

The Referential/Attributive Distinction:
Its Status and Scope

Atheer Al-Khalifa
B.A. Hons. (Philosophy)

Department of Philosophy, School of Humanities
The University of Adelaide



THE UNIVERSITY
of ADELAIDE

Thesis submitted in fulfilment of the requirements for the degree of Master
in Philosophy

November 2020

Contents

Abstract	3
Declaration	4
Acknowledgments	5
Introduction	6
Chapter 1: Preliminaries	9
1.1 Characters, Truth Conditions, Propositions and Possible Worlds	9
1.2 The Meaning of Russell's Definite Descriptions	12
1.3 Object-Independent vs Object-Dependent Terms	16
1.4 Weakly vs. Strongly Directly Referential Terms	20
1.5 The Meaning of Frege's Definite Descriptions	21
Chapter 2: Donnellan's Constitutive R/A Distinction:	23
2.1 Donnellan's Observation	23
2.2 Correctness and Completeness	24
2.3 Donnellan's Attributive Uses	26
2.4 Donnellan's Referential Uses	28
Chapter 3: Absorbing the Distinction Pragmatically	35
3.1 Neale's Basic Case	35
3.2 Conversationally Implicating the Referential Use	37
3.2a Cancellability and its Significance	40
3.2b Derivation and its Significance	43
Chapter 4: The Systematicity of the R/A Constitutive Distinction	50
Chapter 5: The Pervasiveness of the R/A Constitutive Distinction	56
5.1 Demonstratives	56
5.1a King's Theory of Complex Demonstratives	59
5.2 Pronouns	63
5.3 More Data	67
Chapter 6: The Causal R/A Distinction	72
Chapter 7: Two Semantic Treatments	83
7.1 Approach 1: Intention Sensitive Indexicals	84

7.2 Approach 2: Syntactic Ambiguity	90
Chapter 8: The Problem of Misdcriptions	94
Conclusion	100
References	101

Abstract

In this thesis I investigate Donnellan's Referential / Attributive distinction (R/A distinction); a distinction about using (say) definite descriptions in two truth-conditionally different ways. I propose an argument *from non-misdescriptions to the R/A distinction* which does not utilise those cases of misdescriptions that Donnellan focused on. After arguing that the distinction should not be captured via conversational implicature, I point out a certain systematicity between the two relevant uses which tells against the view that the R/A distinction arises due to *lexical ambiguity*. I then extend the R/A distinction to demonstratives like 'that F' as well as some pronouns and suggest that it may be even more pervasive. Given this systematicity and pervasiveness of the R/A distinction, I propose two unified semantic treatments. The first is a treatment under which terms that exhibit the R/A distinction (R/A terms) are intention-sensitive indexicals and the second is a treatment under which R/A terms induce a syntactic ambiguity in the sentences they are embeded in. I conclude by distinguishing my argument *from non-misdescriptions to the R/A distinction* from the (perhaps) more familiar argument *from misdescriptions to the R/A distinction*, to which I adopt the well-known Kripkean (1977) position.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Atheer Al-Khalifa

4th November 2020

Acknowledgments

This thesis would not have been possible without the insight, guidance and encouragement of some people; I wish to thank them personally here. First and foremost, I wish to thank Antony Eagle, my primary supervisor. He is a fantastic person to talk philosophy with. There is not much that he said that did not, at once, make me re-evaluate my position, yet, he could not have been more approachable or more willing to hear me out and see where things go. I also wish to thank my secondary supervisor Jordi Fernández. It is due to Jordi's penetrating and fresh eye that some important sections and clarifications owe their existence – especially §7.2. I also wish to thank all staff in the philosophy department at the University of Adelaide for putting so much of their time and energy in creating such a philosophically lively, friendly and warm environment. A very special thank you goes to Jon Opie whose suggestion before commencing my candidature that I may be able to retract my rejection of scholarship saved my life! I wish to thank my fellow postgraduates who made each day of research so much fun: Edward Burston, Rebecca Randell, Brigitte Everett, Danny Wardle, Robert Farquharson, James Vlachoulis, Michael Lazarou, Henry Phillips, Jessica Pohlmann. A special thank you goes to Riley Harris and Matthew Nestor for some notational guidance pertaining to §7.1. I also wish to thank my family for their financial assistance when things got a little tough - a special thanks to my sister Sarah for being my primary life consultant. And last but not least, I wish to thank my partner Noor for all her care and cheering as I worked through this thesis.

Introduction

Since its publication in 1966, Donnellan's *Reference and Definite Descriptions* has been the cause of a longstanding debate. This is partly because, for better or worse, it wraps together various interesting issues into one. There are, as far as I can tell, three separate issues in that paper. The first is what I am going to call the 'Constitutive Referential/Attributive Distinction' or 'Constitutive R/A Distinction' for short. This is a claim about definite descriptions being associated with two different types of *linguistic content*. The second is what I am going to call the 'Causal Referential/Attribution Distinction' or 'Causal R/A Distinction' for short. This is a claim about the *mechanisms* at play in the *determination* of the previously mentioned linguistic contents. The third is the semantic status of *misdescriptions*. Early writers on the topic focused on this third issue of misdescriptions as central to Donnellan's distinction. They argued that his observations regarding this phenomenon are no threat to the classic unitary analysis given by Russell, who espouses no duality in the linguistic content associated with definite descriptions. Over time, however, writers appreciated that some of Donnellan's central observations can be freed from his more controversial claims about misdescriptions and they began to recognise that those observations can be plausibly upheld while setting aside the issue of misdescriptions. These days, if misdescriptions are brought up at all, it is almost always in an auxiliary way.

I think this trend is justified. Misdescriptions, as I will go on to emphasise, are a special case of the more general phenomenon of *misdesignation*. This latter phenomenon should be investigated independently of the other issues presented in Donnellan's paper. In Chapter 1, I begin this thesis by introducing the framework I'll be adopting, some terminology and some helpful distinctions. I then introduce Donnellan's distinction in Chapter 2 and distinguish between two independent routes to the constitutive R/A distinction: *the argument from misdescriptions* and *the argument from non-misdescriptions*. I will not have much to say about the argument from misdescriptions until the final chapter (Chapter 8), where I adopt the well-known Kripkean (1977) position. However, in this second chapter, I emphasise that there is a route through *the argument from non-misdescriptions* to a misdescription-free constitutive R/A distinction whose status is independent from the status of the argument from misdescriptions. My R/A constitutive distinction is based on uses of definite description that contribute a property-based condition to the proposition, as opposed to the object satisfying this property-based condition. But the property-based condition associated with both uses is just the familiar Russellian one.

A lot of the writing on the topic has been about the *status* of this misdescription-free R/A distinction. That is, about whether the duality of meaning associated with the two uses of definite descriptions is a result of only the information encoded in definite description terms (i.e., their semantics) or a result of the broader idiosyncratic conversational situation in which the speaker and audience find themselves when uttering those definite description terms (i.e., their pragmatics). Most of this literature can be viewed as a response to the question of whether definite descriptions are *ambiguous* between those referential and attributive readings. But this trend is curious. Most of the attention has been on whether the attributive analysis (which most people take to be the Russellian analysis) of definite descriptions can somehow absorb the referential use (which I will be attributing to Frege and not Strawson). If it can, then it is concluded that definite descriptions are not ambiguous; if it can't, then it is concluded that they are. But Donnellan's criticism was against both the Russellian and Fregean analyses. Donnellan's accusation of ambiguity was due to neither analysis being, allegedly, adequate. Why isn't as much attention given to whether the Fregean (or referential) use can absorb the Russellian (or attributive) use? Maybe it is because most writers favour the Russellian analysis, and maybe that's because they think that the Russellian analysis is in fact explanatorily more powerful. Still, it will be instructive to look at how the Fregean might capture the R/A constitutive distinction (i.e., how the Fregean absorbs or explains away the Russellian use). This is not because the Fregean analysis is better, but because going down that path exposes a limitation with a certain popular defensive tactic that Russellians usually employ to explain away Donnellan's observations regarding those divergent referential uses. My main focus in Chapter 3 will be on the role that Gricean conversational implicature is supposed to play in the debate and whether we can use it to diffuse Donnellan's criticisms. My conclusion will be that we can, but in an unsatisfactory way.

In the second part of the thesis, starting in Chapter 4, I take a closer look at the distinction as it applies to the special case of definite descriptions. I observe that there is a systematic relation between the two types of linguistic content that the R/A constitutive distinction isolates. That is, there is a *truth-value dependence* (a somewhat *logical systematicity*) between the relevant meanings. This, for me, is a reason *against* the view that the R/A constitutive distinction arises due to a lexical ambiguity. Typical cases of lexical ambiguity (like homonymy or polysomy) do not exhibit this systematicity. This will not be a knock-down argument against the lexical ambiguity proposal; it just means that those defending this way of capturing the R/A constitutive distinction will need to say that the R/A distinction doesn't just point to a lexical ambiguity, but to a *special* kind of lexical ambiguity. This, I think, weakens their thesis. In Chapter 5, I look beyond definite description and point out that the same misdescription-free duality which Donnellan observed (or could have observed) in definite descriptions is also observed in other terms like demonstratives,

pronouns and perhaps even more quantifiers. In Chapter 6, I investigate the causal R/A distinction. This is an investigation into the *mechanisms* involved in the *determination* of the two types of linguistic content that the R/A constitutive distinction isolates. I conclude that there are two candidates for this mechanism, depending on how we capture the R/A distinction. According to the proposal under which R/A terms (these are terms that exhibit the R/A constitutive distinction) are intention-sensitive indexicals, this causal role has to be played by intentions and other candidates like *having in mind*, *grounds for assertions* or the *common ground* will not do. And according to the proposal under which the R/A distinction is a product of a syntactic ambiguity associated with the sentence the R/A terms are embedded in, this causal role is also an intention directed at a certain syntactic arrangement which the audience then recognises in order to disambiguate the sentence. In Chapter 7 I finally present these two proposals. I conclude that given the observations relating to the systematic relation between referential and attributive uses and their pervasiveness, one of these proposals is called for. In the final chapter (Chapter 8), I return to the issue of misdescriptions and situate it within the broader issue of misdesignation. I finish by emphasising how my discussion is motivated independently of any issues relating to misdesignation.

Chapter 1: Preliminaries

Before we begin discussing Donnellan's distinction, I need to present the framework I am adopting (i.e., what I mean by 'meaning') and introduce some terminology and distinctions that I will appeal to throughout the thesis. After doing this I will then present Russell's analysis of definite descriptions and contrast it with a different kind of analysis: the directly referential analysis which, as I will go on to claim, is what definite descriptions are analysed as according to Frege^{1,2} In the later chapters I will be connecting these two analyses to Donnellan's distinction.

1.1 Characters, Truth Conditions, Propositions and Possible Worlds

In this thesis I am going to assume a *truth-conditional* framework for doing semantics. This is an approach in the study of meaning that states that the meaning of a sentence (or at least a statement) is given by its truth conditions; the conditions under which the sentence or statement would be true.³ To explain Russell's analysis of definite descriptions (which is more or less the starting point for many discussions of Donnellan's distinction) and its *essential* feature, I am going to contrast it with the directly referential analysis of certain terms. In my explanation I will be appealing to *propositions* and *possible worlds*. This is how I will cash out the notion of truth conditions. For example, take 1:

1. Tom Cruise is happy

Standardly, we say that 1 is true iff Tom Cruise is happy. This truth-schema tells us that 1 is true iff Tom Cruise is happy. This is a condition that needs to be satisfied for 1 to be true. Satisfied by what? A possible situation in which 1 is true. Abstractly, we can take these conditions on possible situations to be modelled by a set of such situations or possible worlds in which a sentence is true, or alternatively, a *function* from possible worlds to truth values. That is, instead of the above underlined *wordy* representation of truth conditions, we can represent the truth conditions of an utterance as a function that takes some possible world as its argument and maps it to a truth value. In the case of 1, this function will take some arbitrary possible world as argument and output T (for *true*) iff Tom Cruise (the very man himself) has the property of being happy and F (for *false*)

¹ It is usually Strawson's analysis that is contrasted with Russell's when presenting Donnellan's distinction. But in this thesis, I ignore Strawson completely. I touch on this later, but nothing substantive turns on this choice.

² As Almog (2012) points out, it was in fact Donnellan's famous distinction that went some way to laying the foundations of the *directly referential* conception of certain terms.

³ I acknowledge that it is usually propositions that are said to be either true or false, but I am postponing talk of propositions until a bit later.

otherwise. More concretely, I will say that an utterance of 1 expresses (or would express when uttered appropriately) 2:

2. $\langle \text{Tom Cruise, being happy} \rangle$

2 will be our way of representing the *structured proposition* that 1 would express; I will be assuming a structured view of propositions throughout my thesis.⁴ Once we have this structured proposition however our function from possible worlds to truth values gets off the ground. This function will take some possible world and check if Tom Cruise (the first item in the above pair) has the property in the second position of the above pair (i.e., *being happy*) to output a truth value. This will give us a set of ordered pairs $\langle w, t \rangle$, where ' w ' stands for some possible world like w_1 , w_2 or w_3 and ' t ' which stands for some truth value; either T for true or F for false. Suppose we only have three worlds. Let's then say that we get the set in 3:

3. $\{\langle w_1, T \rangle, \langle w_2, T \rangle, \langle w_3, F \rangle\}$

3 is a representation of a function from possible worlds to truth values; it is a set of possible worlds in which the utterance is true/false. 3, then, is a formal representation of the meaning of the utterance in 1; it represents the truth conditions of 1. In other words, the linguistic information contained in 1 is represented by the possibilities that the utterance includes (i.e., w_1 and w_2) and excludes (i.e., w_3). According to Stalnaker (2014, p. 16) this fits well with a popular overall picture of language: *we use language to convey information*. 3, then, is a formal representation of an item of information. So, utterances express structured propositions, and structured propositions go into generating functions from possible worlds to truth values. But we cannot identify the meaning of 1 with the set in 3, for surely another expression, in that same context, like 'that guy running really fast is happy' uttered while pointing at Tom Cruise will have expressed the same proposition in 2, which in turn would have generated the same function in 3. But surely these utterances ('Tom Cruise is happy' and 'that guy running really fast is happy') have different meanings. But we can now relativise meanings to context and, therefore, we need two conceptions of meaning. The first we have already discussed; 2 and the function it generates in 3 turns out to be a representation of the meaning of 1 in a particular context and, therefore, it is what 1 means in *context*. However, the second aspect of meaning we will call, following Kaplan (1989), the 'character'. The character of an utterance is a rule that determines the structured proposition that an utterance expresses (which in turn

⁴ The proposition expressed by an utterance is sometimes called the 'content' of the utterance and the propositional contribution of a term is sometimes called the 'content' of a term or the 'semantic value' of that term. I have assumed a structured view of propositions. Some see this as an alternative to the view under which propositions are sets of possible words and truth values, but I see them as complimenting each other as suggested by Kaplan (1978).

generates the set of possible worlds and truth value pairs) for *each context*. More formally, the character is a function from contexts to propositions. The character of an utterance is a set of context and proposition pairs.⁵ It maps each context of utterance to the proposition expressed by that utterance in that context.⁶ This way, we distinguish nicely between ‘Tom Cruise is happy’ and ‘that guy running really fast is happy’ by saying that although they express the same structured proposition in 2 (and therefore generate the same set in 3) in context *c*, there will be another context *c'* in which the proposition expressed by those sentences will diverge.⁷ That is, their characters are different; the set of context and proposition pairs they generate is different.⁸ To put it simply, according to this view, the reason why ‘Tom Cruise is happy’ and ‘that guy running really fast is happy’ have different meanings is because the set of context and proposition pairs they generate is different and therefore the range of sets of possible world and truth value pairs they generate is different.

But there are other ways that people helpfully conceive of the character of an expression. Suppose you are walking down the street and you hear someone screaming, but you don’t know who because you cannot identify them. Suppose this someone is screaming the following:

4. I am angry

If understanding is in any way indicative of meaning, then in one sense of ‘understand’, you clearly *did not* understand the meaning of 4. You don’t know who it is that is angry and, therefore, you cannot evaluate the utterance for truth; we don’t know what was *expressed*. But in a different sense of ‘understand’, you clearly did understand something, something along the lines of *the speaker (whoever it is) is angry*. To distinguish between these two ways of understanding, Neale (1990, p. 70) thinks (and I think that he is right) we need two notions of meaning. This enables us to say (in the case of 4 above) although you may have (partly) grasped the character of 4, you did not grasp the proposition expressed by 4. Contrastingly, in 1 (above), we grasped both the character and proposition (this is why I chose ‘Tom Cruise’, the name of a famous person). Alternatively, we may think of the character of an expression as the kind of meaning that

⁵ Characters have been called the ‘linguistic meanings’, ‘senses’ or ‘modes of presentation’ of an expression. Some think of the character, informally, as the meaning of expression at the level of *type* (i.e. type truth conditions).

⁶ I say ‘utterance’, but perhaps simple terms have a character too (i.e. functions from context to propositional contributions). There is a debate about what things possess a character: simple or complex expressions. See Higginbotham (1998) and King and Stanley (2005).

⁷ Suppose that in context *c'* I utter, ‘That guy running really fast is happy’ while gesturing at Usain Bolt.

⁸ I am further assuming that *all* utterances have characters, even utterances that contain what some might call ‘context insensitive’ terms. The character for utterances involving such terms will be a trivial constant function that maps each context to the same proposition.

is constant throughout all uses of that expression. For example, when I, you, or Tom Cruise utter, 'I am angry', that utterance always means something like *the speaker* (in some context) *is angry*. This is the constant meaning of the expression throughout all its uses. However, depending on who utters the expression, different things will be expressed. When I utter it, I express that *Atheer is angry* and when Tom Cruise utters it, he expresses that *Tom Cruise is angry*. These propositions have different truth conditions. So we have two sorts of meanings: the character of an utterance and the proposition expressed by an utterance in context. I assume throughout this thesis that although (as we will come to see) one and the same utterance can express different propositions in various contexts (like 'I am angry', for example), that is not enough to deem that utterance *ambiguous*. I am going to assume that it is duality of *character* that is sufficient for ambiguity. I do not intend to go into detail about why I make those assumptions. I take them as standard/implicit in most discussions on the topic.⁹ This is all I have to say about the general framework I'll be adopting. Next, I introduce Russell's classic analysis of definite descriptions.

1.2 The Meaning of Russell's Definite Descriptions¹⁰

Definite descriptions are phrases of the form *the so and so*. 'the tallest man', 'the table', 'the happy customer', 'the show that's playing now', 'the missing car', 'the biggest number' and 'the first human born in the 22nd century' are all examples of definite descriptions. Other examples are 'his first car', 'John's girlfriend', 'Smith's murderer' or 'my last vacation'. These are, apparently, logically equivalent to 'the first car of his', 'the girlfriend of John', 'the murderer of Smith' and 'the last vacation of mine' respectively. According to Russell, definite descriptions are one type of a *denoting phrase*.¹¹ They are so, Russell (1905, p. 479) says, solely in virtue of their logical form. These phrases may denote an object so that one may talk or think about it. A simple subject-predicate statement of the form 'The F is G', containing a definite description in the subject place, roughly means:¹²

There exists exactly one object that is F and it is G

More precisely, definite descriptions, according to a standard interpretation of Russell,

⁹ A parallel point can be made about the character of terms. We can say that although a term can make different propositional contributions in different contexts (like 'I', for example), that is not enough to deem that term *lexically ambiguous*. It is duality of *character* that is sufficient for lexical ambiguity.

¹⁰ Strictly speaking, Russell (1905, p. 480) thought that definite descriptions, in themselves, have no meaning, but the sentences they are embedded in are meaningful. However, I talk as if he thought they had meaning merely because it is easier for my contrastive purposes. Nothing substantive hinges on this decision.

¹¹ Others are indefinite descriptions such as 'a young man' or 'an old apple'.

¹² Where 'G' stands for some property that is predicated on what is denoted by 'the F'. And 'F' here stands for some property or other given by the noun-phrase directly accompanying the definite article 'the'.

are *quantificational terms*. A statement containing them in subject place has the following logical form:

$$1. \exists x((Fx \wedge Gx) \wedge \forall y(Fy \rightarrow x = y))$$

In semi-English, 1 says that *there exists at least one x such that x is F and G and for all y if y is F then x is identical to y*. 1, then, makes three different claims:

- 1a. There exists at least one x such that it is F. [This claim involves *existence*]
- 1b. There is at most one x such that it is F. [This claim involves *uniqueness*]
- 1c. Everything which is F is also G.¹³ [This claim involves *predication*]

Recall from the last section that we are working within a truth-conditional semantic framework. This means that 1a, 1b, and 1c above provide the meaning (in quasi-formal language) of a statement containing a definite description in the subject place *because* they specify the conditions under which that statement would be true. 1a and 1b are also the conditions some object must satisfy in order for the definite description to successfully designate it.¹⁴ Let us call definite descriptions analysed in this way ‘Russellian’ or ‘quantificational’.

Now, following Russell, we will call the mechanism by which a Russellian definite description designates an object ‘denoting’. We can think of denoting as *application* or *satisfaction*. When we say, for example, that Brian May is denoted by 2:

2. The lead guitarist of Queen

we are saying that 2 applies to Brian May uniquely. It applies to him uniquely because he, alone, has the property of being a lead guitarist of Queen. Brian May then, is the denotation of 2. This way, we can say that 2 is *about* Brian May. Denoting definite descriptions are about some object because they encode object-independent, *general, descriptive* or *property-based* information (under some quantificational constraint) that fits that object (it doesn’t matter which) uniquely. If Brian May were, all of a sudden, to get replaced by (say) Slash, then it is Slash that will be denoted by 1. Application or satisfaction is one way in which we can linguistically talk and think about objects in the world. I have never seen or been in contact with the biggest grain of sand in the Sahara Desert, but I can certainly talk about it if it exists. I can do that using a definite description because, presumably, there is a particular object out there that would fit my description

¹³ Or property overlap; that the Fs are among the Gs.

¹⁴ I use the term ‘designate’ to stay neutral between ‘refer’ and ‘denote’ because these last two terms have come to carry certain semantic connotations as we will see.

uniquely.¹⁵ Russell would say that 2 is only about Brian May *indirectly*. What Russell means to say, I think, is that linguistically, 2 has no more to do with Brian May than anything else. Surely some of us think about the denotation of 2 (i.e., Brian May) when reading and thinking about 2, but the denotation of the definite description (the very object that satisfies it) is not part of the linguistic information encoded in 2. It is (only sometimes) part of the extra-linguistic information associated with 2 in a given utterance situation; or so Russell would claim. That the denotation of a Russellian definite description is *not* part of its linguistic meaning is crucial for Russell's above analysis, for one of the things he aimed to do with his theory of definite descriptions was to provide an analysis of definite descriptions under which the object denoted did not feature in their meaning at all. For Russell, this solved certain well-known metaphysical, semantic and epistemic problems. 1, then, is the logical representation of Russell's analysis of the meaning of expressions of the form 'The F is G' that occur in our natural and spoken English language. In addition to this logical representation, we also saw the wordy representation of this expression given in conditions 1a, 1b and 1c. But we can also represent the proposition expressed by the 'The lead guitarist of Queen is happy' in the following way:

3. ⟨⟨'The', lead guitarist of Queen⟩ being happy⟩¹⁶

Russell, in analysing the meaning of sentences of the form 'The F is G', also gave us what we (the speakers) express when we competently and sincerely utter a statement containing a definite description as subject. When we say, for example, 'The author of Waverley is sad' we are not linguistically referring or *pointing* to a particular individual and saying *of* them that they are sad. Rather, we are saying that Waverley was written by exactly one thing and that thing is sad. This is what I meant earlier by saying that according to Russell, definite descriptions are *quantificational*; they specify in a given domain, that the quantity of things having a certain property is at least one and at most one (i.e., exactly one). The definite description quantifier functions in the same way as (say) an indefinite quantifier like 'a' or 'some', except that in the case of 'the', the upper limit of what things must fit the noun phrase (e.g., 'writer', 'singer') is also specified; it is also one. Now straight off the bat, Russell's analysis faces some difficulties, as is widely recognised:

4. The dogs are barking

¹⁵ Vagueness worries may mean this example is not a case of determinate denotation.

¹⁶ Don't get thrown off by this bizarre representation. This is the way Kaplan (1978, p. 227) represents the structured propositional contribution of 'the'. He says that the fact that the logical word 'the' is '... just carried along reflects our treatment of them [definite descriptions] as syncategorematic, i.e. as having no independent meaning but as indicators of how to combine the meaning-bearing parts in determining the meaning of the whole'.

- 4a. #There exists exactly one dogs that is barking
 5. The water is dirty.
 5a. #There exists exactly one water and it is dirty

Applying Russell's analysis to 4 we get 4a, which sounds clearly infelicitous. Applying Russell's analysis to 5, we get 5a. But a natural reading of 5 does not imply the existence of exactly one thing that is water; 5a seems to get things wrong. To get around these problems, some have tried to revise Russell's analysis to suit. For example, one may logically analyse 4 as:

$$4b. \exists x \exists z ((Fx \wedge Fz) \wedge x \neq z) \wedge (\forall y)(Fy \rightarrow Gy)$$

This says two things: *that all Fs are Gs* (this is given by $(\forall y)(Fy \rightarrow Gy)$) *and that there exists more than one F* (this is given by $\exists x \exists z ((Fx \wedge Fz) \wedge x \neq z)$). 4b says *that there are at least two things that are a dog (and those two things are not identical) and for everything, if it is a dog then it is barking*. Similarly, one might just revise 5 as (Sharvy, 1980):

$$5b. \exists x (Fx \& \forall y (Fy \rightarrow y \leq x) \& Gx)$$

Where F can be either a singular count noun (like 'dog'), a plural count noun (like 'dogs') or a mass noun (like 'water'), and the symbol ' \leq ' indicates a parthood relation. This accommodates natural kind definite descriptions like the one in 5 as well as plural definite descriptions like the one in 4. As Ludlow (2018) points out, 5 says that an F that every F is part of satisfies G. It follows then, that a definite description 'the F' ('the Fs') denotes the maximal sum or mass in the extension of F. So, for example, when we utter, 'the dogs', we are denoting the maximal set of dogs in the domain of discourse. When we utter, 'the water', we are denoting the maximal mass of water in the discourse. In the case of non-plural definite descriptions (like 'the dog') we are still denoting the maximal set of dogs in the domain of discourse (since in this case there is only one dog).

