

Compositionality, innocence and the interpretation of NPs^{*}

GEORGE POWELL

Abstract

Neale (1999) outlines a theory of meaning which he claims can account for a range of problematic semantic data while respecting two intuitively plausible constraints on semantic theory: compositionality and semantic innocence. In this paper I argue that Neale's theory, hindered as it is by its underlying conception of the relation between language and thought, fails to satisfy the criteria it sets for itself. Once this conception is abandoned, a new analysis becomes available which accounts for the same semantic phenomena within a more parsimonious framework, as well as respecting both compositionality and semantic innocence, properly understood.

1 Introduction

There are two constraints on the design of a semantic theory which, while both apparently desirable, have proved difficult to reconcile: on the one hand, the Principle of Compositionality requires that the meaning of a sentence should be a function of the meanings of its component lexical items and their syntactic combination; and, on the other, according to the Principle of Semantic Innocence, a particular linguistic expression should have the same meaning regardless of the linguistic context in which it appears. In this paper I want to examine and assess one attempt to account for a range of traditionally problematic meaning phenomena within a semantic theory that respects both of these principles, that of Neale (1999). Neale focuses on two separate areas of semantic difficulty: those aspects of meaning which Frege grouped under the *colouring* banner, and which others have thought of as non-truth conditional meaning, and the interpretation of certain types of NP. In this paper I want to focus on the second of these areas: I aim to argue that Neale's account of the interpretation of NPs is the product of a

^{*}I would like to thank Robyn Carston for her patience and support, as well as for many helpful comments on earlier drafts of this paper.

particular conception of the relation between language and thought, and that, once this conception is abandoned, a more theoretically elegant account begins to emerge, one which respects both compositionality and semantic innocence, properly understood, as well as the *efficiency* of language in the sense of Barwise & Perry (1983/1999):

By the *efficiency* of language we mean this: expressions used by different people, in different space-time locations, with different connections to the world around them, can have different interpretations, even though they retain the same linguistic meaning.

Barwise & Perry (1983/1999, p. 5)

It's worth stating at the outset, however, that I don't intend to pick over the minutiae of Neale's account, except inasmuch as they cast light on his overall conception of meaning. As he himself makes clear, his aim is not to present a fully articulated analysis of the expressions he tackles, but rather the outline of a fruitful approach to the semantics and pragmatics of those expressions. Given this, there may well be shortcomings in the detail that could be changed without affecting the overall framework. My interest is in the framework itself.

2 Frege: the clash between compositionality and semantic innocence

Compositionality has been one of the primary desiderata of more or less all semantic theorising since Frege placed the compositionality of both sense and reference at the heart of his semantics. But why should we want a semantic theory to be compositional? Firstly, compositionality ties in with some deeply-felt intuitions: sentence (1), for instance,

(1) Cats like plain crisps

appears to mean what it does because the words 'cats', 'like', 'plain' and 'crisps' mean what they do and because they're combined in the way they are. It's just because of this, that sentences (1) and (2) and (3):

(2) Plain crisps like cats

(3) Dogs like plain crisps

are similar in some semantic respects but different in others.

Beyond this intuition, there is another frequently cited argument for the compositionality of meaning: we are all, as native speakers of a natural language, capable of understanding an infinite number of the sentences of our language. A non-compositional semantic theory, the argument goes, would be unable to account for this fact: we cannot store an infinite number of whole sentence meanings in a finite brain, so we must instead store a finite number of meaning atoms and a finite number of recursive rules for combining those atoms. The meaning of a sentence, therefore, must be made up from the meanings of its lexical items/morphemes and their syntactic combination; it must, in other words, be compositional.

A strict adherence to compositionality within a semantic theory, however, seems to throw up some unwanted results, particularly when it comes to the analysis of belief contexts. The problem can be approached from various angles. Consider, for instance, the following sentences:

- (4) The moon is made of cream cheese
- (5) Bill believes that the moon is made of cream cheese

As is well known, belief contexts such as (5) present a particular problem for a straightforwardly Fregean semantics, on which sentences express propositions and refer to truth values: (4), as a sentence, must express (i.e. have as its sense) a proposition, and refer to a truth value; (4) appears to be a constituent of (5), so, armed with our commitment to compositionality, we must want to say that the proposition expressed by (5) is partly determined by the proposition expressed by (4), and the truth value of (5) is partly determined by the truth value of (4). But that can't be so: the truth values of (4) and (5) are independent. That is to say, it is possible for (5) to be true when (4) is true, but also when (4) is false: Bill may believe that the moon is made of cream cheese whether he is right to do so or not.

How might this difficulty be resolved? There appear to be two possibilities: either we can sacrifice our commitment to compositionality (we can, in other words, say that the reference of (5) is not, as we had believed, partly determined by the reference of the expression 'the moon is made of cream cheese'), or we might, as Frege did, claim that the referent of the expression 'the moon is made of cream cheese' as it appears in (5) is not the same as the referent of that expression when it appears alone, as in (4). For Frege, a sentence embedded in a belief context refers not to a truth value, but to the proposition it expresses when it stands alone, what Frege (1892) calls its *customary sense*.

