

# **LANGUAGE, THOUGHT AND REFERENCE**

**GEORGE POWELL**

**Thesis submitted in partial fulfilment  
of the requirements for the degree of Ph.D.**

**Department of Phonetics and Linguistics  
University College London**

**2003**



## Abstract

How should we best analyse the meaning of proper names, indexicals, demonstratives, both simple and complex, and definite descriptions? In what relation do such expressions stand to the objects they designate? In what relation do they stand to mental representations of those objects? Do these expressions form a semantic class, or must we distinguish between those that are referential and those that are quantificational? Such questions have constituted one of the core research areas in the philosophy of language for much of the last century, yet consensus remains elusive: the field is still divided, for instance, between those who hold that all such expressions are semantically descriptive and those who would analyse most as the natural language counterparts of logical individual constants.

The aim of this thesis is to cast new light on such questions by approaching them from within the cognitive framework of Sperber and Wilson's Relevance Theory. Relevance Theory offers not just an articulated pragmatics but also a broad conception of the functioning of natural language which differs radically from that presupposed within (most of) the philosophy of language. The function of linguistic expressions, on this conception, is not to determine propositional content, but rather to provide inferential premises which, in parallel with context and general pragmatic principles, will enable a hearer to reach the speaker's intended interpretation. Working within this framework, I shall argue that the semantics of the expressions discussed should best be analysed not in terms of their relation to those objects which, on occasions of use, they may designate, but rather in terms of the indications they offer a hearer concerning the mental representation which constitutes the content of a speaker's informative intention. Such an analysis can, I shall claim, capture certain key data on reference which have proved notoriously problematic, while respecting a broad range of apparently conflicting intuitions.

## Contents

<b>Abstract</b>	<b>2</b>
<b>Acknowledgements</b>	<b>6</b>
 <b>Chapter 1: Introduction</b>	 <b>7</b>
1.1 Reference: a brief history	10
1.2 Language, mind and reference	16
1.3 Structure	17
 <b>Chapter 2: The Theoretical background</b>	 <b>18</b>
2.1 A general view of mind and cognition	18
2.2 Thoughts of individual particulars	21
2.2.1 Individual concepts – an outline	21
2.2.2 Two types of individual concept	22
2.3 Relevance-theoretic pragmatics	31
2.4 Relevance-theoretic semantics	33
2.5 Explicature vs. what is said	35
2.6 Semantic theory and intuition	37
 <b>Chapter 3: Proper Names</b>	 <b>43</b>
3.1 The central data on proper names	45
3.1.1 Co-reference	45
3.1.2 Emptiness	47
3.1.3 Rigidity	48
3.2 Proper names and individual concepts	50
3.2.1 The bare bones	50

3.2.2	Names, individual concepts and questions 1 to 3	56
3.2.3	Individual concepts and the core data	65
3.2.3.1	Co-reference	65
3.2.3.2	Empty names	76
3.2.3.3	Individual concepts and rigidity	93
3.3	Names, individual concepts and alternative accounts	98
3.3.1	Neo-descriptivism on names	99
3.3.2	Proper names as indexicals	110
3.3.3	Neo-Millianism on names	116
<b>Chapter 4: Indexicals</b>		<b>118</b>
4.1	The central data on indexicals	120
4.1.1	Indexicals and co-reference	121
4.1.2	Indexicals and emptiness	125
4.1.3	Indexicals and rigidity	127
4.1.4	Indexicals, answer machines and post-it notes	129
4.2	Indexicals and individual concepts	131
4.2.1	The bare bones	131
4.2.2	Indexicals, individual concepts and the core data	147
4.2.2.1	Indexicals, individual concepts and co-reference	148
4.2.2.2	Indexicals, individual concepts and emptiness	161
4.2.2.3	Indexicals, individual concepts and rigidity	162
4.2.2.4	Indexicals, individual concepts, answer machines and post-it notes	169
4.3	Conclusion	175
<b>Chapter 5: Definite Descriptions</b>		<b>178</b>
5.1	The debate on descriptions	179
5.1.1	Donnellan's position	180



5.1.2	Yes-theory vs. no-theory	183
5.2	The semantics and pragmatics of definite descriptions	192
5.3	Conclusion	205
<b>Chapter 6: Complex demonstratives</b>		<b>207</b>
6.1	Referentialism	208
6.2	Quantificationalism	214
6.3	Hybrid accounts	222
6.4	Complex demonstratives and communication	224
6.5	A semantics for complex demonstratives	230
6.6	Defending the analysis	233
<b>Chapter 7: Conclusion</b>		<b>237</b>
7.1	The class of singular expressions	237
7.2	What kind of account is this?	240
7.3	Singular expressions, individual concepts and semantic theory	243
<b>Bibliography</b>		<b>245</b>

## Acknowledgements

My first thanks go, as ever, to my supervisor, Robyn Carston. Over the last few years, she has offered not just stimulating, challenging discussion on the topic of this thesis, and many other subjects besides, but also endless support, encouragement and time. She has shown a faith in my work that I have not always felt was well-placed but I have always felt was very welcome. Thank you Robyn.

I would also like to thank Deirdre Wilson, who has always been unstintingly generous in her willingness to offer help and advice. Her patience in the face of ceaseless questioning while stuck in economy class on the way to this or that conference is greatly appreciated.

Thanks also go to the members of the pragmatics reading group at UCL who have offered me both the opportunity to discuss my own work in a stimulating environment and an introduction to fascinating areas of pragmatics of which I would otherwise have been unaware. I thank in particular Richard Breheny, Richard Horsey and Tim Wharton.

There have been times during the writing of this thesis when I have felt very close to losing my way. At those times, Corinne Iten and, in particular, Milena Nuti, have offered the friendship, support and insight that have helped me find my way again, for which I thank them both from the bottom of my heart.


And finally, a very big thank you to Laura Winton, who has put up with living with an overstressed and undercommunicative Ph.D. student with extraordinarily good humour. Without her support this thesis would never have been finished. Thank you.

## Chapter one

### Introduction

#### 1.0

Around 150AD Claudius Ptolemy published a book that was to define astronomy for the next 1400 years, a book that has come to be known as the *Almagest*. In the *Almagest*, Ptolemy set out to provide a mathematical model of the motion of the Sun, the Moon and the five known planets, taking as his starting point two of the key tenets of Aristotelian cosmology: that the Earth was at the centre of the universe and that the motion of the heavenly bodies was perfectly circular, i.e. that heavenly bodies moved in a circular orbit at a constant rate around the centre of that orbit. Hampered from the start by these premises, Ptolemy devised, in his attempt to represent the observed motion of the planets, a model of spectacular sophistication but also of spectacular complexity. In the system of the *Almagest*, the planets are taken to move in circular orbits, with the centre of their orbits themselves orbiting around points which are



slightly removed from the position of the Earth. The centres of the planetary orbits do not, however, move at a constant rate around the centre of their orbits, but around a point, known as an *equant*, which is at the same distance from the centre of their orbit as the Earth on a line passing through the Earth and the centre of their orbit (a typical Ptolemaic planetary orbit is illustrated by the diagram on the left)<sup>1</sup>.

