

*The referential-attributive distinction - a cognitive account**

GEORGE POWELL

Abstract

In this paper I approach the referential-attributive distinction in the interpretation of definite descriptions, originally discussed by Donnellan (1966), from a cognitive perspective grounded in Sperber and Wilson's Relevance Theory. In particular, I argue that definite descriptions encode a procedural semantics which is neutral as between referential and attributive readings, with each reading arising as a result of the differing links that exist between different mental representations and the world, rather than as a result of the differing links between language and mental representations.

1 Introduction

The aim of this paper is to take a new look at an old problem in the philosophy of language. The problem, as commonly formulated, is whether the referential and attributive uses of definite descriptions are semantically significant. What I want to do is twofold: first I want to look at two of the main strands in the debate that this question has given rise to and suggest that some of the apparent differences between these strands may not be as they at first appear, and second I want to sketch a possible account of the interpretation of descriptions, an account from which an answer to the original question (or, rather, a reformulated version of the original question) may fall.

The account I propose approaches the questions surrounding descriptions not directly from a truth-conditional stance, but from within a general theory of communication and cognition, grounded in Sperber and Wilson's relevance theory (see e.g. Sperber and Wilson (1986)). I will, however, have a certain amount to say, towards the end of the paper, on how my account ties in with the truth-conditional issues central to the philosophical debate on descriptions. In particular I shall argue that bringing mind into

* I would like, once more, to thank Robyn Carston for her unending support and patience during the writing of this paper. I would also like to thank Deirdre Wilson for her helpful comments on an earlier version of this paper.

the picture, rather than simply addressing the language-world relationship, clarifies where the key issues lie for the truth-conditionalist: it seems to me that questions on the referential-attributive distinction are more satisfactorily treated at the level of the mind-world relationship than at the level of the language-mind relationship.

Let me start by laying out the central problem. Donnellan (1966) pointed out that there are two apparently different ways in which definite descriptions (expressions typically of the form ‘the F’) can be used. On the one hand they can be used to talk about whoever or whatever satisfies their descriptive content, and on the other they can be used to pick out a particular individual. Take, for instance, the sentence in (1):

(1) The Ferrari driver has an unfair advantage

Now imagine two contexts in which this sentence might be uttered. In the first context, speaker and hearer are discussing a forthcoming motor race; neither of them knows who will be driving any given car but they both know that the Ferrari is going to be allowed to start the race in front of all the other cars. The speaker then utters (1). In this context the speaker would, on Donnellan’s analysis, mean something like (2):

(2) The Ferrari driver, whoever he is, has an unfair advantage

Now consider another situation. This time hearer and speaker are wandering around the pit lane before the race when they overhear a conversation between the race organiser and one of the drivers, Jones, who they take to be the Ferrari driver. The race organiser is telling Jones that he will do all he can to help him win. The speaker then utters (1). In this context Donnellan’s analysis would suggest that the speaker will mean something like (3):

(3) Jones has an unfair advantage.

If, in uttering (1), the speaker means (2), then she is using the description ‘the Ferrari driver’ attributively, in Donnellan’s terms; if she means (3) she is using the description referentially. This, then, in pre-theoretical terms, is the referential-attributive distinction. But what is its significance to an understanding of how we use natural language to convey meaning?

2 The debate on descriptions

2.0 There are certain key questions that have guided the debate on referential and attributive uses of definite descriptions and I would like to lay them out clearly before getting into the debate itself. As I said above, the central question, as commonly formulated, is whether the referential-attributive distinction is semantically significant. But what does this mean? There are two ways that it is commonly cashed out in the literature on descriptions:

- a) Does the referential-attributive distinction equate to a difference in truth conditions? In other words, are two utterances which differ only in that one contains a definite description referentially used and the other contains the same description attributively used truth-conditionally distinct?
- b) Is 'the' lexically ambiguous? (or, sometimes, is the definite description as a whole ambiguous?)

As I hope to make clear in the first half of this paper, many contributions to the referential-attributive debate have assumed a particular relation between these two formulations: they have assumed, in essence, that if you answer 'yes' to (a), you are thereby committed to answering 'yes' to (b), and if you answer 'no' to (a) you are committed to answering 'no' to (b). Part of my aim in this paper is to argue that this assumption is false. Moreover, once it's appreciated that this assumption is false, a great number of the apparently thorny issues which have surrounded descriptions dissolve.

The other question central to the debate on descriptions is, whether they're ambiguous or not, what semantics do we want to give for definite descriptions? Again there are two possible ways of cashing this out, parallel to those for the last question:

- c) What contribution do definite descriptions make to the propositions expressed by utterances in which they appear?
- d) What is the encoded meaning of definite descriptions?

I want to tackle these two questions in much the same way as the two above: again it seems to me that a relation that has been taken to hold between the answers to these two questions does not necessarily hold. It will be my claim that a particular univocal semantics can lead to two different types of proposition expressed.

Those, then, are the central questions that any contribution to the debate on descriptions must aim to answer. I next want to look at some of the answers that have been given to these questions, in particularly drawing out two main strands in the debate. But, since the referential-attributive distinction derives from Donnellan (1966), I will start by taking a brief look at how he himself analyses it. In particular, I want to show that the answers Donnellan gives to these questions foreshadow much of the most recent research in the area, and point in the direction in which I aim to go.

2.1 Donnellan's position

What is Donnellan's answer to question (a)? It seems clear to me (although not to Kripke (1979)) that for Donnellan there is a truth-conditional difference in the proposition expressed between referential and attributive uses. An utterance of a sentence containing an attributively-used definite description in subject position expresses an object-independent proposition: in Donnellan's terms the description occurs *essentially*, in that its descriptive content appears in a specification of the truth conditions of this proposition. An utterance of a sentence containing a referentially-used definite description in subject position, by contrast, expresses an object-dependent proposition: the descriptive content of the description plays no part in a specification of truth conditions, serving merely to guide the hearer towards a particular individual, and then falling out of the picture.

This truth-conditional difference is, in fact, at the heart of Donnellan's account. Donnellan sees the referential-attributive distinction as a challenge to Russell's semantics for definite descriptions (see, for instance, Russell (1905)), on which any utterance containing a definite description has object-independent truth conditions; 'the F is G', for instance, is true iff there is exactly one F and that one is G. Donnellan is prepared to concede that Russell's semantics may provide an adequate account of attributive uses, but his argument is that it simply fails to account for referential uses. Now if the distinction is to stand as a challenge to Russell, it must affect the truth conditions of utterances in which descriptions appear. Donnellan must, therefore, be taking the position that the referential-attributive distinction corresponds to a difference in the truth conditions of the proposition expressed or, in his terms, of the 'statement made'¹.