But, in adopting this solution, Frege is abandoning semantic innocence, a guiding principle that seems, to many, to have as much intuitive support as compositionality: on Frege's account, the meaning of an expression changes from one linguistic context to another. Much the same problem arises for referring expressions within belief contexts:

- (6) Cary Grant was born in London
- (7) Archie Leach was born in London
- (8) Mary believes that Cary Grant was born in London
- (9) Mary believes that Archie Leach was born in London

In (6) and (7), the co-referring proper names 'Cary Grant' and 'Archie Leach' seem to be substitutable *salva veritate*. This kind of evidence (among much else) has led many theorists within what has become known as the Direct Reference tradition to conclude that the semantics of a proper name is exhausted by its referent. Yet, as is well known, this substitutability seems to disappear in belief contexts: (8) may well be true at the same time as (9) is false, although they differ only in the substitution of one name for another co-referential name. Again we seem to be faced with the same choice: if we want to preserve compositionality, i.e. if we want to claim that 'Cary Grant' contributes to the meaning of (8), then we seem to be forced to the conclusion that 'Cary Grant' has somewhat different meanings in (6) and (8); we're forced, in other words, to abandon semantic innocence¹.

But why is this a problem? Why do we want a theory of meaning to respect semantic innocence? The fundamental intuition behind semantic innocence is most clearly expressed by Davidson (1969) (from whom Barwise & Perry (1981) originally took the term 'semantic innocence'):

If we could but recover our pre-Fregean semantic innocence, I think it would be plainly incredible that the words "the earth moves", uttered after the words "Galileo said that", mean anything different, or refer to anything else, than is their wont when they come in other environments. Language is the instrument it is because the same expression, with semantic features (meaning) unchanged, can serve countless purposes.

Davidson (1969/1984, p.108)

¹ For a version of this view, see Loar (1972).

A theory which violates semantic innocence, then, runs contrary to some fundamental notions about the nature of language: language operates in the way it does precisely because we associate fixed meanings with particular linguistic expressions; we wouldn't expect the meaning of an expression to change according to the syntactic context in which it appears². So there seem to be good reasons to hold to each of our two principles: a semantic theory which respects both is, *ceteris paribus*, preferable to a theory that ditches one principle to save the other.

3. Neale and the multiple proposition framework

3.1 Multiple propositions and contextual weighting

Neale (1999) expressly sets out to develop a semantic framework which, while respecting both compositionality and semantic innocence, can account for a range of problematic phenomena, in particular the place within a semantic theory of Fregean *colouring* (colouring is, for Frege, the sort of thing that differentiates the meaning of 'and' from that of the truth-conditionally equivalent 'but') and particular types of NP, including definite and demonstrative descriptions and proper names. As I said earlier, I don't intend to examine Neale's account of colouring, except inasmuch as it casts light on the rest of his account: I suspect that much of what he has to say on colouring may be unobjectionable. The extension of his account to NPs, by contrast, seems to me to be largely a matter of sticking square pegs into round holes (although, as I hope will become clear, I'm not claiming, even in this area, that Neale has nothing of interest to offer).

At the heart of Neale's account is the idea that one utterance may express more than one proposition³. Take, for instance, the sentence in (10):

² Neale talks of the reference of expressions remaining constant across different *linguistic* contexts, but he clearly intends the principle of semantic innocence to apply only to linguistic contexts *syntactically* defined. In (i) and (ii), for instance:

- (i) The king of France was executed during the French Revolution
- (ii) The king of France built Versailles in the seventeenth century

the definite description 'the king of France' refers to different individuals, and, since all we have to go on is the surrounding linguistic context, it must be this which effects the change in reference. Neale, however, would not want to see this as a threat to semantic innocence.

³ Neale is not the only theorist to make use of multiple-propositions; for another recent account of this kind, see Bach (1999).

(10) Napoleon, who recognised the danger to his right flank, personally led his guards against the enemy position.

Frege (1892), who originally discussed the example, suggests that (10) may express a conjunction of two propositions, those in (11) and (12):

(11) Napoleon personally led his guards against the enemy position.

(12) Napoleon recognised the danger to his right flank.

For Neale, this idea can, with one alteration, shed light on some famously thorny semantic issues. The alteration that Neale proposes is that, rather than expressing a conjunction of propositions, a sentence such as (10) expresses a *sequence* of propositions. In other words, (10) no longer expresses a single proposition, with a single set of truth conditions; rather it expresses two discrete propositions, with two discrete sets of truth conditions. For Neale, the shift from conjunction to sequence is necessary if we are to account for intuitions on the truth or falsity of sentences such as (10) assessed against different circumstances. Consider, for instance, the circumstance that Napoleon did lead his guards against the enemy position, but did not realise the danger to his right flank. Would an utterance of (10) be true or false in this circumstance? For Frege it must be false, since it expresses a conjunction of propositions one of which is false. But does that tie in with our intuitions? Would it not be more natural to claim that something false has indeed been said, but that something true has also been said? This intuition falls out naturally on a multiple-proposition approach: two separate propositions have been expressed, one true, one false.

Neale also claims for his multiple proposition framework the advantage of being able to account for some notoriously divided intuitions: consider, for instance, the following sentence, adapted from an example originally discussed by Grice (1967a).

(13) Bill is a philosopher, therefore he is brave

For Grice, a speaker uttering (13) will be *saying* that Bill is a philosopher and that Bill is brave but will merely be *indicating* that Bill's being brave follows from his being a philosopher. The intuition underlying this analysis is that an utterance of (13) will be true so long as Bill is brave and Bill is a philosopher, regardless of whether the one follows from the other. But, as Neale points out, this intuition is far from universal: for McCawley (1993), for instance, an utterance of (13) is false if the causal connection between being a philosopher and being brave does not hold. How might we account for

two such different intuitions? For Neale, the multiple-proposition framework gives rise to a natural explanation: any utterance of (13) will express three propositions, that Bill is a philosopher, that he's brave and that being brave follows from being a philosopher. In different contexts, however, these three propositions will receive different weighting: in some contexts the main force of the speaker's communicative intent will lie with the first two propositions, in others it will lie with the third. The difference in intuitions can, then, be put down to the different contexts of utterance imagined: those who, with Grice, hold that a failure of the causal connection makes no difference to the truth value of an utterance of (13) are imagining a context of the former type, those who go the other way with McCawley are imagining a context of the latter type.