This model represents the motion of the planets with reasonable accuracy; it has, however, certain problems on its own terms: although close to the centres of the orbits of the centres of their orbits, the Earth is not itself taken to be at the centre of the planets' orbits, thus violating one of the key Aristotelian tenets. Furthermore, planets are not taken to move constantly around the centre of the centre of their orbits, thus violating the other key Aristotelian tenet. Despite its enormous complexity and its failure to respect, at least to the letter, its own foundational principles, it would not be until 1543, with the publication of Copernicus' *De revolutionibus orbium coelestium*,

---

<sup>1</sup> If you find this exposition of Ptolemaic cosmology less than entirely clear, <http://www.history.mcs.st.ac.uk/story/Mathematicians/Ptolemy.html> offers a lucid presentation of the key elements of the *Almagest*.

that any serious challenge would be mounted to the Ptolemaic system and the assumptions on which it is built. Copernicus' key insight was that it was possible to do away with Ptolemy's complex theoretical machinery, his eccentrics, epicycles and equants, by the simple step of abandoning one of his foundational principles, the principle of geocentricity. If one places the Sun, rather than the Earth, at the centre of one's cosmology, one no longer needs to postulate planets moving constantly around points which are not only not the centre of their orbits, but not even the centre of the orbit of the centre of their orbits.

In 1892 Gottlob Frege published *Über Sinn und Bedeutung*, a paper that has been one of the foundation stones of the philosophy of language for the last century. In the paper, Frege mounts an argument that looks something like this: it seems clear that expressions such as proper names refer to things; why not, then, assume that all linguistic expressions refer? If that is the case, however, then whole sentences must refer; but what sort of thing might they refer to? We can assume that whatever serves as the reference of a sentence must be determined by the references of that sentence's constituent parts. Now this offers us a strategy for finding out what sentences refer to: it should be the case, if the reference of a sentence is determined by the references of its parts, that, if we take one expression in a sentence and swap it for another with the same reference, the reference of the sentence as a whole should remain unchanged. But what remains unchanged in such circumstances? Nothing but the sentence's truth value. Therefore, sentences must refer to their truth values. This argument, and later more sophisticated versions of it, have had a profound and lasting effect on the philosophy of language: they have located truth fairly and squarely at the centre of theories of linguistic meaning.

This thesis will be concerned with many of the questions that have been at the heart of the philosophy of language, and, in particular, the philosophy of reference, since Frege placed truth at the centre of meaning: what do singular expressions, expressions such as proper names, pronouns and definite descriptions, mean? What do they contribute to the truth conditions of sentences in which they appear? How do they make this contribution? Must we distinguish between those singular expressions which denote and those which refer? If we answer 'yes' to this question, must we make a parallel cognitive distinction between ways of thinking of entities? If, on the other hand, we answer 'no', how can we account for intuitive differences between, for

instance, proper names and definite descriptions? I shall also explore the issues which provide a background to these questions: how might we want to individuate propositional content? What sorts of evidence should we rely on in the search for a theory of meaning? What role should be assigned to speaker intuition?

My aim will be to address these questions from within a particular theoretical framework, the communication-based framework offered by Sperber and Wilson's Relevance Theory (see, for instance, Sperber and Wilson (1986/95)). Sperber and Wilson offer not just an articulated pragmatics, but a conception of meaning which is fundamentally different from that presupposed within much of the truth-conditional literature: on this conception, the role of language is not to encode (or determine) propositional content, but rather to provide clues as to the content of a speaker's informative intention. Thus the hearer of an utterance will retrieve that utterance's propositional content not merely by decoding the sentence uttered, but by a combination of decoding and pragmatic inference. By approaching the meaning properties of singular expressions from within this framework, I hope to suggest that *Über Sinn und Bedeutung* has had a similar effect on twentieth century philosophy of language as the *Almagest* had on mediaeval cosmology: putting truth at the centre of theories of linguistic meaning has led to the construction of a great number of semantic eccentrics, epicycles and equants. This is not to say that truth has no place in theories of meaning, and, in particular, in theories of the meaning of singular expressions. It is clearly the case that utterances have truth conditions, and that tokens of singular expressions contribute to the truth conditions of utterances in which they appear. But this fact no more shows that truth must be at the centre of theories of meaning than the fact that the Earth has a place in any satisfactory model of the solar system shows that we should return to Aristotelian geocentricity. One of the key claims of this thesis, then, will be that we should oust truth from the central place it has held in those theories of meaning that have held sway for the last century. But what should we replace it with? The theory which I shall outline over the following chapters suggests that we should, in essence, replace truth with mind. My claim will be that the eccentrics of twentieth century semantics are the product of a failure to appreciate the nature of the mediatory role that mind plays between language and the world.

## 1.1 Reference – a brief history

It should already be clear that an examination of the notion of reference will be central to this thesis. But what are we talking about when we talk about reference? The fundamental intuition is that certain types of expressions, expressions such as proper names, demonstratives, definite descriptions and so on, pick out particular individual entities in order that we may talk about them. There is an intuitive contrast between these sorts of expression and those which seem to be meaningful not by dint of picking out individuals. So, if we consider sentence (1):

### 1) Bert is happy

it looks like ‘Bert’ and ‘happy’ are doing fundamentally different kinds of thing: ‘Bert’ picks out a particular individual, whereas ‘happy’ picks out a property.

This then is the starting point for discussions of reference: an intuitive distinction between the kind of work done by different types of expression. Over the course of the last century in particular, this distinction has been subject to extensive refinement. To set the scene for what follows, and to give a sense of what is at stake in debates on the nature of reference, I want to examine the most significant of those refinements. This examination will, however, be brief: my aim is not primarily to present the refinements themselves, since they are familiar, but rather to draw out their underlying concerns and to focus on those concerns that are common to all; we shall also have occasion to return to some of the most influential accounts of reference in later chapters.

Although Frege established the terms in which the modern discussion of reference is conducted, the origin of that discussion can be traced back beyond Frege, at least to Mill (1843). The intuitive distinction between those expressions which serve to pick out individuals and those which do not finds a corollary in Mill’s work in the distinction between *individual names* and *general names*:

A general name is ... a name which is capable of being truly affirmed,  
in the same sense, of each of an indefinite number of things. An

individual or singular name is a name which is only capable of being truly affirmed, in the same sense, of one thing.

Mill (1843, p. 131)

Mill draws another distinction, within the class of individual names, that has been pivotal to the modern philosophy of language. For Mill, proper names are distinct from other individual names, such as definite descriptions, in that their meaning properties are exhausted by their referents, i.e. by that which they name. Put in Mill's own terms, all individual names other than proper names have both *connotation* and *denotation*, i.e. both descriptive meaning and referent, whereas proper names have only denotation. Given what we shall have to say about later accounts of reference, it is perhaps unfortunate that Mill talks of the referent of a proper name as its denotation; its referent, however, is what he is talking about. This distinction between connotative and non-connotative names prepares the ground for much that follows: the question of how one takes the distinction can be seen as one of the key areas of disagreement between opposing camps in the debate on reference.