There is, however, an obvious difficulty in deciding whether different propositions really are expressed on referential and attributive uses: in most contexts, the truth values of the two interpretations will covary. Taking (1) again, if the Ferrari driver, whoever he is, has an unfair advantage and Jones is the Ferrari driver, then Jones has an unfair advantage; and equally, if the Ferrari driver doesn't have an unfair advantage and Jones is the Ferrari driver, then Jones doesn't have an unfair advantage. This difficulty has, it

¹ If you are still not convinced that Donnellan is in direct opposition to Russell, you need only consider his views on the statement made when the existential presupposition of a definite description fails. For Russell, if the existential presupposition fails then the statement made must be false; for Donnellan it may still be true. This is closely linked to the issue of misdescription discussed in the next section.

seems to me, led to a great deal of the confusion that surrounds the discussion of referential and attributive uses.

In an attempt to demonstrate that propositions with different truth conditions really are expressed on referential and attributive uses, Donnellan turns to contexts in which truth values for the different uses do not coincide; in other words, he turns to cases of misdescription. To see how this works, consider (1) again but, this time, imagine that the speaker and hearer have made a mistake: Jones is not, in fact, the Ferrari driver at all, he drives for McLaren. Brown, the real Ferrari driver, far from having an unfair advantage, is being schemed against by the race organisers. What is the truth value of (1) now? For Donnellan, in this context the proposition expressed, or 'statement made', by an utterance of (1) in which the description is referentially used is true: although Jones is not, in fact, the Ferrari driver, the proposition expressed makes no reference to the property of being the Ferrari driver, it is a singular proposition about Jones himself, and, since Jones does have an unfair advantage, it is true. On the other hand, an utterance of (1) in which the description is attributively used would clearly be false in this context: the utterance predicates the property of having an unfair advantage of whoever happens to be the Ferrari driver, and Brown, who is the Ferrari driver, does not have an unfair advantage. Donnellan's argument, then, runs like this: in cases of misdescription, referential and attributive uses can have different truth values in the same context; any two utterances which can have different truth values in the same context must have different truth conditions; therefore referential and attributive uses must be truth-conditionally distinct.

This reliance on cases of misdescription has opened Donnellan to some of the most forceful criticism his account has had to face, in particular that of Kripke (1979). The central problem is that this phenomenon is not peculiar to definite descriptions. It is, for instance, possible, in an appropriate context, to achieve reference to Jones by using the proper name 'Smith'. This should not lead us, so the argument goes, to posit an ambiguity in proper names. Many of those who argue in favour of a Donnellan-type position now accept that, at the very least, the referential/attributive distinction is not best illustrated by recourse to misdescription (see, for instance, Recanati (1993) and Wettstein (1981). Wettstein even goes so far as to agree that, in cases of misdescription, truth conditions will not depend on the intended referent. As I shall discuss later, I am inclined to accept this conclusion.

So Donnellan answers 'yes' to question (a); how about question (b)? As I mentioned earlier, there has been a default assumption in much of the literature on descriptions which has followed Donnellan's paper that answering 'yes' to (a) commits you to answering 'yes' to (b). Donnellan himself, however, makes no such assumption; in a much-quoted passage he says:

[It does not] seem at all attractive to suppose an ambiguity in the meaning of the words; it does not appear to be semantically ambiguous. (Perhaps we could say that the sentence is pragmatically ambiguous: the distinction between roles that the description plays is a function of the speaker's intentions.)

(Donnellan 1966, p. 59)

It's clear that Donnellan does not see an ambiguity in the encoded meaning of 'the'; what is not so clear, however, is what he might mean by 'pragmatically ambiguous'. It seems to me (and also to Recanati (1993)) that an appropriate cashing out of this claim may well point in the right direction.

As for questions (c) and (d), I have already, in laying out Donnellan's answer to (a), pointed towards his view on (c): an attributively used description contributes some kind of Russellian complex to the proposition expressed by an utterance in which it appears, something like the property of being a unique F; a referentially used description simply contributes an individual. On question (d) Donnellan remains silent: the above comment shows him to believe that descriptions are univocal in that they encode only one semantics, but what that semantics may be he does not make clear.

2.2 Yes-theory vs no-theory

On, then, to the two main strands in the debate that followed the publication of Donnellan's paper. I will start with a few general introductory comments. It seems to me that these strands are based on two good insights and one bad assumption: the one bad assumption, as I have already mentioned, is that answering 'yes' to question (a) above commits you to answering 'yes' to question (b); the two good insights are that, on the one hand, descriptions are univocal and, on the other, that the referential-attributive distinction corresponds to a difference in truth conditions. Anyone who relies on the bad assumption, however, can clearly not consistently maintain both the good insights: if you answer 'yes' to both questions (a) and (b), then you are denying the univocality of descriptions, if you answer 'no' to both (a) and (b) then you are denying that the referential-attributive distinction is truth-conditional. And these are just the positions that the two main strands take.

Let's start by looking at the second of these two possibilities: that the answer to both (a) and (b) is 'no', that the referential-attributive distinction is not truth-conditional and that descriptions are not ambiguous. This is, broadly speaking, the position adopted by a group of theorists including Grice (e.g. (1969)), Neale (e.g. (1990)), Kripke and others.

From now on I shall refer to those who advocate versions of this view as no-theorists. For a no-theorist the picture looks something like this²: descriptions encode a Russellian semantics; on the basis of the subsidiary bad assumption that, with slight caveats, the same answer must be given to question (c) as to question (d), the no-theorist therefore holds that the proposition expressed by any utterance containing a description must have Russellian truth conditions; in other words, for the no-theorist, even where ‘the F’ is used referentially, the proposition expressed by ‘the F is G’ will be true iff there is exactly one F and that one is G.

But how, then, can the no-theorist accommodate intuitions about the reality of the referential-attributive distinction? Typically he splits them into two: he separates the general intuition that there is some significant difference between referential and attributive uses from the specific intuition that the truth conditions of an utterance containing a definite description will vary according to whether the description is used referentially or attributively. He then agrees with the first intuition and denies that he has the second intuition. The second intuition is the one that lies behind Donnellan’s claims about misdescription: that an utterance containing a referentially used definite description will be true so long as the predication is true of the intended referent, even although that intended referent does not, in fact, fall within the denotation of the description. This intuition has come under heavy fire from no-theorists; Wiggins (1975), for instance, says:

Donnellan’s [account] ... depends on the, for me, incredible idea that if I say ‘The man drinking champagne is *F*’ and the man I mean, although drinking water, is *F*, then what I say is true.

But if the referential-attributive distinction makes no difference at the level of truth conditions, how does the no-theorist accommodate the first intuition: that there really is some significant distinction to be accounted for here? He turns to the familiar distinction within the Gricean pragmatic tradition between *what is said* and *what is communicated* (or *meant*). The former equates to something like the proposition ‘literally’ expressed by an utterance, the latter includes propositions which are not ‘literally’ expressed but arise as implicatures. So, for instance, to repeat a well-worn Gricean example:

² Clearly there are significant differences between the accounts that I am grouping together here; I believe, however, that there are sufficient similarities for it to be valid for my purposes to treat them as one.