For Neale, the combination of a multiple-proposition framework with the notion of contextual weighting can be used to account for many of the most contentious issues in the current philosophy of language. He initially motivates the framework with examples such as (10) and (13), for which it does indeed seem potentially fruitful. However, he extends the framework to an analysis of the interpretation of NPs, and it is to his claims in this area that I now want to turn.

3.2 Multiple propositions and the interpretation of NPs

Neale claims that only within a multiple-proposition framework is it possible to account satisfactorily for the semantic behaviour of many types of NP. Consider, for instance, the following sentence:

(14) That man by the door is Dave's boss

What contribution does the demonstrative description 'that man by the door' make to the truth conditions of an utterance of (14)? A consideration of counterfactual circumstances suggests that 'that man by the door' is referential: the proposition expressed by an utterance of (14) would seem to be true in any circumstance in which the person indicated, call him a, is, in fact, Dave's boss, regardless of where he happens to be standing in that counterfactual situation. The property encoded by the nominal, the property of being a man by the door, seems to make no contribution to truth conditional content, simply serving to guide the hearer towards the speaker's intended referent. But now consider (15):

(15) That man standing by the door should close it

Here something rather different is going on: an element within the nominal seems to be entering into a semantic relation with another element in the sentence: ‘the door’ seems to be serving as an antecedent for ‘it’. And, on an account that sees demonstrative descriptions as straightforwardly referential, that, it has been argued, is hard to explain. Maybe the descriptive content of the nominal does contribute to propositional content after all⁴.

Neale’s solution to this apparent conflict is to claim that utterances with demonstrative descriptions in subject position express two propositions: an utterance of (14), for instance will express firstly a descriptive proposition:

(16) [the *x*: the speaker is indicating *x* and *x* is a man standing by the door] *x* is Dave’s boss

and then, parasitically, a singular proposition:

(17) *a* is Dave’s boss

This is so, Neale claims, because demonstrative descriptions have the meaning they do: the meaning of a demonstrative description is a set of instructions telling the hearer to build just these two propositions.

If demonstrative descriptions give rise to two propositions, how about definite descriptions? How do the attributive and referential uses of definite descriptions, in the sense of Donnellan (1966), fit into this picture? For Neale, whereas the meaning of a demonstrative description instructs the hearer that he *must* construct the second, singular proposition, the meaning of a definite description tells him that he *may* have to do this; it opens up the possibility, in other words, that it may be the speaker’s intention for him to build such a proposition. In Donnellan’s attributive uses, therefore, there will simply be one proposition expressed, which will be straightforwardly Russellian, whereas in referential uses there will be two propositions expressed, one Russellian, one singular⁵.

⁴ I only use this kind of anaphoric evidence for illustration; there is much other evidence that has been taken to suggest that the descriptive content of demonstrative descriptions contributes to truth conditions (see Lepore & Ludwig (1998) for some of the more convincing evidence).

⁵ As Neale himself seems to recognise, placing the contribution of referential uses within the truth conditional content of an utterance is a radical departure from his earlier position (see, for instance, Neale (1990)).

This multiple-proposition account of the semantics of definite descriptions provides, for Neale, an explanation of the notoriously split intuitions over cases of misdescription. Is what a speaker uttering (18) has said:

(18) The man drinking a martini lives next door to me

true or false in the circumstance that the man whom the speaker intends to talk about, although indeed the speaker's neighbour, is not in fact drinking a martini (and no-one else happens to be drinking a martini)? For some it is straightforwardly true, for others straightforwardly false. Neale takes this split itself to be a significant datum for a semantic theory, and one that falls out naturally from his account: in a context in which the speaker does intend to express a singular proposition with an utterance of (18), she must also express a Russellian proposition. Whether the utterance is taken to be true or false in the circumstance imagined above simply depends on which of these two propositions is taken to receive greater contextual weight.

Finally, Neale extends the multiple-proposition analysis to proper names. Consider again examples (6) to (9) (repeated here as (19) to (22)):

- (19) Cary Grant was born in London
- (20) Archie Leach was born in London
- (21) Mary believes that Cary Grant was born in London
- (22) Mary believes that Archie Leach was born in London

The problem, you'll remember, is that, although 'Cary Grant' and 'Archie Leach' seem straightforwardly referential in (19) and (20), in (21) and (22) they seem to be behaving descriptively: if they're not substitutable, that must mean that something other than their referents is getting into the truth conditions. For Neale, the multiple-proposition framework offers a way out of this difficulty: in both cases (i.e. in (19) and (20) and in (21) and (22)) two propositions are expressed, one descriptive and one referential. So, abstracting away from certain details of Neale's account, an utterance of (19) would express the propositions in (23) and (24):

- (23) [the x : x is called 'Cary Grant'] x was born in London
- (24) a was born in London

But if the *same* two propositions are expressed by 'Cary Grant was born in London' regardless of whether it is embedded in a belief context or not, why do such *different*

things seem to be going on in the two cases? Again, contextual weighting provides Neale with his answer: in a straightforward utterance of (19), Neale claims, it will standardly be the proposition in (24) that receives greater contextual weight, whereas in an utterance of (21), it will be (23). And why might this be so? Although Neale does not address this question directly for proper names, he makes some suggestions in his discussion of demonstrative descriptions which are presumably intended to carry over to the other types of NP discussed. The weighting bias towards general propositions in epistemic contexts and towards singular propositions in modal contexts is a product, Neale suggests, of the more general distinction between epistemology and metaphysics:

When we investigate our thoughts about things we are interested in the properties or features that we use to identify them and the concepts under which we take them to fall; when we investigate the nature of things themselves, we are interested in the things themselves and the properties they actually, necessarily, and accidentally possess.