As the discussion above may already have made clear, the intuitive distinction between those expressions which refer and those which do not is somewhat blurred in Frege's semantic system. For Frege all meaningful expressions may refer, not just those which we would normally think of as doing so. This does not mean, however, that the intuitive distinction outlined above is entirely lost for Frege. In his semantic system it corresponds not to a division between those expressions which refer and those which do not, but rather to a division between those expressions which have saturated meanings and those which have unsaturated meanings<sup>2</sup>. Without going into detail on this distinction, it corresponds closely to the idea that some expressions serve to pick out individuals, while others serve to pick out properties which may be true or false of those individuals. The intuitive distinction we started off with survives, then, albeit in altered form, in Frege's framework. Given the underlying principles which inform Frege's semantic system, however, Mill's distinction between connotative and non-connotative individual names is lost. For Frege, one of the key facts about meaning is that linguistic expressions present what they refer to in a particular way. This idea ties in with the foundational epistemological notion that

---

<sup>2</sup> For a presentation of this distinction, see, for instance, Frege (1891).

we cannot think of an individual other than as the satisfier of certain properties: I cannot entertain a bare thought about the desk I am working at, for instance; I must instead think of it as that which is presented to me in such-and-such a way. Given this epistemological slant there can be, for Frege, no non-connotative names: if a name were non-connotative, it would seem that to grasp the meaning of that name would be simply to think of its referent, without thinking of that referent in any particular way. All singular expression must therefore have descriptive meaning, i.e. sense<sup>3</sup>. Not only does this view tie in with the fundamental tenets of Frege's epistemology, it also offers a way out of a puzzle that has dominated the philosophy of reference for a century, a puzzle that has come to be known as *Frege's puzzle*. We shall have much to say about Frege's puzzle in the coming chapters, particularly when we come on to discuss the semantics of specific types of singular expression. In its simplest form it looks like this: if Mill is right, then proper names lack any meaning beyond their referent. But if that is so, then there can be no difference in meaning between, for instance, (2) and (3):

- 2) Batman is Bruce Wayne
- 3) Bruce Wayne is Bruce Wayne

After all, (2) and (3) differ only in the substitution of one name for another name with the same reference. On the assumption that names have no meaning other than reference, (2) and (3) must therefore have the same meaning. Yet this is clearly wrong, according to Frege: (2) and (3) must have different meaning properties for the simple reason that (2) is potentially informative, i.e. is cognitively significant, whereas (3) is not. This puzzle has been largely responsible for setting the parameters within which discussion of reference has been conducted; in particular it has placed the relation between meaning and cognitive significance centre stage. Frege's assumption is that a plausible semantic theory must be answerable to facts concerning cognitive significance. Whether this is so, or whether, as some recent contributors to

---

<sup>3</sup> In identifying Fregean sense with descriptive meaning I believe that I am not only being faithful to Frege's intention, but I am also aligning myself with the majority interpretation of Frege's work. My interpretation is, however, at odds with that of Evans (see, in particular, Evans 1982) who sees the possibility in Frege's work of non-descriptive senses.



the debate would have it<sup>4</sup>, semantic facts and cognitive facts are independent, will be a theme that we will return to on more than one occasion.

A system such as Frege's, on which to be meaningful is to have descriptive meaning, conforms closely with intuition in its analysis of expressions such as definite descriptions and complex demonstratives. But then, as far as these expressions are concerned, Frege's semantics have much in common with Mill's: on both accounts, the meaning of these types of expression must be analysed along two dimensions, for Mill, connotation and denotation, for Frege, sense and reference. The interesting case is clearly going to be the analysis of proper names. The Millian intuition, that names lack descriptive meaning, seems robust: on the face of it, there seems to be no descriptive condition which an individual must satisfy in order to be the referent of 'Condoleezza Rice', she simply has to be Condoleezza Rice. Yet this intuition is beyond the grasp of a semantic system such as Frege's: meaningful as they are, expressions such as 'Condoleezza Rice', must indeed have a sense, maybe the same sense as the description 'George W. Bush's National Security Advisor'. I don't want to go any further into this issue here; we shall pick it up again when we come to discuss the semantics of proper names in chapter 3. The point I want to stress, since we are examining the intuitive basis for the debate on reference, is that in this respect, Frege's semantic system seems, at least, counterintuitive. Once we do return to this question in chapter 3, however, it will turn out that there is a way of looking at Frege's account on which it seems less counterintuitive.

Mill's distinction between connotative and non-connotative individual names resurfaces in Russell's account of reference, based, as it is, on an epistemology which differs from Frege's in fundamental respects<sup>5</sup>. I shall have a great deal to say about Russell's theory of descriptions when I come on to look at the semantics of definite descriptions in chapter 5. Here I shall restrict myself, therefore, to the briefest discussion. Underlying Russell's semantics is a distinction between two ways in which we can think about individuals: we can think of an individual either in an essentially Fregean way as the satisfier of some property or set of properties, or in a much more immediate way which Russell terms *knowledge by acquaintance*. This distinction tallies with some robust intuitions: there seems to be a fundamental

---

<sup>4</sup> See, for instance, Wettstein (1986).

<sup>5</sup> For Russell's semantic and epistemological views on reference see, in particular, Russell (1905) and Russell (1911).

difference between the way you can think, for instance, about your next door neighbour and the way you can think about the oldest woman in the world (whoever she may be); while you can entertain thoughts about both, that, for instance, your next door neighbour left his rubbish on your doorstep again and that the oldest woman in the world can't be looking forward to the future much, your thoughts about your next door neighbour seem, on an intuitive level, to be linked to their object in a way that your thoughts about the oldest woman in the world are not. For Russell, this cognitive distinction has a linguistic corollary: some expressions, which Russell terms *logically proper names*, are anchored to their referents in the same way as acquaintance-based thoughts: they do not present their referents as the satisfiers of any set of properties but rather simply label those referents. Given their lack of any mode of presentation, Russell's logically proper names can be identified with Mill's non-connotative individual names. We can thus see the re-emergence of a distinction that has since become fundamental to the philosophy of language between, on the one hand, certain classes of singular expression which refer directly to their designata, and, on the other, classes of singular expression which operate on something like a Fregean model. As I shall discuss at some length in the chapters that follow, the assumption embodied by Russell's account, that truth-conditional distinctions between types of cognitive entity must carry over to types of linguistic entity, has become part of the bedrock of modern truth-conditional semantics.