A: I am out of petrol

B: There is a garage round the corner

Grice (1967, p32)

What B *says*, in this example, is just that there is a garage around the corner from where the conversational exchange is taking place; what B *communicates* includes, according to Grice, a proposition to the effect that ‘the garage is, or at least may be open’.

It is this distinction between what is said and what is communicated that the no-theorist uses as his main weapon against a Donnellan-type position. Sure enough, he may concede, a speaker who uses a definite description referentially *conveys* an object-dependent proposition, and that accounts for intuitions about the difference between referential and attributive uses. But that’s no reason to believe that the referential-attributive distinction corresponds to a difference in the truth conditions of the proposition expressed. The object-dependent proposition conveyed by a referential use, so the no-theorist argues, arises not at the level of what is said (i.e. not at the level of the proposition expressed by an utterance), but at the level of what is communicated.

What evidence is there for this claim, beyond the Wiggins-type intuition on cases of misdescription? The most heavily-used weapon in the no-theorist’s armoury is methodological rather than empirical: the no-theorist argues that his account is more parsimonious than the alternative; he need only posit one semantically encoded meaning for ‘the’ and the rest of the work is done by independently motivated pragmatic machinery. Of course any account which argues for a truth-conditional difference between referential and attributive uses must, so long as it’s based on the bad assumption, posit two semantically encoded meanings for ‘the’ if it’s going to maintain that the referential-attributive distinction corresponds to a difference in the truth conditions of the proposition expressed. And that, so the no-theorist argues, violates Modified Occam’s Razor, Grice’s principle that ‘senses are not to be multiplied beyond necessity’ (Grice (1967)). Of course, as Recanati (1993) points out, this argument only carries any weight so long as the bad assumption is maintained; so long, in other words, as you equate pragmatically derived meaning and implicated meaning. This equivalence has come under heavy attack over the last twenty years: Wilson and Sperber (1981, 1986) and Carston (1988, 1998), among many others, have shown that pragmatic inference is intimately involved in the retrieval of propositional content. Once this is appreciated, it becomes possible to argue that a univocal semantics for definite descriptions leads, via inference, to differing truth conditions in different contexts. Although such an account allows for a truth-conditional distinction between referential and attributive uses, it is no less parsimonious than the standard no-theory account, since it posits only one encoded semantics for descriptions, with the rest of the work being

done by the same sort of independently motivated pragmatic machinery that the no-theorist relies on.

The no-theorist is left, then, with the Wiggins-type intuition as the bedrock of his claim that referential interpretations are communicated as implicatures; this intuition is the heart of his answer to (a) and, as I shall argue later, it seems to me that this intuition is not the evidence for the no-theorist's claims that it has been taken to be. The no-theorist has rather more arguments at his disposal when it comes to answering (b), the question of whether 'the' is lexically ambiguous or not (and remember that, for the no-theorist who holds to the bad assumption, a 'no' answer to (b) is every bit as good as a 'no' answer to (a), since the one is supposed to follow from the other). Kripke (1979 fn 28), for instance, shows that in some perfectly intelligible cases it is hard to say whether a description is being used referentially or attributively, an unlikely state of affairs for a true lexical ambiguity.

Rouchota (1994) points out that the referential-attributive distinction far from exhausts the possible uses of definite descriptions. Consider, for instance, sentences (4)-(6):

- (4) The man I live next door to has some very strange habits
- (5) The president of the United States changes every five years
- (6) The koala is a native of Australia

Imagine (4) uttered in a context in which the speaker has a particular person in mind and yet has no intention that, in understanding her utterance, the hearer should think of that particular person. Such a use seems neither truly referential nor truly attributive; Ludlow and Neale (1991) call such uses *specific*. Again, in so-called *functional* uses such as (5), the definite description seems to function neither to pick out a particular individual nor to talk about whichever individual happens to satisfy the description; in this case it seems to be used to talk about the role of president independently of any of the individuals who have filled that role. And finally in (6) the definite description is used generically, as equivalent to the bare plural 'koalas' (with appropriate syntactic tinkering). As Rouchota argues, anyone who wanted to claim that the referential-attributive distinction derives from a lexical ambiguity would have to posit a many-more-than-two-way ambiguity; from the above examples alone, plus the original referential and attributive uses, it would, presumably, have to be at least a five-way ambiguity.

Another argument against the ambiguity thesis produced by many no-theorists goes something like this: referential uses do not just arise for definite descriptions, but for most quantified phrases. Consider sentence (7):

(7) Every student currently in my office did well on the test

A speaker uttering (7) could certainly intend to convey an object-independent quantified proposition. But she could also intend to convey an object-dependent proposition about the particular students in her office. Given that referential and attributive uses can be shown to arise with quantified expressions generally, it becomes a very unattractive option, so the no-theorist argues, to claim that each of these expressions is lexically ambiguous. And for the no-theorist, holding as he does to the bad assumption, that's strong evidence in favour of the claim that the referential-attributive distinction is not truth-conditional.

The no-theorist, then, believes the following: the referential-attributive distinction is non-truth-conditional, and all utterances with a definite description in subject position, however that description is used, have Russellian truth conditions; the referential use of descriptions arises not at the level of what is said but at the level of what is communicated, i.e. it arises as an implicature. And I have pointed towards a few of the stronger arguments that the no-theorist relies on in support of his claim. I next want to turn, briefly, to the other major strand in the debate over definite descriptions: the yes-theoretic position³. At the heart of this position are the claims that the referential-attributive distinction does correspond to a difference in truth conditions and, therefore, on the basis of the bad assumption, that 'the' is lexically ambiguous⁴.

Again I want to look at some of the main arguments used in defence of the yes-theoretic position. The most familiar (see, for instance, Wettstein (1981)) concerns the use of incomplete descriptions, descriptions that fail to pick out any single individual in the world. The argument goes like this: most descriptions, as commonly used in everyday speech, are incomplete; consider the following sentence:

(8) The policeman is pulling cars over

In the appropriate context, for instance, if speaker and hearer are driving along and see a policeman ahead, an utterance of (8) could clearly be used to convey a determinate proposition. It is, however, so the yes-theorist argues, ludicrous to believe that that proposition has Russellian truth conditions; if that were the case it would be true iff there

³ Versions of this position are taken by Wettstein (1981) and Peacocke (1975), among others.