Neale (1999, p. 67)

3.3 Compositionality, innocence and multiple propositions

Among the claims, then, that Neale makes for his multiple-proposition framework are the following:

- (i) It can be used to account for a range of problematic semantic phenomena, including the interpretation of demonstrative descriptions, definite descriptions and proper names;
- (ii) It respects compositionality
- (iii) It respects semantic innocence

In the last section we saw, at least in outline, what claim (i) amounts to, but how about claims (ii) and (iii)? So far I've been more or less assuming that the content of these claims is unproblematic, that we know what it means for a semantic theory to respect compositionality and semantic innocence. But an attempt to spell these claims out reveals that things are not quite so straightforward. As Neale himself points out, it is now standard in the philosophical/semantic literature to distinguish between two distinct levels of meaning: sentence meaning and propositional/truth-conditional content. This distinction is motivated, if by nothing else, by the gap between linguistic meaning and propositional content which characterises indexicals.

This distinction between sentence meaning and propositional content immediately gives rise to (at least) two different versions of the principles of compositionality and semantic innocence:

Compositionality_{LM}: The linguistic meaning of a sentence is determined by the linguistic meanings of its component parts and their syntactic combination

Compositionality_{TC}: The proposition expressed by (an utterance of) a sentence is determined by the linguistic meanings of its component parts and their syntactic combination

Semantic innocence_{LM}: A particular expression has the same linguistic meaning regardless of the linguistic context in which it appears

Semantic innocence_{TC}: A particular expression corresponds to the same propositional content regardless of the linguistic context in which it appears

The question then arises: which versions of these two principles does Neale's multiple-proposition framework respect? I want here to focus on compositionality, although much the same points can be made about semantic innocence, which I'll return to later. Neale appears to believe that his framework respects compositionality_{LM}:

The linguistic meaning of 'I am here now' is determined by, and only by, the meanings of the lexical items of which it is composed and the syntactic structure that here holds them together.

Neale (1999, p. 43)

But, as he goes on to point out, the notion of sentence meaning is, on this picture, 'epiphenomenal' (Neale 1999, p. 43): the instructions that constitute the linguistic meanings of particular lexical items project, in line with their syntactic combination, not onto anything like sentence meaning, but directly onto propositions; sentence meaning simply plays no theoretical role in the multiple-proposition framework. Neale cannot, then, be claiming merely that sentence meaning is compositional; he cannot, in other words, be intending his theory only to respect compositionality_{LM}.

So is he claiming that the multiple-proposition framework respects compositionality_{TC}? It would seem that he's not claiming this either, given the intimate role which he accepts is played by pragmatic inference in the retrieval of propositional

content. He considers, for instance, examples originally discussed by Searle (1980), Sperber & Wilson (1986/95) and Carston (1988):

(25) I have cut the grass

(26) I have cut the cake

(27) I have had breakfast

(28) I have been to Tibet

(29) She gave him her key and he opened the door

(30) Mr Jones has been insulted and he's going to resign

Searle points out that 'cut' makes different contributions to truth-conditional content in (25) and (26): the speaker of (26), for instance, will not be speaking truly if she has trimmed the top of the cake with a lawn mower. Similarly, Sperber & Wilson show that the past tense marker gives rise to truth-conditionally distinct interpretations in (27) and (28): an utterance of (27) will be true iff the speaker has had breakfast *this morning*, whereas an utterance of (28) will be true so long as the speaker has been to Tibet at any time in the past. Finally, Carston shows that 'and' contributes more than simple logical conjunction to the propositions expressed by utterances of (29) and (30): an utterance of (29) will be true iff she gave him her key and *then* he opened the door, while (30) will be true iff Mr Jones is going to resign *as a result of* the insult. Of examples such as (25) to (30), Neale says:

The morals that are very rightly drawn from such examples are (i) that linguistic meaning radically underdetermines the proposition expressed, and (ii) that the same sorts of principles that play a role in theories of pragmatic implications ... and nonliteral meaning ought to have an equally important role in theories that purport to characterise the proposition, or propositions straightforwardly expressed by an utterance.

Neale (1999, p. 44)

This viewpoint is clearly incompatible with compositionality_{TC}: the heart of compositionality_{TC} is precisely that linguistic meaning does not underdetermine the proposition expressed.