While Russell drew a hard and fast distinction between logically proper names, on the one hand, and definite descriptions on the other, applying this distinction to the semantics of natural language proved problematic. As we shall see when we come on to look at the semantics of proper names in chapter 3, the assumptions upon which his epistemology was based forced Russell to some notoriously unpalatable conclusions, conclusions that undermine the intuitive basis for his distinction between logically proper names and descriptions. Despite the unpalatability of these conclusions, the fundamental distinction drawn by Russell between those singular expressions which serve as no more than labels for their designata and those which designate via the satisfaction of a property, between, in other words, those expressions which refer and those which denote, has become one of the main planks of the currently dominant school of thought on reference. As we saw above, a Fregean model of reference seems to work fine (speaking roughly) for

expressions such as definite descriptions and complex demonstratives. It seems to work less well, or, at least, is less in tune with pre-theoretical intuition, when it comes to an analysis of expressions such as proper names which seem, on the face of it, to lack descriptive meaning. As Kripke has shown (Kripke 1972 etc), this is far from being the only problem we face if we apply a Fregean model to proper names. A Fregean story on names depends on the claim that, when we successfully refer by use of a proper name, this must be because our use of the name corresponds to a mode of presentation, a sense, which uniquely picks out that to which we have referred. Yet, as Kripke shows, we may refer to a particular individual using a proper name even if every piece of information we associate with the name is false of that individual, or, indeed, where we have no uniquely identifying information concerning that individual at all.

Such observations have heralded a return to the Millian model over the last thirty years or so. The key tenets of the new version of Millianism, espoused by Kripke, Kaplan, Perry et al., are that we must distinguish between those expressions which refer by virtue of their linguistic meaning and those which denote by virtue of their linguistic meaning, and that those expressions which are referential contribute nothing to truth-conditional content but their referents. Even in the neo-Millian analysis of referring expressions, however, we can still make out the reflection of a Fregean two-tier semantics. For Kaplan (in particular 1989a), the key distinction between referring and denoting expressions does not lie in whether or not their linguistic meaning is descriptive; both, on his semantic system, may have descriptive linguistic meaning. Rather it depends on what role that descriptive meaning plays with regard to propositional content. What distinguishes referring expressions, as a class, from denoting expressions, as a class, is that whatever descriptive meaning the former have drops out of the picture prior to the level of propositional content, serving only to constrain reference. The latter by contrast contribute their descriptive meaning to propositional content. Given this analysis, Kaplan distinguishes between the linguistic meaning of a referring expression, its *character* in his terms, and its *content*, i.e. that which it contributes to propositional content which, on his account, is nothing

but its referent (or, more accurately but less perspicuously, a constant function from circumstances of evaluation to its referent)<sup>6</sup>.

Neo-Millianism can be viewed as an attempt to develop the Russellian distinction between logically proper names and denoting expressions into a semantic theory for natural language. Any such attempt will, of course, have to confront those data around which the Fregean model was designed, and in particular Frege's puzzle itself. After all, if Mill's original story fell foul of Frege's puzzle, what reason might there be to suppose that a new version will not also come to grief on it? Much ink has been spilled in answering just this question, and many variants of the basic Millian picture have been developed as a result.

## 1.2 Language, mind and reference

Regardless of the differences between different flavours of neo-Millianism, more or less all accept the fundamental Russellian assumption that there is a direct correlation between type of expression linguistically individuated and type of propositional contribution truth-conditionally individuated. Put another way, it is assumed that one type of expression must always give rise to the same type of propositional contribution. It is this assumption, above all, that has led, I shall claim, to the eccentricities, epicycles and equants of modern semantic theory. Why should this be so? There is, I shall argue, extensive evidence that all types of singular expression can be used either referentially or descriptively; they can, that is, give rise to either singular or general truth conditions. If, however, one works on the assumption that expression types must always give rise to the same type of propositional constituent, truth-conditionally individuated, then some sophisticated theoretical machinery is going to be called for to explain this apparent data. You have two choices: either you claim that all uses of your chosen singular expression give rise to singular propositions, or you claim that they all give rise to descriptive propositions. Either way, you need some way of explaining how it is that there are apparent examples in which your singular expression gives rise to the wrong kind of proposition. Much use has been made of the Gricean notion of implicature in this cause, particularly in an analysis of

---

<sup>6</sup> The identification of character with linguistic meaning is clearest in Kaplan (1989a).

what Donnellan<sup>7</sup> terms the *referential use* of definite descriptions, of which we will see much more in chapter 5.

Why, rather than employ this theoretical machinery, should we not simply accept, in light of the strong evidence available, that singular expressions can make different types of propositional contribution in different contexts? There are two key answers that are either explicit or implicit in the literature: firstly, any account on which singular expressions can be either referential or descriptive is thereby an ambiguity account, and we should, on theoretical grounds, prefer accounts that can do without positing ambiguities. I shall argue in the final chapter of the thesis that to view the account I favour as an ambiguity account would be to misunderstand why avoiding ambiguity is a goal worth pursuing for the semantic theorist. The second reason why the one-expression-one-propositional contribution assumption has remained so pervasive is this: if we abandon it, then we abandon the hope that we might be able to explicate the linguistic meaning of singular expressions in terms of the notion of truth, and, since Frege placed truth at the centre of meaning, it has largely been taken on faith that truth offers our best hope of a satisfactory theory of meaning. I shall argue that, at least as far as singular expressions are concerned, we must bite the bullet, accept that truth cannot lie at the heart of theories of linguistic meaning, and thus excise the Ptolemaic machinery which has become a familiar feature of the modern philosophy of language.

### 1.3 Structure

I propose to structure the thesis as follows: in chapter 2 I shall lay out the theoretical machinery which I intend to use in building my theory of the semantics of singular expressions. This machinery will come in three flavours: cognitive, pragmatic and semantic. Once I've laid out the theoretical machinery, I shall then, in chapters 3 to 6, examine the semantics and pragmatics of different types of singular expression, outlining the account of each that I favour and comparing it with the main accounts currently on the philosophical market. Finally in chapter 7 I shall address how the accounts I will by then have developed fit in with broader semantic concerns.

---

<sup>7</sup> See, in particular, Donnellan (1966).

## Chapter 2

### The theoretical background

#### 2.0

In this chapter I want to undertake some preliminary groundwork, groundwork upon which the account of the semantics and pragmatics of singular expressions which I shall develop in chapters 3 to 6 will build. At the heart of my thesis is a reassessment of the relations between semantic, pragmatic and cognitive facts in an account of singular expressions. It will thus be necessary to provide both background and technical machinery for each of these dimensions of the account. I shall not, however, address them in this order: the view of semantics which I shall be working with depends upon a proper understanding of the role played by pragmatic inference in the determination of propositional content, while the nature of our pragmatic machinery depends on prior facts about human cognition. I shall thus start with some cognitive scene-setting, and, once that is in place, move on to look at the pragmatic and semantic issues that underlie the claims I shall be making in the rest of the thesis.

#### 2.1 A general view of mind and cognition

As discussed in chapter 1, the account of the semantics and pragmatics of singular expressions which I aim to defend will make extensive use of the relevance-theoretic framework developed by Sperber and Wilson (e.g. 1986/95). Given this, I am committed from the outset to certain claims on the nature of mind which are embodied in that framework; in particular the account I shall develop presupposes a representationalist view of mind. That is to say, I shall take the contents of propositional attitudes to be representations in something like a language of thought. As Fodor (see, in particular Fodor (1975)) has convincingly argued, all vaguely plausible views on the operation of mind presuppose computation, and computation presupposes something like a machine code. It thus seems that all plausible views on the mind presuppose some form of language of thought.