⁴ I suspect that I am doing many of those who might think of themselves as yes-theorists a disservice: their thesis is, in most cases, based exclusively on the first of these two claims, with little or no mention of whether 'the' is ambiguous or not.

were exactly one policeman and that one was pulling cars over. Since there exists more than one policeman, such a proposition would be false, whereas it is easy to picture a situation in which what is conveyed by (8) is true. There have been various attempts by neo-Russellian no-theorists to defend their position against this objection. Generally speaking, there are two possible approaches: on one approach, the incomplete descriptive content of a definite description is supplemented with contextually available information so as to make it uniquely denoting; on the other, the domain of the definite description is constrained, on the basis of context, so as to allow only one interpretation⁵. However, both of these approaches face fundamental difficulties. The key problem for the first is that there is no principled way to decide which uniquely denoting description appears in the proposition expressed, given that there will standardly be more than one available. The second solution is undermined by cases in which there is no satisfactory way to restrict the domain so as to give rise to the right interpretation (for a discussion of the shortcomings of these approaches see Breheny (1999)⁶).

Ramachandran (1996), raises another interesting argument in favour of the yes-theoretic position. He attacks the claim, outlined above, that referential uses arise for other quantified expressions just as they do for definite descriptions, suggesting that all such uses are non-truth-conditional. Consider, Ramachandran asks, the following sentences:

- (9) The table is broken
- (10) There is exactly one table in this room and it is broken

(9) is, obviously enough, an incomplete description sentence; (10) is its supposed contextually-completed Russellian paraphrase. Imagine now that a speaker utters each of (9) and (10) while gesturing at a particular table; in both cases she may succeed in communicating an object-dependent proposition, a proposition which could equally well have been communicated by an utterance of (11):

- (11) That table is broken

⁵ Whether these two approaches really are distinct is not entirely clear. Neale (1990) suggests that 'when all is said and done, the explicit [contextual supplementation] and implicit [domain restriction] methods might turn out to be notational variants of one another' (Neale (1990, p115, fn 48).

⁶ Breheny (1999) proposes an interesting alternative to these approaches to incompleteness, an approach on which the determiner itself is marked as requiring contextual supplementation in line with the descriptive content of the description.

However there is, so Ramachandran argues, a tension in the interpretation of (10) between what is ‘literally asserted’ and what is ‘conveyed’, a tension which is lacking in the interpretation of (9): it is quite natural to interpret (9) as conveying an object-dependent proposition, whereas in the interpretation of (10) there appears to be a gap between what is said and what is communicated. Ramachandran calls this, ‘the prevalence of intuitively correct and wholly proper uses of incomplete descriptions’ (Ramachandran (1996, p.379), the *KosherRef phenomenon*. It seems to me that the *KosherRef* phenomenon, while providing no evidence for a lexical ambiguity in ‘the’ (Ramachandran does not intend it to be so taken), does offer some evidence that the referential use of definite descriptions affects propositional content, rather than just arising at the level of what is communicated, as the referential interpretation of uncontroversially quantified expressions appears to.

These, then, are a couple of arguments in favour of a ‘yes’ answer to question (a). How about question (b)? What arguments does the yes-theorist have to support a lexical ambiguity for ‘the’? Rouchota (1994) points to some of the better-known arguments, of which I will look briefly at just one: the argument from anaphora. The argument looks something like this: indexicals can operate either as genuine referring expressions or as bound variables; consider, then, an example such as (12), taken from Rouchota (1994, p.195):

(12) The girl in the pink suit is one of my students. She is clever

in which the indexical ‘she’ is not syntactically bound by a quantified antecedent; if ‘she’ is not acting as a bound variable then it must be acting as a genuine referring expression; ‘she’ receives its interpretation from the definite description ‘the girl in the pink suit’; therefore ‘the girl in the pink suit’ must be a genuine referring expression; so definite descriptions must be ambiguous as between quantified and referential senses. However, as Neale (1990) points out, free indexicals can receive their interpretation from definite descriptions which are clearly being used attributively; consider (13), taken from Neale (1990, p.175):

(13) The inventor of the wheel was a genius. I suspect (s)he ate fish on a daily basis

Evidence such as this appears to rob the argument from anaphora of all its force. It seems to me that all other such attempts to employ syntactic tools to prise apart two encoded senses for definite descriptions fall at the same hurdle: what holds for referential uses holds just as well for attributive uses.

The two main positions on definite descriptions should now be clear: yes-theorists believe that the referential-attributive distinction is truth-conditional (and maybe that ‘the’ is lexically ambiguous); no-theorists believe that ‘the’ is univocal and that the referential-attributive distinction is non-truth-conditional. I hope, however, that in outlining some of the arguments used on both sides of this debate, I have at least hinted that each camp has better arguments in favour of one of its two claims than in favour of the other. The yes-theorist seems to have some good arguments in favour of the claim that the referential-attributive distinction is truth-conditional; he seems rather less well equipped with arguments to support the view that ‘the’ is lexically ambiguous to the extent that many yes-theorists simply do not address the issue. The no-theorist, on the other hand, has some very good arguments to show that ‘the’ is univocal; whereas his arguments for referential uses arising as implicatures are more or less entirely theory driven: for the no-theorist the proposition expressed by an utterance is essentially defined as the product of semantic decoding (with a certain amount of give and take to allow indexicals and ambiguity into the picture); pragmatic inference occurs, for the no-theorist, only in the retrieval of implicated meaning. Since the only way to get from a univocal semantics to two distinct interpretations is via pragmatic inference, one of the interpretations of descriptions must, for the no-theorist, arise as an implicature⁷. Neale (1990), in particular, seems firmly committed to this assumption as it applies to definite descriptions. In a discussion of the account of Recanati (1989), on which a univocal semantics leads in some contexts to an object-dependent proposition and in other contexts to an object-independent proposition, Neale remarks:

Since two utterly distinct types of proposition may be expressed, I fail to see how a theory with such flexibility can fail to be a theory that is postulating a semantic ambiguity.

Neale (1990, p. 112, fn 36)

Neale seems to be rejecting out of hand the possibility that the gap between semantically encoded meaning and propositional content can, here, be filled by pragmatic inference. This is particularly puzzling in the light of his willingness to accept the proposals of Carston (1988), who argues that the temporal and causal implications associated with certain uses of ‘and’ are, while not the product of a lexical ambiguity, nevertheless part of the meaning explicitly communicated by an utterance (see Neale (1990, p 108, fn 27)). Why Neale should be prepared to accept the one proposal while rejecting the other is unclear.

⁷ This is really no more than a reformulation of the bad assumption.

The difference between these two camps is, then, as much a matter of the theoretical machinery with which they approach the questions surrounding descriptions, as of the substantive claims they are making about the meaning and interpretation of descriptions. Each starts from one good insight into the workings of descriptions and is then forced, by an unwillingness to reject the bad assumption, to deny the other good insight.