Another problem case comes from what have been called 'incomplete definite descriptions':

6. The book is on the table
 6a. There exists exactly one book and it is on the table.

Examples like 6 have also been thought to be counterexamples to Russell's analysis for in uttering 6, we would not be typically implying 6a (there is more than one book in the universe and most people know this!). Again, some (like Stanley & Szabó, 2000) discuss contextually restricted domains of quantification. Roughly speaking, they propose inserting a placeholder into the logical form of the expression to be filled in by context

that restricts the domain of quantification to things like Adelaide or room 720. Others have claimed that utterances of 6 (and not the sentence in 6 mind you) are elliptical for some more complete expression with a fuller description that denotes the object being spoken about. Now, all this is great; however, in this thesis I am not interested in seeing how Russell's analysis might absorb issues like the ones illustrated in 4, 5 and 6. This is because these types of revisions *preserve* the essential feature of Russell's analysis. Many names have been thrown around to name this feature in the context of Donnellan's discussion, such as 'Russellian', 'quantificational', 'general' or 'object-independent'. I will pick the term 'object-independent' from Neale (1990) to name this feature that is essential to Russell's analysis of definite descriptions. I will explain this feature in the next section when I contrast it with the *object-dependent* feature of directly referential terms. In any case, I take it to be non-controversial that at least some definite descriptions (or at least, for now, some definite description *uses*) have this object-independent feature, even though the details of that feature (as I suggested in this section by looking at some revisions of Russell's analysis) may not be set in stone. In this thesis, rather, I am investigating Donnellan's distinction and, as will become clear in the second half of the thesis, it just so happens that definite descriptions are one place this distinction shows up. What I intend to do in the first half of the thesis when focusing on definite descriptions is to look at only one allegedly non-Russellian use: the so called 'referential' use. And the reason why this particular use is of special interest is because it seems to resist an object-independent analysis; what is essential about Russell's analysis.

1.3 Object-Independent vs Object-Dependent Terms

In order to explain the object-independent feature of Russell's analysis, it will be helpful to contrast it with the object-dependent feature of what are called 'directly referential terms'. Kaplan (1977/1989) proposed that some terms are used to express object-dependent propositions. In the last section, I gave the following example:

7. Tom Cruise is happy

Many people think that 7 would express an object-dependent proposition when uttered appropriately. This is because 7 contains 'Tom Cruise' – a name and names are frequently conceived as directly referential terms. Previously, I represented the proposition that 7 would express as the following ordered pair (a structured proposition):

8. ⟨Tom Cruise, being happy⟩

I also said that this ordered pair (consisting of Tom Cruise and the property of being happy) generates a function from possible worlds to truth values. This function

represented for us an item of information, namely truth conditions. This function required us to locate Tom Cruise himself in some possible world to check whether he is happy or not. Now, we are not literally going to do this; this is just a colourful way of saying something we said earlier. We said that the meaning of an expression in context is given by the conditions under which it is or *would* be true (or just the conditions under which it would be true). Would 8 be true in a world (or counterfactual scenario) in which Tom Cruise himself is happy? Yes, it will be. The proposition in 7 consists of Tom Cruise (the very man himself) and it depends on him for truth evaluation in any possible world, including the actual one. 7 then exhibits this object-dependent feature because it contains a term like 'Tom Cruise', which contributes an object itself to the proposition it expresses; it contributes the object it designates in the actual world. 8, then, is an object-dependent proposition. It is a proposition that is *directly* about Tom Cruise and so it has Tom Cruise as a *direct constituent*, as some writers usually put it.

In the last section, I talked about Russell's view that definite descriptions are terms which are about some individual but only *indirectly*. I said that he treated definite descriptions as quantifiers, which are things that specify the quantity of things with certain properties in a given domain. Take the following sentence:

9. The lead guitarist for Queen is happy

Earlier we said that we can represent the proposition that an utterance of 9 would express under the Russellian analysis as 10 below:

10. ⟨⟨'The', lead guitarist of Queen⟩ being happy⟩

10 goes into determining that function from possible worlds to truth values. 10 requires us to locate one and only one thing (in some contextually determined domain) that has the property of being the lead guitarist of Queen as well as the property of being happy. If such an individual is located in some world then the proposition is true in that world. If not, then the proposition is false. But how does 8 differ *in kind* from 10? Well, 10 does not consist of Brian May or any other object as a *main* constituent. Notice here that I say 'main constituent', because it is plainly obvious that the definite description in 3 contributes the object designated by 'Queen' (assuming 'Queen' is term that contributes an object to the proposition), but that object is not the main constituent of the first item in the pair in 10; it is a constituent *of* a constituent. The definite description in 9 is, strictly speaking, not about Queen but about the lead guitarist of Queen, whoever that happens to be. And *who* that happens to be is irrelevant to the truth evaluation of 10, so long as it is exactly one who is a lead guitarist for Queen and that thing is also happy. This is what I mean by the *essential* object-independent feature of Russell's analysis of definite descriptions. It is that feature that makes it the case that it is not an object that goes (as a

main constituent) into the relevant proposition. Russellian definite descriptions like 'the lead guitarist of Queen' exhibit this feature and they are used to express the object-independent proposition in 10. It is a proposition that is about some particular thing but only indirectly, that is, only insofar as that thing satisfies some properties uniquely. Now, one might say that there is really no difference in meaning between 11 and 12 below:

11. Brian May is happy

12. The lead guitarist of Queen is happy

A person might say that if 11 is true, then 12 is true, and if 12 is true then 11 is true. They might add that if someone says 11, then it is as if they had said 12, and vice versa. But this is not right. Recall that the meaning of an expression is given by the conditions under which it is true *and* the conditions under which it *would* be true. We said that we are going to assume that sentences express structured propositions in various contexts which in turn generates a function. This function gives us a set of possible world and truth value pairs for each context of utterance. Now, of course, if we limit ourselves to only the actual world then the meanings of 11 and 12, one might insist, are the same.¹⁷ That is, this person might say, although 11 and 12 express different structured propositions, they nevertheless generate the same function from possible world to truth values at each context. But recall that we have assumed a truth-conditional semantics which just commits us to the view that the meanings of sentences is given by the conditions under which they *would* be true had the world been a different way. For example, it seems that if things had gone differently for Slash and he ended up as a very happy lead guitarist for Queen then 12 would be true, whereas 11 would be false in this counterfactual world where Brian May had ended up being a checkout clerk with only unfulfilled dreams of being in a rock band. This way, the truth conditions of 11 and 12 diverge, and this is why 11 & 12 differ in meaning. There are a few tests that allegedly confirm this.

13. Brian May is happy, and it is not the case that the lead guitarist of Queen is happy

If names are synonymous with Russellian definite descriptions, we should expect 13 to be a contradiction. But it does not seem to be a contradiction, although it may be false. That is, it seems that 13 can indeed be true (as our little imaginative story above involving Slash illustrates). This point shows that the name 'Brian May' and the definite description don't have the same *range* of propositional contributions as 'Brian May' and therefore they have different characters. Also, take 14 below, disambiguated with the scopes in 15:

14. Brian May might not have been the lead guitarist of Queen

¹⁷ For the sake of argument let us ignore the complication that the designation of 12 can change over time in the actual world and the fact that 11 and 12 express different structured propositions.

15. There exists some possible world in which Brian May is not the lead guitarist of Queen

Surely 15 is true. This point shows that 'Brian May' and 'the lead guitarist for Queen' don't have the same range of propositional contributions or character. If they did, then 15 should be false. Again, the Russellian will say, the fact that we may *think* of Brian May when we utter or hear, 'the lead guitarist for Queen' is truth-conditionally inert. But still, you might object and respond by saying that although 11 is not synonymous with 12, 11 is still synonymous with 16 below:

16. The *actual* lead guitarist for Queen is happy

This is because, you might say, 11 and 16 although they express different structured propositions, they generate the same function from possible worlds to truth values at each context. This is correct, they will, but that just means that for each name there may be a definite description that can *mimic* its modal profile. But this does not mean that definite descriptions in *general* have the same range of propositional contribution as names. This is because there are obviously *non-rigid* definite descriptions. But still, once the modal profile of rigid definite descriptions is recognised, you might then notice that *every* definite description can be rigidified. Therefore, you might say that once we rigidify definite descriptions, I cannot distinguish between definite descriptions and names by pointing out their modal profiles. My response to this is that although we can rigidify all definite descriptions, that does not mean that all definite descriptions are synonymous with their rigidified variants. For otherwise, we lose the modal significance of scope interactions (I will say more about this in Chapter 2). I take it for granted that non-rigid definite descriptions must be possible. Now, there is plenty of data that is supposed to justify the truth-conditional differences between object-dependent and object-independent terms. I have no intention to survey it here. I just assume that these two kinds of analyses (the object-dependent and object-independent analyses) of terms exist. What will be a matter of debate in this thesis is *how* these different analyses are distributed within and across various designational terms. Note, however, that I didn't have to use definite descriptions and names to distinguish between these two kinds of analyses, but I wanted to initiate you into the coming debate. Donnellan's discussion will challenge some of the claims I made above. So we have distinguished object-dependent propositions from object-independent ones. In the former, the main constituent is an object, and in the latter, the main constituent is some property that is, in the case of Russellian definite descriptions, accompanied by a quantificational structure (or whatever it is that corresponds in reality to *uniqueness*).

1.4 Weakly vs. Strongly Directly Referential Terms

Before I end this chapter, I want to discuss those directly referential terms in more detail. Directly referential terms come in two importantly different varieties which I did not distinguish between in the last section. I said that directly referential terms exhibit that object-dependent feature which allows them to contribute objects to the proposition these terms are used to express: object-dependent propositions. But *how* they do this depends on whether the directly referential term is *weakly* referential or *strongly* referential. Names are usually thought to be strongly directly referential (or at least we will assume this for the purpose of our contrast) whereas demonstratives and indexicals (e.g., 'he', 'she', 'you', 'this', 'that', 'I', 'you') are weakly directly referential (at least this is how they are treated classically). Now to bring out this difference between the two kinds of directly referential terms intuitively, I appeal to Martí's (1995, pp. 281-282) discussion. Consider the following two questions:

- a. What is the mechanism by which a name contributes an object to a proposition?
- b. What is the mechanism by which a demonstrative or pronoun contributes an object to a proposition?

In (a), we may be asking one of two questions. The first is a historical question about how a name came to be linked to an object. And in the second, we may be asking a special instance of a more general question regarding how a thing comes to stand or represent another thing. This question is one that needs to be answered in the general framework of a theory of what it is for a symbol (linguistic or otherwise) to stand for something. The important point here is that neither are questions that should especially be answered by a semanticist. Names are directly referential because they are stand-ins, ad-hoc marks or tags for objects. A noise is just noise, but if you name something with that noise, you can use it to say things that are true or false, to utter commands or to make promises. All it takes to incorporate a noise into a language is to succeed in giving it to an individual as a name. On the other hand, (b) is certainly a question that a semanticist should be able to answer. If Kaplan's account of the semantics of demonstratives or indexicals is correct, then they are very different to names. This is because what a demonstrative or indexical contributes to the proposition is in part uniformly determined by some object satisfying a certain descriptive condition and the choice of this condition is not done by mere stipulation. I'll illustrate this with an example I used previously:

17. I am angry

I said that the indexical 'I' in 17 can contribute to the proposition different persons in different contexts. If we ask question (b) about 'I', an important component of the answer

will consist in there being some kind of descriptive condition that some object must satisfy in order for that object to be contributed to the proposition by 'I' in some context. I hinted at this previously when I spoke of the 'character' of a term. The character of a term was a function from contexts to propositional contributions that involves some descriptive condition. So 'I' doesn't just stand for an object simpliciter, in the way that a name allegedly does. It stands for an object (given some context) partly by virtue of that object satisfying a certain descriptive condition. In the case of 'I', this would be something like, *being the speaker* in that context. This information is supposed to be linguistically encoded in the indexical. In contrast to this, a name like 'Tom Cruise' (assuming it is strongly directly referential) designates its object not by virtue of it satisfying any descriptive condition but by virtue of it (perhaps) being causally or historically connected in some appropriate way to that object.¹⁸

1.5 The Meaning of Frege's Definite Descriptions

Now that we have distinguished strongly directly referential terms from weakly directly referential terms, we can say what a Fregean definite description is. A Fregean definite description is a weakly directly referential term that contributes to the proposition the object that is in fact the satisfier of the description in a given context. So, according to Frege (and contrary to Russell), the propositional contribution of the definite description in 18 below is Brian May and 18 as a whole expresses the structured object-dependent proposition in 19. In a nutshell, under the Fregean analysis, the propositional contribution of a definite description is the very Russellian denotation of the definite description.

18. The lead guitarist of Queen is happy

19. ⟨Brian May, being happy⟩

Note, however, that the Fregean analysis is not subject to the criticisms I raised (at the end of §1.3) against the proposal under which a definite description like 'the lead guitarist for Queen' just means 'Brian May'. Fregean definite descriptions do not, in general, have the same range of propositional contributions as names that name the denotation of the definite description.¹⁹

¹⁸ One might say that this can't be what distinguishes names from indexicals. For in the case of names too the relevant object that is contributed to the proposition has to satisfy some uniform descriptive condition like (say) *being causally (or historically) connected to object o*. But this is the wrong picture of what is going on with names. The *causal / historical* aspect relevant to names is *relational* not *satisfactional*.

¹⁹ An example will illustrate the point. Suppose I say, 'the dog is lost' (using the definite description in the Fregean way) and suppose that the dog I am talking about is called 'Silver'. In this context both utterances 'the dog is lost' and

Before moving on to the next chapter to present Donnellan's distinction, let's summarise the main points in this chapter. We have so far presented at least three categories of linguistic analysis, that, for the time being, are best exemplified in three kinds of terms. We have names, which are strongly directly referential because the object they contribute (as a main constituent) to the proposition is not determined via some descriptive condition.²⁰ We have demonstratives like 'I' or Fregean definite descriptions which are weakly directly referential, as the object they contribute to the proposition is determined via a descriptive condition that is linguistically encoded in the term. Finally, we have terms like Russellian definite descriptions that do not really refer at all, but may merely *apply* to some object that satisfies its descriptive information uniquely. But the object that definite descriptions apply to (the denotation) is not part of their meaning. The propositional contribution of names and demonstratives is an object, whereas the propositional contribution of a definite description is not an object, but some sort of quantificational structure accompanied by properties or, let us say, a *property-based condition* or just *condition* for simplicity. In the next section, I am going to link Donnellan's attributive uses to Russellian definite description and his referential uses to Fregean definite descriptions. I choose the name 'Fregean' and not Strawsonian (as many writers on the topic do) because Strawson rejected Russell's uniqueness condition. Frege, on the other hand, thought that the contribution which a definite description makes to the truth-value of a sentence is determined by the fact that it stands for a certain object: the unique satisfier of the description. In our way of speaking then (which is different to Frege's), he held that the propositional contribution of a definite description is the object which it stands for (Miller, 2007, p. 65).²¹ Now, Frege might disagree with this characterisation, but it doesn't matter. I choose to call it 'Fregean' because the Fregean analysis (as will become clear in the next chapter) is one that captures what is plausible and essential about Donnellan's observations in a way that does the least amount of violence to Russell's analysis. Of course, the Fregean analysis is still in some way wildly different from Russell's, but in some other way it is not so wildly different. Russell and Frege agree about what condition an object must satisfy to be the denotation of the definite description, they would just disagree about what goes into the proposition. Russell would say that the propositional contribution is a condition, while Frege would say that it is the denoted object itself. Of course, Donnellan's observations are as much a challenge to Frege as they are to Russell. This is what I shall turn to now.

'Silver is lost' express the same structured proposition which determines the same set of possible world and truth value pairs. But if we vary the context, Silver will still name Silver, but the 'the dog is lost' may denote 'Foxy'. This will go into generating two different functions.

²⁰ Although that object may be fixed by some description, it never becomes relevant to both, the truth conditions and the character that determines the truth conditions.

²¹ See also Elbourne (2003, pp. 42-46) for a survey of passages from Frege's writings that are highly suggestive of the analysis that I am attributing to Frege.

Chapter 2: Donnellan's Constitutive R/A Distinction:

In this chapter, I will present Donnellan's famous distinction between referential and attributive uses of definite descriptions. I will first taxonomize the various kinds of definite descriptions that can be used either referentially or attributively and focus on one; those that are assumed to be *complete and correct*. I want to see how one may motivate Donnellan's distinction using this class of definite descriptions. Donnellan's distinction is supposed to account for some data that, apparently, Russell and Frege did not when giving their semantic analyses of definite descriptions. The upshot of Donnellan's observations is supposed to be that neither analysis is adequate for all uses of definite descriptions, and, therefore, definite descriptions call for some kind of ambiguity.

2.1 Donnellan's Observation

In *Reference and Definite Descriptions* (1966) Donnellan points out that on one occasion, a speaker may *use* a definite description in an expression to state something about whoever or whatever is the so-and-so; this will be an *attributive* use. On another occasion, a speaker may use that same definite description term in the same expression to enable his audience to pick out whom or what he is talking about and state something about that person or thing; this will be a *referential* use. Take the following classic example:

1. The murderer of Smith is insane

Suppose, first, that a speaker comes upon poor Smith, who has been foully murdered. Based on the brutal manner of the killing (and given that Smith was just the most lovable guy), the speaker utters 1. This is an attributive use of the definite description in 1. The contrast with such a use is an utterance of 1 in a situation in which the speaker expects and intends their audience to realise whom they *have in mind* when they speak of the murderer of Smith and, most importantly, to know that it is *this* person about whom they are going to say something. For example, suppose that Jones has been charged with Smith's murder and has been placed on trial. Imagine that there is a discussion of Jones' odd behaviour at his trial. The speaker might sum up their impression of Jones' behaviour by uttering 1. If someone asks to whom the speaker is referring to by using this description, the answer here is Jones. This is the referential use of the description in 1. Donnellan adds that in the attributive case, the definite description occurs *essentially*, for the speaker wishes to assert something about *whatever or whoever* fits that description. However, in the referential case, the definite description is merely one tool for calling attention to a person or thing and any other device for doing the same job; another

description or name would do as well (Donnellan, 1966, pp. 285-286).

Now, it is going to be very hard to try and come up with a Donnellian analysis of these attributive and referential uses. This is mostly because, unlike Russell, Donnellan relies very heavily on the actions and psychologies of speakers. For example, he relies heavily on phrases like 'using an expression', 'having things in mind', 'speakers asserting', 'speakers wishing' and 'speakers using a definite description to call attention to'. He does not give us an analysis of definite descriptions as used in either of these ways, but merely observes what speakers are *doing* when they utter these things. Still, I think Donnellan's two uses are intuitively real and it will be important to try to get to the bottom of what the difference between them amounts to. One thing I am *not* going to do in this first part of my thesis is to attempt to give necessary and sufficient conditions for the mechanisms involved in giving rise to those uses. At this early stage of the thesis, we are only trying to work out what propositional contribution a referential and attributive definite description makes and not how these contributions are produced by speakers or ascertained by hearers. But, of course, this later issue will be extremely relevant in the second part of the thesis when I explore the *causal* R/A distinction. I am here distinguishing between two *in virtue of what* questions: a constitutive one and a causal one.

- Constitutive Question: In virtue of what propositional contribution is a definite description referential/attributional?
- Causal Question: In virtue of what does a referential/attributional definite description make the relevant propositional contribution?

The constitutive R/A distinction distinguishes two different propositional contributions; at this stage I am merely claiming this, but shortly I will be arguing for it. But the causal R/A distinction distinguishes two mechanisms that determine those propositional contributions. I am going to pay little attention to the causal side of things in the first half of the thesis. In this chapter, I will be concerned with the constitutive question. Before I go through what the answer is, I need to say something about the various subcategories of definite descriptions *within* those referential and attributional uses themselves.

2.2 Correctness and Completeness

There are at least two sub-categories of definite descriptions that we need to discuss to

prevent ourselves from tangling up separate problems.²²

	Referential				Attributive	
Correctness	Correct		Incorrect		Correct	
Completeness	Complete	Incomplete	Complete	Incomplete	Complete	Incomplete

Table 1: Subcategories of definite descriptions

The above table indicates that referential uses of definite descriptions can be either *correct* or *incorrect*. It indicates that those correct and incorrect referential definite descriptions can be either complete or incomplete. And it also indicates that attributive uses cannot be incorrect and, when they are correct, they are either complete or incomplete. But what does all this mean? The sub-category of correctness tracks the descriptive or predicative component of the definite description (i.e., the F component of ‘the F’). Whether a definite description is correct or incorrect depends on whether the object designated, if there is one, instantiates the properties attributed to it by the descriptive component of the definite description or not. For example, take 1 from the previous section (repeated here):

2. The murderer of Smith is insane

I can use 2 to talk about someone who is either in fact a murderer of Smith (this can be a referential or attributive use) or to talk about someone who in fact is not the murderer of Smith (this only arises in the referential case).²³ Suppose, as we did earlier, that Jones has been charged with Smith’s murder and has been placed on trial. Imagine that there is a discussion of Jones’ odd behaviour at the trial. We might sum up our impression of his behaviour by uttering 2. But now suppose that it turns out that Jones is innocent and never committed a murder in his life. The intuition here is that although, in some sense, we referred to Jones, we incorrectly described him and, therefore, our referential definite description is incorrect of him. This is a case involving *misdescriptions*. The category of completeness, on the other hand, tracks the *definiteness* (i.e., uniqueness) component of the definite description indicated by the definite article ‘the’. Whether a definite description is complete or not depends on whether the descriptive material in it uniquely singles out the designated object. An example of an incomplete definite description (that can be attributive or referential) is ‘the murder’. This definite description can be used to designate an object that is obviously not alone in being a murderer. However, a definite description such as ‘the short murderer of Smith’ may very well be a complete definite

²² There is, in fact, at least one more sub-category which we will call ‘empty / non-empty’. This is a subcategory of definite descriptions which tracks whether the definite description denotes some object in the first place or not. This can be a referential or attributive definite description. This sub-category will not be very important for our purposes.

²³ In Donnellan’s words, when I use a description attributively, I use it to talk about *whoever or whatever* fits the description. If some object fits the description, then it must be correct of them.

description that doesn't just apply to that short murderer in question but applies to them uniquely. Now the point of all this is to emphasise that one can motivate Donnellan's constitutive R/A distinction using any one of the following combinations of types of definite description:

- i. Definite description uses that are **complete and correct**
- ii. Definite description uses that are **incomplete**
- iii. Definite description uses that are **incorrect**

The R/A constitutive distinction, as I will end up conceiving it, is a distinction between referential definite descriptions that are used to express object-dependent propositions and attributive definite descriptions that are used to express object-independent propositions. There is a route or argument from each of the above categories of definite descriptions to an R/A constitutive distinction. As is well known, Donnellan focused on cases of type (iii); on incorrect uses of definite description or misdescriptions to motivate his R/A constitutive distinction. Others, like Wettstein (1981), focus on cases of type (ii); cases of incomplete definite descriptions to motivate the R/A constitutive distinction. Now, the route or argument from misdescriptions to the R/A distinction is perhaps the one that is mostly associated with Donnellan's R/A distinction. I am going to set aside the misdescription route to the R/A distinction until the last chapter in the thesis, but only to explicitly distinguish it from the route I will end up pursuing. The route I will pursue is through type (i) cases; cases of complete and correct uses of definite descriptions. I am also going to completely ignore the route to the R/A distinction from incomplete uses of definite description. In the remainder of this section, I want to see how one can motivate Donnellan's R/A distinction while ignoring issues of completeness and correctness. That is, let us only consider definite descriptions that are *assumed* to be complete and correct and ask ourselves the following question: if we restrict our attention to this category of definite descriptions, how, if at all, could Donnellan have motivated his constitutive R/A distinction? This is my task in the next two sections.

2.3 Donnellan's Attributive Uses

Recall that Russell analysed definite descriptions as quantificational. Quantificational definite descriptions are terms that merely specify the quantity of something (in a given domain) as having certain properties. Is that what a definite description used attributively is, or is it a name or demonstrative? Donnellan did not actually commit himself regarding whether those attributive uses can be treated in a Russellian way or not, but most writers take it for granted that those uses are correctly analysed as Russellian (or in the spirit of the Russellian analysis, given its essential object-independent feature). The question we are asking in this section is whether the

propositional contribution of an attributively used definite description is, as Russell would have it, a condition in which no object features as a main constituent. There are plenty of uses or constructions that favour this Russellian analysis, but the question here is whether those attributive uses are Russellian. There is, as far as I can see, one reason to think that those attributive uses are in fact Russellian. It can be brought out by looking at Table 1 above. Note that attributive uses cannot be incorrect, just like Russellian definite descriptions. I think what explains this is that they both share the same analysis (at least in its essential object-independent feature). That is, the reason why they cannot be used incorrectly is that whatever the definite description designates is completely constrained by the descriptive material in the description. There is no possibility of conflicting designation – there is just one kind of designation; the satisfactoral kind. This is what I think Donnellan is getting at when he says that in the attributive case, the definite description occurs *essentially*, for the speaker wishes to assert something about whatever or whoever fits that description. Some writers have been uncharitable to Donnellan by pointing out that this notion of *whatever or whoever* does not distinguish referential from attributive uses. For example, Yagisawa says:

‘The speaker using ‘the murderer of Smith’ [referentially] in the courtroom example does wish to assert something about whoever fits the description. To wish to assert something about whatever or whoever fits a particular description is not at all incompatible with the intention to call attention to a particular object. If one believes that whatever or whoever fits a given description is in fact the object one intends to call attention to, then one can intend to call attention to that object and at the same time wish to assert something about whatever or whoever fits the description’. Yagisawa (1985, p. 111)

Of course, Yagisawa is right. But, as I said, I think what Donnellan is getting at here by talk of *whatever or whoever* is that in those attributive uses, there is only one kind of designation: the satisfactoral kind. But in those referential uses there is the satisfactoral kind of designation as well as the *speaker reference* kind of designation (more on that later) and those may conflict in cases of incorrectness or misdescription.²⁴ This understanding of *whatever and whoever* distinguishes attributive from referential uses and it *doesn't* distinguish attributive from Russellian uses, which is just what we need here. I am not going to have more to say about why those attributive uses are Russellian. Maybe they ultimately aren't, but I don't think that it matters, for it is a lot more interesting and challenging to pair those Russellian uses with referential uses for the purpose of this

²⁴ Donnellan (1966, p. 291) explicitly says, ‘When a definite description is used attributively, however, there is not the same possibility of misdescription’.

thesis. We certainly need to pair the Russellian uses with the referential uses in order to engage in the relevant debates.