Placing my account within a relevance-theoretic framework also commits me to certain claims about human mental architecture; in particular I am committed to the

view that mind is modular in something like the sense of Fodor (1983)<sup>1</sup>. The details of what such a claim must amount to have been the subject of much recent discussion, within the relevance-theoretic framework and elsewhere. It is beyond the scope of this thesis to argue for any one particular take on the notion of mental modularity. I want, however, to lay out those features of the modular view which will be significant for the kind of account of singular expressions which I wish to develop.

Fodor (1983) presents a view on which the architecture of the mind is divided between input modules and central systems. Among the hallmarks of modular processes is that they are domain specific, mandatory and fast; central processes, by comparison, are domain general, optional and relatively slow. Given these typical feature sets, language processing seems, for Fodor, to be a good candidate for modularity. This view, that the mind contains a language processing input module which takes as its input linguistic stimuli and, having processed those stimuli, passes its output on to other mental components, dovetails elegantly with the relevance-theoretic view of language and communication. I don't want to say much more about this here, since I will return to the interrelations between relevance and modularity when I lay out the semantic and pragmatic background to my account later in this chapter. In brief, however, relevance theorists take the distinction between those interpretive processes which are conducted within the language module and those which are conducted at a higher cognitive level to correspond to a distinction between linguistic decoding and pragmatic inference. As I shall discuss below, the view that the encoded content of linguistic stimuli goes no further than whatever the output of a dumb language processing module may be has profound implications for the semantics-pragmatics interface.

In addition to these general claims about the nature of mind and mental architecture, the relevance-theoretic analysis of communication rests on a specific claim about human cognition: that it is guided by the search for relevance. For Sperber and Wilson (Sperber and Wilson (1986/95), Wilson and Sperber (2002) etc.), relevance is taken to be a property of inputs to cognitive processes, be those inputs external (environmental stimuli) or internal (e.g. thoughts, memories etc.). As far as the determination of relevance is concerned, such inputs have two key features: on the

---

<sup>1</sup> It might, perhaps, be better to say that I am committed to the view that mind is computationally modular in the sense of Segal (1996).

one hand they may offer an agent positive cognitive effects – they may combine with other available information to provide new information, they may reinforce pre-existing information or they may contradict and eliminate pre-existing beliefs – while on the other they will cost the agent effort to process: accessing the available positive cognitive effects will place a lesser or greater demand on the agent's cognitive resources. It is the balance between these two factors that, for Sperber and Wilson, determines the relevance of an input:

#### **Relevance of an input to an individual**

- a. Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.
- b. Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time.

Wilson and Sperber (2002, p. 252)

Sperber and Wilson's claim is that human cognition is geared towards optimising the balance between positive cognitive effects and processing effort. This claim is formalised in their *cognitive principle of relevance*:

#### **Cognitive principle of relevance**

Human cognition tends to be geared to the maximisation of relevance.

Wilson and Sperber (2002, p. 254)

The significance of this principle to our current concerns will not become clear until we examine the relevance-theoretic account of utterance interpretation later in this chapter. As we shall see then, however, utterance interpretation, on the relevance-theoretic view, depends on the exploitation of the cognitive principle of relevance.



## 2.2 Thoughts of individual particulars

### 2.2.1 Individual concepts – an outline

While I shall presuppose, in this thesis, the sort of picture of mind and cognition outlined above, my main focus, as far as cognition is concerned, will be on mental representations of individuals. I shall, throughout, refer to such representations as *individual concepts*. This term seems to me to have the advantage of offering an intuitive way to talk about conceptual representations of individuals. There is, however, a caveat that needs to be raised: I am not using the term in line with the semantic tradition on which an individual concept is taken to be a function from indices to individuals, i.e. the intension of a singular expression. The sense in which I use the term should become clear as this section progresses. I want to underline here, however, that I take individual concepts to be subjective cognitive entities; to be, in other words, whatever representation a cognitive agent entertains when she thinks of a particular individual.

I want to start out, then, with the following idea: that, on a representationalist view of the mind, what it is to have a thought about an individual is to entertain a propositional mental representation one of the components of which is a representation of the individual in question, i.e. one of the components of which is an individual concept of that individual. But what do such concepts look like? What, in other words, makes a concept an individual concept? The answer I want to suggest is broadly familiar within the philosophical literature<sup>2</sup>. I want to identify individual concepts with dossiers containing information all of which is taken by the holder of the concept to be satisfied by the same individual. Thus, for instance, you may have a concept of Humphrey Bogart which contains the information that he was a film star (which I shall henceforth represent as *x was a film star*), that *x acted in *The Big Sleep**, that *x was married to Lauren Bacall* and so on. On the kind of story I am proposing, to say that you have placed these pieces of information in the same dossier, i.e. in the same individual concept, is just to say that you take them all to apply to a single individual. Looking at the same point from a different angle, individual concepts are

---

<sup>2</sup> See, for instance, Grice (1969), Recanati (1993), Larson and Segal (1995).

taken by their holders to license particular inference schemas: a cognitive agent who has placed the information that *x is F* and that *x is G* into the same individual concept takes this to license the inference to  $\exists x (Fx \ \& \ Gx)$ .

Although this is only, as yet, the barest outline of an anatomy for individual concepts, there might already seem to be an objection: while it is clearly the case that we do have individual concepts of actual individuals, we also seem capable of entertaining individual concepts of non-actual individuals. I can certainly entertain a concept of Humphrey Bogart, but I can also entertain a concept of Philip Marlowe, a concept that may contain information such as *x is the protagonist of The Big Sleep, Farewell my Lovely etc., x was created by Raymond Chandler, x is a detective in Los Angeles* and so on. But, so the objection might go, on the story I have told so far, my concept of Philip Marlowe cannot be viewed as an individual concept: I take none of the information in my Philip Marlowe concept to be satisfied by a particular individual, since Philip Marlowe is fictional. I cannot therefore, *a fortiori*, take *all* the information in my Philip Marlowe concept to be satisfied by the same individual. I don't want to say a great deal on this question here; I will return to the issues raised by non-denoting individual concepts at some length in the next chapter. What I will suggest there is, in essence, that the apparent problems posed by empty individual concepts are not best addressed by altering our story on the anatomy of individual concepts, but should rather be addressed via a distinct theory of fictional contexts (or, more broadly, of non-actual contexts).

### 2.2.2 Two types of individual concept

What I have claimed so far, then, is that when we think of individuals we do so by entertaining individual concepts which comprise dossiers of information taken to be true of a single individual. This, however, leaves open a key question: what makes an individual concept a concept of a particular individual? In answering this question we will need to distinguish between two types of individual concept, a distinction that harks back to the work of Russell (e.g. 1905, 1911) .