3 The semantics and pragmatics of definite descriptions: a proposal

3.1 Preliminaries

Before getting into the details of my account I want to draw out the general considerations which will guide it. As I have already made clear, it seems to me that both of the strands of thought above capture certain key facts about the interpretation of definite descriptions. In particular, the no-theorist is right to claim that definite descriptions are univocal and the yes-theorist is right to claim that the referential attributive distinction is truth-conditional (I am, of course, not the first to make this claim; see, for instance, Recanati (1993), Rouchota (1992, 1994), Bezuidenhout (1997)). By taking this position, I am committing myself to a further claim: that the bad assumption is wrong; that it is entirely consistent to give different answers to questions (a) and (b) (and also to questions (c) and (d)). But this is not the controversial claim that it might once have been. Sperber and Wilson (1981), Carston (1988, 1998) and Recanati (1993), among others, have argued convincingly that the role of pragmatic inferencing in the retrieval of propositional content is much more pervasive than the standard truth-conditional approach allows. Carston (1988), for instance, has shown that, from a relevance theoretic perspective, there are very good reasons to suppose that certain elements of meaning which have traditionally been thought of as implicated, such as the temporal sequencing and causality commonly carried by ‘and’, are best treated as pragmatically inferred parts of the explicature of (or proposition expressed by) an utterance. I do not propose to re-present Carston’s arguments here, merely to accept that they provide very strong evidence for the central claim that inference is involved in the retrieval of propositional content. Once this claim is accepted, the need to give the same answers to (a) and (b) (or (c) and (d)) evaporates. It is now possible to answer ‘yes’ to (a) and ‘no’ to (b): referentially and attributively used definite descriptions may well make different contributions to propositional content, but that is no longer any reason to suppose that ‘the’ is ambiguous; the truth-conditional difference may be the result not of semantic encoding, but of pragmatic processes operating in the retrieval of the proposition expressed. And it is, essentially, a version of this view that I aim to advocate.

I therefore intend to build on the growing understanding, amongst pragmatists and philosophers, of the extent to which linguistically encoded meaning typically underdetermines propositional content. There is also a second burgeoning area of research on which I aim to draw: within the relevance theoretic framework there has been much work conducted over the last few years into a distinction originally proposed by Blakemore (1987) between two fundamentally different kinds of meaning. Blakemore's original insight was, roughly, this: that, since utterance interpretation involves two radically different kinds of process, first the retrieval of a logical form and then the manipulation of that logical form through pragmatic inference, we might well expect to find two distinct kinds of meaning corresponding to these processes. And this, for Blakemore, is just what we do find. On the one hand there is conceptual meaning, which is straightforwardly representational; on the other there is procedural meaning, which lays constraints on the retrieval of implicatures.

Over the last few years the application of this distinction has been steadily extended within the relevance-theoretic framework. In particular, the assumption that procedural meaning is involved only in the retrieval of implicatures has been challenged. Wilson and Sperber (1993), among others, have suggested that indexicals may encode procedures rather than, or as well as, concepts, procedures that are therefore involved in the retrieval of explicatures. It seems to me that the distinction between conceptual and procedural information can cast a useful light on the questions surrounding definite descriptions⁸.

Before going on to the details of my proposal I want to put in place one further piece of cognitive groundwork. I take it as relatively uncontroversial that entertaining thoughts about individuals involves entertaining mental representations of those individuals; to think about something, in other words, you must have a concept of that thing. I propose to call such concepts, concepts which are taken by their holders to represent an individual (or set of individuals)⁹, *individual concepts*. What do individual concepts look like? What must we have in order to be able to think about something in the world? Crucially what we need is uniquely identifying information, information which applies to one and only one thing. In cognitive terms, we need to have information which allows us to distinguish one individual concept from all other individual concepts; no two individual concepts can contain the same information. The next question is, what different kinds of uniquely identifying information might there be? Or, looking at the

⁸ The conceptual-procedural distinction has been applied to the interpretation of descriptions before (see, for instance, Kempson (1986), Breheny (1999)).

⁹ I leave aside, here, questions concerning concepts of fictional entities.

same question from another angle, in what kind of ways can we come across uniquely identifying information? It seems to me that in answering this question, we are going much of the way towards disentangling the issues surrounding descriptions.

Following the work of Russell (see, for instance, Russell (1911)), much research has been devoted to investigating two fundamentally different sorts of route via which we can come upon uniquely identifying information, routes dubbed by Russell *knowledge by acquaintance* and *knowledge by description*. Russell, himself, presents the difference between these two forms of knowledge thus:

I say that I am *acquainted* with an object when I have a direct cognitive relation to that object, i.e. when I am directly aware of the object itself...I shall say that an object is “known by description” when we know that it is “the so-and-so,” i.e. when we know that there is one object, and no more, having a certain property.

Russell (1911, pp 108 & 113)

Russell’s determination not to allow into his analysis the possibility of referring without there existing anything to which reference is made, however, forced him into a position which is now widely regarded as untenable: that the only things of which one can have knowledge by acquaintance, and therefore the only things to which one can truly refer, are one’s own sense data (although not the distal objects which the sense data represent) and oneself.

Following Russell, many attempts have been made to free the distinction between acquaintance-based thoughts (or *de re* thoughts) and description-based thoughts (*descriptive* thoughts) from the straitjacket of Russell’s epistemology (see, for instance, Bach (1987) and Recanati (1993)). I don’t propose to commit myself here to any one detailed analysis of the nature of *de re* thoughts; I’m sure that Bach’s and, in particular, Recanati’s accounts are very much along the right lines. I do, however, want to draw out the main features of *de re* thought that are common to these two (and other) accounts. The central feature of *de re* thought, or rather of the *de re* individual concepts which serve as constituents of *de re* thoughts, is that they derive, in one way or another, from perception. A *de re* concept must, in other words, be anchored to its referent by some kind of perception-based causal chain. I do not want to delve into the very fraught details of exactly what kinds of causal chain will do for *de re* thought. The main point, as far as this paper is concerned, is that it is a causal chain which links a *de re* concept to its referent, not the satisfaction of any particular descriptive content. *De re* thought, in other words, is relational, not satisfactual.

Descriptive thought, by contrast, is satisfactional, not relational. There is no causal chain which anchors a descriptive individual concept to any particular individual in the world; rather, to think of an individual under a descriptive concept is to think of that individual as the unique satisfier of a collection of descriptive information. A descriptive individual concept, therefore, is no more than a collection of non-relational information which the holder of the concept believes to apply, in its entirety, to one and only one individual.

The distinction between *de re* and descriptive thought is clearly of fundamental truth-conditional significance: *de re* thought is, by its very nature, anchored to particular individuals in the world; descriptive thought is, by its very nature, anchored to no particular individual, but to whoever satisfies the relevant descriptive content. A brief consideration of the truth values of *de re* and descriptive thoughts in counterfactual situations will, I am sure, convince you of this point: say that, unbeknownst to me, my friend John is the author of *Love in the Morning*. I may entertain two beliefs: one a descriptive belief that the author of *Love in the Morning* cannot write for toffee and the other a *de re* belief that John cannot write for toffee (on the evidence, say, of a letter he has written me). Now it's clearly possible to imagine a counterfactual situation unlike the actual world in that the first of these beliefs is not about John, a counterfactual situation in which John is not, in fact, the author of *Love in the Morning*. But it's hard to see what it might mean for the second thought not to be about John in a counterfactual situation. These are, essentially, the same intuitions that underlie Kripke's claims about rigid designation and Kaplan's claims about direct reference.