2.4 Donnellan's Referential Uses

We have now identified Russellian uses with attributive uses. What about referential uses? How do they square with the Russellian analysis? In §2.2, we said that we were going to ignore issues of correctness and completeness, and I said that there is a route from complete and correct uses of definite descriptions to an R/A distinction. To anticipate, this route is going to be through the following argument:²⁵

- I. Some uses of definite descriptions are Russellian. [Assumption]
- II. Some uses of correct (and complete) definite descriptions can be used to express a proposition that is true in w iff the object designated in the actual world is G , regardless of whether it is F in w . [Claim]
- III. Those uses cannot be Russellian and must contribute an object to the proposition. [Claim / Best Explanation]
- IV. There is an R/A distinction for definite descriptions. [Conclusion]

The purpose of this section is to support claims II and III which will allow me later to naturally interpret referential uses as Fregean. Now Donnellan emphasises the point that a referential definite description is being used by the speaker to refer to some object o . Since Kripke's (1977) paper, 'Speaker's Reference and Semantic Reference', the phrase 'speaker reference' has come to carry certain pragmatic or non-semantic connotations. You might think that there is just no kind of reference (semantic, speaker or otherwise) that does not involve speakers; words (whatever they are) do not refer to things on their own. But even if that's what you think, we will still be able to distinguish two different kinds of reference that involve speakers. The first is the one *required* by the information encoded in certain terms for the determination of their propositional contribution – this is Neale's (2007, p. 359, n.7) metaphysical sense of *content determination*.²⁶ This is the kind of speaker-involving reference that the phrase 'semantic reference' is trying to isolate. The other kind of reference also involves speakers, but it serves a different purpose. Its

²⁵ I call these 'arguments' but they are more like *strategies* that involve reasonable moves rather than logical based inferences.

²⁶ Various people put this point in different ways. For example, Recanati (1993, p. 300) says that directly referential terms (that are only directly referential) are linguistically (i.e., pre-propositionally) *marked* as referential, that is, they are marked for reference at the level of *type* (the level of expression abstracted away from any particular context). That is, if their semantic value is not an object then they cannot be used appropriately. Stokke (2008, p. 403) says that the character of a directly referential term is a rule or function that maps certain occurrences of indexicals to an object. Figuratively speaking, he says that the character of (say) 'he' or 'now' *looks* for an intention to fix its reference. That is just part of their meaning (abstracted way from token utterances).

best formulation, I think, is given by Bach (2007, p. 518). This is supposed to be a four-place relation between speaker, object, term and audience, in which a speaker has object *o* in mind (or has thoughts about *o*) and uses a term to get their audience to think about *o* or to refer their audience to *o* – this relation is relevant for successful communication.²⁷ So although reference may always involve speakers, we can still distinguish between the semantically relevant kind (the former kind) and the not semantically relevant kind (the latter one). Moreover, these two kinds of reference usually accompany each other. These issues will be relevant later in Chapter 5; for now, I just want to emphasise that when I talk about speakers referring to something, this should not indicate that I have taken a position on the semantic *status* of the relevant term. I have not yet said anything regarding the semantic status of those referential (or attributive) uses. I remain neutral on this issue until I actually start discussing it in the next chapter. With this out of the way, take again the following example:

3. The murderer of Smith is insane

Suppose the definite description in 3 is used by the speaker to refer to Jones (the actual Russellian denotation of the definite description). In this case, as Donnellan would say, the definite description is being used to merely pick out an individual for their audience whom they have in mind. He adds, although not using the same terminology, that had the speaker used some other referential device instead (like a name or a demonstrative) they would have expressed the same proposition as the one they would have expressed by uttering 3. Recall that Donnellan (1966, p. 285-286) says that, in the referential use, the definite description is merely one tool for calling attention to a person or thing and any other device for doing the same job; another description or name would do as well. There is an underlying truth conditional intuition here (that I share) which is that 3 expresses a proposition that *would* be true in a counterfactual scenario in which Jones (the actual satisfier of the definite description) is insane regardless of whether he is F or not in that counterfactual scenario. Call this intuition the ‘counterfactual intuition’ or ‘CI’ for short. I claim that if a use of 3 is staged in the referential way, then we would get this CI. One might wonder whether we have intuitions about the truth value of propositions in different possible worlds. I think we do; we just don’t express these intuitions in those terms. For example, suppose I say to my friend, ‘This priest is kind’. Thinking (mistakenly) that I am attributing kindness to the priest in question because of him being a priest, my friend (who knows this priest well) responds by saying, ‘You know, he didn’t have to be a priest’. I then say, ‘Still, what I said would be true’. This sounds natural enough. It can be rephrased in our technical terminology as meaning that what I said

²⁷ Bach does not use the phrase ‘have in mind’. This is my addition, although still, he thinks that this second kind of speaker reference involves some psychological relation to this object *o*.

would be true in a counterfactual scenario in which the person we are actually talking about is not a priest but is kind in that scenario. Note also that this counterfactual intuition I am speaking of is *not* the intuition that Donnellan explicitly shares in his paper. Donnellan's original intuition is that an expression like 'The murderer of Smith is insane' (with the definite description used referentially) is true even if Jones, the person referred to by the definite description, is not *actually* a murderer. I do not want to discuss this controversial claim relating to misdescriptions. I am supposed to be carving out a misdescription-free pathway route to the R/A distinction. But you might now also say that a CI accompanies attributive uses just as well as those referential uses, so I cannot appeal to CI to distinguish referential and attributive uses. But this cannot be true. There are uses of definite descriptions in which you are forced to deny having a CI, and those have to be referential. Take 4 below, disambiguated with the scopes in 5:

4. The murderer of Smith might have been Atheer
5. There exists some possible world in which the unique murderer of Smith is Atheer

5 seems true, but if you agree with this (which I think you should), then you could not have that CI in relation to the truth conditions of 5. This is because if Jones is the actual satisfier of the definite description in 5 then surely 5 will be false; Atheer may never be Jones. The definite description used in 5 must be Russellian, and recall that in the previous section we more or less stipulated that Russellian uses are identical to attributive uses. Of course, this reading in which we don't get a CI is no problem for Russell; this reading confirms his theory. The reading is just that *wide scope* reading in which the Russellian definite description is not modified by the modal operator (more on that shortly). But no issues of scope arise in 3 in which the definite description is used referentially although we have that CI – this is the challenge for Russell.²⁸ So let us take CI as distinguishing the referential from the attributive uses. Now, let us ask: *how come we have this counterfactual intuition in the case of 3?* I think we would only have this intuition if:

- a) The definite description in 3 is Russellian but contained some rigidifying term like 'actual', or
- b) The definite description in 3 is Russellian, but somehow syntactically took *widest* scope, or

²⁸ Take the expression 'the murder is a murderer'. If definite descriptions can really sometimes contribute objects to propositions, then this expression can express a contingently true proposition. Does this reading intuitively exist? Admittedly, the scenario in which this expression will be an appropriate thing to say is hard to stage and so I cannot bring out this contingency intuitively. In any case, my point here will be comparative. If, as some people think, the analogous expression involving a demonstrative (i.e., 'that murderer is murderer') can be uttered falsely, then I see no barrier to making the analogous claim in the case of definite descriptions, especially if we have that CI.

- c) The definite description contributed an object (namely Jones) to the proposition that 3 expresses. That is, if 3 expresses an object-dependent proposition with Jones as a main constituent.

Why? Well, I already discussed what happens if (c) is the case. Object-dependent propositions, I said in the previous chapter, are constituted of the object referred to in the world of utterance. An object-dependent proposition requires the location of that very object in every counterfactual scenario for truth evaluation. And so if the propositional contribution of the definite description in 3 is Jones (the person who actually satisfies the definite description), then what is expressed by 3 *would* be true in a counterfactual scenario in which (say) there are no murderers but Jones is insane. Therefore, we would have captured CI. That is, Jones himself (and not any of his properties) enters into the structured proposition expressed by 1. Now, this same evaluatory effect is achieved if (b) is the case. That is, if the definite description took widest scope. Let me quickly introduce this concept because we will need it later. Russell (1905, pp. 489-490) showed us via the notion of scope that the following expression is ambiguous:²⁹

6. Atheer asserted that the murderer of Smith is insane

This is because the definite description may take wide or narrow scope in relation to some scope inducing operator (like **[assert]** in this case). 6 permits two readings:

- 6a. Atheer **asserted** that the murderer of Smith is insane
6b. The murderer of Smith is such that Atheer **asserted** of him that he is insane

When I say that a definite description *takes wide scope relative to a propositional attitude verb (or operator)*, I mean that the definite description is not modified by that verb; when I say that it *takes narrow scope relative to a propositional attitude verb*, I mean that it is. 6a and 6b are two different syntactic arrangements of 6. In 6a, the definite description is modified by the propositional attitude verb and in 6b it is not. In 6a, the speaker expresses the proposition that *Atheer asserted that the murderer is insane*. Contrastingly, in 6b, the speaker expresses the proposition that *the murderer of Smith is such that Atheer asserted of him that he is insane*. The effect of this syntactic arrangement in which the definite description takes wide scope (in 6b) is that what is relevant for truth evaluation in any possible world is fixed in the actual world. It is the actual denotation of the definite description, namely Jones. This is unlike the syntactic arrangement in 6a, in which what is relevant for truth evaluation is merely whether some object in that world of evaluation is indeed the

²⁹ Russell (1905, p. 489) called these 'primary' and 'secondary' occurrences of definite descriptions.

murderer of Smith and it is insane. However, if we want to pursue this approach in explaining our CI we confront a problem. The problem is that there is no visible scope-inducing operator in 3, and even if there was, although Jones would be relevant for truth-evaluation, he would be relevant in the wrong way. In 3, the expression is about Jones being insane and not about Jones being such that (say) Ather asserted of him that he is insane. Inducing scope does not come free, it brings its own truth conditional contributions. 3 is true iff Jones is insane, but 6b is true iff Jones is such that Ather asserted of him that he is insane! It might appear then that ‘option’ (b) is not even an option to begin with, so why did I include it? I included it because there is a somewhat elusive proposal by Mendelsohn (2008) under which the referential and attributive uses are simply the result of a syntactic ambiguity involving a covert but (apparently) truth-conditionally inert hyperintensional operator of assertion. I think his proposal has promise, but it needs to be justified properly and modified. I return to this much later on in §7.2.

Finally, option (a) seems to also be out, since there are no overt rigidifying terms in the description, such as, for example, ‘actual’. But perhaps the description contains a covert rigidifying term, and so 3 (with the definite description used referentially) is really an abbreviation for the following (which involves a rigid Russellian definite description):

7. The *actual* murderer of Smith

Regardless of the plausibility of this proposal, our problem still persists, for now we need to say whether definite descriptions are ambiguous between their rigid and non-rigid variants. In some sense, then, we haven’t achieved anything.³⁰ By process of elimination, then, and given that we are going to ignore Mendelsohn’s syntactic proposal until much later, we are now left with option (c). That is, we are left with the option that 3 must express an object-dependent proposition in order to account for our CI; according to which 1 *would* be true in a possible world *w* in which Jones (the satisfier of the definite description in the actual world) is insane, even though he is not the murderer of Smith in *w*. In other words, we are left with the option that referential definite descriptions are Fregean – they contribute the very object that is the satisfier or denotation of the definite description in a certain context. But, as I hinted at the end of the previous chapter, the Fregean too is challenged by Donnellan’s observations. His challenge however are those attributive uses of definite descriptions. This is because in the classic attributive scenario in which the detective utters 8 below:

³⁰ However, this proposal preserves the object-independent feature of Russell’s analysis and so it might have more promise. Perhaps we may say that there are contexts (which Donnellan calls ‘referential’) in which definite descriptions are accompanied by an actuality operator and others (which Donnellan calls ‘attributive’) where that is not so. Unfortunately, I won’t have the space to do justice to this proposal and, therefore, I shall ignore it.

8. The murderer of Smith is insane

the intuition is that the utterer would express something that is true in a counterfactual scenario in which John, Michael, James or anyone else is insane so long as they are unique in being a murderer. The question, then, is as follows: since 3 can make different propositional contributions (the very object which satisfies some condition vs. that condition), doesn't that mean that definite descriptions are ambiguous? This claim assumes something like this: if the same term can be used in different contexts to make different propositional contributions then that term is ambiguous. But this conditional is false. Counterexamples are terms like indexicals (such as 'I'). These terms make different propositional contributions depending on who is uttering them in context. But there is a reason why some people (Neale, 1990; Elbourne, 2013) don't think that 'I' and other indexicals are ambiguous. First, both of these terms make the same *kind* of propositional contribution: an object. And second, that object is determined by virtue of it satisfying a certain uniform descriptive condition in all contexts of utterance. That is, in all their uses, indexicals can only semantically contribute an object that fits a certain semantic profile (in the case of 'he' for example, it is the object that is uniquely male, in some context). And if meaning is something like a rule that determines propositional contribution (a character) then there is no reason to think that indexicals are ambiguous – they utilise the same rule in each context to output the same kind of semantic value. But Donnellan's claim is that definite descriptions make different *kinds* of propositional contributions, so you would think that in each case, there is a different rule involved; a different character. If such a claim is true, then, as I have assumed in Chapter 1, this is sufficient for ambiguity. So, is there an ambiguity? At this point, it is too early to tell. One might think that we were too quick on Russell (and Frege of course). That is, can the Russellian or the Fregean really not resist this duality of propositions while still accounting for our CI? Can they each really not absorb the other use? This is what I will explore in the next chapter. Before I do that, I want to re-hash, in short form, the first version of the argument from non-misdescriptions to the R/A distinction. I say 'short version' because we will be adding more premises to it later.

- I. Some uses of definite descriptions are Russellian. [Assumption]
- II. Some uses of correct (and complete) definite descriptions can be used to express a proposition that is true in w iff the object designated by the 'the F' in the actual world is G regardless of whether it is F in w . [Premise from CI]
- III. Those uses cannot be Russellian and must contribute an object to the proposition. [Premise II, Best Explanation]
- IV. There is an R/A distinction for definite descriptions. [Conclusion]

Note this argument is *silent* on the status of the R/A distinction. That is, it says nothing about whether those referential and attributive uses are semantic or pragmatic. But this is what we shall turn to next.

Chapter 3: Absorbing the Distinction Pragmatically

In this chapter, I want to explore how one may pragmatically absorb Donnellan's distinction and then talk about the significance of these pragmatic accounts in resolving claims of ambiguity of the sort that have been so far suggested. The upshot of this chapter will be that the necessary and sufficient conditions for conversational implicature on their own do not provide a compelling reason for thinking that referential uses arise due to conversational implicature. And, in fact, availing ourselves of Gricean conversational implicature denies us of any principled grounds for favouring the Russellian analysis over the Fregean analysis of definite descriptions.

3.1 Neale's Basic Case

I'll start with Neale's popular pragmatic account of referential uses. Neale (1990, pp. 65-102) defends a unitary Russellian analysis of the semantics of definite descriptions. In discussing referential uses, he distinguishes between what he calls 'basic cases' and 'non-basic cases'. I am going to ignore these non-basic cases because these are cases involving issues of correctness. In this chapter, I will be concerned with the basic cases which more or less involve referential uses of definite descriptions that are correct (and complete). Neale (1990, p. 65) says that in those referential uses involving no incorrectness or misdescription such as 1 below:

1. The murderer of Smith is insane

a referentialist needs to say that whatever is expressed by 1 is true iff it is true on the Russellian interpretation too. That is, whatever proposition is expressed by 1 is true iff *there exists one and only one thing that is a murderer of Smith and they are insane*. But that, of course, is just false iff referential definite descriptions are identified with Fregean definite descriptions. Under the Fregean analysis, the condition of there existing only one thing that is a murderer of Smith is only relevant to determining which object goes into the proposition. Once this has happened, it is the very denotation (or object) that becomes relevant for truth-evaluation. Neale adds that unless the referentialist can provide an argument that demonstrates beyond any doubt that one must entertain an object-dependent proposition about some object *o* in order to grasp the proposition expressed by 1, she is forced to argue for referential uses using other cases, perhaps cases involving incorrectness or incompleteness. Neale's observation here is connected to a reason that Russell (1919, p. 171) provides against the view that indefinites like 'a woman' are referential. Suppose that I met some woman who happens to be Lilly. Russell says that it

is clear that when I utter, 'I met a woman' my audience would *not* be required to identify Lilly if they were to understand my utterance. Similarly, Neale is saying that in uttering 1 the audience is not required to grasp an object-dependent proposition. But my argument from non-misdescriptions is a reason to think just that; that one must entertain an object-dependent proposition about object *o* in order to grasp the proposition expressed by 1. Given our CI, if the audience does not entertain that object-dependent proposition associated with referentially used definite descriptions then they have not understood the speakers' utterance.³¹ Under the Fregean interpretation, Russell's existence and uniqueness condition are only required to identify an object, and then that object itself goes into the proposition. So unfortunately, Neale's and Russell's point about understanding does not serve the debate very much. In any case, Neale still tries to give a detailed pragmatic account of how a speaker may use a definite description to communicate an object-dependent proposition. He thinks that definite descriptions only have one semantic analysis: a Russellian one (or an attributive one; remember that for us, these are identical). But sometimes, he thinks one might use 2 below to conversationally implicate an object-dependent proposition and not semantically express it:

2. The murderer of Smith is insane

That is, we use 2 (which just semantically expresses an object-independent proposition) to merely conversationally implicate an object-dependent proposition. I assume some elementary understanding of Grice's mechanism of conversational implicature, however I will just say a few words to refresh the readers' memory. Grice suggested that there is an underlying principle that determines the way in which language is used (effectively and with maximum efficiency) to achieve rational interaction in communication. He called this principle the 'co-operative principle' which he divided into nine *maxims of conversation* classified into four categories: Quality, Quantity, Relation and Manner. The co-operative principle and its component maxims ensure that in an exchange of conversation, the right amount of information is provided and the interaction is conducted in a truthful, relevant and perspicuous manner (Huang, 2007). Given this, we are able to use words to convey information that is not linguistically encoded in those words because we rely on the common knowledge (between speaker and audience) of that co-operative principle and maxims of conversation. In a nutshell, conversational implicature describes the phenomenon by which we use words with a certain meaning *p* to arrive at a meaning *q* via a certain two-way rationality-based procedure. In those interesting cases, the proposition that is conversationally implicated *q* is that proposition

³¹ Just like if one utters in the appropriate context 'that taxi driver is going fast'; the audience is required to understand an object-dependent proposition in order to fully understand the utterance. I say 'fully' because recall in the first chapter that we distinguished between two ways of understanding. One way of understanding was relevant to what we called the 'character', and the other was relevant to what we called the 'proposition'.

that the audience must assume the speaker to think (and intend to convey) over and above *p* in order to maintain the assumption that the speaker is being cooperative and, therefore, observing the maxims of conversation. For example, we might on some occasion utter 'war is war'. Confronted with this blatant breaking of the maxim of *quantity*, the addressee assumes that the speaker is actually co-operative and has to work out why he or she has made such an apparently uninformative utterance. The only way to do this is to interpret it as, in fact, highly informative. As an example, this utterance may conversationally implicate something like 'terrible things always happen in war'. This is a case that involves flouting a maxim. But we could also conversationally implicate a proposition while *observing* the cooperative principle and maxims. Suppose one utters 'Michael Jackson is dead'. Given the sub-maxim of *quality*, when someone utters this, they are taken to conversationally implicate that they believe it to be true.

3.2 Conversationally Implicating the Referential Use

Before getting into how the Russellian may explain away Donnellan's observations, let me get clear on what the dialectic is supposed to be. Neale thinks there is plenty of data that requires the positing of a Russellian analysis. But this is something that a Donnellian welcomes given their acceptance of the R/A distinction and the associated duality of use. Neale would also probably say that a CI also accompanies uses of other quantifier phrases like indefinites and universals. He will then go on to point out that it is very implausible to think that other quantifier phrases are also ambiguous between attributive and referential readings.³² But note that this is something the Donnellian is also happy to accept.³³ But now Neale will put forward the following argument:

- I. Much data requires the positing of a Russellian analysis.³⁴
- II. Referential uses cannot be treated in a Russellian way because they violate Russell's object-independency feature.
- III. Other quantifiers can also be used referentially (i.e., they are also accompanied by that CI).
- IV. We already have at our disposal the mechanism of conversational implicature; it is independently motivated.
- V. There is a Gricean derivation for the referential use; it is also cancellable without contradiction.

³² Although, as is well known, a Donnellian will point out that we would only have CI with much more stage setting; see the *argument from convention* in Devitt (2004). It is suggested in more recent literature that this argument has been effective.

³³ Note also, that a Donnellian so far is merely committed to the R/A distinction and not its status. All that the Donnellian has argued for is that we can *use* definite descriptions in two ways.

³⁴ This premise is *not* saying that certain data requires a *unitary* Russellian analysis to be posited, just a Russellian analysis.

- VI. Methodological Principle: do not multiply meanings beyond necessity.
- VII. Conclusion: we should treat the referential use pragmatically.

This is an argument against the claim that those definite descriptions are ambiguous. But at this point, the Donnellian will say that she accepts all premises of the above argument except the conclusion. This is because she will say that premise (III) is in fact a reason to accept the view that the duality we observe is due to the semantics of the term involved (but they will be neutral on whether that is due to ambiguity or something else)! And premise V, although true, is (as I will go on to explain in §3.2a and §3.2b) not very compelling at all. Finally, the Donnellian will also point out that just as what are standardly conceived as descriptive or quantificational terms can be used referentially, what are standardly conceived as referential terms can be used descriptively or quantificationally. So Donnellan's observations are not just about a peculiar kind of pervasive use in certain quantifiers; they are about a certain systematic duality of two kinds of meanings (more on that in Chapter 4) in many terms that have been traditionally conceived as unitarily of one kind. So Donnellan will now provide the following update to the argument from non-misdescriptions to the R/A distinction in which he is committed to it being a semantic distinction:³⁵

- i. Some uses of definite descriptions must be Russellian. [Assumption]
- ii. Some uses of correct (and complete) definite descriptions can be used to express a proposition that is true in some arbitrary world w iff the actual Russellian denotation of 'The F' is G regardless of whether they are F in w . [Premise from CI]
- iii. Those uses cannot be Russellian and must contribute an object to the proposition. [Premise ii, Best Explanation]
- iv. There is an R/A distinction for definite descriptions. [Conclusion]
- v. Russellian ways of explaining the referential use via conversational implicature are not compelling. [Claim 1]
- vi. There is a systematic relation between these two meanings. Although they depend on the context, they are not arbitrary. [Claim 2]
- vii. Other terms, such as quantifiers and directly referential terms, exhibit the R/A distinction. [Claim 3]
- viii. The R/A distinction should receive a unified semantic treatment. [Conclusion].

Call this 'the argument to a semantic R/A distinction'. I will not go into what it is that distinguishes *semantics* and *pragmatics* but what I mean by a semantic R/A distinction is just that the duality we observe in definite descriptions arises due to the information

³⁵ Again, no claim about *ambiguity* is being made here.

encoded in the words themselves *and* whatever that information is sensitive to in the extra-linguistic environment. From here until the end of Chapter 6, everything I say will be a defence of *Claims 1, 2 and 3* of this argument. Let us start by defending Claim 1. To anticipate, the reasoning behind this claim is that the necessary and jointly sufficient conditions for conversational implicature being met, on their own, is not a compelling reason for thinking that Russellian definite descriptions conversationally implicate an object-dependent proposition. But, more importantly, the Russellian availing themselves of conversational implicature does not give them any principled grounds for accepting the Russellian theory of definite descriptions over the Fregean theory of definite descriptions. Ultimately, my point in Claim 1 is that the debate becomes just entrenched opinion, where the Russellian and Fregean can defend their preferred view using Gricean tools. But what we want isn't this kind of defensive move; we need a reason to opt for one theory in the first place.

Let us now look closer at Neale's (1990) claim that one can conversationally implicate an object-dependent proposition using a Russellian definite description. Take 3 below, analysed in the classical Russellian way:

3. The murderer of Smith is insane

Let us see how one may go about defending such a claim. One might simply claim that Grice's conditions for conversational implicature are satisfied in this case. With the help of Neale (Neale, 1992, pp. 527-529), we can extract a set of five necessary and sufficient conditions from Grice's discussion. These conditions can be used to determine whether a meaning counts as a case of conversational implicature (where 'p' stands for the meaning literally expressed and 'q' stands for the alleged conversationally implicated meaning):³⁶

- A. The speaker is presumed to be observing the above maxims, or at least the cooperative principle.
- B. The supposition that he is aware that, or thinks that, q is required in order to make the speakers' saying or making as if to say p (or doing so in those terms) consistent with this presumption.
- C. The speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that

³⁶ Note that the familiar non-detachability condition is not even necessary for conversational implicature. To illustrate, suppose that the participants of a conversation (who are surgeons) don't think that Sally (who is a dentist, and everyone knows her to be) has a very challenging or complicated job. In defence of her, I utter, 'Well, Sally is a professional who specialises in the diagnosis, prevention, and treatment of diseases and conditions of the oral cavity' thereby implicating that Sally indeed *does* have a challenging and complicated job. But replacing 'dentist' with the underlined segment (which is what 'dentist' means) will detach the implicature.

the supposition mentioned in B is required.

- D. A conversational implicature is cancellable either explicitly or contextually without contradiction.
- E. There exists a derivation procedure for q , one has to produce an account of how it could have arisen and why it is there.

I am going to tentatively grant it for Neale (or the Russellian), that conditions A, B, and C are satisfied. This is because there is not much that can be said to argue for them directly. These are conditions based on the speakers' intentions towards their audience and their audience recognising these intentions. And while in certain cases of conversational implicature (involving, say, sarcasm), conditions A – C really do seem to be satisfied, it is just not obvious in many cases of plausible and alleged conversational implicature whether they are. So, dialectically, they will just be asserted by the Russellian and denied by the Donnellian. Writers in the debate usually focus on condition E, but I will look at conditions D and E respectively.

3.2a Cancellability and its Significance

Condition E in the previous section requires that a conversational implicature be cancellable either explicitly or contextually, without contradiction. For if the conjunction of p and $\neg q$ is a non-contradiction then p does not semantically entail q . This is supposed to be interesting for our discussion because if proposition q is not semantically entailed by p , a necessary condition would have been met for the truth of the claim that q is no part of the information encoded in p . Therefore, a necessary condition would have been met for the truth of the claim that q is merely conversationally implicated by p . So now take 4 (the Russellian object-independent proposition) and 5 (the Fregean object-dependent proposition).