The theory of meaning developed by Russell depends, as was briefly discussed in the last chapter, on some fundamental epistemological claims. Underlying the distinction, central to Russell's semantics, between logically proper names and

denoting expressions, is the parallel distinction between knowledge by acquaintance and by description. Russell (1911) sketches the distinction thus: 'I say that I am *acquainted* with an object when I have a direct cognitive relation to that object, i.e. when I am directly aware of the object itself'(p. 108); 'I shall say that an object is "known by description" when we know that it is "the so-and-so," i.e. when we know that there is one object, and no more, having a certain property'(p. 113). As mentioned in chapter 1, such a distinction seems to have a robust intuitive basis: it seems, for instance, that my thoughts about my next door neighbour, call him 'Bob', are of a fundamentally different kind from my thoughts about the oldest woman in the world (whoever she may be). As is well known, however, Russell came up against some formidable problems in applying these distinctions, between logically proper names and denoting expressions and between knowledge by acquaintance and knowledge by description, to an analysis of natural language. A brief examination of these problems will, I believe, be instructive.

Russell's account approaches meaning from a strictly realist stance. As Neale (1990, p. 15) puts it, 'every meaningful item of language stands for something real. The meaning of an expression is simply that entity for which it stands.' This semantic view has a necessary epistemic corollary: if the meaning of an expression is just that entity for which it stands, then to grasp the meaning of an expression must just be to grasp the entity it stands for. One implication of this is that, should an expression fail in fact to stand for any entity, then there should be nothing which would count as understanding that expression. The mental objects which constitute our understanding of referential expressions must therefore be such that they cannot exist in the absence of something for them to stand for. Put another way, Russell requires mental modes of presentation that would simply not be available if that for which they stand did not exist. Acquaintance is, for Russell, the relation in which a cognitive agent must stand to an entity, in order to be able to think of it in this object-dependent way. The difficulty is that such mental modes of presentation, such concepts, are very hard to come by. We could be mistaken about the existence of more or less any entity which we believe ourselves to have thoughts about. Thus, I currently believe that I am having a thought about a particular mug to the effect that it is empty but would be better full. It is entirely possible, however, that I am merely hallucinating, that there is in fact no mug. Given this, Russell must conclude that my thought about the mug,

although seemingly as direct as any thought could be, is in fact descriptive; I am thinking of the mug not by acquaintance, that is, but by description. The difficulty of finding any concepts which are truly immune to error, and thus of finding linguistic expressions which might express such concepts, led Russell down some dark alleys. The best known of these alleys is to be found in his account of the semantics of proper names, an account which will receive more detailed attention in the next chapter. For the time being I just want to point to Russell's central thesis on names. As we have seen, to understand an expression is to grasp that which it stands for. It seems intuitively clear, moreover, that I can understand the sentence:

1) Santa Claus is coming to town

despite the failure of 'Santa Claus' to pick out any actual individual<sup>3</sup>. If Russell accepts that I can understand the name 'Santa Claus', as he must, he is forced to the following conclusion: that proper names do not, in fact, stand for individuals. And this is, of course, just the view that Russell held: that proper names are semantically equivalent not to his logically proper names but rather to definite descriptions; that they stand, in other words, not for individuals but for complex quantified structures.

The requirement, imposed on him by his realist framework, that the semantic value of a concept, and thus of any expression which expresses that concept, can only be identified with an individual if the concept is immune to error, robs Russell's epistemology of its intuitive support. Consider once more the intuition we originally used to motivate a distinction between different ways of thinking of things, the intuition that my thoughts about my neighbour Bob are of a different kind from my thoughts about the oldest woman in the world. On Russell's picture this isn't so: my thoughts about both must be descriptive, since I certainly have no concept of Bob which is entirely immune to error. If we want to capture this intuition, then, we need to look for an alternative to Russell's notion of acquaintance. This is not to say that Russellian acquaintance will have no role to play: as I shall come on to discuss shortly, there do seem to be certain concepts which are immune to error, and which may or may not correspond to particular linguistic expressions. However, if we are

---

<sup>3</sup> I'm assuming that the name 'Santa Claus' as we use it should not be taken to refer to Saint Nicholas, a 4<sup>th</sup> century bishop from Asia Minor. Even if we do so take it, however, the same point can be made, since there seems to be some evidence that Saint Nicholas did not, in fact, exist.

interested in capturing the intuitions that support Russellian acquaintance, Russellian acquaintance itself will be too inflexible a theoretical notion. My aim in what follows is to outline an alternative to Russell's distinction which I believe will do the work we want of it. I shall not here have much to say in defence of this outline; it must stand or fall according to how well it allows us to account for the semantic behaviour of singular expressions, and we will not come on to discuss that question until I present my account of singular expressions starting in the next chapter.

The lesson we should learn from the problems facing a Russellian epistemology is that infallibility is a dead end when it comes to capturing the intuitive distinction between these different ways of thinking of individuals. An alternative basis for the distinction has been advocated, in one form or another, by a number of recent philosophers of language (see, for instance, Recanati (1993), Bach (1987)). On this approach, the distinction between, for instance, the way I can think of my neighbour Bob and the way I can think of the oldest woman in the world comes down to a difference between the sort of facts that make each of these ways of thinking a way of thinking of a particular individual. It seems reasonably clear that my concept of the oldest woman in the world attaches to its designatum satisfactorily; this is just to say that it is a concept of whoever satisfies its descriptive content, i.e. of whoever is the oldest woman in the world. It is equally clear, however, that my concept of Bob does not attach to its designatum satisfactorily: making use of the sort of intuitions which Kripke (e.g. 1972) marshals as evidence for the rigidity of proper names, it seems that my concept of Bob may contain information which is false of Bob or may fail to contain information which Bob uniquely satisfies without thereby ceasing to be a concept of Bob. If, then, my concept of Bob is not of Bob because Bob satisfies some or all of its descriptive content, what is it that makes my concept of Bob of Bob? The answer which is now familiar within the philosophy of language and mind is that my concept of Bob is of Bob because of certain causal relations in which Bob and my concept stand; because Bob is, in some appropriate sense, the causal origin of my Bob concept. I don't propose to spend time discussing just which causal relations are appropriate for this corollary of Russellian acquaintance: interesting though this question is, there are many brands of answer currently available on the market which I am sure are along more or less the right lines, and a proper discussion of the question would require an extensive digression

from our central concerns. The point I wish to focus on is that a distinction between those concepts which pick out their designatum satisfactionally and those which pick out their designatum relationally may well offer us an appropriate alternative to Russell's distinction between acquaintance and description.