The upshot of this is that *de re* concepts are anchored to their referents in such a way that the truth conditions of any *de re* thought are also anchored to that referent (I shy away from the suggestion that the referent itself appears in a specification of the truth conditions of a *de re* thought since it is hard to see what this might mean within a cognitive account). A descriptive concept, on the other hand, is not so anchored; it is satisfactional, rather than relational. The truth conditions of a descriptive thought will therefore be more or less Russellian, although much more remains to be said about the precise details of the truth conditions of descriptive thought (for further discussion of the distinction between *de re* and descriptive concepts, see, for instance, Bach (1987), Recanati (1993) and Powell (1998)).

3.2 The proposal

I think that's probably all the groundwork I need to do. So, bearing semantic underdeterminacy, the conceptual-procedural distinction and the distinction between *de re* and descriptive concepts in mind, I shall attempt to lay out an analysis of the interpretation of definite descriptions. Probably the clearest way to present my account is by answering the four questions set out at the beginning of the paper. It makes more sense, however, not to answer them in the order in which they appear: I therefore intend to answer first questions (b) and (d) together and then questions (a) and (c) together. The answer I propose to (b) and (d) is this: definite descriptions are univocal and encode procedures for the retrieval of propositional constituents; they instruct a hearer to access (or initiate) an individual concept, whether *de re* or descriptive, which includes the descriptive content of the description. 'The F', for instance, instructs the hearer to access an individual concept which includes the information 'is F'¹⁰. Indefinite descriptions, by contrast, do not instruct the hearer to access or initiate an individual concept. While definite descriptions are used where the hearer has access to uniquely identifying information about the description's referent, indefinite descriptions are used where the hearer has access to no such information¹¹.

How about questions (a) and (c)? I hope that by now it's reasonably clear what my answer to question (a) is going to be: it seems to me that the arguments put forward by the ambiguity theorist in favour of his contention that the referential-attributive distinction corresponds to a difference in truth conditions are very convincing. There is, however, still the issue of the Wiggins-type intuition to deal with, the intuition that in cases of misdescription, even if the predicate is true of the intended referent, so long it is false of the actual denotation of the description used, the proposition expressed by an utterance containing a referentially-used definite description will be false. I intend to address this issue again once I have laid out my proposal a bit more fully: it's my belief that a reasonable explanation for this intuition should fall out from the account. My

¹⁰ I use the word 'includes' here very loosely. I am inclined to accept a view of concepts along the lines proposed by Fodor or Sperber and Wilson, on which concepts are simple, unanalysable units. Recast within such a theory of concepts it might be more accurate to replace 'includes' with 'is associated with', however, as long as this is borne in mind, the terminology shouldn't lead to too much confusion.

¹¹ What constitutes access to uniquely identifying information, particularly in terms of linguistic context, is a more complicated question than I have allowed for. However, there is not the space to discuss it in detail here.

answer to (a), then, is that there is a truth-conditional distinction between referential and attributive uses.

How about question (c), the question of what contribution definite descriptions make to the propositions expressed by utterances in which they appear? I believe, along with Donnellan and Recanati, among others, that the proposition expressed by an utterance containing an attributively-used definite description is object-independent, whereas the proposition expressed by an utterance containing a referentially-used definite description is object-dependent. The next question is, how do these two different interpretations arise? I want first to give the outline of an answer before working briefly through a couple of examples to show how the whole thing might go in practice. In outline the answer looks like this: as I've already suggested, what is semantically encoded by definite descriptions does no more than constrain the hearer's search for an interpretation to an individual concept containing the description's descriptive content. The gap between this encoded meaning and the propositional constituent corresponding to the description is filled by inferential pragmatic processes guided by considerations of relevance. The difference between referential and attributive interpretations of definite descriptions amounts to no more than this: in some contexts the optimally relevant interpretation of a description will be a *de re* individual concept, in other contexts it will be a descriptive individual concept.

As I suggested, this might all be a bit clearer if I work briefly through a couple of examples. Let me remind you of sentence (1) and its two proposed interpretations, repeated here as (14), (15) and (16):

- (14) The Ferrari driver has an unfair advantage
- (15) The Ferrari driver, whoever he is, has an unfair advantage
- (16) Jones has an unfair advantage.

Let's deal with the referential interpretation, (16), first. Remember that, in the imagined context for (16), speaker and hearer are standing in the pit lane looking at a driver, Jones, who they believe to be the Ferrari driver and who is scheming with the race organiser. In that context it is mutually manifest to speaker and hearer that both have (or are capable of acquiring) a *de re* concept of Jones, since, however one defines the epistemic relationship requisite for *de re* thought, one is going to want to say that the speaker and hearer in this case both have sufficiently close epistemic contact with Jones to be able to entertain a *de re* concept of him. The speaker then utters (14) and sets the hearer on a hunt for an individual concept which includes the information 'is a Ferrari driver'. I presume that the speaker may well have various individual concepts which include this information, some of them *de re* and some of them descriptive, i.e. some of them linked

to individuals in the world via a causal chain and some of them not; he may, for instance, have a *de re* concept of Nicky Lauda and a descriptive concept of whoever won the 1998 Monaco Grand Prix, both of which contain the information ‘is a Ferrari driver’. It is clear, however, that in the context suggested these are not going to be the most salient, the most easily accessible Ferrari-driver individual concepts; there are two much more obvious candidates, both of which are likely to be reasonably salient: one is a descriptive concept something like ‘the Ferrari driver in this race’ and the other is a *de re* concept of Jones, standing in front of him. The question is, which one does he accept? And the answer to this, on a relevance theoretic picture, is going to depend on which concept offers him more contextual effects. It’s not hard to see, then, which this will be in the proposed context: since the hearer believes Jones to be the Ferrari driver, any contextual effects carried by the descriptive interpretation come along for free with the *de re* interpretation. However, the *de re* interpretation also has additional contextual effects: it might, for instance, contain the information that Jones is wearing a red helmet, in which case the hearer could infer that the driver with a red helmet has an unfair advantage and so on. In this context, then, the *de re* interpretation is going to be optimally relevant since it offers the hearer more contextual effects for no extra processing effort; the description will, therefore, be interpreted referentially.