- 4. $\langle\langle$ 'The', murderer of Smith) insane \rangle
- 5. \langle Jones, is insane \rangle

Surely a speaker can assert 4 and deny 5 without contradiction (although she would have asserted a falsity). So the necessary condition in D is met for the Russellian who is claiming that the object-dependent proposition is merely conversationally implicated. But note that if Donnellan's counter claim is right (that 4 and 5 diverge in truth conditions) then this is precisely what we should expect. Donnellan motivated the semantic challenge to Russell by stating that 4 is what is sometimes expressed using a definite description. But Donnellan's observation would hardly be a challenge if 5 did not have different truth conditions to 4. So the condition of cancellability without contradiction is required to be met by the Russellian just as much as the Donnellian if

each of the their respective claims is true. The condition of cancellability being met does not, on its own, advantage one position over the other. The general point was made by Sadock (1978, pp. 293-295); cancelability does not distinguish cases of ambiguity from cases of univocality plus conversational implicature. If a term is alleged to have two different meanings, then it should always be possible to assert one and deny the other. If a term is alleged to be used to conversationally implicate a meaning other than its conventional meaning, it should always be possible to assert one meaning and deny the other.

But you might say this is not how the cancellability test is supposed to work. You might say that we should apply the cancellability test at sentence level from *within* the staged referential scenario (inside, say, the trial room in which we find Jones behaving strangely). Imagine that situation again and ask if you can utter 6 below without contradiction:

6. The murderer of Smith is insane, but Jones/this guy/he isn't

And, of course, 6 too is no contradiction. Maybe this is no surprise given all I have said trying to distinguish names, Russellian definite descriptions and Fregean definite descriptions. But note that here, too, we observe the same limitation pointed out by Sadock (1978). Cancelability without contradiction does not distinguish cases of alleged ambiguity from cases of alleged univocality plus possible conversational implicature; both views of definite descriptions (the Russellian and Donnellian) *predict* what we observe in 6. The Russellian will say that in 6, the second segment of the utterance proceeding the comma cancels the object-dependent proposition that would have been conversationally implicated had the speaker stopped immediately after '...insane'. And if Donnellan is right and definite descriptions are ambiguous, then we should expect the speaker, in a context, to always supplement his utterance with more words to help their audience work out what propositional contribution they intended. The observation that the expression 'The bank is high, but the riparian zone isn't' is not a contradiction is surely no threat to someone who claims that 'bank' can have multiple propositional contributions (due to ambiguity). So all that this cancellability test has shown is that definite descriptions cannot always mean (say) *Jones*; a claim no one is making. So again, condition D is met for the Russellian just as much as for the Donnellian; it does not advantage one position over the other. To be vindicated, both positions require condition D to be satisfied.

But someone might say that if definite descriptions can really make the relevant Russellian *and* Fregean propositional contributions, then the following should not be a contradiction:

7. *The F is G, but it's not the case that the F is G*

This is because each occurrence of the definite description may have a different propositional contribution. But 7 is surely a contradiction, so it will be said that Donnellan's claim can't be right. But this sort of observation/response is confused for take 8 below:

8. *I am happy but I am not happy*

which is clearly a contradiction, but one that is perfectly compatible with the widely held view that 'I' can have different propositional contributions in different contexts. It is true that 7 and 8 cannot be uttered truly in the same context (they are a contradiction) but this does not speak to the issue of the range of propositional contributions in different contexts or utterance situations of a given term, which is Donnellan's semantically relevant issue. Here, it becomes important to start separating *semantic entailment* from *logical entailment*. To check for logical entailment, we fixate on a world, and check if there is some interpretation of the non-logical terms that makes the first sentence true but the second (underlined sentence) false *in that* world.³⁷ But when checking for semantic entailment, we fix the proposition at the world of utterance and check for truth in all possible worlds. Observations of contradiction speak to what is logically entailed by what, but do not speak to what is semantically entailed by what. These observations push us to the core of what I take Donnellan's distinction to be. Designational terms like 'the F', 'that F', 'me' and 'he' encode certain information. Tests of contradiction at sentence level (like 7 and 8) are not sensitive to *what* role this information plays in content determination. Sometimes, that information itself goes into the proposition. Sometimes, that information goes into determining some particular object *o*, and then that object itself goes into the proposition. But condition D may not even be necessary for conversational implicature. Take 9 and 10 below:

9. That cop is nice.

10. There exists a cop that is nice.

³⁷ We also need to allow hypothetical variations in domain size, otherwise e.g., 'There are seventeen things' turns out to be a logical truth because it is actually true and contains no non-logical terms. Now, if we logically interpret 'the F' in 7 as a logical constant that stands for the thing that is uniquely F in the context of utterance then 7 is a logical contradiction, similarly with 8. But this has nothing to do with the range of propositional contributions that 'I' and 'The F' are alleged to make in various contexts

9 logically entails 10, so the conjunction of 9 and the negation of 10 is a contradiction, but surely 9 still conversationally implicates 10!³⁸ 10 must be assumed to be intended to be communicated by the speaker so as to preserve the assumption that they are being cooperative by observing the conversational maxims. So my conclusion about condition D is that it either does not advantage the Russellian over the Donnellian, or that it can be ignored because it is not even necessary. The purpose of this section is to emphasise the (perhaps already obvious point) that the object-dependent proposition being cancellable without contradiction is *not* a point in favour of *just* the Russellian position. Let us now look at derivation.

3.2b Derivation and its Significance

Let us now move to condition E regarding *derivation*. Take the following two passages from Grice in relation to this condition:

- ... the final test for the presence of conversational implicature has to be, as far as I see, a derivation of it. One has to produce an account of how it could have arisen and why it is there. And I am very much opposed to any kind of sloppy use of this philosophical tool, in which one does not fulfil this condition. (Grice, 1981, p. 187).
- The presence of a conversational implicature must be capable of being worked out; for even if it can be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature... (Grice, 1989, p. 31).

Here we see that, according to Grice, for a piece of communication to count as an instance of conversational implicature, there must exist a reasoning procedure between the speaker and hearer relating to how this proposition could have been communicated and why it is there. But before getting into Neale's account for how this derivation is supposed to go, I want to address Devitt's argument that there is something deeply wrong with relying on derivation from one meaning to another to explain away the referential use in

³⁸ Perhaps you might say that this point merely shows that some logical entailments are *degenerate* conversational implicatures. That is, they are cases that fit the definition of what conversational implicatures are but not because they only become available via the mechanism of conversational implicature, rather they are available by logical entailment. If so, then we may persevere the claim that the condition in D is necessary (in the non-degenerate cases). But, on the other hand, you may question whether 12 is indeed a logical entailment. You may say that the Fregean proposition is represented logically as ' Gb ', where b is the F. But since b is a non-logical expression; we can interpret it so it does not denote the F. But then ' Gb ' will not logically entail 'the F is G '. This is because there is a uniform interpretation of ' b ' on which it doesn't denote an F. Now I do not think that a Fregean definite description should be represented logically as a constant, but for now this doesn't matter.

a Gricean way. Devitt (2004, p. 284) is against explaining the referential use by merely providing a Gricean derivation for its origins. But although derivation is not the only necessary condition for conversational implicature (as Devitt seems to assume), let us go through his argument anyway, because I think it is flawed. He thinks this strategy (of using a Gricean derivation to explain away a certain use) is based on the following false claim:

C1. Where an utterance has a conventional meaning and we can derive a different speaker meaning from this conventional meaning using appropriate assumptions about the context and mind of the speaker, there is no need to suppose that this derived meaning is another conventional meaning.

Devitt claims that going further down C1's path leads to a contradiction. He invents a character that he calls 'the fundamentalist Gricean' who will go further down C1's path but arrives at a contradiction by arguing that

... there are no conventional meanings at all: it is pragmatics all the way down. Consider the time before there were languages. Then there were clearly no conventional meanings to feature in derivations of speaker meanings. Yet, presumably, there came a time when people spoke and others derived the speaker meanings of these utterances simply from appropriate assumptions about the context and the mind of the speaker (just as we still often have to in a foreign country). There came to be regularities in the associations of speaker meanings with sounds and, we are inclined to say, over time some of these 'caught on' and became the first conventional meanings of a language (Devitt and Sterelny 1999: 7.6). But our fundamentalist is not inclined to say this at all, he will say: 'There is no need to take these first meanings to be conventional because they can all be derived from assumptions about the context and other minds. We must always remember the word of Grice. His "Modified Occam's Razor" tells us "Senses are not to be multiplied beyond necessity" (1989: 47). I say unto you that no senses are necessary (Devitt, 2004, p. 284)

Let's try to make this more precise. C1 can be simplified as C2 below:

C2. If U has conventional meaning p, and q can be derived from p, then q is not another conventional meaning of U.

The assumption behind Devitt's argument is that language originated via a Gricean style system of speaker meaning. Devitt and Sterelny (1999, p. 155) state that '[our ancestors] grunted or gestured, *meaning something by* such actions. There was speaker meaning without conventional meaning. Over time the grunts and gestures caught on: linguistic conventions were born'. Take the term 'run' for example. Presumably (according to Devitt and Sterelny), at some point in the distant past, this was a mere grunt or sound with no conventional meaning. Hearers then tried to make sense of it, allegedly, via a Gricean style derivation procedure and arrived at something like *to move quickly in a particular direction*. This is its speaker meaning. Over time, that grunt (which we today write as 'run') started to regularly mean what people derived it to mean, and that meaning (which was merely its speaker meaning in the past) stuck as conventional. But our fundamentalist Gricean, according to Devitt, does not need to say anything like this; he can just deny that 'run' has a conventional meaning at all. The only meaning it has, our fundamentalist will claim, is what speakers use it to mean on a particular occasion. Why? Because he subscribes to C2 above and will therefore reason as follows:

- a) 'Run' has conventional meaning p
- b) p can be derived from q
- c) \therefore p is not conventional

which is clearly a contradiction. This is supposed to be why C2 is false, because it leads to a contradiction. But that's just wrong. The above reasoning is not licensed by C2; C2 only licenses the following reasoning:³⁹

- d) 'Run' has conventional meaning p
- e) q can be derived from p
- f) \therefore q is not conventional.

And this is not a contradiction. So I don't see what the reductio is supposed to be. This leads me to think that Devitt's issue is really about our principled grounds for accepting the derivation in one case and rejecting it in the other. That is, he might ask: *what principled basis does one have for accepting the derivation from one direction (i.e., the one in d – f) and not the other (i.e., the one in a – c)?* But here an obvious reply is available. The derivation involved in a – c is just not a conversational implicature derivation. The derivation involved in conversational implicature is a derivation *from* a conventional meaning (and not *to* it); this is just what it means for a conversational implicature derivation to be available. So Devitt's fundamentalist Gricean is an impossible (not a hypothetical)

³⁹ Note also that if Devitt is right, there are no conventional meanings for the Fundamentalist Gricean. And hence C2 is vacuously true: true because its antecedent is always unsatisfied.

creature. Devitt says that to see what is wrong with C2, we may go further down its path. But we are not going further down C1's path, we are going in its opposite direction. And this choice of direction is not arbitrary or merely stipulative. One of the aims of the theory of conversational implicature is to separate the meaning encoded in a sentence from what speakers use it to mean on different occasions. These grunts are not words and have no meaning in and of themselves, and while people can grunt to mean things via certain principles of rationality, this is not the phenomenon that conversational implicature was designed to capture. So despite what Devitt says, there is nothing principally wrong with the mere availability of a derivation from p to q counting against the view that q is conventional. Now, Neale offers the following derivation as a candidate for the kind that will be involved in conversationally implicating an object-dependent proposition using a Russellian definite description (Neale, 1990, pp. 89-90).

Derivation:

- a) S has expressed the proposition that $[\text{the } x: Fx](Gx)$.⁴⁰
- b) There is no reason to suppose that S is not observing the CP and maxims.
- c) S could not be doing this unless he thought that Gb .⁴¹ Gloss: on the assumption that S is observing the maxim of relation, he must be attempting to convey something over and above the object-independent proposition that whoever is uniquely F is G . On the assumption that S is adhering to the maxim of quality, he must have adequate evidence to think that what we are willing to entertain as the F is G . It is not plausible to suppose that he has just general grounds for this belief; therefore, he must have object-dependent grounds. I can see that there is someone b in the perceptual environment who I and the speaker are willing to entertain as the unique F . Therefore, the grounds for his assertion that the F is G are plausibly furnished by the belief that Gb .
- d) S knows that I am willing to entertain that b is the F , that I know that S is willing to entertain that b is the F , and that I can see that S thinks the supposition that he thinks that Gb is required.
- e) S has done nothing to stop me thinking that Gb .
- f) S intends me to think, or is at least willing to allow me to think, that Gb .
- g) And so, S has implicated that Gb .

⁴⁰ This is just a semi-logical way of representing the object-independent proposition, the one that is expressed using a Russellian definite description.

⁴¹ Again, this is just a semi-logical way of representing the object-dependent proposition, the one that is expressed using a Fregean definite description.

First, that the possibility of a derivation along the above lines can be identified is not so impressive. Pick any term that is alleged to have two meanings, stipulate (as Neale does due to his pre-existing commitments about the semantics of definite descriptions) that one is the conventional one, put it in a context (or stage it in a way) that naturally permits the other one, and the audience will be forced to reason (via a derivation like the above) to that other meaning via some maxim (e.g., relevance or truth). You could do that with 'bank' if you weren't committed to anything about its lexical ambiguity. You could do that with dead metaphors (terms that now conventionally mean what they, in the past, were only used to conversationally implicate) and you can do that with polysemous terms (which are arguably ambiguous terms). Moreover, since the meanings associated with these terms are cancellable without contradiction, the Gricean tool of conversational implicature provides us with a swift means of showing that what is going on in all these cases is simply conversational implicature. But I don't think anybody is going to seriously think that just because the meanings associated with these terms are cancellable and calculable, we have sufficient reason to think that duality of meaning associated with these kinds of terms is a result of conversational implicature.

So given that 1) we merely assumed Grice's conditions A-C are met without argument, 2) cancellability (on its own) is either not a necessary condition for conversational implicature or it confirms the Russellian *as well as* the Donnellian thesis and 3) derivation comes too cheaply, what reason is there to think these object-dependent propositions are merely conversationally implicated? Well, as we saw, at the beginning of §3.2, Neale justified his strategy on the methodological grounds that a combination of unitary meaning and conversational implicature is always better than ambiguity. But the problem now is that once the Russellian avails herself of conversational implicature to capture those referential uses, she has availed herself of too much in a way that denies her any principled basis for favouring the Russellian theory of definite descriptions over the Fregean theory. That is, if we assume that definite descriptions are all Fregean, and that they contribute the object that uniquely satisfies the description (if there is one) to the proposition, then whenever we are confronted with a use that favours the Russellian interpretation, we will be able to say that those are merely conversationally implicated. And methodologically, our theory would stay just as simple and economic in such cases. Let me say more. Suppose we assume that definite descriptions are Fregean. This way, those referential uses of definite descriptions are captured straightforwardly. And those attributive uses can be captured, too, by showing that they are conversationally implicated from the Fregean meaning. Imagine that same attributive scenario in which

the detective stumbles across Smith's badly mutilated body, and take the following derivation from the Fregean meaning to the Russellian meaning:⁴²

- a') S is attempting to say *of* a contextually/historically salient object that is F that it is also G.
- b') There is no reason to suppose that S is not observing the CP and maxims.
- c') S could not be doing this unless he thought that [the x: Fx](Gx). Gloss: on the assumption that S is observing the maxim of relation, he must be attempting to convey something other than an object-dependent proposition because there is no contextually/historically salient object that is F in my immediate environment.⁴³ Also, on the assumption that S is adhering to the maxim of quality, since there is no contextually/historically salient object that is F, the speaker must be trying to communicate an object-independent proposition about whoever fits this description.
- d') S knows (and knows that I know that he knows) that I know that there is no b in my immediate contextual/historical environment that is F, and that I can see that S thinks the supposition that he thinks that [the x: Fx](Gx) is required.
- e') S has done nothing to stop me thinking that [the x: Fx](Gx).
- f') S intends me to think, or is at least willing to allow me to think, that [the x: Fx](Gx).
- g') And so, S has implicated that [the x: Fx](Gx).

You might say that this derivation does not even get off the ground, because in this case no proposition was actually conventionally expressed. But to say this would be to place an unfair constraint on the Fregean. Surely, I can conversationally implicate a proposition without having expressed a proposition. An example will illustrate the point. Suppose you and I are in the middle of a conversation and as I am asking you a question, you seem to have become distracted. You notice my reaction to your distracted state, and ask, 'Sorry, are you talking to me?' I reply sarcastically, 'No, I am talking to him', while gesturing to the wall behind you. Since there is no male person behind you, you realise that I am being sarcastic and say, 'Sorry, I just zoned out for a minute'. Here, the audience would have implicated something like, 'C'mon, of course I am talking to you, there is no one else here'. The reason why I think it is possible to implicate a proposition from an utterance that does not express a proposition is that the audience relies on their

⁴² Note, there are two possible styles of derivations depending on whether the definite description has a semantic value or not (i.e., whether it has a propositional contribution or not). I present the derivation in which no proposition is expressed by the Fregean definite description. The other kind is also available along similar lines.

⁴³ This is just one candidate gloss, but there are others that don't depend on there being no salient individual that satisfies the definite description. For example, suppose that after seeing Jones in the trial room and thinking that anybody in his position would have done the same, I say, 'The murderer of Smith might have been me'. The Fregean will say that this expresses the object-dependent proposition that Jones might have been Ather. But since this is clearly false, the audience will realise that the speaker intends to communicate the object-independent proposition according to which the definite description takes narrow scope relative to the modal operator.

knowledge of the character of the utterance. That is, the speaker relies on the audience's knowledge of the *type truth-conditions* of the utterance, and the hearer recognises this. I do not mean to suggest that one can linguistically implicate something without uttering anything; that would be absurd. I just think it is very plausible that just as all sorts of facts about the external environment can be an informational source for the audience to draw from in arriving at an implicature (e.g., temperature, birds chirping, etc.) surely the character of an expression is relevant too; otherwise, how could we make sense of the above example?

But there may be a bigger worry with this Fregean strategy. It may be pointed out that those Russellian meanings that are allegedly conversationally implicated from the Fregean meaning are also logically entailed by the Fregean meanings, and so they will not be cancellable without contradiction. However, as I pointed out in the previous section, conversational implicatures can also be logical entailments. And although some logical entailments are degenerate cases of conversational implicature the above isn't. This is because although the Russellian meaning, one might say, can be derived logically from the Fregean meaning, it is only with the aid of conversational implicature that it becomes truth-conditionally relevant. The Fregean analysis of definite descriptions under which the Russellian meanings are merely conversationally implicated is a plausible candidate. Now, the point of all this is to emphasise that once we recognise that there is a parallel defensive tactic available for the Fregean under which those Russellian meanings are conversationally implicated, then what principled grounds does the Russellian have for favouring one theory over the other? None! Of course, this might not be worrisome; you might ask, 'how can two successes be considered a failure?' That is, if two theories can handle the same equally well, how does this count against either of the theories? But note that I am not saying this is a failure. I merely want to conclude that in following the path of conversational implicature in Neale's strategy to its end, we reach a deadlock. I take this to be sufficient reason to momentarily set aside this sort of discussion and explore the distinction further and wherever else it appears. As I said at the start of §3.2, Donnellan's observations do not just describe a peculiar kind of pervasive referential use in certain quantifiers. Rather, they describe a certain systematic duality of two kinds of meanings in many terms that have been traditionally conceived as unitarily of one kind. This chapter has been a defence of *Claim 1* of Donnellan's counterargument (on page 39). I now move to *Claim 2* and 3. If the distinction is quite and systematic and pervasive as these claims suggest, then I take it that a semantically unified way of capturing it will be welcomed and therefore Donnellan's counter conclusion is justified. Let us now turn to this alleged systematicity in *Claim 2*.

Chapter 4: The Systematicity of the R/A Constitutive Distinction

We have seen how the R/A distinction plays out in the case of definite descriptions. We have also seen that when a Russellian points out that Grice's conditions for conversational implicature are satisfied, this does not (on its own) provide a compelling reason to think that those referential uses are conversationally implicated. Moreover, we have seen that where Russellians avail themselves of conversational implicature, they deny themselves a principled basis for favouring the Russellian theory of definite descriptions over the Fregean theory of definite descriptions. This was concerning since Neale's strategy was supposed to give us reason to opt for *one* theory in the first place. I now want to set aside this debate regarding whether or not Donnellan's observations are a challenge to Russell (or Frege), and look more closely at the distinction itself. In this chapter, I want to note some interesting features of the R/A distinction. Due to limitations of space, I am going to use definite descriptions again to expose these features, but I expect that what I say is generalisable to the other kinds of terms that exhibit the R/A distinction. I will be surveying some of these in the next chapter.

We have observed that in the referential use of a definite description, the propositional contribution is an object that satisfies some property-based condition. Further, in the attributive uses of definite descriptions, the propositional contribution is that very same property-based condition that some object or other may or may not satisfy. Under both analyses, the condition of *there being exactly one F* is essential.⁴⁴ In the referential or Fregean use, this condition plays a role in determining an object as the propositional contribution (and then it is no longer relevant for truth evaluation). In the attributive or Russellian use, *that same condition* goes into the proposition. Mendelsohn (2008, pp. 169-170) writes that if the object-independent proposition *that the murderer of Smith is insane* is true, then the object-dependent proposition *that he is insane* will be true of the murderer of Smith. Conversely, if the object-dependent proposition *that he is insane* is true of the murderer of Smith, then the object-independent proposition *that the murderer of Smith is insane* will be true. He takes this observation as evidence that there is a logical dependence between the two uses. But if there is a logical dependence, it cannot be two-way logical entailment or logical equivalence. This is because although the truth of one sentence may guarantee the truth of the other, the falsity of one sentence doesn't guarantee the falsity of the other. To illustrate, take the following one-premise argument (subscript 'R' for Russellian analysis and subscript 'F' for Fregean analysis):

1. The_RF is G

⁴⁴ Of course, Donnellan claimed that in those referential uses, the description is *not* essential. But I am using 'essential' in a different sense, as will soon become clear.

Conc: The_F F is G

The first (and only) premise or sentence is true iff there exists some object x such that it is unique in being F and it is G . Can the first premise be true while the conclusion false? No, that is impossible! This is because the conclusion is true iff some object x that is uniquely F is also G . The truth of the first premise guarantees the truth of the conclusion. Now, take 2:

2. The_F F is G

Conc: The_R F is G

The first (and only premise) is true iff there exists some object x that it is unique in being F and it is G . Can the first premise be true while the conclusion is false? No, that is impossible! This is because the conclusion is true iff some object x such that it is uniquely F is also G . The truth of the first premise guarantees the truth of the conclusion in both arguments. In both arguments there is no valuation of the non-logical terms that makes the first premise true but the conclusion false. But this is not enough to say that premise 1 and 2 *logically* entail each other. This is because although the falsity of 2 guarantees the falsity of 1, it is not obvious that the falsity of 1 guarantees the falsity of 2; that would depend on whether 2 is truth-value gappy or false when the relevant condition is not satisfied, something I will not be discussing.⁴⁵ So to describe this systematicity, we need to move to something weaker than logical entailment.⁴⁶ Mendelsohn (2008, p. 171) metaphorically brings out this systematicity by saying that in one use (the attributive or Russellian use) the proposition expressed *includes* the procedure for picking out an object (i.e., a condition), and in the other use (the referential or Fregean use) the proposition includes only the output of that procedure (i.e., the object). In other words, there is a satisfactoriness condition involved in both uses. But in one use, the condition is a *truth condition*, and in the other it is merely a *designational condition* (a condition whose satisfaction is relevant to merely fixing the propositional contribution which, in the referential use, is the object itself if there is one). That is, in the attributive use, the condition is a truth condition in the semantic sense of what is relevant for truth evaluation in counterfactual scenarios. But in the referential use, the truth of the condition is relevant merely for determining which *object* is relevant for truth evaluation in counterfactual scenarios. This semantically relevant issue is not brought out at the level of the logic of

⁴⁵ Sentence 1 (the Russellian sentence) can be false if there exists no unique F or if there exists a unique F but it is not G . If the former, then sentence 2 (the Fregean sentence) maybe truth-value gappy.

⁴⁶ Of course, the logical relationship between these two meaning depends on how we logically treat the Fregean definite description. I already hinted in footnote 38 on page 43 that I don't think they are logical constants which means we cannot interpret them as objects that are *not* uniquely F . But as we shall now see, we are going to move to something weaker than logical entailment to describe this systematicity anyway.

these terms (as I emphasised in §3.2a, although we may say that both arguments are *truth preserving in every context of utterance* which might amount to a kind of validity, something that Stalnaker (1975, p. 271) calls a ‘reasonable inference’.⁴⁷ Recall from the previous chapters that this systematicity cannot be described as *semantic entailment* either. There is no relation of semantic entailment going from the object-dependent proposition to the object-independent proposition and vice versa. That is, utter 1 and 2, fixing an object-independent proposition with the first and an object-dependent proposition about *a* (the thing which is the actual unique F) with the second and consequently, a world where there is a unique F which is G but *a* is not G will make the proposition expressed by 1 true but that which is expressed by 2 false. Similarly, a world where *a* is G but there is no unique F will make the proposition expressed in 2 true but that expressed by 1 false.

If we turn to formal semantics, the systematicity of the relationship between these two kinds of meaning has also been recognised. The Fregean inspired denotation of ‘the’ is as follows (Heim & Kratzer, 1998, p. 75):

- $\lambda f : f \in D_{\langle e, t \rangle}$ and there is exactly one x such that $f(x) = 1$. The unique y such that $f(y) = 1$.

This is a function from some function denoted by some noun phrase like ‘murderer’ (an $\langle e, t \rangle$ type function to an object. This is the object that is in fact alone in being a murderer. So here, the definite article denotes the following function: $\langle \langle e, t \rangle, e \rangle$; a function from $\langle e, t \rangle$ type functions to entities. And the definite description as a whole denotes an entity and is, therefore, of type e . On the other hand, under the Russellian analysis, the denotation of ‘the’ is as follows:⁴⁸

- $\lambda f : f \in D_{\langle e, t \rangle} . [\lambda g \in D_{\langle e, t \rangle} . \text{there exists only one } x \text{ such that } f(x) = 1, g(x) = 1]$

This is a function from some function denoted by a noun phrase like ‘murderer’ to another function that is a function denoted by a predicate like ‘insane’ (an $\langle e, t \rangle$ type function) to truth values. The truth value will be generated depending on whether the above underlined condition is met or not. So here, the definite article denotes the following function: $\langle \langle e, t \rangle, \langle \langle e, t \rangle, t \rangle \rangle$ and the definite description as a whole denotes the

⁴⁷ Stalnaker (1975, p. 271) informally defines a *reasonable inference* as ‘...an inference from a sequence of assertions or suppositions (the premisses) to an assertion or hypothetical assertion (the conclusion) is reasonable just in case, in every context in which the premisses could appropriately be asserted or supposed, it is impossible for anyone to accept the premisses without committing himself to the conclusion...’ This is in contrast to his informal definition of *entailment* in which ‘... a set of propositions (the premisses) entails a proposition (the conclusion) just in case it is impossible for the premisses to be true without the conclusion being true as well.’