How will this idea look when recast in the language of individual concepts? What it would suggest is this: we should distinguish between two dimensions of individual concepts, what I shall call their *internal* and *external* dimensions. All individual concepts have structurally similar internal dimensions along the lines suggested above; their internal dimensions will essentially be repositories of open sentences structured in such a way that they support particular inferences. It is thus in terms of facts concerning their internal dimensions that the class of individual concepts as a whole is to be individuated. It is only when we come to look at their external dimensions that we start to see distinctions within the class of individual concepts. For those concepts which pick out their designatum satisfactionally, henceforth, following standard usage, *descriptive* individual concepts, the relation between internal and external dimensions is intimate: essentially one or more elements of the internal dimension are co-opted to do work as the external dimension. In the limiting case, the internal dimension of a descriptive individual concept may comprise a single piece of information; I may, for instance, know nothing at all about the oldest woman in the world. On the view I want to propose, my concept of the oldest woman in the world would thus have identical internal and external dimensions. Given that this concept, as it stands, licenses no inferences, why should I want to say that the unique piece of information that *x is the oldest woman in the world* enters the internal dimension at all? Why should I not rather say that this individual concept has an external dimension which anchors it to whoever is the oldest woman in the world, but has an empty internal dimension? The answer to this concerns the inferences I will be disposed to draw once new information is added. Say, for instance, that I learn that the oldest woman in the world lives in Asia. On the picture painted above, this should lead me to add the information *x lives in Asia* to the internal dimension of my oldest-woman-in-the-world concept. If this were the only piece of information within the internal dimension, I should still not be in a position to draw inferences from the concept. But clearly I am in a position to do so; I will, for instance, be disposed to infer that  $\exists x(\text{Oldest-woman-in-the-world } x \ \& \ \text{Lives-in-China}$

x). Since this is an instantiation of just the inference schema which internal dimensions of individual concepts characteristically support, there seems to be good reason to suppose that we should place a copy of the condition which comprises the external dimension of such a concept in its internal dimension.

What of the more complex case, in which the internal dimension of a descriptive individual concept comprises more than one piece of information? What I want to suggest is that any uniquely identifying condition (or conjunction of conditions) within the internal dimension can serve as external dimension<sup>4</sup>. Thus, consider a concept the internal dimension of which contains the information that *x is the oldest woman in the world* and that *x is the oldest woman in China*. Either of these conditions can serve as external dimension; the bearer of the concept can, in other words, entertain the thought that, for instance, the oldest woman in the world is G, from which she should be inclined to infer that the oldest woman in China is G, as well as the thought that the oldest woman in China is G from which she should be inclined to infer that the oldest woman in the world is G. It seems to me to come down to little more than preference whether we are inclined to say that we have here two distinct individual concepts or one individual concept with distinct external dimensions on distinct occasions of use. Given what I shall come on to say about relational individual concepts, I am inclined to view the two concepts as distinct although, given that they share an internal dimension, intimately linked.

What I have suggested, then, is that there is a close link between the internal and external dimensions of descriptive individual concepts. One upshot of this is that a cognitive agent may not rationally revise all the information within the internal dimension of a descriptive individual concept without also revising its external dimension; she cannot, for instance, jettison the information that *x is the oldest woman in the world* from the internal dimension of a descriptive concept while retaining this condition as the concept's external dimension at pain of entertaining the contradictory belief that the oldest woman in the world is not the oldest woman in the world. This shows up one of the key distinctions between descriptive concepts and those individual concepts which pick out their designata relationally, henceforth (again following common usage) *de re* individual concepts. As suggested by the claim

---

<sup>4</sup> I shall have more to say on the requirement that the condition serving as the external dimension should be uniquely denoting below.

that such concepts pick out their designata relationally, the satisfaction of none of the conditions which comprise its internal dimension is criterial to determining the reference of a *de re* concept<sup>5</sup>. On the view I want to suggest, the descriptive conditions within the internal dimension of a *de re* concept are overridden, as far as the external dimension is concerned, by a causal relation in which the concept stands to a particular individual. Thus my concept of my neighbour Bob may contain the information *x is a teacher*, *x is called 'Bob'*, *x is male*. I can, nevertheless, come to revise all of these pieces of information without rendering my concept contradictory; that is to say that, for any condition *x is G* which appears within the internal dimension of my Bob concept, I can rationally come to believe that Bob is not G.

I have so far laid out a reasonably simple ontology for individual concepts: on the one hand there are descriptive individual concepts which pick out their designata satisfactorily and thus contribute a complex condition to the truth conditions of thoughts in which they occur, while on the other there are *de re* individual concepts which pick out their designata relationally and thus contribute nothing but their designata to truth-conditions. On this picture, all individual concepts are of a kind as far as their internal dimensions are concerned, the difference between the two distinct types arising only at the external level. As it turns out, however, things are not quite as straightforward as this. Firstly, although the causal link that determines the reference of a *de re* concept is a matter of that concept's external dimension, the *fact* of the link must be internally represented. Recanati (1993) gives an analysis of the semantics of what he takes to be directly referential expressions on which these expressions encode, as part of their linguistic meaning, a feature which Recanati terms *REF*. This *REF* feature serves to mark an expression as directly referential, i.e. limits the truth-conditional contribution of the expression in question to its referent. To give a flavour of what is to come, I shall argue throughout this thesis that singular expressions are not marked by dint of their linguistic meaning as either referential or descriptive. They may, however, given appropriate context, give rise to referential or descriptive readings. Individual concepts, by contrast, are marked as either referential or descriptive: that is to say, we have intuitions concerning whether particular concepts are linked to their reference in the intimate way which Russell identified with knowledge by acquaintance, or the indirect way identified with knowledge by

---

<sup>5</sup> The following owes much to the two-component picture defended by Recanati (1993).



description. It thus seems a natural move to transfer Recanati's REF feature from the encoded meaning of particular linguistic expressions to the internal dimension of particular individual concepts, specifically those individual concepts which are *de re*. Why should we want REF to be a constituent of the internal dimension of *de re* concepts? Firstly to account for our intuitions concerning the truth conditions of certain thoughts. But secondly, this move allows us to capture the intuitions underlying Russell's distinction between knowledge by acquaintance and by description without the problems inherent in Russell's epistemology. Recall the intuition with which we originally motivated some such distinction, the intuition that my thoughts about my next door neighbour are of a fundamentally different kind from my thoughts about the oldest woman in the world. Now we cannot rely on actual causal chains to capture this distinction for just the reasons that so constrained Russell's account: my next-door-neighbour concept seems contentful, whether or not there is a neighbour for it to be a concept of. In the terms just developed, however, what distinguishes my next-door-neighbour concept, whether or not it is in fact causally linked to any individual, from my oldest-woman-in-the-world concept is the presence of REF within its internal dimension. What distinguishes the two concepts, in other words, is that I believe one but not the other to be causally linked to an individual in the way appropriate for *de re*-ness. I am not, of course, intending to suggest that this will be the only difference standardly to be found between the two types of concept: *de re* concepts will, for instance, standardly be weighted towards perceptual information in a way that descriptive concepts will not. It is on the presence or absence of the REF feature, however, that, so I want to claim, the intuitive distinction between these different types of concept is based.

There is a further refinement I wish to make to the conceptual ontology so-far sketched. As things stand there are two truth-conditional possibilities for individual concepts: either the external dimension of an individual concept is satisfactional and that concept gives rise to general truth conditions, or the external dimension is relational and the concept gives rise to singular truth conditions. There is, however, a third truth-conditional possibility: that the external dimension is descriptive but the truth conditions are singular. What does this possibility amount to? It essentially equates to placing the descriptive condition which constitutes the external dimension of a descriptive individual concept within the scope of a rigidifying operator

something like Kaplan's dthat operator (see, for instance, Kaplan (1978)). The result of this move is to anchor a concept truth-conditionally to whatever actually satisfies its external dimension. The significance of this third possibility, which I shall label *stipulative de re*-ness, will become apparent in the next chapter.

