Schema is valid—has been somewhat more controversial. But this consensus is mistaken. Even if we accept the Attitude Transparency Schema, there might well turn out to be good reason to favour clausalism over the other accounts of the source of the context-sensitivity of attitude reports, in which case the Russellian account of propositions is not adequate. The philosophy of language still has plenty to contribute to the metaphysical debate between Russellian and Fregean theories of propositions. But the relevant issue in the philosophy of language is not that of “transparency” versus “opacity”, but the question which of the constituents of attitude-reports are responsible for their context-sensitivity.

References

- Bach, Kent (1994). 'Conversational Implicature'. *Mind and Language* 9, pp. 124–162 (cit. on pp. 42, 44).
- (1997). 'Do Belief Reports Report Beliefs?' *Pacific Philosophical Quarterly* 78, pp. 215–241 (cit. on p. 10).
- Braun, David (1988). 'Understanding Belief Reports'. *Philosophical Review* 107, pp. 555–595 (cit. on p. 6).
- (2002). 'Cognitive Significance, Attitude Ascriptions, and Ways of Believing Propositions'. *Philosophical Studies* 108, pp. 65–81 (cit. on p. 6).
- Cappelen, Herman and Ernest Lepore (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden, Mass.: Blackwell (cit. on pp. 12, 40).
- Cohen, Jonathan (forthcoming). 'Indexicality and the Puzzle of the Answering Machine' (cit. on p. 35).
- Crimmins, Mark (1992). *Talk About Beliefs*. Cambridge, MA: MIT Press (cit. on p. 10).
- Crimmins, Mark and John Perry (1989). 'The Prince and the Phone Booth: Reporting Puzzling Beliefs'. *Journal of Philosophy* 86, pp. 685–711 (cit. on pp. 10, 44).
- Dorr, Cian (2005). 'What We Disagree About When We Disagree About Ontology'. In *Fictionalist Approaches to Metaphysics*, ed. Mark Kalderon. Oxford: Oxford University Press, pp. 234–286 (cit. on p. 3).
- (2011). 'De Re A Priori Knowledge'. *Mind* 120 (cit. on p. 10).
- Dorr, Cian and John Hawthorne (MS). 'Semantic Plasticity' (cit. on p. 40).
- Field, Hartry (2008). *Saving Truth from Paradox*. Oxford: Oxford University Press (cit. on p. 20).
- Fiengo, Robert and Robert May (1994). *Indices and Identity*. Cambridge, MA: MIT Press (cit. on p. 26).
- Fine, Kit (2007). *Semantic Relationalism*. Malden, Mass.: Blackwell (cit. on p. 26).
- Hofweber, Thomas (2000). 'Quantification and Non-existent Objects'. In *Empty Names, Fiction and the Puzzles of Non-Existence*, ed. Anthony Everett and Thomas Hofweber. Stanford: CSLI Press (cit. on p. 3).
- Kaplan, David (1989). 'Demonstratives'. In *Themes from Kaplan*, ed. Joseph Almog, John Perry and Howard Wettstein. Oxford: Oxford University Press, pp. 481–566 (cit. on pp. 32, 34).
- Lewis, David (1979). 'Attitudes De Dicto and De Se'. *Philosophical Review* 88, pp. 513–543 (cit. on p. 10).