⁴⁸ This is just the function type of a generalised quantifier term like ‘every’, ‘no’, or ‘some’ found in Heim and Kratzer (1998, p. 146), although with a different constraint - a constraint about existence and uniqueness.

following type of function: $\langle\langle e, t \rangle, t \rangle$. Under both denotation types of the definite description (e or $\langle\langle e, t \rangle, t \rangle$) the same underlined condition of *there being exactly one x that $f(x) = 1$* plays a role. Under the Fregean analysis of 'the' the satisfaction of the condition is required to output the value of the function: an entity or object that goes into determining a truth value for the whole sentence. And under the Russellian analysis, that same condition is required for the outputting of a truth value for the whole sentence. In her paper about type-shifting principles, Partee (1986, p. 365) emphasises the close relationship between these two kinds of denotation types and introduces a function called 'lift' (which maps an object to the set of its unique properties) and an operation called 'lower' (which maps a set of properties to the unique object that has them, if there is one). She says these meanings (the Fregean and Russellian) are related to each other by the fact that whenever the Fregean function of the definite article has a value, then $lift(The_F F) = The_R F$ and $lower(The_R F) = The_F F$. Furthermore, whenever the Fregean function is valueless, $The_R F$ is vacuous in that it denotes the empty set of properties. This is on the formal semantics side.

On the propositional side, this systematicity is brought out by observing that in the actual world, the truth of the object-independent proposition expressed by an attributive used definite description *coincides* with the truth of the object-dependent proposition expressed by a referentially used definite description, and vice versa. But note that this coincidence, though it may not be logically based, is not accidental. There is nothing accidental about Jones being the denotation or referent of the condition *there exists exactly one murderer of Smith* if Jones is indeed unique in being the murderer of Smith! We may say that the truth values will always coincide in the following sense: for whatever property that is predicated and actually instantiated by the unique thing that is a murderer of Smith, it will also be instantiated by Jones. Now, this systematicity explains why both uses of definite descriptions seem (loosely speaking) *conventional, literal and straightforward, requiring no special stage-setting* as Devitt (2003, 2007) has claimed. That is, this sense that both uses are conventional is due to the fact that in both uses the condition of their existing exactly one F is relevant for truth evaluation, albeit, as I have been emphasising, in two different and subtle ways. But what is interesting about this systematicity I have described is that it is just not present in typical cases of lexical ambiguity. Terms that are lexically ambiguous typically contribute different things to the propositions they are used to express (even in one and the same context), but they do so without the sort of systematicity described above. For example, take the following expression and its disambiguation:

3. The bank is nice.
 - a. The financial institution is nice.
 - b. The riparian zone is nice.

There is no reason to think that whatever property (in this case the property of *being nice*) that is instantiated by the financial institution will also be instantiated by the riparian zone. In the actual world (a) may be true in context *c* but (b) might be false in that same context. The above is an example of a *homonymy* type ambiguity. The other type of lexical ambiguity is *polysemy*, exhibited in a term like ‘mouth’, which is allegedly ambiguous between different but very related meanings – the mouth of a human, the mouth of a bottle, the mouth of a cave, and so on. Based on Donnellan’s observations, Amaral (2008) has proposed that the definite article is an instance of polysemous lexical ambiguity. But the systematicity I have pointed to above tells against this kind of story too, for take the following utterance 4 and two of its disambiguations:

4. We can’t use it. The mouth is too small.
 - a. We can’t use it. The cave’s mouth is too small.
 - b. We can’t use it. The bottle’s mouth is too small.

There is no reason to think that in the same context of utterance whatever property (in this case the property of *being too small*) that is instantiated by the cave’s mouth will also be instantiated by the bottle’s mouth in the same context. In the same context, (a) may be true but (b) might be false. Of course, it might be said that the alleged ambiguity associated with definite descriptions is just a *very special kind* of lexical ambiguity. But if by this claim we mean something like indexicality, we can set aside the discussion until Chapter 7. But if we mean ‘typical lexical ambiguity but very different’ then proponents of the view that definite descriptions are lexically ambiguous (like Devitt, 2004) will be advancing an unattractive position. In fact, the systematicity I have described above is very suggestive of a syntactic ambiguity. I will be offering a proposal in Chapter 7, under which expressions containing definite descriptions (or other R/A terms) are internally syntactically ambiguous. This will be inspired by a proposal from Mendelsohn (2008); however, it does not postulate an *invisible* hyperintensional operator of *assertion*. For now, let us still just ignore this syntactic ambiguity proposal.

We have seen that the conditions of conversational implicature, on their own, don’t provide a compelling reason to think that one use is conversationally implicated in the case of definite description. Furthermore, these uses are systematically related in the above way, which is not typical of standard lexical ambiguity. Finally, the R/A distinction (as we shall see in the next chapter) is quite pervasive. With all this in mind, one might assume the need for a special semantic treatment, that is, a treatment under which those propositional contributions are directly predicted by the semantics of the terms. Such a treatment would not utilise lexical ambiguity or conversational implicature, but rather the information encoded in the terms themselves, which may be sensitive to aspects of

the extra-linguistic environment. Before I do this let me finally justify *Claim 3* regarding the pervasiveness of the R/A constitutive distinction.

Chapter 5: The Pervasiveness of the R/A Constitutive Distinction

In this chapter I intend to show that other terms (not just definite descriptions) exhibit the R/A distinction. I will first look at cases in which demonstratives seem to be used referentially and attributively. Then, I will briefly outline King's treatment of these cases to show how it *mimics* the effects of our R/A distinction and why I don't think there is a sharp contrast between definite descriptions and demonstratives as he claims. After that I turn to pronouns and point out more data in which the pronouns are not plausibly seen as contributing an object to the relevant proposition. I will then finish by distinguishing this data from other data that has been offered in support of the view that the R/A is pervasive.

5.1 Demonstratives

The more familiar uses of demonstratives are those that involve a kind of *speaker reference* to an object that typically is (or was) available in the perceptual environment. In such uses, the speaker, in some way, demonstrates this object (e.g., via pointing). Take the following use of 'that', in which the speaker is pointing at some object:

1. That cop is nice.

In this situation 1 is plausibly semantically associated with an object-dependent proposition: that proposition which contains the very cop that the speaker intends to refer to as a constituent, and the property of being nice. This proposition is true in some world w iff Nancy (the very cop in the actual world) is nice, regardless of whether she is a cop in w or not. This is a case involving a *referential* demonstrative. But there are uses that intuitively don't quite fit this picture, as King (2001, p. 9) points out. King calls them 'complex demonstratives' but I will just call them 'attributive uses of demonstratives'. Take 2, for example:

2. That hominid who discovered how to start fire was a genius.

2 can be uttered felicitously without the speaker having anyone in mind and, therefore, without any speaker reference or demonstration. These uses, according to King, involve *no demonstration* and *no speaker reference* (in the semantic or pragmatic sense identified in §2.3). In this situation, 2 is plausibly semantically associated with an object-independent proposition; that proposition which is constituted merely of a property-based condition. And this proposition is true in some arbitrary world w iff whoever is unique in being a

hominid who discovered fire in w is indeed a genius in w , regardless of whether that hominid is in fact the actual one or not. But there are more reasons to think that some uses of demonstratives cannot fit the classical, directly referential picture of demonstratives. Take the following utterance:

3. That woman with the highest salary is the most powerful.

What seems to be expressed by 3 is a rule-like expression; something along the lines of *for everything, if it's a woman with the highest salary, then it is the most powerful*. The truth of this expression does not depend on some particular woman (the same woman) for truth evaluation. Indeed, 3 can be true in a world in which there are no women. This is easily explained if the propositional contribution of the demonstrative is not an object, but some property-based condition. Take also the following modal embedding inspired by King (2001, pp. 91-92):

4. That last number I picked may have been odd!

King thinks that 4 has a false as well as a true reading, and that seems to be correct. The number itself (if it was even) would not have been odd, but I could have picked an odd instead of an even number. Suppose 4 is uttered after my audience and I discover that that last number I picked was in fact the number 8. Upon hearing someone say about me, *'See, Atheer was always going to lose,'* I angrily utter 4 in response, implying that I only lost due to bad luck and not due to some inherent mediocrity. Now here, although the audience and I know that the last number I picked was 8, it seems that 4 expresses a truth in this situation. But this is not what we should expect if the demonstrative in 4 is directly referential and contributes the number 8. If the demonstrative in 4 is directly referential, then 4 would be used to express *that it is possible that 8 is odd*, which is false! Hence, we will not get that plausible true reading we are after because directly referential terms are scope invariant. However, if here the demonstrative merely contributes a property-based condition, then the true reading can easily be captured. This is because there would now exist a narrow scope reading (the true reading) under which the demonstrative is within the scope of (and therefore modified by) the modal operator. This is the reading according to which 4 is uttered to express that *in some possible world, there exists a unique last number that Atheer picked, and that number is odd*. There will also be the reading according to which the demonstrative is outside (and therefore not modified by) the modal operator (the false reading).

As King also points out, some demonstrative uses (those that I am alleging to be attributive) exhibit the expected behaviour when embedded in hyperintensional environments. To see how this is, take a case also inspired by King (2001, pp. 2-8). Imagine

that Greg is telling someone that a student failed the exam (so, by stipulation, he asserts an object-independent proposition). Assume that Greg has come to know this by checking a system-generated report that does not disclose the identity of students. I overhear Greg's statement and I go to report it to my friend (my friend and I do not have the relevant object in mind):

5. Greg believes that that student who failed the exam is stupid.

It seems that 5 is true. But if the demonstrative contributed an object to the relevant proposition, then 5 should be false, because Greg does not believe the relevant object-dependent proposition (we stipulated this). But 5 being true is explained easily if what I am attributing to Greg, using the demonstrative, is the object-independent belief that there exists some unique student who failed the exam, and that student is stupid. Moreover, it seems that 5 will be true even if Martin is that student who failed the exam (unbeknownst to Greg), but is such that Greg thinks he is a genius. That is, it seems that 5 and 6 below will be true in this scenario:

6. Greg believes that Martin is a genius.

But how can that be on the directly referential account of demonstratives? According to this account, 5 and 6 attribute contradictory beliefs to Greg. If the demonstrative in 1 really did contribute Martin to the proposition, these observations (i.e., that 5 and 6 are true) cannot be explained. This is because directly referential terms contribute objects to the proposition and are scope invariant. These observations are easily explained, however, if the demonstrative contributes a property-based condition to the proposition. Moreover, there are also hyperintensional embeddings, like 7 below:

7. Greg believes that that student who failed the exam is not Martin.

Here, it does not seem to me that we are saying that Greg believes that Martin is not Martin. Here, I and the speaker know that Martin is that unique student who failed the exam (so we have him in mind), but I utter 7 to communicate to my audience that Greg doesn't know this. It seems that I can use 7 felicitously to attribute to Greg the belief that there exists a unique student who failed the exam, and it is not Martin. Now, I think the above examples indicate that there are attributive/non-referential uses of demonstratives. If so, then there is also an R/A distinction in demonstratives. In the next section I want to go through King's quantificational account of demonstratives that is intended to capture this duality of uses of demonstratives, which we have surveyed above. It is important to note at the outset that although King recognises that there are indeed two different uses of demonstratives, he offers a semantically unified (non-ambiguity-based) account of

these different uses under which they are both quantificational, although they are quantificational in a way that captures the truth-conditional differences we have so far been observing.

5.1a King's Theory of Complex Demonstratives

King (2001, pp. 24-25), who is working within a structured propositional framework, proposes that demonstratives like 'that' should receive a quantificational semantic analysis. Just like 'every', for example, 'that' is to be analysed as contributing to propositions a two-place relation between properties. For example, 'every' contributes to a structured proposition a relation **R** between the property of being a woman and being smart. Thus, the sentence 'every woman is smart' expresses a proposition that is true in some world iff the property expressed by 'woman' stands in the relation **R** (the relation expressed by 'every') to the property expressed by 'smart'. That is, it will be true in *w* iff every instance of the property of being a woman is an instance of the property of being smart in *w*. Secondly, King (2001, pp. 28, 150) claims that what is relevant to the semantics of 'that' phrases is that they allow supplementation by speaker intentions. These intentions, according to King, uniquely determine various properties which then figure in the semantics of such phrases. It is this feature (the feature of sensitivity to intentions) that King claims is common to all these uses of 'that' phrases (whether they are referential or attributive), and it is what distinguishes them from the other more familiar quantifiers. King distinguishes mainly between two types of intentions that allow 'that' phrases to function in a way very similar to definite descriptions used referentially and attributively function. These two intentions uniquely determine different properties, which then figure in the semantics of 'that' phrases. As we will come to see, King develops the semantics of demonstratives in a way that *mimics* the truth-conditional effects of our R/A distinction. Now, the first sort of intention is one that occurs in the familiar referential cases of demonstratives. King calls this intention a 'perceptual intention'. This intention uniquely determines a property – the property of *being identical to object o* (where *o* is the object of a current or past perceptual intention). This perceptual intention also determines another property, the property of *being jointly instantiated in w, t* King (2001, p. 42).⁴⁹ This kind of perceptual intention is exemplified in the following use of 'that', in which the speaker is observing some tree and wants to point it out to their audience:

8. That tree is old.

The demonstrative in 8 is referential; naturally it would involve a demonstration or speaker reference. But, according to King, there is another type of intention which occurs

⁴⁹ Where 'w' here stands for the world of utterance and 't' stands for the time of utterance.

when a speaker believes that something possesses certain properties uniquely and intends to say something about whoever or whatever possesses those properties. King calls this a ‘descriptive intention’. Given this descriptive intention, there is a property or conjunction of properties C such that the speaker intends to say something about whatever uniquely possesses C . This intention determines a unique property: *the property of the descriptive intention in question*. This descriptive intention also determines another property, the property of *being jointly instantiated*. This kind of intention is exemplified in the following use of ‘that’, in which the speaker wants to talk about whoever it is that is unique in being a hominid who discovered how to start fire. This will be an attributive use:

9. That hominid who discovered how to start fire was a genius.

Now, as I said at the outset of this section, ‘that’ phrases are (according to King) to be treated as quantifiers contributing a two-place relation to the structured propositions they are used to express. But this is true *only in the context* in which ‘that’ is used. However, according to King, ‘that’, taken outside of any context, expresses a four-place relation. The idea is that when a speaker uses a ‘that’ phrase in a particular context, his intentions determine properties that saturate these additional argument places in the relation expressed by ‘that’. This saturation results in a two-place relation between properties. This two-place relation between properties is then contributed to the proposition expressed in that context by the sentence containing the ‘that’ phrase. So, according to King, ‘that’ (when taken out of any context), expresses a relation with more than two places. But when used in a context, it always ends up contributing a two-place relation, just as other quantifiers do (2001, p. 31). The lexical meaning of ‘that’, then (outside of any context), is the following four-place relation:

10. $_1$ and $_2$ are uniquely $_3$ in an object x and x is $_4$.

But depending on which one of our two intentions are in place, different things will saturate the second and third argument places. Since each of our two intentions determine two properties, it is those two properties that fill the second and third argument places above. If our use is referential like the use in 8, then (after saturation of the two properties by the relevant perceptual intention) the ‘that’ phrase contributes the two-place relation in 11 (below) to the structured proposition. But if our use is attributive then (after saturation by the relevant descriptive intention properties) the ‘that’ phrase contributes the two-place relation in 12 (below) to the structured proposition:

11. $_1$ and = object o are uniquely jointly instantiated in w, t in an object x and x is $_4$.
12. $_1$ and = C are uniquely jointly instantiated in an object x and x is $_4$.

Further, the sentence as a whole (in the case of 12) expresses the proposition that *F and C are uniquely jointly instantiated in an object x and x is G*, or (in the case of 11), the proposition that *F and = object o are uniquely jointly instantiated in w,t in an object x and x is G*; where F and G are the properties expressed by the predicates 'F' and 'G' (King, 2001; pp. 45-46).⁵⁰

So, King (2001, p. 31) claims that the meaning of 'that' phrases when used either way (referentially or attributively), outside of any context, is 10 above. He adds that what is distinctive about the semantics of 'that' phrases is that they require extra-linguistic supplementation by intentions. The only thing that changes from one use to the next (be it a perceptual or descriptive demonstrative) is the sort of intention that is in place which determines different properties that saturate two of the four place relations. This analysis attempts to unify the two different uses of definite descriptions (those exemplified by 8 and 9 above). But, as I said at the beginning of this section, although these are two different uses, they are so only because there are two different kinds of intentions at play. Under both uses, the demonstrative is to be analysed as a quantifier. It is easy to see what King's proposal is supposed to achieve. It allows 'that' phrases to be rigid in certain contexts of utterance and non-rigid in others. When one's intention is perceptual, what is relevant for evaluation in some arbitrary possible world is object *o*, the object of my perception. King's theory achieves this by making the property of *being identical to object o in w, t* necessary for the truth evaluation of the relevant referential utterance. On the other hand, when one's intention is descriptive, what is relevant for truth evaluation in some arbitrary world is merely some property (like *being a guy who discovered fire*) which is not tied to the world and time of utterance. This allows the 'that' phrase to be non-rigid. It also allows the phrase to have flexible designation across worlds. So, we see that King develops the semantics of demonstratives in a way that *mimics* the truth-conditional effects of our R/A distinction. Our R/A distinction was given in terms of different kinds of propositional contributions, which achieves this rigidity/non-rigidity slightly differently, but his proposal is still in the spirit of our distinction. That is, although what saturates the second argument place in both uses is a property, the saturated property in those referential uses is the property of *being object o* (unlike attributive uses).

King wants to use this feature of *sensitivity to intentions* at utterance type level (the pre-propositional semantic level) to draw a sharp contrast between definite descriptions and demonstratives. He says (2001, p. 67) that if we assume that intentions are relevant to the semantics of 'that' phrases but not to definite descriptions, then we can explain why the phrase 'that F' can be used several times in a situation in which there are many Fs to talk

⁵⁰ Recall that 'C' is the conjunction of properties possessed by whatever object the speaker intends to talk about.

about distinct Fs, whereas 'the F' often cannot be used this way. Thus, looking at cars in a lot filled with new cars, one can say:

13. That car is nicer than that car

To talk about distinct cars by means of the distinct occurrences of 'that car'. In contrast, one could not felicitously say:

14. #The car is nicer than the car

With or without accompanying demonstrations. According to King, if demonstrations or speaker intentions are relevant to the semantics of 'that' phrases, different demonstrations or intentions associated with the distinct occurrences of 'that F' would allow them to be used to talk about distinct Fs. This would explain the felicity of 13. By contrast, the semantics of 'the', according to King, does not allow a role for demonstrations or speaker intentions. Thus, there is nothing that allows the distinct occurrences of 'the F' to be used to talk about different Fs, and so the attempt to do so results in infelicity (as in the case of 14). But I think King's claim that the felicity of 13 and the infelicity of 14 is somehow evidence that intentions are relevant for the semantics of 'that' and *irrelevant* for the semantics of definite descriptions is incorrect. For someone could maintain that the felicity of 13 and the infelicity of 14 just shows that intentions are relevant for both terms, but in different ways. That is, the meaning of these terms places different constraints on intentions, but both terms are sensitive to intentions but in different ways. For example, if we individuate definite descriptions (and demonstrative phrases) according to the information encoded in the accompanying noun phrase (e.g., 'car' in 'the/that car') then it may be that the semantics of definite descriptions are such that they only allow the supplementation of at least one (and at most one) intention per definite description. As such, the semantics of definite descriptions would be different to the semantics of 'that' phrases, which, following King, might allow multiple intentions for each 'that' phrase.

Once we recognise this possibility, we can maintain that it is this that explains the difference in felicity between 13 and 14, and not the insensitivity of definite descriptions to intentions. Moreover, this doesn't seem like a merely *possible* response, but a plausible one. For while both 'that F' and 'the F' allow the speaker to talk about distinct or unique Fs, the main difference in their semantics stems from the fact that the definite article signals the speaker's intention to exhaust the range of a certain predicate (Vendler, 1967,

p. 51).⁵¹ For example, in 14, the first definite description indicates that the designated object (if there is one) exhausts the predicate. This further constrains uses of the same definite description in the same sentence. If this is true, then it is the reason for the difference in felicity between 13 and 14, contrary to King's claims. This is in contrast with 'that' phrases, where it seems obvious that the object designated by a 'that' phrase leaves it open as to whether the object exhausts the predicate or not and, consequently, does not constrain further uses of 'that' phrases in the way that is apparent in definite descriptions. This is true regardless of whether the term is being referentially or attributively. So given that demonstratives and definite descriptions allow supplementation of intentions, and given that they both exhibit the R/A constitutive distinction, we can probably modify King's proposal and extend it to definite descriptions too. Doing this will not mean that definite descriptions and demonstratives have the same meaning. This is because each kind of term places a different contrast on the *number* of intention-saturated properties in context. But I do not intend to extend King's proposal to definite descriptions. In Chapter 7, I propose indexical and syntactic treatments for these terms. But let us now move to pronouns and see how they exhibit the R/A constitutive distinction.

5.2 Pronouns

The discussion about pronouns sometimes starts by pointing out that they can be used *deictically* and *anaphorically*. A pronoun is used deictically when it receives its reference from the extralinguistic utterance context, and it is used anaphorically when it picks up its reference from another phrase in the surrounding text. Paradigm cases of deictic uses are those that involve some sort of demonstration or speaker reference. But the term 'deictic' also applies to cases where the intended referent is sufficiently salient in the utterance context, even without the help of pointing (Heim & Kratzer, 1998, p. 239). For example, 15 below might be uttered immediately after a certain man has left the room:

15. I am glad he is gone.

Thus, in a deictic use or a referential use of the pronoun 'he', the speaker has the relevant object in mind, and she uses it to express the proposition that *she is glad that a particular object o is gone*. That is, the propositional contribution of 'he' in the above utterance is an object. But there are also anaphoric uses of pronouns. In 16 below, 'he' is anaphoric on 'Atheer', and in 17, 'she' is anaphoric on 'the/that murderer of Smith'. Some call these 'lazy pronouns' because they abbreviate some earlier (sometimes descriptively richer)

⁵¹ This is, according to Vendler, what unifies non-plural definite descriptions (of the sort we have mostly been dealing with in this thesis) and plural definite descriptions.

subject term. Subscript 'A/R' indicates that the definite description or demonstrative can be used attributively or referentially:

16. A_{A/R}theer was reading. He is asleep now.

17. The_{A/R}/that_{A/R} murderer of Smith is happy. She is insane.

But there are also those uses of pronouns in which a term like 'her' occurs as a *variable* bound by (say) a universal quantifier like 'every':

18. Every man put a screen in front of her.⁵²

Now, although they can be used to motivate the R/A distinction, I am going to put cases like 16 and 17 aside. Anaphoric uses of terms, you may think, are just a different phenomenon to the one I am investigating. You might say that anaphoric terms just inherit the semantic value (or propositional contribution) of a certain term that precedes them (i.e. the name in 16 and the definite description/demonstrative in 17, be it attributive or referential). Maybe definite descriptions and demonstratives themselves can function as anaphora. Take the following as examples:

19. A_{A/R}theer thinks he is a good philosopher. The guy is insane.

20. [In reference to a large pink diamond] Everyone touched it. That thing is mesmerising.

Moreover, cases like 18 above are also not typically seen as a threat to the directly referential analysis of pronouns. Perhaps this is because it was always only certain (and not *all*) uses of pronouns that were supposed to be directly referential.⁵³ Or perhaps you might say that none of the cases involving pronouns I have presented so far pose a threat to directly referential analysis, because all of these uses can be reconciled under the conception of pronouns as variables at the level of logical form or syntax. But although I think this claim is orthogonal, I shall put the cases in 16, 17 and 18 to one side. Instead, to begin with, the uses with which I want to contrast the deictic uses in 15 with are cases that (in King's words) don't involve any speaker references or demonstrations. Still, they do seem to express a proposition. The following example from Schiffer (1995, p. 123) is sometimes presented:

⁵² Notice that in 18, the pronoun 'her' cannot just be a variable bound by the universal quantifier; it also encodes some descriptive information (i.e., female-hood).

⁵³ Kaplan (1989, pp. 489-490) says of 'he': 'These words have uses different from those in which I am interested [directly referential indexicals] (or perhaps, depending on how you individuate words, we should say that they have homonyms in which I am not interested).'

21. He is a giant,

uttered in the presence of a huge footprint in the sand. I am going to view cases like 21 as involving an attributive use of the pronoun 'he' (analogous to those paradigmatic attributive cases of definite descriptions and demonstratives in which you don't have the relevant individual in mind. These attributive cases will be in contrast to cases like 15 above, which I will view as involving a referential use of the relevant pronoun. I take these to be analogous to the paradigmatic examples of referential cases in which one does have the relevant object in mind. Similar observations can be made about other pronouns like 'you' and 'we'; their referential uses are familiar. But take an example of an attributive use of 'you' from Bezuidenhout (1997, p. 384)⁵⁴:

22. You are entering Grand Rapids.

Suppose this is printed on a board at the side of a road leading into Grand Rapids. In this context, the pronoun in 22 does not refer to 'you' or to any specific individual. Rather, 22 means something like: *the appropriately positioned reader(s) of this message is/are about to enter Grand Rapids*. Another example of an attributively used pronoun, also from Bezuidenhout (1997, p. 384), is the following, involving 'we': suppose you are standing around with a group of colleagues after a talk by a visiting speaker. I turn to the visitor and say 23 below:

23. We traditionally go for drinks after the talk.

Here you are not referring to the group including yourself and your colleagues who are standing around the visitor at this moment. Rather, you are saying something like: *there is a tradition for the people present at the talk (whoever they are) to go for drinks after the talk*. I now intend to present examples analogous to those I presented using demonstratives and definite descriptions. These will be cases in which a given pronoun is used to express a rule-like expression, cases in which the pronoun is modified by a modal operator and cases in which the pronoun is modified by a hyperintensional operator. These uses seem to express an object-independent proposition even when we do have the relevant object in mind. Take first the following rule-like expression:

24. She that works hard always wins.

Now it might be said that the pronoun here is bound by a covert universal quantifier, if so then we may put 24 aside. But note that claiming this maybe beside the point. The claim that I am making is that the pronoun in 24 does not contribute an object to the

⁵⁴ Her example is inspired by Numberg (1991).

proposition, but rather some property-based condition (if pronouns are directly referential then they cannot be bound by a quantifier). 24 seems to convey a rule-like expression; its truth is not apparently dependent on any particular female. Indeed, it seems that 24 can be true even if there are no women. Cases in 24 are analogous to the rule-like expressions using demonstratives that I presented in the previous section (e.g., ‘that woman with the highest salary is the most powerful’). In any case, take now a modal embedding of the pronoun ‘she’.