What, then, does Neale's claim that the multiple-proposition framework respects compositionality amount to? A clue is given by the definition he proposes for

semantics: ‘the study of propositions whose *general form and character* is determined by word meaning and syntax’ (Neale 1999, p. 36). Neale’s claim is not that his framework respects either compositionality_{LM} or compositionality_{TC}, but rather a third formulation of the principle of compositionality, which I’ll call compositionality_{FC}:

Compositionality_{FC}: the form and character of the proposition(s) expressed by a sentence are determined by the linguistic meanings of its component parts and their syntactic combination

So what does this claim amount to? Among other things, it amounts to positing an extra level of semantic representation lying between sentence meaning and truth-conditional content, the level of propositional form and character. But what constitutes propositional form and character? Neale does not, himself, have anything explicit to say on this, but at the heart of his notion lie, I believe, two separate issues: on the one hand, propositional form must reflect syntactic structure, and, on the other, the linguistic meaning of a sentence must determine whether the propositions expressed are singular, general or both. But why would Neale want to posit an extra level of representation, given that a theory which posits fewer levels of representation must be preferred, on grounds of parsimony, to one that posits more levels? The answer, given the above discussion, is that he needs somewhere for the principle of compositionality to hold. We already know why he cannot claim that compositionality_{TC} holds: there’s just too much evidence of pragmatic input to propositional content. But why should he not be satisfied with compositionality_{LM}? I suspect the answer will look something like this: if your semantic theory respects only compositionality_{LM}, then it is no longer able to explicate the notion of meaning via the notion of truth, i.e. it abandons the fundamental tenet of all semantics in the truth-conditional tradition. The reason for this is that, if your compositional semantic theory only respects compositionality_{LM}, it will not deliver anything that is truth-bearing. I will eventually suggest that this is not the problem that it may, at first, seem.

Much the same observations can, I think, be made about the place of semantic innocence in the multiple-proposition framework, although I shall not be focusing on this. The claim that Neale is making is that a given expression will make the same contribution to the form and character of the proposition expressed by an utterance in which it is contained regardless of the linguistic environment. His account is intended, in other words, to respect semantic innocence_{FC}:

Semantic innocence_{FC}: A particular expression corresponds to propositional content of the same form and character regardless of the linguistic context in which it appears.

3.4 Compositionality_{FC} and multiple propositions

I think it's clear that Neale does intend his multiple-proposition framework to be taken to respect compositionality_{FC}. But does it do so? There is one obvious aspect of the multiple-proposition story which seems, at the least, to stretch compositionality_{FC}: Neale's analysis of definite descriptions. Remember that, for Neale, an utterance containing a description attributively used in subject position will express just one proposition, along Russellian lines, whereas an utterance containing a description referentially used will express two propositions, one Russellian and one singular. It seems, then, that, in different contexts, utterances of one and the same description sentence can express very different propositional structures; the form and character of the singular proposition expressed in referential uses is nowhere reflected in the proposition expressed in attributive uses. If, as suggested above, the encoded meaning of a sentence should, on Neale's picture, determine whether the propositions expressed by an utterance of that sentence are singular, general or both, this inbuilt optionality seems at least unfortunate (apart from being *ad hoc*).

But there is a rather more revealing threat to compositionality_{FC} within Neale's account. In his discussion of Fregean colouring, Neale touches upon the semantics and pragmatics of 'and' (although 'and' seems like a somewhat unlikely candidate for inclusion under the colouring banner). His view seems to be essentially that of Carston (1988), mentioned earlier, on which the temporal and causal (and other) elements standardly associated with 'and' are part of propositional content, although not semantically encoded. On this story, the propositions expressed by utterances of (31) to (33):

- (31) London is in England and Paris is in France
- (32) Mary took off her socks and jumped into bed
- (33) John ate poisoned mushrooms and died

will look like (34) to (36):

- (34) London is in England & Paris is in France
- (35) Mary took off her socks & then jumped into bed

(36) John ate poisoned mushrooms & then, as a result of eating poisoned mushrooms, died

But crucially, on Carston's story, these truth-conditional differences are not down to lexical ambiguity. Where, then, do they come from? Carston (2000) provides strong arguments against the idea that they arise from anything like the assigning of reference to covert indexical elements. One of the characteristics of overt indexicals, for instance, is that a failure to assign reference results in a failure to retrieve full propositional content. Consider the sentence in (37):

(37) She lives in Sacramento

A hearer who fails to assign reference to the indexical 'she' will fail to retrieve the proposition expressed by an utterance of (37). But the same is clearly not the case when it comes to the retrieval of the temporal and causal elements in the interpretations of (32) and (33), as can be seen from (31): in (31) 'and' seems to be straightforwardly truth-functional; if there is a covert indexical element at LF, then it has not been assigned reference in the interpretation of (31), yet (31) appears to express an entirely determinate proposition.

If the temporal and causal elements in the interpretations of (32) and (33) are not represented at any syntactic level of representation, where do they come from? The answer must be that they are the product of pragmatic inference. But, if this is so, and, by accepting Carston's story, Neale appears to be accepting that it is indeed so, then that poses a serious threat to compositionality_{FC}: part of the *form* of the propositional content of utterances of (32) and (33) is determined by non-compositional factors. Accepting the idea that the form (and character) of propositional content can be influenced by non-compositional, pragmatic factors opens the way to an alternative framework, one that respects compositionality and semantic innocence to the degree that they need respecting, and which can account for the same semantic phenomena as Neale's account with greater theoretical parsimony.

4. NPs, underdeterminacy and a return to single propositions

4.1 Underdeterminacy and the interpretation of NPs

If we accept, as I think we must, that underdeterminacy is a pervasive feature of natural language use, i.e. that the linguistic meaning of a sentence standardly underdetermines

not just the detail, but the form and character of the proposition expressed by an utterance of that sentence, then the goalposts for an analysis of the semantics and pragmatics of NPs move somewhat. We need no longer claim that the varying interpretations available for particular NPs must be reflected in their linguistic meaning: so long as there is an inferential path available from a unitary semantics to each of those interpretations, we can let pragmatics do the job for us. The idea that I want to put forward is that the NP types discussed by Neale, far from encoding a multi-level set of instructions, have a relatively minimal semantically encoded meaning. This minimal semantics, combined with context and pragmatic principles can then lead to a variety of different interpretations.