- Lewis, David (1980). 'Index, Context and Content'. In *Philosophy and Grammar*, ed. Stig Kanger and Sven Öhman. Dordrecht: Kluwer Academic Publishers (cit. on p. 34).
- (1986). *On the Plurality of Worlds*. Oxford: Blackwell (cit. on p. 23).
- (1993). 'Many, But Almost One'. In *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*, ed. Keith Campbell, John Bacon and Lloyd Reinhardt. Cambridge: Cambridge University Press, pp. 23–42 (cit. on p. 18).
- Ludlow, Peter (1995). 'Logical Form and the Hidden Indexical Theory: A Reply to Schiffer'. *Journal of Philosophy* 92, pp. 102–107 (cit. on p. 46).
- Oppy, Graham (1992). 'A Semantics for Propositional Attitude Ascriptions'. *Philosophical Studies* 67, pp. 1–18 (cit. on p. 10).
- Predelli, Stefano (2005). *Contexts: Meaning, Truth, and the Use of Language*. Oxford: Oxford University Press (cit. on p. 35).
- Quine, Willard van Orman (1940). *Mathematical Logic*. Boston: Harvard University Press (cit. on p. 29).
- Rabern, Brian (MS). 'Monsters in Kaplan's Logic of Demonstratives' (cit. on p. 43).
- Recanati, François (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell (cit. on p. 10).
- (2011). *Truth-Conditional Pragmatics*. Oxford: Oxford University Press (cit. on p. 44).
- Richard, Mark (1983). 'Direct Reference and Ascriptions of Belief'. *Journal of Philosophical Logic* 12, pp. 425–52 (cit. on p. 23).
- (1990). *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*. Cambridge: Cambridge University Press (cit. on pp. 10, 23, 45).
- (1993). 'Attitudes in Context'. *Linguistics and Philosophy* 16, pp. 126–148 (cit. on p. 10).
- Salmon, Nathan (1986). *Frege's Puzzle*. Cambridge, MA: MIT Press (cit. on p. 5).
- Saul, Jennifer (2007). *Simple Sentences, Substitution, and Intuitions*. Oxford: Oxford University Press (cit. on p. 6).
- Schiffer, Stephen (1979). 'Naming and Knowing'. In *Midwest Studies in Philosophy II: Contemporary Perspectives in the Philosophy of Language*, ed. Peter French, Theodore E. Uehling Jr. and Howard K. Wettstein. Minneapolis: University of Minnesota Press, pp. 28–41 (cit. on pp. 10, 46).
- Sidelle, Alan (1991). 'The Answering Machine Paradox'. *Canadian Journal of Philosophy* 21, pp. 525–39 (cit. on p. 35).
- Soames, Scott (1986). 'Incomplete Definite Descriptions'. *Notre Dame Journal of Formal Logic* 27, pp. 349–75 (cit. on p. 14).

- Soames, Scott (1987). 'Direct Reference, Propositional Attitudes, and Semantic Content'. *Philosophical Topics* 15, pp. 47–87 (cit. on p. 5).
- (2005). 'Naming and Asserting'. In *Semantics Versus Pragmatics*, ed. Zoltán Gendler Szabó. Oxford: Oxford University Press, pp. 356–382 (cit. on pp. 10, 42).
- (2008). 'Drawing the Line Between Meaning and Implicature — and Relating both to Assertion'. *Noûs* 42, pp. 529–554 (cit. on p. 39).
- (2009). 'The Gap Between Meaning and Assertion. Natural Language: What It Means and How We Use It'. In *Philosophical Essays, vol. 1*. Princeton: Princeton University Press (cit. on pp. 39, 40).
- Sosa, Ernest (1970). 'Propositional Attitudes *De Dicto* and *De Re*'. *Journal of Philosophy* 67, pp. 883–896 (cit. on p. 10).
- Stalnaker, Robert C. (1978). 'Assertion'. In *Syntax and Semantics* 9, ed. P. Cole. New York: New York Academic Press, pp. 312–32 (cit. on p. 35).
- (1997). 'Reference and Necessity'. In *A Companion to the Philosophy of Language*, ed. Bob Hale and Crispin Wright. Oxford: Basil Blackwell (cit. on p. 2).
- Stanley, Jason (2003). 'Context, Interest-Relativity, and the Sorites'. *Analysis* 63, pp. 269–81 (cit. on p. 47).
- (2007). *Language in Context*. Oxford: Oxford University Press (cit. on pp. 26, 45).
- Stanley, Jason and Zoltán Gendler Szabó (2000). 'On Quantifier Domain Restriction'. *Mind and Language* 15, pp. 219–261 (cit. on p. 15).
- Szabó, Zoltán Gendler (2008). 'Compositionality'. In *The Stanford Encyclopedia of Philosophy (Winter 2008 edition)*, ed. Edward N. Zalta (cit. on p. 30).
- Yli-Vakkuri, Juhani (MS). 'Propositions and Compositionality' (cit. on p. 43).