25. She might not have had green eyes.

If ‘she’ can be attributive, then the truth value of 25 should vary depending on whether it is modified by the modal operator or not. If we can stage a scenario in which 25 is false with ‘she’ (say) taking widest scope relative to both operators (‘might’ and ‘not’), but true with (say) ‘she’ taking intermediate scope relative to ‘might’ and ‘not’, then that would be evidence in favour of an attributive reading of ‘she’. We need to find a property that a particular human female has necessarily that human females, in general, may or may not have. In 25, I have chosen the property of *having green eyes*.⁵⁵ Suppose that my colleague and I are very bored at the office and we decide to make a bet on the eye colour of people entering the office. I make my bet by uttering, ‘green’ and my colleague makes her bet by uttering, ‘brown’. We wait a few minutes and a female comes in with dark brown eyes. I lose the bet and my colleague laughingly says, ‘You are always doomed to fail’ I respond by saying, ‘No I am not. She might have had green eyes.’⁵⁶ It seems that I have uttered a truth. This is the reading in which it is stated that there exists some possible world such that the unique female (who just walked through the door does) not have green eyes. This reading is just not available if ‘she’ contributed an object to the proposition. If it did, then 25 will be false. This is because, assuming that eye colour is a feature that humans necessarily have, no human would have a different eye colour than the one they have.

Now take a hyperintensional embedding. Suppose me and my friend Jane both work at the local op shop. Jane and I empty hauls of what are usually second-hand clothes. We notice that we are receiving male attire that is brand new and which is all of the same brand. Jane believes that these clothes belong to one person, and expresses this object-independent belief by making the utterance in 26 below, which expresses an object-independent proposition:

26. Some guy is donating a lot lately.

⁵⁵ Perhaps a better property can be found to illustrate the point.

⁵⁶ Implying that I only lost due to bad luck and not due to some inherent mediocrity.

Jane walks away and another colleague of mine, June, comes to help me unload that same collection of clothes. However, she reports to me that everyone (i.e., all the employees) are so frustrated because they are not receiving enough donations. As a result, she reports to me that the staff are now thinking that people in general are just very stingy these days. As I hang up that new shipment of clothes on the rack, I look at them and then back at my colleague (who is looking at me) and utter:

27. That's not right. Jane thinks he is donating a lot.

It seems to me that 27 is true. But how can that be if the classical picture of directly referential pronouns is correct? There might be some sort of reference to a shirt, but this shirt is not the appropriate object to be a referent of 'he'. Also, the occurrence of 'he' in 27 cannot be anaphoric on Jane's earlier discourse in this situation; the two are disconnected. But even if 'he' somehow contributed the right individual, 27 should be false since Jane does not believe the relevant object-dependent proposition (we stipulated this) by making her utter 26. However, this is easily captured if what I am attributing to Jane is the object-independent belief that the male owner of these clothes is donating a lot. Moreover, it seems that 27 is true even if (unbeknownst to any of us) the actual person donating these clothes is Kevin who is such that Jane thinks he is stingy because he never donates anything. That is, it seems that 28 below as well as 27 above are true in this scenario:

28. Jane believes that Kevin is not donating a lot.

But this is problematic on the directly referential account of 'he'. According to this account, 27 and 28 attribute contradictory beliefs to Jane. This result is avoided if the 'he' in 27 contributes a property-based condition about being a unique male (that donated those clothes). Let us now look at more data.

5.3 More Data

Now some writers have claimed that the R/A distinction is even more pervasive than the above section indicates. For example, Gómez-Torrente (2015, pp. 102-103) points out that Donnellan's observations regarding the duality in use can be extended to many more quantifiers. Consider the following sentences in which the (a) variants are uttered from the analogous attributive scenario in which the detective stumbles across Smith's badly mutilated body and the (b) variants are uttered from the analogous referential scenario in which the speaker and audience are looking at the convict(s) standing trial in the court dock (where subscript 'A' and 'R' stand for 'Attributive' and 'Referential' respectively):

29. (a) Every_A murderer of Smith is insane.
 (b) Every_R guy in the dock is insane.
30. (a) Most_A murderers of Smith are insane.
 (b) Most_R guys in the dock are insane.
31. (a) Many_A murderers of Smith are insane.
 (b) Many_R guys in the dock are insane.
32. (a) Several_A murderers of Smith are insane.
 (b) Several_R guys in the dock are insane.
33. (a) A_A murderer of Smith is insane.
 (b) A_R guy in the dock is insane.
34. (a) A_A few murderers of Smith are insane.
 (b) A_R few guys in the dock are insane.

Gómez-Torrente thinks that it is easy to imagine utterances of 29-34 by which an utterer would be attempting to communicate object-dependent contents about particular persons. That is, it is easy to see how, if some of the detainees in the dock behaved in suitable ways, the speaker could use the quantifier phrases in all of the (b) variants intending to communicate, and successfully communicating, a variety of contents involving some particular detainees. This means that in each case, those particular detainees would enter as constituents into the proposition and, hence, it would be mandatory to view the corresponding utterances as containing referential uses of the corresponding quantifier phrases (unlike the (a) attributive variants). Similarly, Bezuidenhout (1997), inspired by some data surveyed by Nunberg (1993 & 2002), points out (as I have been in §5.2 & §5.3) that the R/A distinction appears in many indexicals. I agree with this, but her R/A distinction is more inclusive than mine. She appeals to the following kind of data. Suppose that Bill Clinton is giving a press conference to introduce his latest appointee to the Supreme Court. He utters the following sentence (Bezuidenhout, 1997, p. 384):

35. The Founding Fathers invested *me* with the power to appoint Supreme Court justices.

Here, it does not seem that Clinton is saying that the Founding Fathers had Bill Clinton in mind in particular when they wrote the Constitution. Obviously, they could not have done so. For this reason, it is plausible that what is said by an utterance of 35 in the circumstances imagined is something like the following:

36. The Founding Fathers invested *the US president* with the power to appoint Supreme Court justices.

Or suppose you are alone in the house and not expecting your spouse to return for many hours. You hear a loud crash as the door bangs open, and you think that an intruder has entered your house. Just then, your spouse walks into the room and you say (Bezuidenhout, 1997, p. 384, fn. 5):

37. I thought *you* were a burglar.

Here, it is plausible to say that what you thought was that the person who crashed through the door was a burglar. You did not entertain the thought that your spouse was a burglar. Or take the following example from Nunberg (2002, p. 5):

38. Condemned Prisoner: *I* am traditionally allowed to order whatever I want for my last meal.

Here, it does not seem that the person who utters 38 (say it is John) is saying that there is a tradition according to which he is allowed to order whatever he wants for his last meal. Instead, 38 seems to communicate something like the following:

39. *The inmate on death row* is traditionally allowed to order whatever they want for their last meal.

Now I don't have much to say regarding the quantifiers in examples 29-34. I think that the counterfactual intuition (CI) that applied to definite description might apply to them just as well, but I do not have the space to argue for this. It is certainly not as well established that those quantifier terms really do function as referential devices sometimes. I merely mention those examples to suggest that the distinction might be

much more pervasive than is typically assumed, and than my survey indicates. In any case, we shall see in Chapter 7 a syntactic treatment inspired by Mendelsohn (2010) that *would* capture these uses (those in 29-34) naturally if they really do exhibit the R/A constitutive distinction.

And finally, I have no doubt that expressions 35, 37 and 38 communicate the relevant object-independent propositions above. However, I don't want to say that the occurrence of 'me' in 35 (for example) represents an attributive use in my particular sense. According to my R/A distinction, the descriptive condition that is associated with an attributively or referentially used term is the *same*. Moreover, this condition is the one typically associated with what writers informally call the 'character'. For example, in the case of 'the murder', whether it is being used referentially or attributively, the descriptive condition involved in both cases is something about unique murderers. In the case of 'that window', it is something about some unique window. In the case of 'he', it is something about some unique male. In the case of 'me', it is something about the unique speaker, and so on. My R/A distinction is not just a distinction between the relevant term being used to express an object-dependent proposition at times and an object-independent proposition at other times – this is Bezuidenhouts' R/A distinction. My R/A distinction, in contrast, is between two specific roles that the descriptive information encoded in the relevant term can play; either that information itself goes into the proposition, or an object that satisfies this descriptive condition goes into the proposition. Recall how Mendelsohn (2008, p. 171) put it: in one use (the attributive use) the proposition expressed *includes* the procedure for picking out an object (i.e., a condition). In the other use (the referential use), the proposition includes only the output of that procedure (i.e., the object). In both cases the descriptive information is the same. But the descriptive condition relevant to the uses in 35, 37 and 38 does not fit this picture. What seems to be going on with these cases is actually something more sophisticated than what I have been exploring so far. That is, the condition that is being contributed to the proposition in each of those attributive cases is a property-based condition *of the object that is the referent* (in some informal sense) of the relevant term. For example, in 35, the property-based condition is something about presidents, a property that Bill Clinton (the utterer of 'me' in this context) possesses.

Now I do not know exactly what to think about the cases in 35, 37 and 38. There seems to be at least two distinct contextual processes going on, as Nunberg (1993) recognises. The first process is one that goes into determining the object (i.e., Bill Clinton) which is the satisfier of the standard descriptive condition associated with the term (e.g., the speaker in the case of 'me'). The second process is the one that goes into determining a contextually relevant or *salient property* of that object (i.e., being the US president). This is

not exactly the R/A distinction I have been looking at. But I mention these cases merely to distinguish my data from the R/A distinction I have been discussing, and to set them aside. I take it that what is at the heart of the R/A constitutive distinction, as Donnellan suggests, is the truth-conditional role of the descriptive information encoded in the term. That is, does this information go into the proposition, or does it determine an object and then fall away? The data of which 35, 37 and 38 are a sample seems to identify yet a third role for this descriptive information, and my R/A distinction might be a special case of this more general R/A distinction. Alternatively, it might be that this more sophisticated distinction is some other phenomenon altogether. In any case I have not tried to justify *Claim 3* in Donnellan's counter argument (on page 38). So we have now justified *Claim 1, 2 and 3* which, I think, push us to conclude that the R/A constitutive distinction should receive a unified semantic treatment. This semantic treatment cannot be lexical ambiguity, so what can it be? I think there are two plausible possible treatments. One based on syntactic ambiguity of the sentence the (say) definite description is embedded in and the other is based *narrow* context sensitivity. These treatments will be presented in Chapter 7. In the next chapter I want to investigate the causal R/A distinction. It will become clear shortly why I want to do this now before looking at those proposed treatments in Chapter 7.

Chapter 6: The Causal R/A Distinction

I now want to investigate the mechanisms that give rise to the R/A distinction. That is, I want to look at the question of: *in virtue of what does a referential/attributive definite description make the relevant propositional contribution?* Note that our discussion so far has already excluded certain answers to this question. We set aside an answer according to which the distinction is due to conversational implicature and lexical ambiguity. In §2.4 and in the last chapter, I very briefly hinted at Mendelsohn's scope-based syntactic ambiguity proposal and I said that I would propose a more straightforward version of it in Chapter 7. Under this proposal, the R/A causal distinction is just given in terms of intending one syntactic structure as opposed to the other. The audience then relies on certain pragmatic process for disambiguating which syntactic structure the speaker intended. We will return to this proposal later, but if we want to avoid ambiguity completely what can the mechanism that gives rise to the R/A distinction be? In this chapter I am going to claim that it is *intentions*. Again, I will be focusing on definite descriptions, but what I say should be generalisable to other terms that exhibit the R/A constitutive distinction that I have already surveyed. Let us make our discussion more precise. Take the following question:

Metaphysical Causal Question (MCQ):

- What features of context or use determine whether a definite description contributes an object or condition to the proposition expressed?

This is the causal question that is relevant to our discussion if the R/A distinction really does arise solely due to information encoded in terms that exhibit the R/A distinction (R/A terms). There have been many attempts aimed at distinguishing referential and attributive uses by getting precise on *what it is that makes a term attributive or referential*. However, here, writers have not been so careful to distinguish between two kinds of answers to this question: constitutive answers and causal answers. I have already given the constitutive answer in terms of propositional contributions in Chapter 2, and now I want to focus on the causal one. But even here an additional distinction needs to be made. Guided by Neale's discussion (2007, p. 359, n.7) we can distinguish between the metaphysical causal question above (MCQ) and an epistemic causal question below (ECQ):

Epistemic Causal Question (ECQ):

- What features of context or use determine whether a definite description is interpreted by the audience as having contributed an object as opposed to a condition to the proposition expressed?

We need only concern ourselves with MCQ – this is the content determination question that is relevant to the semantics of the relevant term. Some writers (like Donnellan himself) appeal to things like *having in mind* and *intentions* and provide the following necessary and sufficient feature for what makes a definite description referential as opposed to attributive (Donnellan 1979, pp. 119-120):⁵⁷

- Donnellan’s Causal Answer (DCA): In those referential uses, the speaker has the relevant object *o* in mind and intends their audience to recognise that it is *o* that they are talking about (partly) by virtue of their use of the definite description. In the attributive cases, there is no such intention.

But DCA appears to be the start of an answer to ECQ. That is, the feature in DCA (regarding the referential use, for example) resembles that phenomenon that Bach called ‘speaker reference’. Recall from our discussion in §2.4 that according to Bach (2008, p. 518), this is at least a four-place relation between speaker, object, term and audience, in which a speaker has *o* in mind (or has thoughts about) *o* and uses a term *t* to refer their audience to *o*. Bach points out that this is a pragmatic phenomenon in which the speaker intends their audience to interpret their utterance in a certain way. DCA does not quite involve the kind of intention that is relevant to fixing the semantics of the term; rather, it may just be relevant to the audience ascertaining what the proposition expressed is, and not to actually expressing that proposition. But the intention-relevant observations in DCA regarding referential and attributive uses would typically accompany those uses *if* these uses were a product of the semantics of the terms anyway. There is a dependence (assuming cooperativeness by the speaker) between 1) the intentions that go into determining what the semantics of a term are and 2) intentions that go into determining the audience’s interpretations of those semantics. And since we are going to assume a semantic view of the R/A distinction, we can transpose DCA to a semantic version like the following (which is an answer to MCQ):

- Metaphysical Causal Answer Referential (MCQR): In the referential uses, the speaker intends to refer to particular object *o*.

⁵⁷ Donnellan (2012, p. 145, n. 5) says that this is in fact only a sufficient condition for what he calls ‘speaker reference’. This is due to anaphoric uses of definite descriptions. But here, Donnellan is not thinking about our R/A distinction. He says in footnote 10 of that same paper: ‘*In this paper I use the term “reference” and variants on it in connection with definite descriptions both in referential contexts and in attributive. This is not intended to imply that there may not be a big difference between the two situations, as “Reference and Definite Descriptions” suggested.*’

- Metaphysical Causal Answer Attributive (MCQA): In the attributive cases the speaker does not intend to refer to a particular object *o*.

This kind of speaker reference does not need to pertain to audiences and their interpretation. That is, it need not be Bach's speaker reference. But you might think that MCQR/A simply boils down to the claim that referential uses are those in which the speaker intends to contribute an object to the proposition, which are contrasted to uses in which the speaker has no such intention. If so, you might say that this answer is on the one hand not illuminating and on the other not even possible, because we are just not able to have such technical intentions pertaining to propositional contributions. But this latter complaint is nothing but a terminological quibble; we can always spell out the intention in MCQR/A using a less technical notion like *intentions to talk about* or *intentions to speak of*. I am using the term 'refer' in an informal, everyday sense without any theoretical commitments. And the former complaint regarding non-illumination would be worrying if we weren't presenting a theory of definite descriptions under which they are semantically intention-sensitive terms, but that's just what we are going to do.

Now you would have noticed the notion of *having in mind* (which appeared in DCA) no longer appears in MCQR/A. This is because the notion is only relevant to whether or not one can have an intention in MCQR. Donnellan, like Bach and many others, thinks that in order to express an object-dependent proposition about an object *o*, you must have the relevant object in mind. Let us say that you have an object in mind when you are cognitively or perceptually connected to that object; when you are in some way acquainted with it.⁵⁸ Perhaps this involves an interesting causal link between the speakers' mind and the relevant object or a causal link between the object and your perceptual system that leave impressions on your perceptual system that allow you to have object-*involving* thoughts about that object, as opposed to general or property-based thoughts. Or perhaps 'having in mind' involves uniquely identifying or discriminating knowledge of an object. What I mean when I say that the speaker has the *relevant* object in mind (with respect to, say, definite descriptions) is that the speaker has the object that they know/believe/pretend to satisfy the definite description being used. In any case, this is where I will leave the notion of 'having in mind' – vague enough that it can be understood as a weak or strong relation between the speaker and the relevant object. But the important point here is that the condition of *having in mind* (whatever its precisification is) is only a *pre-condition* for MCQR; a condition that may be necessary in order to have the intention in MCQR. In other words, it is a pre-condition of managing to use 'the F' referentially, not part of what makes one's use referential. Now, in the

⁵⁸ Not necessarily in the strict Russellian sense of 'acquaintance'.

remainder of this chapter, I want to survey some examples. The purpose of this survey is to appreciate why, if there is a correct answer to MCQ, it can't be via the notion of 'having in mind' or something like *grounds for assertion* (which is sometimes appealed to), but it must be given in terms of the intentions in MCQR/A if we are to propose a theory under which definite descriptions (and other terms that exhibit the constitutive R/A distinction) are narrow context sensitive terms.

The Missing Student

Suppose I have been tasked with identifying the missing student in my school who I know to be Sam. When I finally observe Sam sitting on the school bench, I call the principle and, to put her mind at ease, I say, 'The missing student has been found.' Suppose the principal doesn't know who the missing student is and has never had him in mind, and the speaker knows this.

My utterance here seems attributive, and the principal understands it as such. I (the speaker) have the relevant object in mind, yet my use of the definite description seems attributive. That is, what seems to be expressed is an object-independent proposition. MCQA explains this by saying that the speaker has no intention to refer to any particular object *o* by virtue of their use of the definite description.⁵⁹ But perhaps you might say that what this example shows is that the speaker *and* audience must have the relevant individual in mind. Now, I think such an appeal to the audience for metaphysical content determination is wrong, but in any case, take the following counterexample to this suggestion: the case of Vladimir (Donnellan, 2012, p. 119):

Vladimir:

I and my audience see Vladimir (and therefore we both have him in our mind), someone who we know to be the strongest amongst men. We see him lifting 450lbs, and I utter, 'The strongest man can lift 450lbs' not intending to communicate anything about Vladimir but intending to communicate something general about strong people and their lifting limits.

Here it seems that the definite description in my utterance is not referential, although I do have the relevant individual in mind (recall that the relevant object in mind is the

⁵⁹ Let me say why I think this case is attributive. I assume that if the audience grasps the same proposition that the speaker expressed then this will be an instance of successful communication. In the case of The Missing Student, the principal is blocked (by stipulation) from entertaining the relevant object-dependent proposition constituted of Sam as a main constituent. But if in this case there is successful communication (as I assume there is), then the speaker must have expressed the relevant object-independent proposition, because that is the only proposition (out of the two) that he can grasp.

object that the speaker knows/believes/pretends to be the satisfier of the definite description in question). This is because I am trying to communicate an association or a certain pattern of co-instantiation. That is, I am trying to communicate a universal or rule-like expression, something along the lines of ‘for everything, if it is the strongest among men, then it can lift 450lbs.’ Now of course, this is exactly what I have stipulated to be the case in the above example (by saying that I am intending to communicate something general and nothing about Vladimir). But that was just the quickest way to get you to *feel* the rule-like tone of the example. Now take a modified version of the case of Smith’s murderer, mentioned in Yagisawa (1985, p. 117):

Smith’s Murderer V2

You are that same detective (from Donnellan’s classic attributive case) but after the conviction of the murder you go inside the courtroom and look at the convicted murderer, but you don’t see him acting strangely. You look over to your friend (who is also looking at Jones) and you utter, ‘Smith’s murderer is insane,’ but you attribute insanity to the murderer precisely because of the condition of Smith’s body and the fact that Smith was the most lovable person in the world.

I think this case is attributive, although my audience and I *do* have the relevant object in mind. And I think this is because I am, like in Vladimir case, trying to communicate a rule-like expression by saying that there is an association or a certain pattern of co-instantiation between the property of being Smith’s murderer (the property in ‘F’) and the property of insanity (given in ‘G’). This is the effect of stipulating that you are attributing insanity to the murderer due to him being Smith’s murderer – it guarantees generality. Now, I am not quite sure what is going on in these cases (*Smith’s Murderer V1* – the classical case, *Smith’s Murderer V2* and *Vladimir*). I am sure that they are all attributive, but are they cases in which the description is inside the scope of a covert universal quantifier? If so, perhaps, the effect of this covert universal quantifier is that it always forces an attributive reading of the embedded subject term that is accompanied, only sometimes, by an existential commitment.⁶⁰ Alternatively, it might be that these cases are simply analysed in the standard classical Russellian manner, though they may be justified by an antecedent commitment to a claim of property inclusion: that the property of being the strongest man is included in the property of being able to lift 450lbs. In this case, the rule-like feature of the utterance holds because of a relation between the properties and has nothing to do with ‘the’. Note that these cases are analogous to those rule-like expressions I presented in the last chapter involving pronouns (e.g. ‘She that

⁶⁰ By ‘existential commitment’, I mean that in some rule-like expressions (such as the one in *Smith’s Murderer V2*), although are to be treated as universal expression, they still require the existence of a unique murderer for their truth. But in cases like *Vladimir*, our utterance seems true even if there is, in fact, no strongest man (or men for that matter).

works hard wins') and demonstratives (e.g. 'that woman with the highest salary is the most powerful). In any case, the last two examples show that we (the speaker *and* audience) may have the relevant object in mind although our use is attributive. Now, you might say that once we ignore those attributive uses of definite descriptions in alleged universal or rule-like environments, all the attributive uses that are left are those in which the speaker does *not* have the relevant object in mind. This is not true – take the following case.

Hyperintensional Embedding – Jones' Mother:

I turn to my friend in the courtroom while observing Jones acting strangely – the actual murderer of Smith and utter the statement below.

- Sally believes that the murderer of Smith is me.

Sally is Jones' mother who is still in denial about who the murderer is. My audience and I both know this and know that about each other.

It seems that I have uttered a truth and I have faithfully reported Sally's belief, although it does not seem that I am attributing to Sally the belief that Jones is Atheer, which is what I would be attributing to her if the definite description contributed an object to the proposition. So although my audience and I have the relevant individual in mind (we are both looking at him standing trial), I do not intend to talk about the individual by virtue of my use of the definite description. And that is because I am trying to communicate the content of Sally's belief, and Jones is no part of it. So much for *having in mind*.

Now, motivated by some of the above examples, Donnellan (2010, pp. 119-120) appeals to the notion of *grounds for assertion*. Ludlow and Neale (1991, p. 177) also appeal to this notion to explain what is going on in certain referential cases. Others like D.E. Over (1985) appeal to a very similar notion – that of *constructive vs. non-constructive inference* – to distinguish referential from attributive cases. This notion (of grounds for assertion) hasn't received much in the way of analysis, but on first inspection it might seem to nicely distinguish Donnellan's classic referential cases from those attributive cases. For example, in the classic attributive cases, Donnellan has us inferring insanity from the condition of Smith's body, whereas in the referential case, Donnellan has us inferring insanity directly or perceptually by seeing Jones, the actual murderer of Smith, acting strangely in the court room. Ludlow and Neale (1991, p. 177) say that referential uses are those in which the speaker's grounds for assertion are furnished by an object-dependent

belief concerning some particular object *o*.⁶¹ Applying this to our causal R/A distinction investigation, we may say that what makes a definite description referential is its being used to express an object-dependent belief which we might contrast to those attributive cases in which the definite description is used to express an object-independent belief. But this claim just pushes the need for explanation deeper; surely, we are rightly tempted to ask *what allows those attributive/referential definite descriptions to function in this way – that is, to be used to express object-dependent vs. object-independent beliefs?* This is the causal question we started with! It is not like *all* linguistic terms can be used in this way. Or maybe this is not how the dialectic is supposed to go. Perhaps the dialectic starts by pointing out that definite descriptions are only Russellian, but they are used referentially only when we want to communicate an object-dependent belief. But here, we are not assuming that the correct unitary analysis of definite description is Russellian. Moreover, again (given that we have already excluded the answer according to which those referential meanings are conversationally implicated), we are tempted to ask the same question we started with: *what allows those attributive definite descriptions to function in this way – that is, to be used to communicate object beliefs?* Is it something about having the relevant object in mind, or something else? We've made no progress.

But maybe what is being claimed or suggested by Ludlow and Neale (1991) is that what makes a definite description referential is its being used to express a belief that is inferred perceptually. This is in contrast to those attributive cases in which the definite description is used to express a belief that is inferred via general, property-based or descriptive information. But this can't be right. Suppose I believe (on completely general grounds) that recently imported furniture is of very poor quality due to some international trade tensions. As I arrive at a furniture store with my partner, the store clerk tells me that everything at the back of the store was imported years ago. After looking at a bunch of tables, she asks me which one I prefer; I respond by uttering the following:

- The old table at the back is of better quality.

Here, it seems that I have expressed an object-dependent proposition; that is, my use seems to be referential, but the belief that motivated my assertion was inferred generally or descriptively and not by directly inspecting the quality or the age of the table in question. And in the case of *The Missing Student*, my use seems attributive although the belief that motivated my assertion was inferred perceptually (by observing Sam, sitting on the bench). But maybe the notion of grounds for assertion is supposed to signal to the *way* in which the descriptive information contained in the noun phrase of the

⁶¹ Ludlow and Neale (1991) focus on referential uses of *indefinite* descriptions (not definite descriptions) when discussing the notion of grounds for assertion, but they suggest (in footnote 12) that the notion of 'speaker's ground' in relation to the discussion of definite descriptions ought to carry over (with modification) to the study of definites.

designational term is *inferred* (i.e., the 'F' in 'the F'). But this, too, doesn't seem right. For example, suppose I am reading an encyclopaedia article about microwaves and I read that they were invented by exactly one person in 1945 (assume I have never met this person and, therefore, do not have him in mind). Suppose further that I read that this person is very tall, and so then I utter the following to my friend:

- The inventor of microwaves is quite tall.