It's not my aim here to provide detailed analyses of the semantics and pragmatics of these expressions but, as an illustration of how this story might go, let's first consider the interpretation of definite descriptions⁶. Take, for instance, the following sentence:

(38) The Ferrari driver has an unfair advantage

Compare two contexts in which this sentence might be uttered: in one context speaker and hearer are discussing a forthcoming Grand Prix and, although they don't know who is going to be driving any car, they do know that the Ferrari is the best car in the race. In the other context, speaker and hearer are in the pit lane before the race and overhear a man in red overalls, who they take to be the Ferrari driver, scheming with the race organiser. In the first context the speaker uses the description 'the Ferrari driver' attributively, in Donnellan's sense, and in the second she uses it referentially. On Neale's story, in the first context the hearer follows an instruction to build a Russellian proposition and then follows another instruction to decide whether, given contextual factors, the speaker intends him to build a second, singular proposition. In this case he will decide, on the basis of his inferential abilities, that no such proposition is intended and stop. In the second context, by contrast, he will decide that the speaker does intend him to build a second proposition and will therefore do so.

On the analysis I'm proposing, things look rather different. In very rough terms, the semantically encoded meaning of the definite description will tell the hearer that he must find a conceptual representation of a Ferrari driver. After that, everything's down to pragmatics. So how might the pragmatic story go? I propose to outline a story which makes use of the relevance theoretic framework of Sperber & Wilson (1986/95). In this

⁶ For a more detailed analysis of the interpretation of definite descriptions within this framework, see Powell (1999).

framework, relevance is taken to be a property of stimuli such that the more contextual effects a stimulus offers, the greater its relevance, and the less processing effort required to achieve these effects, the greater the relevance. According to Sperber & Wilson's Communicative Principle of Relevance, every ostensive stimulus communicates a presumption of its own optimal relevance:

Presumption of optimal relevance

- (a) The ostensive stimulus is relevant enough for it to be worth the addressee's effort to process it.
- (b) The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences

Sperber & Wilson (1986/95, p. 270)

The presumption of optimal relevance licenses a particular interpretation strategy: what a hearer will do is retrieve the most accessible interpretation first, check whether it satisfies his expectation of relevance and, if it does, accept it. If, however, it doesn't meet his expectation of relevance, then the hearer will start the procedure again by retrieving the next most accessible interpretation, and so on. How, then, will this strategy apply to the interpretation of 'the Ferrari driver' in (38)? In context one, the hearer will hunt around for the most accessible Ferrari-driver individual concept in the context, i.e. the most accessible individual concept associated with the information 'x is a Ferrari driver'. He may well have a number of different such concepts: one corresponding to the Ferrari driver who ran over his foot last week, another corresponding to Michael Schumacher and so on. However, given the context in which he's processing the utterance, a conversation about a particular race in which a Ferrari is involved, one Ferrari-driver individual concept is likely to be highly accessible: the non-referential concept of *the Ferrari driver in this race*. The hearer will then try this interpretation out for relevance and, finding that it produces sufficient contextual effects for no undue processing effort, is justified in accepting it as the optimally relevant interpretation.

How about context two? The hearer will again start with the same semantic encoding: find a Ferrari-driver individual concept. But he may well, in this case, find that two potential Ferrari-driver concepts are more-or-less equally accessible: one non-referential concept of *the Ferrari driver in this race*, and one referential concept of *that man over there*. Which should he choose? Referential concepts standardly serve to group together a range of distinct but co-referential modes of presentation, each of which will offer a particular contextual effect. So a hearer retrieving a referential

interpretation of (38) may also be able to infer that the man who looks like *that* has an unfair advantage; the man who speaks with a French accent has an unfair advantage, and so on. Given that the referential interpretation offers these extra contextual effects, it is this interpretation that satisfies the hearer's expectation of relevance: it is the most relevant interpretation compatible with what the hearer perceives to be the speaker's abilities and preferences. He should thus accept it as optimally relevant.

And what of demonstrative descriptions? How might we account for examples such as (14) and (15), repeated here as (39) and (40)?:

- (39) That man by the door is Dave's boss
 (40) That man standing by the door should close it

Remember that in (39), the demonstrative description 'that man by the door' appears to be directly referential, i.e. the nominal 'man by the door' appears to make no contribution to truth conditions, whereas in (40), an element within the nominal, 'the door', appears to be acting as an antecedent for the anaphoric 'it'. I suggest that, if we take underdeterminacy seriously, an elegant solution presents itself: since 'that man by the door' appears to be directly referential in (39), let's assume that it is also directly referential in (40); after all, the consideration of counterfactual circumstances that argued for the rigidity of the demonstrative description in (39) should lead us to the same conclusion in (40). But how, then, can 'the door' act as antecedent for 'it'? There is only any problem with this if we assume that anaphoric relations must hold at the level of propositional content. But why make this assumption, when it seems clear that the gap between sentence meaning and propositional content is a wide one? If, instead, we allow for the possibility that anaphoric relations can be resolved at an earlier stage of interpretation, the difficulty evaporates. The story might go something like this: since the speaker of (40) has used the demonstrative description 'that man standing by the door' to refer to her intended referent, it will be mutually manifest to speaker and hearer that the speaker believes her intended referent is standing by a door (why else would she use this particular tool of reference?). Given this, when the hearer comes to interpret 'it', one highly accessible interpretation will be a conceptual representation of the door a is standing by (where a is the speaker's intended referent). This does not entail, however, that a specification of the truth-conditional contribution of 'the man standing by the door' should make any mention of the door; it may simply contribute a non truth-conditional mode of presentation of a to propositional content.