Now let’s consider (15), the attributive interpretation. Remember that in the imagined context for (15), speaker and hearer are discussing a race; they do not know who the Ferrari driver is, they merely know that whoever he is, he has an unfair advantage since he will be allowed to start the race in front of the other cars. Again the speaker utters (14) and the hearer starts his hunt for an individual concept. What options are open to him this time? Again there are the concepts of Nicky Lauda and the winner of last year’s Monaco Grand Prix, but again they are going to be much less salient in the given context than another Ferrari-driver individual concept: the concept of whoever is driving the Ferrari in the race currently under discussion. Notice that the hearer, in this context, has no *de re* concept of the Ferrari driver in question: that is simply what is meant by saying that he doesn’t know *who the Ferrari driver is*. So the first individual concept he will access will be this descriptive concept of the Ferrari driver in the race under discussion, i.e. a satisfactoral concept containing information such as ‘is a racing driver’, ‘drives a Ferrari’, ‘is driving a Ferrari in this race’ and so on. He will then, as ever, assess this interpretation for optimal relevance and, given that it will have enough contextual effects - the race is unfair, the Ferrari driver is likely to be the winner, the other drivers may complain - for no undue processing effort, he should accept it as the intended interpretation.

Does this entail that wherever a referential interpretation is available, a hearer should accept it as the intended interpretation? Although this will standardly be the case, for

reasons along the lines presented above, there are certain situations in which, although speaker and hearer are both capable of entertaining *de re* representations of the individual which falls within the denotation of a description, nevertheless an attributive interpretation of that description will be optimally relevant. The clearest examples are those in which, by using a definite description, a speaker who intended her utterance to express a singular proposition would be putting her hearer to gratuitous processing effort. Consider the following situation: Wimbledon fortnight has just ended and Pete Sampras has won the men's singles title, a fact that is mutually manifest to A and B. A then utters (17):

(17) I bet the winner of this year's Wimbledon men's singles title also wins the US open

How is B going to interpret the definite description 'the winner of this year's Wimbledon men's singles title'? As with all definite descriptions, the hearer's task is to find the optimally relevant individual concept containing the descriptive content of the description. But will this be a *de re* concept or a descriptive concept in the given context? The *de re* interpretation is eliminated on the grounds that, had that been A's intended interpretation, there is at least one alternative utterance available to her which would have put her hearer to less effort for the same contextual effects, namely (18):

(18) I bet Pete Sampras also wins the US open

The optimally relevant interpretation must, therefore, be the attributive one, on which A wins the bet iff whoever won Wimbledon also wins the US open. If, for instance, Sampras is stripped of his Wimbledon title and Andre Agassi is awarded the title in his stead, A's bet will depend on whether Agassi wins the US open or not. It is thus the case that, although an available referential interpretation will normally be optimally relevant, this will not always be so.

This, then, is the outline of my proposal: definite descriptions encode procedures instructing the hearer to access an individual concept compatible with the description's descriptive content. The rest is left to pragmatic inference, guided by considerations of relevance. The distinction between referential and attributive interpretations arises, on this picture, not as a product of the encoded meaning of descriptions, nor, indeed, in some sense, as a product of the interpretation *process*: in all cases the interpretation of definite descriptions proceeds along the same lines, leading the hearer from encoded meaning to the optimally relevant individual concept in the context. The referential-attributive distinction lies, rather, at the level of the mind-world relationship: although all definite descriptions lead to individual concepts, some of those concepts, those which

are *de re*, are anchored in a particular way to objects in the world, and thus give rise to singular truth conditions, whereas others, those which are descriptive, are not so anchored and give rise to general truth conditions.

The main advantage I claim for this proposal is that it has the flexibility to account for the data: I hope with the above example I have convinced you that, in the standard referential and attributive cases, there is a good story to be told about how relevance can lead from the semantics proposed to the two alternative interpretations. It seems to me that this account can also deal with the other uses of definite descriptions mentioned above. Consider, for instance, the functional use, illustrated above by (5), repeated here as (19):

(19) The president of the United States changes every five years

Recanati (1989) gives a story on definite descriptions which is similar to the above account inasmuch as the semantics proposed is neutral as between referential and attributive readings. The basic idea is this: definite descriptions encode a Russellian semantics but are unspecified as to whether the descriptive content of the description is to be interpreted as truth-conditional or as a ‘contextual condition’. Recanati (1993) has since rejected this account, however, on the basis that a Russellian semantics is unable to account for functional uses such as (19). It seems to me that my account can accommodate such uses. As I have already discussed, the key distinction between *de re* and descriptive concepts is that the former are anchored to particular individuals in the world, whereas the latter are not so anchored. One upshot of this is that, whereas a *de re* individual concept will pick out the same individual across possible worlds, a descriptive individual concept may pick out different individuals. My proposal is that functional uses such as the above are, essentially, metarepresentational; in concrete terms, an utterance of (19) will express a proposition to the effect that the denotation of the descriptive individual concept corresponding to ‘the President of the United States’ changes every five years. This interpretation is accessed, as ever, via considerations of relevance: neither the straightforward *de re* or descriptive interpretations achieve optimal relevance, since, for both, the construction of a context in which the interpretation yields sufficient contextual effects puts the hearer to too much processing effort. The proposed interpretation, however, yields sufficient contextual effects - such as, for instance, Presidential elections must take place every five years, the current President has only a year to run in office and so on - and will thus be accepted as optimally relevant.

One further advantage of the account proposed is that it can deal with the data upon which the main alternative to a Russellian semantics relies. While Russell’s account focuses on the uniqueness implication apparently carried by definite descriptions, a

parallel tradition, advocated, in one form or another, by, for instance, Heim and Kempson, focuses on the familiarity or accessibility associated with definite descriptions. Consider the following sentence:

(20) A cat and a mouse are in the hall; the cat is about to eat the mouse

For Heim (1988), such examples suggest that what definite descriptions encode is not uniqueness but familiarity. Building this into her overall *file change* semantics programme, definite descriptions instruct a hearer to update an existing file, whereas indefinite descriptions instruct the hearer to initiate a new file. This view has come under attack on the basis that, although definite descriptions are standardly used to talk about the familiar, they can be used to introduce new referents. Consider, for instance, (21):

(21) The oldest panda in China is coming to London Zoo.

An utterance of (21) can clearly be understood by a hearer who had no pre-existing file on the oldest panda in China. Heim has mounted a defence against this objection on which there are mechanisms which link such descriptions to existing files, but it seems to me that a more natural explanation of this data falls from the account proposed above, along with standard communicative principles. There seems to me no need to propose that definite descriptions encode anything like familiarity. Remember that, on my account, what a definite description does encode is an instruction to access an individual concept. What, then, if the hearer does not have, or is not able to acquire, the intended individual concept? In this case the speaker would be putting her hearer to gratuitous processing effort, thus guaranteeing that her utterance will fail to be optimally relevant. In order to comply with the presumption of optimal relevance, then, a speaker who uses a definite description must have reason to believe that the individual concept which is the intended interpretation of the description is either already in the hearer's conceptual repertoire (i.e. is familiar), or is accessible to the hearer. The familiarity or accessibility standardly associated with definite descriptions therefore fall out naturally from the semantics proposed, in tandem with the presumption of optimal relevance¹².