This utterance seems to express an object-independent proposition (it would to Donnellan, by stipulation, given his underlying assumption that you cannot express an object-dependent proposition about an object that you do not have in mind. But it might be claimed that this is really because the speaker inferred the existence of microwave inventors descriptively and not perceptually. But that can't be right. Suppose I do end up meeting this person. Suppose he is signing microwaves inside a room with a sign at the front that says, 'Meet the man that made cooking a breeze: the inventor of the microwave!' After reading this sign, I walk in and finally meet this man that I now believe to be the inventor of the microwave (due to the sign I saw earlier). I then utter the above expression to my friend because I can see that the man is, indeed, quite tall. Here I think I would express an object-dependent proposition, although I inferred the descriptive information generally or descriptively (i.e. non-perceptually) from that sign. Or take the attributive use of the definite description in the case of *Vladimir*. It could have been that the speaker inferred the existence of a unique strongest man by simply looking at Vladimir's impressive bodily form and size; still, the definite description would not be referential if I intended to communicate a rule-like expression by uttering, 'The strongest man can lift 450lbs.'

But perhaps someone might say that the utility of the notion of *grounds for assertion* is in noting whether or not the evidential basis of the assertion is object-dependent evidence or purely object-independent evidence. That is, sometimes *thinking about* a particular object is an additional unnecessary cognitive task (if the evidence is general) or is something one has already done (in one's consideration of the object-involving evidence). That is, the notion of *grounds for assertion* may be cashed out in terms of whether *justification* for the assertion requires presenting an object or merely general, property-based, descriptive information. The speaker relying on the latter does not typically intend any particular object. They could, but it would be additional cognitive labour to form the intention to think about *o* when *o* is not in the evidence on which the speaker relies. But if *o* were part of the evidence, then there would be no additional work in intending to refer to them, since one is already having object-directed thoughts about *O*. There are two examples that make me doubt whether this proposal is right. The first is, again, the case of the Missing Student. We said that this definite description use is attributive, but surely

my assertion would be justified if I presented the individual Sam (who we know to be the missing student). The other example involves a definite description embedded in a modal environment.

Modal Embedding – Ill Jones:

Imagine the classic referential scenario involving Jones acting strangely in the courtroom, where we (the speaker and audience) know Jones is the unique murderer of Smith:

- The killer of Smith might have been me.

My friend and I are looking at Jones and we both have him in mind as the satisfier of the definite description that appears in the utterance. I turn to my friend and utter the above.

Now, it seems that I have said something true. If the definite description is attributive and is within the scope of the modal operator, then the truth of the above utterance is easily explained. But if the propositional contribution of the definite description is Jones (the person my audience and I have in mind), then I will surely be uttering a falsehood, for it is impossible for Jones to be me. But now suppose that my justification for making the above assertion is seeing Jones on trial with a particular medical condition that makes him wave his hand frantically all over the place. This is my justification because Jones' defence pointed out that on the particular day of the killing, Jones happened to be holding a kitchen knife (he was preparing dinner) when opening the door to welcome his best friend Smith. But unfortunately, it was at this precise moment that Jones started to have a medical fit which resulted in the savage killing of Smith. So basically, Jones could not control what his body was doing, and this is why he savagely killed Smith. What justified me (the speaker) in making the above utterance is the thought that if I had Jones' condition, I would have savagely killed Smith too.

But could I not present ill Jones (the object I am perceiving) as justification for the above assertion? I think so. But that is a problem for the current proposal, because my use of the definite description seems attributive. On the other hand, MCQA easily explains this; although I do have the relevant object in mind, I do not have the intention to refer to him. That is because I am trying to communicate something about who *would* have been the murderer if Jones wasn't.⁶² But again, maybe this is not how the dialectic is supposed to

⁶² Note that some cases of embeddings naturally lend themselves to a object-dependent proposition interpretation. That is, to an interpretation in which the definite description contributes an object. An example like 'The murderer of Smith might have been caught' is one such example in which I am expressing that *Jones might have been caught*.

go. Maybe the dialectic starts by pointing out that all definite descriptions are semantically Russellian, but sometimes, if our justification for the assertion requires the presentation of an object, our use is referential. But this sounds clearly wrong; if our assertion is Russellian, no object is required to justify it. So perhaps what should be said instead is that assuming that all definite descriptions are semantically Russellian, if our justification *for the belief* motivating the assertion requires the presentation of the relevant object, then our use will be referential. But again, the case of the *Missing Student* and the *Modal Embedding – Ill Jones* are counterexamples. The belief that motivated both assertions was partly object-dependant. And if we move to something along the lines that all beliefs that motivate an assertion have to be either object-based or property-based, then our proposal is now capable of explaining nowhere near as much data, assuming the proposal can even get off the ground. Finally, someone might respond to this by saying that we need to make a distinction *within* our collection of justification. There is evidence that is genuinely justificatory and there is evidence that is merely incidental. This person might go on to say that those assertions (or those beliefs that motivate those assertions) involving attributive uses of definite descriptions are only justified by genuine property-based evidence. This is in contrast to those that are justified by genuine object-based evidence. They will say that surely in both cases (referential and attributive), both kinds of evidence are available, but we need not look at merely incidental evidence. For example, it may be pointed out that in the *Modal Embedding – Ill Jones* example, ill Jones is merely incidental. What actually justifies my assertion is being the kind of thing that has a certain illness which, in turn, makes one prone to uncontrollably and savagely killing people. But what are we going to say about cases like the *Missing Student*? Is Sam being found merely incidental in justifying the relevant assertion? I don't think so. This is why I don't think the notion of *grounds for assertion* can distinguish the uses under consideration correctly. It just seems implausible that I am blocked from using a term attributively if my assertion is somehow justified by perceptually based or object-based evidence.

Now the above examples show that the feature of 'having in mind' is not enough to distinguish the two uses (it features in referential uses as well as some attributive uses) and that there is no version of the notion of *grounds for assertion* that is plausible. But perhaps it might be said that something like Stalnaker's familiar notion of *common ground* can do the job. According to Stalnaker (2002), a proposition *p* is part of the common ground in a conversation iff the speaker and the addressee accept *p*, they both believe that they accept *p*, they both believe that they both believe that they accept *p*, ... and so on. Maybe this kind of context is the thing that makes one use referential as opposed to referential. But the common ground that is associated with referential and uses contains all sorts of object-dependent propositions *and* object-independent propositions, just as is the case with attributive uses. But if so, there will be nothing in the common ground that

can be systematically exploited by the semantics of the term (other than intentions) in a way that distinguishes both uses. And if we are not going to distinguish the propositions in the common ground based on whether they are object dependent or object independent how can they instead be distinguished? It is not clear.

This, I think, leaves intentions as the only suitable candidate that can play the causal role in determining whether the definite description in question is referential or attributive. That is, if we are to treat them as context-sensitive terms it seems that our MCQR/A must be the answer to MCQ. But these are not the only reasons why I think intentions are needed. In one and the same utterance situation (or context), one occurrence of definite descriptions can be attributive while the other may be referential. Alternatively, one may be referential (by contributing Jones to the proposition) while the other may too be referential (by contributing James). Take the following examples (where subscript 'R' stands for 'referential' and subscript 'A' stands for 'attributive':

1. The_A inventor of the microwave is the_R guy we saw last night.
2. The_R guy we saw last night is the_R postman.
3. The_R guy we saw last night is the_R mechanic, who is also the_R postman.

So in fact, in the same context, different occurrences of definite descriptions can have different contents. And this might seem to make them much more akin to classically treated pronouns like 'he' and 'you'. I don't see any other feature in the extra-linguistic environment that the semantics of terms that exhibit the R/A constitutive distinction can *systematically* exploit in a way that also *distinguishes* referential from attributive uses. And examples like 1, 2 and 3 above show that definite descriptions can vary their content from within the same common ground, and the only thing that can play this role would be a different intention for each occurrence of definite description. Now that we have seen that the answer to MCQ must be given in terms of intentions if we proposing a treatment under which R/A terms are narrow context sensitive we are finally in a position to present this treatment in addition to the treatment under which R/A induce a syntactic ambiguity in the sentences they are embeded in.

Chapter 7: Two Semantic Treatments

Writers have proposed various ways of dealing with the R/A constitutive distinction. We have already looked at one way of dealing with this duality at the beginning of the thesis when looking at definite descriptions: the way of conversational implicature. In the case of definite descriptions, we saw in Chapter 3 how Neale tried to explain away the object-dependent proposition that was allegedly expressed using referential definite descriptions via conversational implicature. He claimed that definite descriptions are always attributive or Russellian, though we sometimes use them to conversationally implicate an object-dependent proposition, not to literally express an object-dependent proposition. We saw in Chapter 3 how following this line of thought leads to a dead end. Depending on which unitary analysis of definite descriptions you are committed to, you can explain away the other use via conversational implicature.

Some, like Devitt (2004, 2007), proposed to deal with this duality of usage by claiming that the definite article, for example, is lexically ambiguous – that it is a case of either homonymy or polysomy. I gave reasons in Chapter 4 regarding the systematicity between the relevant meanings that tells against this proposal. We also briefly touched on Mendelsohn’s (2010) proposal, under which the R/A distinction is a product of a certain special scope-based syntactic ambiguity due to the terms’ interaction with the invisible hyperintensional operator of assertion. This is the syntactic ambiguity approach which we will return to shortly. But there has also been a proposal under which terms that exhibit the R/A distinction are context-sensitive terms (i.e. narrow context sensitive terms)⁶³. Recanati (1986) first proposed such treatment, and others have followed this path with various modifications. For example, Bezuidenhout (1997) proposes that definite description and the other terms that exhibit the R/A distinction are *semantically underdetermined* and require certain pragmatic processes to determine the relevant proposition. Under this proposal, those pragmatic processes are not the ones involved in (say) conversational implicature but are more akin to those at play in disambiguation and reference fixing. Under this proposal the meaning of terms that exhibit the R/A distinction encodes no rules about what propositional contribution is made; this is left up to the speaker in their context. I take this underdetermination approach to be the kind of approach to fall back on *if all else fails* and, therefore, I will not be discussing it. Given this landscape and what I pointed out regarding the systematicity between these two uses in Chapter 4, as well as the pervasiveness of these two uses in Chapter 5, I find that a semantic treatment is called for along non-lexical ambiguity and non-conversational-implicature lines. My aim in this section is to sketch the beginnings of two semantic

⁶³ This narrow context is just the set of extra-linguistic features that are required to fix the content of the uttered words/sentence. It is supposed to contrast with the *wide context*, which includes the narrow context as well as a whole bunch of other features in the extra-linguistic environment that are relevant for audience interpretation.

treatments of the R/A distinction according to which the R/A distinction arises due to 1) intention sensitivity (much like classically treated narrow context-sensitive indexicals) and 2) syntactic ambiguity. I will be focussing on just definite descriptions but of course, I think that both treatments can be extended to other R/A terms but due to limitations of space I will not be doing this.

7.1 Approach 1: Intention Sensitive Indexicals

We might motivate this kind of approach by an offline discussion between Salmon and Donnellan. Salmon (2004, p. 233, fn. 5) writes that Donnellan confirmed to him during personal correspondence (in October 1993) that he (Donnellan) had always held an *indexicality conception* of the R/A distinction. Donnellan wrote in the following passage:

The grammatical structure of the sentence seems to me to be the same whether the description is used referentially or attributively; that is, it is not syntactically ambiguous. Nor does it 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. 297).

Here is how Salmon expands upon the indexicality conception based on this offline discussion with Donnellan:

It is a thesis to the effect that a definite description is indexical, expressing different semantic contents with respect to different contexts, depending (at least in some instances) on whether it is used referentially or attributively. Salmon (2004, p. 232).

The idea is supposed to be that definite descriptions function in a manner analogous to the way in which classic indexicals were classically conceived to function. That is, classically indexical terms like 'I', 'you', or 'he' have different propositional contributions in different contexts of utterance. But Elbourne (2013, p. 105) thinks that this is a very odd use of the term 'indexical'. Given the different contributions that referential and attributive descriptions are supposed to make, this alleged indexicality would presumably have definite descriptions contributing individuals (e.g. Jones) to propositions on some occasions and generalised quantifier meanings (e.g., *whoever uniquely murdered Smith*) on other occasions. According to Elbourne, this is not how indexicals generally work. Indexicals such as 'I', 'you', and 'today', he says, generally

have a fixed semantic type and contribute various things of that type, subject to any other requirements encoded in their linguistic meaning: various individual people or days, for example. But Elbourne's complaint is ineffective. We have already seen in Chapter 5 that what are called 'classical indexicals' themselves admit of the R/A distinction. So, as Bezuidenhout (1997, p. 393) points out, this complaint begs the question against Donnellan's indexicality proposal. But Elbourne adds that if we allow ourselves the resources that Donnellan needs here, we will quickly come to the conclusion that there is no ambiguity (no polysemy or homonymy) at all in natural language. Bank? No homonymy here! This word is just an indexical which expresses different contents in different contexts, sometimes expressing a content having to do with gently lapping water and sometimes expressing a content having to do with sub-prime loans. But as we have seen in Chapter 4, the meanings in the R/A distinction are not arbitrary; there is a systematic truth-value dependence between them – although this dependence was not quite two-way logical entailment. The propositional contributions given in the R/A distinction are very constrained; one is the descriptive condition typically associated with that term, and the other is the satisfier of that condition in some context. When we propose to treat terms that exhibit the R/A distinction as indexicals, we are not just proposing to treat them as terms that vary their contents depending on context. Rather, we are proposing to treat them as terms that vary their systematically related or tightly constrained contents depending on context, by virtue of some systematic feature in that context. This proposal cannot be extended to lexically ambiguous terms like 'bank'. Still, Neale (1990, pp. 110-112, fn. 36) formulates the challenge for the contextualist approach as follows: if we have no explanation for how such very different types of propositional contributions could result from something with a unitary meaning then the only sensible explanation is that there are two different meanings or characters, a referential and a quantificational meaning (hence a lexical ambiguity) which lead to these two different types of propositions. It is this challenge that I intend to address in this section.

The starting point for my positive approach under which terms that exhibit the R/A distinction are intention-sensitive terms is a discussion by Vignolo (2012). In this discussion, Vignolo attempts to fill this explanatory gap (the one that answers the question of in virtue what is one term referential as opposed to attributive or vice versa) by proposing that in (say) the case of definite descriptions, there is a two-stage process that leads from the contexts of utterance to the semantic content of a definite description. At the first stage, he says, the linguistic meaning determines the *pointer*. At the second stage, a contextually determined function takes the pointer as an argument and yields the semantic content as value (2012, p. 624). That is, once the linguistic meaning of a definite description determines the descriptive content as a pointer, a pragmatic process gives that pointer as an argument to a function yielding the semantic content as a value. The semantic content might be the pointer itself, an enrichment of it, or a particular object that

satisfies or is believed to satisfy it within a restricted domain. According to this proposal the linguistic meaning of definite descriptions operates at a sub-compositional level. Given a definite description 'the F', the description's linguistic meaning yields a sub-compositional representation:

1. $\langle [\text{the } x: Fx], f \rangle$

Which consists of a pair formed from the descriptive content of the definite description – the pointer – and a functional variable. Contextual information provides the value of the functional variable, which, as stated above, is a function from the pointer to the truth-conditional content. Some contexts will be such that they contribute a function that takes the pointer and outputs an object-independent proposition, and some will be such that they output an object-dependent proposition given the same pointer. Now, all this is well and good, but without Vignolo telling us what it is about the context that determines which function gets contributed to the second position in the pair in 3, *no* explanatory gap has been filled; we have simply pushed the need for explanation deeper. This is what I intend to do by modifying a proposal given by Stokke (2008) which utilises intentions as the element of the extra-linguistic narrow context. Recall that I claimed in Chapter 6 that intentions would be a suitable candidate for this role because they distinguish referential from attributive uses in a systematic way (unlike the common ground, having in mind or grounds for assertion). Now Vignolo does claim that intentions (along with beliefs, presuppositions and other things) are what determines this function, but my proposal is different in the following ways: it is more precise, it utilises *just* intentions, and it models R/A terms on already existing conception of indexicals. Now, according to Stokke (2008), classically treated indexicals have an intention-sensitive semantics. Recall our Metaphysical Causal Question from Chapter 6, modified to be asked about indexicals:

Metaphysical Causal Question (MCO):

What features of context or use determine whether an indexical contributes some object *o* to the proposition?

Stokkes' answer to the above question is given in terms of intentions. That is, the propositional contribution of a given occurrence of an indexical term is some object *o* that is determined by the speakers' intention to refer to *o*. Stokkes' (2008, p. 393) main idea is that a referential intention can be modelled as a *function from indices to referents* and that function is contributed to a context. An index, according to this picture, is a number which corresponding to an occurrence of an intention sensitive term. And the narrow context (the part of the context that the semantics of a term is sensitive to) is to be represented as a tuple consisting of various things such as agent, time, world and so on,

but also that very intention determined function. In particular, let $f(i) = o$ be a referential intention that takes an index and maps it to an object o . This function is contributed to the context tuple and the character of an intention sensitive term α is that function ξ such that for any index i and context tuple $c = \langle f(i), \dots \rangle : \xi(c) = f(i)$.

So take the following example:

2. He₁ and he₂ are both logicians

The tuple that represents this context will look like this (the '...' indicates that other things are contributed to the context tuple like agent, time and world):

3. $c_1 = \langle f(1) = \text{John}, f(2) = \text{Mike}, \dots \rangle$

We then apply the character of 'he' twice to this context and get the two intended referents, respectively, as values (Stokke, 2008, p. 394). That is, the character ξ takes that c_1 context as argument and outputs the value (i.e. object) of the intention function for each index. This is Stokkes' proposal regarding the standard referential uses of indexicals. But since I am trying to also capture those attributive uses (of pronouns and all other R/A terms) via speaker intentions, this function that goes into the narrow context tuple and the character function will have to be a bit different. The idea, again, is to model the intention of the speaker in context as a function that goes into the narrow context tuple. This function is, again, sensitive to the intentions of the speaker. Metaphorically speaking, this function, as Stokke would say, *looks* for a speaker intention for each index: is the speaker intending to refer to o (using this token R/A term), or do they have no such intention? If they are, an object is the value of that function, and (this is where my modification comes in) if not, the value of that function is that same index. More formally we will say:⁶⁴

$$\text{Intention Function } f: i \rightarrow \begin{cases} \text{object } o \text{ iff speaker intends to refer to } o, \text{ and to } \\ i \text{ otherwise} \end{cases}$$

This is the function that is contributed to our context tuple. It takes an index of a token term and maps it to some object o iff the speaker intends to refer to o . But if the speaker has no such referential intention the value is not undefined as Stokke's simple intention function would have it, rather the value is that same index. The motivation behind this is

⁶⁴ Where i on the left side of the arrow is the argument of the function. The right side of the arrow is the value of the function given certain conditions.

that we need some value to be returned in cases where the speaker has no referential intention, so I have just chosen that same index to be returned.

Take now the following utterance and the context tuple that is associated with it (where subscript numbers represent indices and subscript letters represent whether the definite description is being used referentially or attributively).

4. The_{1,A} murder is the_{2,R} guy who won the_{3,A} lotto last night
5. $c_1 = \langle f(1) = 1, f(2) = \text{Jones}, f(3) = 3 \dots \rangle$

Notice that the context tuple in 5 partly consists of the value of our intention function for each index. The character of an indexed R/A term α is then that function ξ that takes the context tuple in 5 above and maps it to either an object that is the value of that index in our intention function f or condition d that is associated with the index otherwise. More formally, we will say:

$$\text{Character Function } \xi: c_1 \rightarrow \left\{ \begin{array}{l} \text{object } o \text{ iff } f(i) = o \text{ and } o \text{ satisfies } d_i, \text{ and to} \\ d_i \text{ otherwise} \end{array} \right\}$$

Let me say more. First, what is ‘ d ’? The motivation behind this is to isolate the constant descriptive condition or information that is linguistically encoded in R/A terms (whether they are being used referentially or attributively). In the case of ‘The murderer’ that is something about *being a unique murderer*. This condition d though is not a sentence, but I am neutral on what it should be – perhaps it can be seen as a property (or as a relation between properties).⁶⁵ Vignolo (2012) called this ‘the descriptive content’ of the term or the ‘pointer’. All our R/A terms have this component (in the case of ‘he’ for example it is something about *being a male*). Every indexed term will be associated with a condition d and if the speaker has no intention to refer to o (if the value of our intention function f is not o but that very index) then the output of our character function is that very condition d associated with that index. Notice also that that above character function only outputs the value o iff o is d_i . That is, if our use is referential, the value of the character function is o iff o satisfies condition d (e.g., iff o is the F). Speakers can refer to anything they want to refer to in a context, but the character of a term is supposed to semantically constrain the object outputs such that cases of incorrect descriptions (i.e. misdescriptions) are cases in which no object is contributed to the semantically expressed proposition.⁶⁶ Note also that the above proposal is neutral on the issue of *descriptive completion*. Perhaps descriptive

⁶⁵ If so, then the what I means by ‘satisfies d_i ’ in our character function is that object o instantiates the property d .

⁶⁶ Note that even under Stokkes’ simple picture (regarding univocal directly referential indexicals) something like condition d is also required. In the case of ‘he’ the objects that are the propositional contributions will have to satisfy the property of being (say) *a unique male*. We don’t want our semantic character to make the wrong prediction in cases of misdesignation.

completion (for descriptions, demonstratives and pronouns) is something that is achieved pragmatically, in which case nothing in the above proposal needs to be modified. Or perhaps descriptive completion is something semantic. If so, and supposing we cash out the notion of *condition* in terms of properties, then we can say that *d* in the case of 'the murderer' is the property of *being a unique murderer*. But we can allow this property to be a richer property in context. This is because in certain contexts we may use 'the murderer' to talk about the murderer *of Smith*. If so, then we can say that property *d* can be enriched to property *d'*. If we further cash out property *d* in terms of a set *D* of objects that are *d*, our character would then constrain *D'* such that it is a subset of *D*. The intuitive idea behind this is that if descriptive completion is something semantic our completed description must *logically entail* the more simple description. The semantics of a term should not allow an enrichment of the property *being a murderer* to *being a dolphin that I saw last night!*

But recall also that in §5.1a, when discussing demonstratives, we said that if intentions are relevant to the semantics of definite descriptions it has to be one and only intention per definite description term (i.e. 'The *F*' and 'The *F*' are one type of descriptive term, but 'The *F*' and 'The *N*' are two; recall that we individuated definite descriptions based on their noun-phrase restriction). Let us say that two tokens of definite descriptions are the same type of description iff they have the same type of noun-phrase. This means that each type of definite description will receive the *same* index. In the example of 4 above all our token definite descriptions are a different type. The character of each of the definite descriptions would apply to the narrow context tuple in 5 above to get the two relevant propositional contributions and the sentence in 4 as a whole would then express the proposition that *there exists a unique murderer of Smith which is identical to Jones* (assuming Jones is the guy who won the lotto last night). This is straight forward, but now take the following example similar to the one we encountered in §6.1a:

6. The₁ car is expensive, but it is not the case that the₁ car is expensive

Our semantics should predict that 6 cannot be true, and that is independent of how the definite descriptions in it are being used (i.e. referentially or attributively). We achieve this result by assigning the same index to both definite descriptions as can be seen. Now, the exact details of this kind of approach will differ depending on which R/A term we all are dealing with. Due to limitations of space I have only been able to sketch the beginnings of how this approach would apply to definite descriptions. In any case, does this proposal meet Neale's challenge? Well, we have said what it is about the context that makes one use referential as opposed to attributive. The presence of an intention to refer to *o* as opposed the lack thereof. And we have addressed the challenge of explaining how a unitary meaning produces the different propositional contributions. The meaning of an

R/A term like a definite description is just given in terms of the descriptive condition *d* that Russell and Frege agreed on with a constraint on the number of intentions per definite description. That is their one true meaning – nothing about the kind of propositional contribution they make. But according to my modified proposal their meaning or character operate on intention-parameters of context (very much like classically conceived directly referential indexicals). In a nutshell, the character of an R/A term looks for an intention function that maps each index to an object *o* iff it is the F and to a condition (corresponding to ‘the F’) otherwise. The simple idea is that when the speaker has no intention to talk about object *o* by virtue of some term, the value of the character is not undefined, it is rather the descriptive information associated with that term. This is just *one* rule.

7.2 Approach 2: Syntactic Ambiguity

I first mentioned in §2.4 there is a somewhat elusive proposal by Mendelsohn (2010) under which the referential and attributive readings of definite descriptions (and even other quantifiers like indefinites and universals) are simply the result of a syntactic ambiguity arising from the effects of an invisible hyperintensional operator of assertion. Now, I am confused by much of what Mendelsohn says, but I think he is on to something and so I shall motivate his broad approach by linking a few of his clearer ideas. Think about what it means for a speaker to assert ‘the F is G’. According to Mendelsohn, it means that the speaker uses or utters ‘The F is G’ to express that the F is G. Notice here, crucially, that the propositional attitude of assertion itself *consists* of a *sentence*. As Mendelsohn (2010, p. 174) writes ‘*assertion bestrides two otherwise distinct logical spaces: it involves using and mentioning the very same sentence*’. Now, if as Coulmas (1985, 3-4) says, the *de dicto* vs. *de re* duality that arises in speech reports is due to an interference between the speaker’s speech and the quoted (or reported speech), then (given what it means to assert a sentence) direct assertion is susceptible to the same ambiguity arising in reported speech. This is because it is just not clear *how* the speaker used ‘The F is G’, did he assert *of* that F that it is G or did he assert that the F is G? Perhaps we can put the point differently. Think about what it means to assert something and take 7 below:

7. Atheer asserted that the F is G

It is well established that constructions like 1 are ambiguous. But now let us replace the word ‘asserted’ which what Mendelsohn says it means. We now get 8:

8. Atheer uttered ‘The F is G’ to express the proposition that the F is G.

Now supposedly 8 too is ambiguous, and that is because, one might say, there are two ways one can report what Atheer uttered ‘The F is G’ to express. But now we have just

confirmed Donnellan's observations! Why would one be able to report Atheer's utterance of 'The F is G' as having expressed one of two different propositions unless it is because Atheer may have *in fact* expressed one of two different propositions with 'The F is G'? That is, the ambiguity we find in 7 *just reflects* an already existing ambiguity in 'The F is G'; it doesn't create it. But all this only bring us half-way to the R/A distinction because we have said nothing about object-independent vs. object-dependent proposition. But here is where Mendelsohn's second claim comes in. According to him, the ambiguity reflected in 7 is an ambiguity between object-independent vs. object-dependent propositions. This is because, for Mendelsohn, wide-scopeness in hyperintensional environments and semantic object-dependency comes hand in hand. Take 9 and 10 below.