This is, I admit, a very sketchy outline, but I hope I have made clear the overall shape of the account I favour: single utterances give rise not to many propositions, with one

being weighted, i.e. doing all the work, but only to whichever proposition (or, in some cases, maybe propositions) is/are doing the work. And that is made possible via a combination of minimal semantic encoding, context and pragmatic inference. Although I don't intend to go into the details here, this general framework can be extended to the other NP types discussed by Neale; each simply encodes a constraint on the search for conceptual representations.

4.2 The underdeterminacy account and theoretical parsimony

There are, it seems to me, good reasons to prefer the account sketched above to Neale's multiple-proposition story. In two fundamental respects it is more parsimonious than Neale's story: it posits fewer propositions expressed and fewer levels of representation. But can we really get away with this pared down theoretical machinery? Yes, we can. After all, we have at our disposal all the resources available to Neale. Once we've accepted, as Neale himself seems drawn to do in his analysis of 'and', that pragmatic inference can affect the form and character of the proposition expressed, then we can adopt the machinery that Neale uses to effect contextual weighting, i.e. the standard machinery of pragmatic inference, and use it instead to determine which proposition(s) is/are expressed. This, in fact, seems the natural way to go, given that, in most of the examples Neale discusses, the contextually weighted proposition is not doing *most* of the work, it is doing *all* of the work. Consider an example that Neale uses in his analysis of proper names:

(41) Vivlos is in Greece

On Neale's analysis, an utterance of (41) will express two propositions:

(42) [the x : x is a place and x is called 'Vivlos'] (x is in Greece)

(43) a is in Greece

But under what circumstances will what is expressed by an utterance of (41) be true? Consider a counterfactual circumstance in which, by some strange quirk of fate, the town that is, in the actual world, called 'Vivlos' has been translated to Turkey, while another town in Greece is named 'Vivlos'. Will the proposition expressed by an utterance of (41) in the actual world be true or false in this counterfactual circumstance? I think it will clearly be false: in that circumstance, Vivlos is in Turkey. Of course, an utterance of (41) in the counterfactual circumstance will be true, but that is an entirely

different matter. So, for a (standard) utterance of (41), all the work will be done by a singular proposition. The general proposition, if one is expressed, seems entirely semantically inert in this context⁷.

This is, I think, no more than Neale is accepting with the notion of contextual weighting: in some contexts it's one proposition that does all the work, in other contexts it's the other. But why, then, would Neale not accept an analysis on which only one proposition is expressed? I have already hinted at one fundamental answer: on an underdeterminacy account, i.e. on an account in which the form and character of the proposition expressed by an utterance can be influenced by pragmatic inference, meaning can no longer be directly approached with truth, and it has long been an accepted axiom of truth-conditional semantics that the solid notion of truth is the only useful tool we have for getting a grip on the slippery notion of meaning. I shall, however, again delay saying more about this.

The reduction in the number of propositions expressed is not the only way in which the account I have sketched is more parsimonious than Neale's. As I have argued above, if we are to take Neale's claim to respect compositionality seriously, it can only be at a level intermediate between linguistic meaning and full propositional content. On the underdeterminacy story, there is no such level of representation: there is just sentence meaning and propositional content, with pragmatic inference leading from one to the other. But, given that compositionality cannot hold at the level of full propositional content, doesn't the underdeterminacy account fall at the very hurdle we initially set out to jump? It may account for the interpretation of NPs, but surely it does so within a framework that respects neither compositionality nor semantic innocence.

⁷ Of course, as Neale shows, a general proposition does seem to carry the communicative burden in other contexts, most notably in belief contexts. Again, I don't intend to propose a detailed analysis of belief contexts; I just want to point out that the pragmatic principles which, for Neale, determine that a general proposition is contextually weighted in belief contexts, can, given the sort of picture painted above, determine that only a general proposition is expressed in such contexts.

4.3 Underdeterminacy, compositionality and semantic innocence

The account I've sketched does violate compositionality_{TC}, compositionality_{FC}, semantic innocence_{TC} and semantic innocence_{FC}. However it respects the LM version of both principles: the linguistic meaning of a sentence is determined by the linguistic meanings of its component parts and their syntactic combination; and a given expression will make the same contribution to sentence meaning, regardless of the surrounding linguistic context. But the question arises: is respecting compositionality_{LM} and semantic innocence_{LM} enough? To find an answer, we must return to the reasons we wanted to build the two principles into our semantic theory in the first place. We wanted our theory to be compositional for two main reasons: compositionality ties in with the intuition that a particular lexical item will make the same contribution to the meanings of the different sentences in which it appears; and only a compositional theory could account for the ability of native speakers to understand an infinite number of novel sentences of their language.