Finally I'd like to tidy up one loose end. I suggested earlier that my account would be able to handle the Wiggins-type intuition that, in cases of misdescription, the truth value of an utterance will depend on whoever the description denotes, not on the individual the speaker is intending to refer to. There are two things I would like to say on this point.

¹² There is more to be said about how my account ties in with the issue of uniqueness. However, limited space again means that I shall have to leave that for another occasion.

Firstly, I agree with Reimer (1998), who argues that the supposedly pre-theoretical intuitions of the opposing camps in this debate are, in reality, anything but pre-theoretical. They are, rather, heavily influenced by the theoretical preconceptions that each camp brings to the debate. But secondly, it does not seem to me to be necessary to oppose Wiggins' intuition in order to maintain a truth-conditionally significant referential/attributive distinction. As I mentioned above, both Recanati (1993) and Wettstein (1981) argue for a truth-conditional distinction between referential and attributive uses, while accepting that the truth conditions of the proposition expressed by an utterance containing a referentially-used definite description cannot make reference to an individual that does not satisfy the description's descriptive content. In other words, an utterance of 'The F is G' can only express a proposition about an individual that is, in fact, F, although it may of course communicate propositions about an individual that is non-F. I am inclined to accept this view, which seems to fall quite naturally from my account: definite descriptions are tools for talking about individuals which satisfy their descriptive content. It is for this reason that their encoded meaning constrains the hearer's interpretive procedure to hunting for individual concepts which contain that descriptive content. Given this, Wiggins-type intuitions on cases of misdescription are just what we should expect: although definite descriptions can be used to express object-dependent propositions, they can only be used to express such propositions about individuals which satisfy their descriptive content.

4 Conclusion

In this paper I have undertaken two more or less separate tasks: firstly I have attempted to lay out the main strands in the ongoing debate on the semantics and pragmatics of definite descriptions, strands that I have labelled the yes- and no-theoretic positions. In the process I hope I have convinced you of the following: what separates the two positions is the product of a theoretical error, more than of any fundamental dispute over the data. Given that the two good insights discussed above are incompatible if one accepts the bad assumption, each camp has taken one of the insights and built it into a theory which excludes the other. Having shown this, I then set about putting forward a proposal on definite descriptions which allows both good insights to be accommodated within one theory. The key to this proposal is abandoning the bad assumption, i.e. accepting that pragmatic inference has a central role to play in the retrieval of propositional content. Once this is accepted an account on which a univocal semantics for definite descriptions leads to different truth conditions becomes available.

The account I have proposed posits a univocal procedural semantics for definite descriptions which is indeterminate between referential and attributive (and functional

and specific etc.) readings. One upshot of the particular account I have proposed is that the distinction between referential and attributive interpretations of descriptions is the product not of the relationship between language and mind, but of the relationship between mind and world: it arises because of a distinction between the kind of link that exists between *de re* concepts and the world and the kind of link that exists between descriptive concepts and the world. This seems to me to be yet more evidence, if evidence is needed, that the real domain of truth-conditional semantics is not sentences of natural language but propositions, i.e. sentences in the language of thought.

References

- Bach, K. (1987). *Thought and Reference*. Oxford: Clarendon Press.
- Bezuidenhout, A. (1997). Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction. *Mind* 106, 375-409.
- Blakemore, D. (1987). *Semantic Constraints on Relevance*. Oxford: Blackwell.
- Breheny, R. (1999). *Context Dependence and Procedural Meaning: The Semantics of Definites*. University of London Ph.D. thesis.
- Carston, R. (1988). Implicature, Explicature and Truth-theoretic Semantics. In R. Kempson. (ed.). *Mental Representations: The interface between language and reality*. 155-181. Cambridge: Cambridge University Press.
- Carston, R. (1998). *Pragmatics and the Explicit/Implicit Distinction*. University of London Ph.D. thesis.
- Donnellan, K. (1966). Reference and Definite Description. *The Philosophical Review* 75, 281-304.
- Grice, H. P. (1967). Logic and Conversation. In *Studies in the way of words*, 22-40. Cambridge: Harvard University Press. 1989.
- Grice, H. P. (1969). Vacuous Names. In D. Davidson and J. Hintikka. (eds.) *Words and Objections*. 118-145. Dordrecht: Reidel.
- Heim, I. (1988). *The Semantics of Definite and Indefinite Noun Phrases*. New York: Garland.
- Kempson, R. (1986). Definite NPs and Context-Dependence: A Unified Theory of Anaphora. In Myers, T. et al. (eds.). *Reasoning and Discourse Processes*. New York: Academic Press.
- Kripke, S. (1979). Speaker's Reference and Semantic Reference. In P. A. French, T. E. Uehling, Jr. & H. K. Wettstein (eds.). *Contemporary Perspectives in the Philosophy of Language*. 6-27. Minneapolis: University of Minnesota Press.
- Ludlow, P. & S. Neale. (1991). Indefinite Descriptions: In Defense of Russell. *Linguistics and Philosophy* 14, 171-202.
- Neale, S. (1990). *Descriptions*. Cambridge: MIT Press.
- Peacocke, C. (1975). Proper Names, Reference and Rigid Designation. In S. Blackburn. (ed.). *Meaning, Reference, Necessity*. 109-132. Cambridge: Cambridge University Press.
- Powell, G. (1998). The deferred interpretation of indexicals and proper names. *UCL Working Papers in Linguistics* 10, 143-172.
- Ramachandran, M. (1996). The Ambiguity Thesis Versus Kripke's Defence of Russell. *Mind and Language* 11, 371-387.

- Recanati, F. (1989). Referential/Attributive: A Contextualist Proposal. *Philosophical Studies* 56, 217-249.
- Recanati, F. (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Reimer, M. (1998). The Wettstein/Salmon Debate: Critique and Resolution. *Pacific Philosophical Quarterly* 79, 130-151.
- Rouchota, V. (1992). On the referential/attribution distinction. *Lingua* 87, 137-167.
- Rouchota, V. (1994). *The Semantics and Pragmatics of the Subjunctive in Modern Greek - a Relevance-Theoretic Approach*. University of London Ph.D. thesis.
- Russell, B. (1905). On Denoting. *Mind* 14, 479-93.
- Russell, B. (1911). Knowledge by Acquaintance and Knowledge by Description. *Proceedings of the Aristotelian Society* XI, 108-28.
- Sperber, D. & D. Wilson. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge: Harvard University Press. Second edition 1995.
- Wettstein, H. (1981). Demonstrative Reference and Definite Descriptions. *Philosophical Studies* 40, 241-257.
- Wiggins, D. (1975). Identity, Designation, Essentialism and Physicalism. *Philosophia* 5, 1-30.
- Wilson, D. & D. Sperber. (1981). On Grice's Theory of Conversation. In P. Werth. (ed.). *Conversation and Discourse*. 155-178. London: Croom Helm.
- Wilson, D. & D. Sperber. (1993). Linguistic form and Relevance. *Lingua* 90.1/2, 1-25.