9. Atheer asserted that *the F is G*

10. The F is such that Atheer asserted of it that *it is G*

According to Mendelsohn, the wide-scopeness of the definite description relative to the operator of assertion in 10 consists, at least partly, in the mere *presupposition* of Russell's conditions. And so these conditions, to speak metaphorically, are screened-off from the propositional attitude. That is, [asserted] in the wide scope construction in 10 is *blind* to the descriptive condition the object must satisfy - it is blind to the way the object is picked out and if this is so then it is not an object-independent proposition that is asserted; it must be an object-dependent one. This last point is generalisable to all hyperintensional attitudes and not just assertion. The reasoning behind this is as follows. Take the sentence 'George IV. wondered whether the F is G'. If George IV. wonders the proposition given by the narrow-scope construction (that's an object-independent proposition) then it would make sense if he (say) started looking for the book titled 'Waverley' in the library to see if Scott is the author, but if he was merely wondering whether Scott is Scott (what is given by the wide scope construction) then he wouldn't do that. That there is a cognitive difference between wondering an object-dependant and an object-independent proposition is obvious but what may not be obvious is that this is also reflected syntactically. If George IV. has gone mad and is wondering whether Scott is Scott, I can faithfully and accurately report this by using the wide scope construction of the expression (i.e. 'The F is such that George IV. wondered whether it is G'). This is because the hyperintensional attitude verb in the wide scope construction, as Mendelsohn emphasises, does not syntactically see the definite description; the way the object is picked out is irrelevant and so it is the object itself that must go into the proposition.

Now, under Mendelsohn's actual proposal, there is supposed to be an invisible (but presumably truth-conditionally inert!) hyperintensional operator of assertion in constructions like 'The F is G'. But I don't think there is a need for that. In other words, if

merely asserting ‘The F is G’ makes it propositionally ambiguous due to the effects of assertion, then why not just say that ‘The F is G’ (by itself) is ambiguous due to an *internal* structural ambiguity arising solely from the overt terms in ‘The F is G’. That is, let us simply just ignore talk of invisible hyperintensional attitude operators and ignore talk about presuppositions. Let us just assume that definite descriptions are unitarily Russellian and provide two phrase structure trees corresponding to each use of 11 below:

11. The murderer is insane




Figure 1: Phrase Structure Tree (Attributive)




Figure 2: Phrase Structure Tree (Referential)

The structure in Figure 1 is the standard phrase structure of 1, and the structure in Figure 2 utilises the standard *quantifier raising* technique (Elbourne, 2011, p. 89) under which a ‘t’ for ‘trace’ is introduced in the place of the definite description, which in turn moves to the upper left of the phrase structure. This forms a new sentence on the right in which the determiner DP (i.e., t) is something like a descriptively empty pronoun (i.e., ‘it’). This pronoun just picks up the very object that is denotation of ‘the murderer’ (i.e. Jones) and combines it with the verb phrase ‘is insane’. That is, what is relevant for counterfactual evaluation under the structure in Figure 2 is the very Russellian denotation (this is just the proposition we want given the Fregean analysis of referential definite descriptions). Let us now embed the syntactically ambiguous ‘The F is G’ in (say) a modal environment and see what the phrase structure tree looks like:

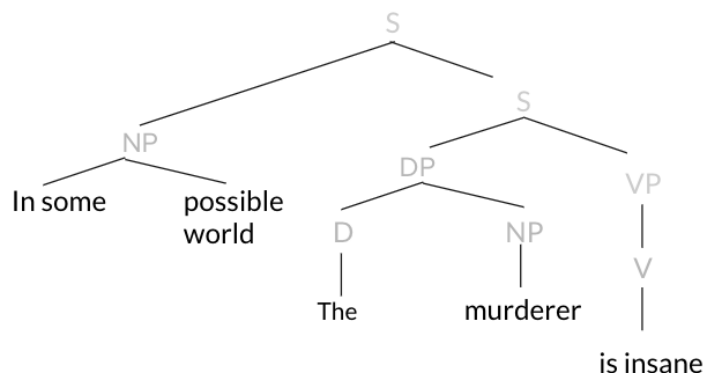


Figure 3: Narrow Scope – Object-independent

The above phrase-structure is one syntactic disambiguation of ‘The murderer might have been insane.’ It corresponds roughly to *there exists some possible world w in which whoever is unique in being a murderer in w is insane in w* . Nothing is supposed to be out of the ordinary here.

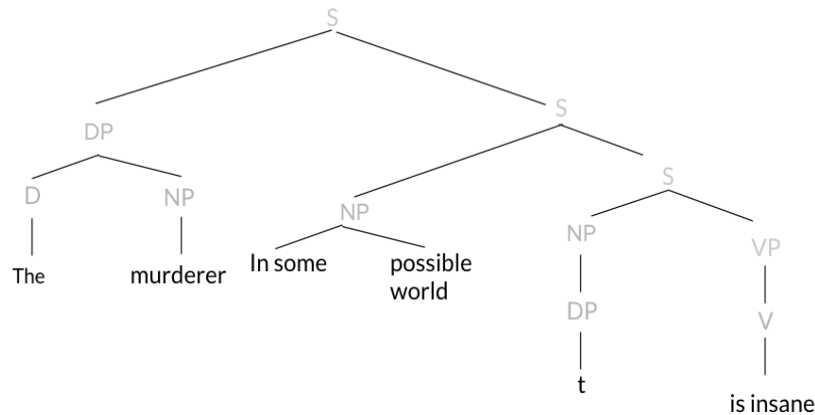


Figure 4: Wide Scope – Object-independent

The above phrase-structure in another syntactic disambiguation of ‘The murderer might have been insane.’ It corresponds roughly to *there exists some unique murderer x (in the actual world) such that there exists some possible world w in which x is insane in w* . Logically, the trace is just a bound variable. Nothing is supposed to be out of the ordinary here either. What is supposed to be new is that yet another syntactic disambiguation of ‘The murderer might have been insane’ is now available which is exactly the same as the one in Figure 4, except that the trace is not a variable but a name for the actual Russellian denotation of ‘the murderer’ which produces the object-dependent proposition we are after. But you might say there is some redundancy here. In Figure 4, whether the trace is a variable or a name does not make a difference to truth-conditions, so you might ask, ‘what is the point of this object-dependent disambiguation?’ The point is that we need to account for Donnellan’s observations which are present at the level where there is *no* scope-inducing operator (i.e. the level of directly asserting ‘The F is G’) and the only way to do that syntactically is to form a new sentence with a name as a trace. Having the trace function as a bound variable in Figure 2 achieves absolutely nothing to that end. But there indeed will always be a redundancy with this proposal; the object-dependent disambiguation (under which the trace names the actual Russellian denotation) will always be truth-conditionally equivalent to the disambiguation under which the definite description takes *widest* scope and ‘ t ’ is a variable (i.e., Figure 4).⁶⁷

⁶⁷ Note that this proposal is not committed to the view that Kripke (1977) criticised by saying is that no two-fold distinction (like the referential / attributive distinction) can replace the wide scope vs. narrow scope distinction which is not a two-fold distinction. The referential / attributive distinction cannot replace the wide scope vs. large scope distinction. The current proposal merely adds a different kind of syntactic arrangement to the already existing array.

Chapter 8: The Problem of Misdescriptions

I have so far not said much regarding what is going in cases of misdescriptions. Donnellan used cases of misdescriptions to argue that definite descriptions have meanings that could not be made compatible with the Russellian analysis. This is because Russell's truth conditions are based on property attributions to some object or other. But Donnellan observed that sometimes what is relevant for truth evaluation is an object that did not have the properties attributed to it by the definite description – this is where misdescriptions or incorrect descriptions come in. Misdescriptions wrap two independent issues into one: the issue of *object-dependency* and the issue of *misdescriptions*. The issue of object-dependency (that definite descriptions can sometimes express object-dependent propositions) we have already discussed. I pointed out that sometimes what is relevant for truth-evaluation is the actual satisfier of the condition associated with definite descriptions, and not the condition itself. We motivated this without appealing to misdescriptions (or incorrect referential definite descriptions). I tried to justify Donnellan's claims in §2.1 by appealing to our intuitions about the counterfactual truth value of the proposition expressed using a referential definite description that is both correct and complete of the object being referred to. It seemed to me that one can use a complete and correct definite description to express an object-dependent proposition about an object that is the unique satisfier of the description. This explains why the following expression (where subscript 'R' stands for 'Referential'):

- The_R murderer of Smith is insane,

Expresses a proposition that would intuitively be true in a counterfactual scenario in which Jones (the unique satisfier of the definite description in the world of utterance) is indeed insane but is not a murderer. I said that Donnellan *could* have motivated his *referential* use in this way, thereby freeing himself from the more controversial claims about the semantic status of misdescriptions. We then went through some data from demonstratives and some pronouns that indicated that the directly referential analysis is not plausible for all their uses. That is, there were some uses of demonstratives and pronouns that did not involve any speaker reference or demonstratives; some of these were embedded in environments that forced a narrow scope interpretation – which is not predicted by a directly referential analysis. All this was done with the assumption that the terms being used involved no misdescription.

But we now have to confront the problem of misdescriptions. The first thing to note about misdescriptions is that they are a special case of the more general problem of *misdesignation*, which occurs knowingly or unknowingly with many designational terms.

For example, I can use the following referential designational terms (in the following expressions) incorrectly:

1. *Jones* is raking the leaves.
2. *He* is quite tall.
3. *That dog* is going to get run over.
4. *The_R guy over there drinking a martini* is sad.
5. *Hitler* has arrived.

In each of the above cases, it seems that the object that would be contributed by the first term's semantics may diverge from the object that is being talked about (and is understood to be being talked about) in conversation. In 1, I use 'Jones' to refer to Smith. In 2, I use 'he' to refer to Sally; a contextually salient female. In 3, I use 'that dog' to refer to a contextually salient small sheep that only looks like a dog. In 4, I use 'the_R guy drinking a martini' to refer to a contextually salient guy who is only drinking water. And in 5, I use 'Hitler' sarcastically to refer to my boss Danny, who isn't Hitler but who is almost just as evil. This phenomenon is orthogonal to the R/A distinction. The R/A distinction is a distinction between condition and satisfier of condition as propositional contribution, while the phenomenon involving misdesignation (at least in the above referential terms) revolves around the distinction between 'semantic referent' and 'speaker referent' (Kripke, 1977). While these two things often converge, sometimes they diverge. Even if we were infallible and incapable of sarcasm and it was impossible that the semantic and speaker referent should diverge, we would still be able to make the R/A distinction. Similarly, if the R/A distinction is impossible, we would still be able to distinguish between the semantic referent and the speaker referent (at least in the case of referential terms) – if we were fallible and capable of sarcasm. In those misdesignation cases, the divergence is not one between two different *kinds* of propositional contributions. That is, it is not a divergence between *two roles* that the descriptive component of a term can play, but between two different objects.

Now, according to Kripke (1977, p. 264), the semantic referent of a referential term is the object determined by the linguistic conventions governing the term (in addition to whatever the semantics is sensitive to in the extra-linguistic environment), but the speaker referent of a term is determined by who the speaker wishes to talk about in context. Now, of course, those two objects are typically identical. For example, if I say under normal circumstances, 'Tom Cruise is happy', then the semantic referent and speaker referent of 'Tom Cruise' converge on the same object – Tom Cruise. That is, 'Tom Cruise' determines Tom Cruise by virtue of the linguistic conventions of that name (i.e., perhaps by virtue of there being the right causal/historical connections) *and* I also use that name because I am thinking about Tom Cruise and I intend my audience to recognise

that it is Tom Cruise that I am talking about by virtue of my use of 'Tom Cruise'. But of course, those two objects (i.e., the semantic/speaker referent) can diverge, and for different reasons. For example, they may diverge due to some *mistake* or perhaps because I am being sarcastic. Let us look more closely at such cases.

I can use the expression in 5, under the right circumstances, to mean that Danny (my evil boss) has arrived. Here, the semantic referent is Hitler (the Nazi German) but the speaker referent is Danny. Here, I am using the name 'Hitler' sarcastically; there are no mistakes. But the semantic and the speaker referent may diverge due to some sort of mistake (knowingly or unknowingly). Example 1 is from Kripke. In this example, both the speaker and the audience think that the person they are looking at is called 'Jones', but in fact it is Smith; they are both mistaken unknowingly. Here, again, the semantic referent of 'Jones' is Jones, but the speaker referent (the object the speaker is thinking about and which is such that the speaker is using a term to get their audience to recognise that it is that object they are talking about by virtue of their use of that term) is Smith. A parallel story could be given about referential uses of the designational terms in 2, 3 and 4 that involve mistakes. I and my audience might think that the contextually salient individual in our environment is a male and, therefore, I might use 'he', intending to refer to it. Here, the same distinction between semantic and speaker referent is called for. There may be no object that is the semantic referent of 'he', because there may be no contextually salient male in the environment, but there will indeed be a speaker referent. And this same distinction appears in cases where the speaker knows that they are mistaken but the audience are not aware of the mistake. This will be a case in which I know that the person I am speaking about is (say) *not* drinking a martini but my audience thinks he is, and so I just use that referential definite description because it may not be cooperative not to, since the person's drinking water might be neither here nor there for the purposes of conversation. Or take the case of 3. Suppose I use 'that dog' to refer to what I think is a dog (which, unbeknownst to me, is in fact a small sheep, as my audience knows). My audience knows that I have made a mistake, but that mistake may be irrelevant for the purposes of conversation, as what truly matters is that an animal is about to die.

Now what exactly is the problem that arises from cases of misdesignation? There is a truth-conditional problem in all the above cases, it seems, that linguistic terms have hooked on to an inappropriate object (inappropriate because the constraints imposed on the object by their semantics are violated). The truth-conditional problem is given in Donnellan's classic observation that you can use 4 to talk about a guy who is only drinking water. But, as I have been saying, the observation can be made about many other designational terms. For example, in Kripke's case, it seems that in some way I would have uttered a truth using 1 if Smith (the person I use 'Jones' to refer to) is raking the leaves. But how can that be when I used the term 'Jones'? Kripke offers a story along the

lines of the story that seems correct in the case of 4: a conversational implicature story.⁶⁸ According to Kripke, we can capture this truth-value intuition by distinguishing between what is strictly *said* vs. *what is meant* (or what is *con conversationally implicated*). That is, in the case of 5 (for example), although I said *Hitler has arrived*, I conversationally implicated *Danny has arrived*. And in the case of 1, although I said *Jones is raking the leaves*, what I meant (and conversationally implicated) was *Smith is raking the leaves*.⁶⁹ Kripke's claim is that the proposition communicated using a misdesignational term is communicated via the action of uttering and not merely by the words themselves (and whatever extra-linguistic features they are semantically sensitive to).

Now, I am perfectly willing to accept the Kripkean story about misdesignational terms. But in the above thesis, I have tried to show that there is still an R/A distinction independently of the phenomenon of misdesignation. But you may now start to question the whole of my project above. Since we, along with Kripke, have now said that actions involving words also have content (via the mechanism of conversational implicature) you might now ask *why then can't this also be what is going on in cases of referential definite descriptions?* You will claim that the action involving a Russellian definite description is used by the speaker to get the audience to think about some object *o*. And of course, then, that object becomes relevant for truth-evaluation and not the condition that it satisfies. But how far can this claim be extended? Will you also say that uses of what are classically conceived as directly referential terms like names and indexicals (involving no mistakes) like 'I' (used to get my audience to think about myself) are merely non-semantically directly referential? No. If you do, you will have no principled basis for accepting the observation in one case and rejecting in the other. But there is a principled basis for thinking that this observation (in the case of misdesignation) is non-semantic because, for example, Sally in the case of 2 is just not a male! I take this to be an obvious, non-disputable violation of the semantics associated with 'he', regardless of whether 'he' is treated in a Russellian or a Fregean way.

Finally, I want to emphasise that there are two routes to the existence of an R/A distinction. The first is the one that Donnellan and some of his earlier critics took, which is based on the phenomenon of misdescription (a special case of the more general phenomenon of misdesignation). I described such a route in this chapter. My view regarding the status of this distinction is that it is pragmatic, as many of those writers would agree (e.g., Kripke, 1977; Neale, 1990; Searle, 1979; Bertollet, 1980). But my view regarding the R/A distinction, arrived at in this first way, is a different matter altogether. That is, there is another independent route to the R/A distinction that I have thus far been

⁶⁸ Kripke (1977, p. 273, fn. 19) seems explicit that it is conversational implicature that he has mind when talking about what is *meant*.

⁶⁹ I presume Kripke would want to extend this kind of response to 2,3 and 4 too.

following. This route, as we saw, does not rely on misdescription or misdesignation. Instead, it relies on our truth-value intuitions of certain constructions involving these terms (some of which are classically treated as object-independent, and some of which are classically treated as object-dependent), the systematicity between the relevant meanings and the pervasiveness of the distinction. To put the point differently, there are two ways in which one can argue that definite description admits of an R/A constitutive distinction. The first (perhaps more familiar approach) is by the following argument from misdescriptions:

1. Some uses of definite descriptions are Russellian. [Premise]
2. Some uses of definite descriptions can be used to utter a truth about some object *o* that does not satisfy the description. [Premise]
3. Those uses cannot be Russellian and must contribute an object to the proposition. [2, Best Explanation]
4. Conclusion: There is an R/A distinction for definite descriptions.

But what I wanted to emphasise in this chapter is that, despite the status of the R/A distinction arrived at via an argument like the above, there is an independent way to arrive at a very similar R/A distinction. This is the path I have followed (we have called this 'the argument from non-misdescriptions):

1. Some uses of definite descriptions must be Russellian. [Assumption]
2. Some uses of correct (and complete) definite descriptions can be used to express a proposition that is true in some arbitrary world *w* iff the actual Russellian denotation of 'the F' is G, regardless of whether it is F in *w*. [Premise from CI]
3. Those uses cannot be Russellian and must contribute an object to the proposition. [2, Best Explanation]
4. There is an R/A distinction for definite descriptions. [Conclusion]

But I have also been presenting an argument from the above conclusion 4 to the semantic status of the R/A distinction, as follows:

1. There is an R/A distinction for definite descriptions. [Assumption]
2. Russellian ways of explaining the R/A distinction via conversational implicature are not compelling. [Claim 1]
3. There is a systematic relation between these two meanings. Although they depend on the context, they are not arbitrary. [Claim 2]
4. Other terms, such as directly referential terms, themselves exhibit the R/A distinction. That is, the R/A distinction is pervasive. [Claim 3]

5. The R/A distinction is semantic but not a product of lexical ambiguity.
[Conclusion]

I justified Claim 1, 2 and 3 and I then went on to present two treatments of the R/A distinction, one based on intention-sensitive indexicality, and the other based on syntactic ambiguity.

Conclusion

I started by linking Donnellan's attributive uses of definite descriptions to Russellian definite descriptions and linking referential uses of definite descriptions to Frege's definite descriptions. This gave us an R/A distinction in terms of definite descriptions that contribute an object (Fregean uses) and those that contribute a property-based condition to a proposition (Russellian uses) This is more or less how many already conceive of the distinction. However, I have tried to argue for a misdescription-free version of this distinction via *the argument from non-misdescriptions*. I argued by process of elimination (given CI) that Fregean definite descriptions contribute an object to the proposition (as opposed to a property-based condition, in the case of Russellian uses). But then I argued that the Russellian way of explaining away the Fregean use pragmatically via Gricean tools gave the Russellian no principled way to favour their theory over the Fregean theory of definite descriptions. Given this deadlock, I explored the systematicity of the relevant meanings, which told against lexical ambiguity treatments. Then, we explored the pervasiveness of the distinction, which made us appreciate that Donnellan's observations were not just very local or idiosyncratic to one class of phrases, but that the R/A distinction is really a *distinction*. I then claimed that this called for a unified treatment across all R/A terms. I presented two such treatments.

References

1. Almog, J., & Leonardi, P. (Eds.). (2012). Referential Uses and the Foundations of Direct Reference. In Almog, J., & Leonardi, P. (Eds.), *Having in Mind: The Philosophy of Keith Donnellan* (pp. 1-256) Oxford University Press USA.
2. Amaral, F. S. (2008). Definite descriptions are ambiguous. *Analysis*, 68(4), 288-297.
3. Bach, K. (2008). What Does it Take to Refer? In *The Oxford Handbook of Philosophy of Language*. Oxford University Press.
4. Bertolet, R. (1980). The Semantic Significance of Donnellan's Distinction. *Philosophical Studies*, 37(3), 281-288.
5. Bezuidenhout, A. (1997). Pragmatics, semantic undetermination and the referential/attributive distinction. *Mind*, 106(423), 375-409.
6. Coulmas, F. (1986). Reported speech: Some general issues. In F. Coulmas, (Ed.), *Direct and Indirect speech* (pp. 1-28). Berlin: Mouton de Gruyter.
7. Donnellan, K. (1966). Reference and Definite Descriptions. *The Philosophical Review*, 75(3), 281-304.
8. Donnellan, K., Almog, J., & Leonardi, P. (2012). Speaker Reference, Descriptions, and Anaphora. In *Essays on Reference, Language, and Mind*. Oxford University Press.
9. Devitt, M. (2004). The Case for Referential Descriptions. In Reimer, M., & Bezuidenhout, A. (Eds.), *Descriptions and beyond*. Clarendon Press.
10. Devitt, M., (2007). Referential Descriptions and Conversational Implicatures. *European Journal of Analytic Philosophy*, 3(2), 7-32.
11. Devitt, M., & Sterelny, K. (1999). *Language and reality: An introduction to the philosophy of language*. Mit Press.
12. Elbourne, P. (2013). *Definite descriptions*. Oxford University Press.
13. Elbourne, P. (2011). *Meaning: a slim guide to semantics*. Oxford University Press.
14. Gómez-Torrente, M. (2015). Quantifiers and referential use. In *Quantifiers, quantifiers, and quantifiers: themes in logic, metaphysics, and language* (pp. 97-124). Springer International Publishing.
15. Grice, H. P. (1981). Presupposition and conversational implicature, in P. Cole (eds.), *Radical pragmatics*, Academic Press, New York, 183-198.
16. Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
17. Huang, Y (2007). *Pragmatics*. Oxford University Press.
18. Kaplan, D. (1989), Demonstratives, in J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan* (481-563), Oxford University Press.
19. Kaplan, D. (1978). Dthat. In Peter Cole (Ed.), *Syntax and Semantics* (221-243). Academic Press.
20. King, J. C. (2001). *Complex demonstratives: A quantificational account*. MIT Press.

21. Kratzer, A., & Heim, I. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
22. Kripke, S. (1977). Speaker's Reference and Semantic Reference. *Midwest studies in philosophy*, 2(1), 255-276.
23. Ludlow, Peter, 'Descriptions', *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/fall2018/entries/descriptions/>](https://plato.stanford.edu/archives/fall2018/entries/descriptions/).
24. Ludlow, P., & Neale, S. (1991). Indefinite descriptions: In defence of Russell. *Linguistics and philosophy*, 14(2), 171-202.
25. Marti, G. (1995). The essence of genuine reference. *Journal of Philosophical Logic*, 24(3), 275-289.
26. Mendelsohn, R. L. (2010). Referential/attributional: a scope interpretation. *Philosophical studies*, 147(2), 167.
27. Miller, A. (2007). *Philosophy of language*. Taylor & Francis.
28. Neale, S. (1990). *Descriptions*. MIT Press.
29. Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and philosophy*, 15(5), 509-559.
30. Neale, S. (2007). On location. In M. O' Rourke & C. Washington (Eds.), *Situating semantics: Essays on the philosophy of John Perry* (pp. 251-395). Cambridge, MA: MIT Press.
31. Nunberg, G. (2004). Descriptive Indexicals and Indexical Descriptions. In Reimer, M., & Bezuidenhout, A. (Eds.), *Descriptions and beyond*. Clarendon Press.
32. Nunberg, G. (1993). Indexicality and deixis. *Linguistics and philosophy*, 16(1) pp. 1-43.
33. Over, D. E. (1985). Constructivity and the referential/attributional distinction. *Linguistics and Philosophy*, 8(4), 415-429.
34. Partee, B. (2002). Noun phrase interpretation and type-shifting principles. *Formal semantics: The essential readings*, 357-381. Wiley.
35. Recanati, F. (1993). *Direct reference : from language to thought* . Blackwell.
36. Russell, B. (1919). *Introduction to mathematical philosophy*. George Allen and Unwin, London.
37. Russell, B. (1905). On denoting. *Mind*, 14(56), 479-493.
38. Salmon, N. (2004). The good, the bad, and the ugly. In Reimer, M., & Bezuidenhout, A. (Eds.), *Descriptions and beyond*. Clarendon Press
39. Saul, J. M. (2002). Speaker meaning, what is said, and what is implicated. *Nous*, 36(2), 228-248.
40. Sadock, J. M. (1978). On testing for conversational implicature. In P. Cole (Ed.), *Syntax and Pragmatics, Vol.9: Pragmatics*. Academic Press: New York.
41. Schiffer, S. (1995). Descriptions, indexicals, and belief reports: some dilemmas (but not the ones you expect). *Mind*, 104(413), 107-131.
42. John R. Searle. (1979). Referential and Attributional. *The Monist*, 62(2), 190-208.

43. Sharvy, R. (1980). A more general theory of definite descriptions. *The philosophical review*, 89(4), 607-624.
44. Stalnaker, R. (2014). *Context*. Oxford University Press.
45. Stalnaker, R. Indicative conditionals. *Philosophia* 5, 269–286 (1975).
46. Stanley, J., & Gendler Szabó, Z. (2000). On quantifier domain restriction. *Mind & Language*, 15(2-3), 219-261.
47. Stokke, A. (2010). Intention-sensitive semantics. *Synthese*, 175(3), 383-404.
48. Vendler, Z. (1967). Singular Terms. In *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.
49. Vignolo, M. (2012). Referential/Attributive: the explanatory gap of the contextualist theory. *Dialectica*, 66(4), 621-633.
50. Wettstein, H. K. (1981). Demonstrative reference and definite descriptions. *Philosophical studies*, 40(2), 241-257.
51. Yagisawa, T. (1985). The Referential and the Attributive: A Distinction in Use?. *The Southern Journal of Philosophy*, 23(1), 109-125.
52. Higginbotham, J. (1988). Contexts, models and meanings: A note on the data of semantics. In R. Kempson (Ed.), *Mental representations: The interface between language and reality* (pp. 29–48). Cambridge: Cambridge University Press.
53. King, J., & Stanley, J. (2005). Semantics, pragmatics, and the role of semantic content. In Z. G. Szabó (Ed.), *Semantics vs. Pragmatics* (pp. 111–164). Oxford: Oxford University Press.