Will compositionality_{LM} satisfy these requirements? On the account sketched above, lexical items do, indeed, make a systematic contribution to sentence meaning, although not necessarily to propositional content. But would we want them to do so? Consider the following two sentences, the latter a standard example of metonymy:

- (44) The mushroom omelette tastes delicious
- (45) The mushroom omelette wants his bill

While we are certainly going to want to say that the description 'the mushroom omelette' has the same *meaning* in both contexts, this intuition must operate at the level of linguistic meaning rather than truth-conditional content: we certainly don't want to say that 'the mushroom omelette' makes the same contribution to truth conditions in both contexts. As for the ability to understand infinite sentences of our native language, compositionality_{TC/FC} isn't necessary to do this job either; compositionality_{LM} will do it just as well. There can be no doubt that we all have complex pragmatic abilities; the retrieval of implicatures, to mention just one phenomenon, demonstrates this. On the picture I have painted the retrieval of propositional content is simply a matter of putting a compositional linguistic meaning together with that complex pragmatic ability and context. So long as sentence meaning is compositional, our ability to understand an infinite number of new sentences is no mystery.

And semantic innocence? Let's return to Davidson's formulation of the intuition underlying semantic innocence:

Language is the instrument it is because the same expression, with semantic features (meaning) unchanged, can serve countless purposes.

Davidson (1969/1984, p.108)

Does semantic innocence_{LM} satisfy this intuition? If anything, the underdeterminacy story takes this idea rather more seriously than does the multiple-proposition account. Unambiguous expressions have the same encoded linguistic meaning regardless of the context in which they appear, defined linguistically or otherwise. Given their basic encoded meaning, however, the use they can be put to is limited only by context and the nature of pragmatic inference. This account, therefore, respects on the one hand, semantic innocence, and, on the other, the efficiency of language.

5 Conclusion

My aim in this paper has been to look at the place of compositionality and semantic innocence in theories of meaning by looking at one account, that of Neale(1999), which claims to build both principles into one unified semantic theory. Although Neale's account makes much of pragmatic input to propositional content, it is, at base level, embedded in a traditional truth-conditional conception of the relation between language and propositional content, a conception on which, since there is a more-or-less direct mapping between language and truth conditions, truth is taken as the basic tool for analysing meaning. Approaching these questions from a perspective grounded in the work of Carston (forthcoming) and Sperber & Wilson (1986/95), I have argued that, once one appreciates the extent to which linguistic meaning standardly underdetermines propositional content, a more parsimonious story is available, a story which can not only account for the phenomena which Neale tackles, but also respects both compositionality and semantic innocence, once they are properly understood within the underdeterminacy framework.

But if we cannot use truth as a tool to analyse meaning, how are we to pin down this notoriously slippery notion? Is the account I have sketched not abandoning the only hope we have of understanding meaning? Not at all: it is simply taking seriously the role played by mind in mediating the relation between language and the world. The path from language to conceptual representations is not as straight as the truth-conditionalist would have it, but that does not mean that it is not mappable.

References

- Bach, K. (1999) The myth of conventional implicature. *Linguistics and Philosophy* 22(4): 327-366.
- Barwise, J. & J. Perry (1981) Semantic innocence and uncompromising situations. *Midwest Studies in Philosophy* 6, 387-403.
- Barwise, J. & J. Perry. 1983. *Situations and Attitudes*. Cambridge, Mass.: MIT Press. Second edition 1999. Stanford: CSLI (references are to this edition).
- Carston, R. (1988) Implicature, explicature, and truth-theoretic semantics. In Kempson, R. (ed.). *Mental Representations*. 155-181. Cambridge: CUP.
- Carston, R. (2000) Explicature and semantics. *UCL Working Papers in Linguistics* 12.
- Carston, R. (forthcoming) *Thoughts and Utterances: the Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Davidson, D. (1969) On saying that. *Synthese* 19: 130-146. Reprinted in Davidson, D. (1984) *Inquiries into Truth and Interpretation*. Oxford: OUP (references are to this edition).
- Donnellan, K. (1966) Reference and Definite Description. *The Philosophical Review* 75: 281-304.
- Frege, G. (1892) Über Sinn und Bedeutung. Trans. as 'On Sense and Reference' in Geach, P. T. & M. Black (eds.). *Translations from the Philosophical Writings of Gottlob Frege*. 56-78. Oxford: Blackwell.
- Grice, H. P. (1967a) Utterer's meaning, sentence-meaning, and word meaning. *Foundations of Language* 4: 225-242. Reprinted in Grice (1989). 86-116. (references are to this edition).
- Grice, H. P. (1967b) Logic and conversation. In Cole, P. & J. Morgan (eds.). *Syntax and Semantics*. Vol. 3. 41-58. New York: Academic Press. Reprinted in Grice (1989). 22-40. (references are to this edition).
- Grice, H. P. (1989) *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Kripke, S. 1980. *Naming and Necessity*. Oxford: Blackwell.
- Lepore, E. & K. Ludwig (1998) The Semantics and pragmatics of complex demonstratives. *Mind* 109: 199-240.
- Loar, B. (1972) Reference and propositional attitudes. *Philosophical Review* 81: 43-62.
- McCawley, J. (1993) *Everything Linguists Always Wanted to Know about Logic*. Chicago: Chicago University Press.
- Neale, S. (1990) *Descriptions*. Cambridge: MIT Press.
- Neale, S. (1999) Coloring and composition. In Murasugi, K. & R. Stainton (eds.). *Philosophy and Linguistics*. 35-82. Boulder: Westview.
- Powell, G. (1999) The referential-attributive distinction: A cognitive account. *UCL Working Papers in Linguistics* 11: 101-125.
- Recanati, F. (1993) *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Searle, J. (1980) The background of meaning. In Searle, J. et al. (eds.). *Speech Act Theory and Pragmatics*. 221-232. Dordrecht: Reidel.
- Sperber, D. & D. Wilson (1986) *Relevance: Communication and Cognition*. Oxford: Blackwell. Second edition 1995. (references are to this edition).