Finally in this section, I just want to point towards a distinction drawn by Recanati (1993), among others, between two types of *de re* individual concept. Recall that we started this section with a discussion of the difficulties Russell faced in applying his distinction between knowledge by acquaintance and knowledge by description to an analysis of natural language. The difficulties, as we saw then, rested on Russell's requirement that acquaintance-based thoughts should be immune to error in that they simply could not exist in the absence of that which they are thoughts about. The types of individual concept which we have looked at so far are certainly not immune to error: I may have just the same thoughts (viewed from an internal perspective) that I have now about, for instance, my neighbour Bob even if Bob does not exist – if I have, say, been the victim of an elaborate hoax. There do seem to be some kinds of thought, however, which are immune to error in just the way required by Russell's epistemology. The clearest and most discussed case is that of first-person thoughts, i.e. thoughts about oneself. It seems intuitively clear that the thought that you might express with the sentence 'I am F' simply could not fail to be about you and therefore could not exist if you did not exist. The same seems to go for the thoughts expressed by 'now is G' and 'here is H': your now-is-G thought simply must be about the time of its own tokening, and your here-is-H thought must be about the place of its tokening. As Recanati points out, the concepts underlying these thoughts are of a very different kind from other *de re* concepts: they do not standardly comprise the sort of beliefs which we have seen so far, but rather comprise predominantly perception-based information. Again, I don't wish to defend any particular view of the nature of such concepts, which, following Recanati, I shall call *egocentric*; I am sure the analysis Recanati proposes is along very much the right lines. I merely want to flag up this distinction, which will become significant when I come on to discuss certain indexical expressions in chapter four.

### 2.3 Relevance-theoretic pragmatics

Now that I have put in place the cognitive machinery which I shall be making use of in the rest of the thesis, I want to turn to the pragmatic background against which my account of the interpretation of singular expressions will be set. I have already laid out, in section 1.2 above, the cognitive assumptions underlying Sperber and Wilson's relevance theory: that the representational view of mind is broadly right, that mind is modular and that human cognition is geared towards the maximisation of relevance. Sperber and Wilson take these assumptions as the foundations upon which their account of communication and the human pragmatic machinery which underlies it is built.

On the relevance-theoretic view, the task facing a hearer is to form and confirm hypotheses concerning speaker intentions. In a case of what Sperber and Wilson (e.g. 1986/95) call *ostensive-inferential communication*, i.e. any case in which we would intuitively want to say that full communication takes place, these intentions will come in two varieties: firstly the speaker will have an intention to inform her hearer of something, her *informative intention*, and secondly, following a broadly Gricean model, she will have the intention to inform her hearer of her informative intention, an intention that Sperber and Wilson term the *communicative intention*. But how might a hearer set about forming and confirming hypotheses concerning a speaker's informative and communicative intentions? In doing so he may rely on the fact that communicators exploit the human cognitive tendency to maximise relevance. Given this tendency, i.e. given the cognitive principle of relevance, a hearer will only attend to a stimulus if it offers him sufficient cognitive effects without undue processing effort, i.e. if it is relevant enough to him to be worth attending to. A speaker who produces an utterance with the intention that a particular hearer should attend to it therefore communicates a presumption that her utterance is at least relevant enough to the hearer to be worth processing. Wilson and Sperber spell out the nature of this communicated presumption in their *communicative principle of relevance*:

### **Communicative Principle of Relevance**

Every ostensive stimulus conveys a presumption of its own optimal relevance.

Wilson and Sperber (2002, p. 256)

They cash out the notion of *optimal relevance* used in this formulation thus:

#### **Optimal relevance**

An ostensive stimulus is optimally relevant to an audience iff:

- a. It is relevant enough to be worth the audience's processing effort;
- b. It is the most relevant one compatible with communicator's abilities and preferences.

Wilson and Sperber (2002, p. 256)

How, then, does the communication of a presumption of optimal relevance help the hearer in his attempt to form and confirm hypotheses on speaker intentions? It does so by justifying a particular interpretive strategy: given the communicative principle of relevance a hearer is justified in testing interpretations in order of their accessibility until he comes across one that satisfies his expectations of relevance. He is then entitled to stop and accept this interpretation. Why does the communication of a presumption of optimal relevance justify this strategy? The key lies in clause (b) of the definition of optimal relevance: a hearer is justified in trying out hypotheses in order of accessibility since a speaker who wants to deliver on the presumption of optimal relevance should make the intended interpretation as accessible as possible in order to minimise processing effort; equally a speaker wishing to maximise the relevance of her utterance should not allow for there to be more than one interpretation which satisfies the hearer's expectations of relevance on pain of putting her hearer to the effort of choosing between competing interpretations. Since, therefore, there can only be one interpretation which will satisfy the hearer's expectations of relevance, as soon as he finds such an interpretation, the hearer can accept it.

Finally in this section I want to flag up a point that I shall return to later in this chapter. In addition to the distinction drawn by Sperber and Wilson between informative and communicative intentions, I shall draw a further distinction between

different kinds of speaker intention which, so I shall claim when I come on to discuss certain types of example later in the thesis, play an important role in determining intuitions concerning the truth-conditional content of utterances. I shall, however, leave the details of this distinction for a later section and turn now to the semantic picture underlying the account of singular expressions which I wish to defend.

## 2.4 Relevance-theoretic semantics

Relevance-theoretic pragmatics, then, takes a hearer's job to be to the formation and confirmation of hypotheses concerning a speaker's informative and communicative intentions. But what role does hypothesis formation and confirmation have to play in grasping speaker meaning? According to a traditional view within the philosophy of language the propositional content of an utterance is determined by the meaning of the sentence uttered. On this view, the sort of inferential processes involved in the relevance-theoretic interpretation strategy may well have a role to play in the retrieval of implicatures, but have no role to play in determining the proposition expressed by an utterance. There are some respects in which linguistic meaning underdetermines propositional content, according to this view; in particular a hearer will have to resolve ambiguities and assign reference to indexical expressions before he can grasp the proposition expressed by an utterance. However, these are not matters of inferential input to propositional content, but are rather matters which are straightforwardly resolved by context.

Work within the relevance theoretic tradition and elsewhere has shown firstly that an analysis of disambiguation and reference assignment requires more than a general appeal to context and that the more required can be provided by the machinery needed to analyse the retrieval of implicatures, and secondly that the underdetermination of propositional content by linguistic meaning goes a great deal deeper than ambiguity and referential indeterminacy<sup>6</sup>. Such evidence supports a conception of the semantics-pragmatics distinction which differs fundamentally from that presupposed by much work within philosophical and formal semantic traditions. On the relevance-theoretic view, the retrieval of truth-conditional content involves

---

<sup>6</sup> For an overview of this work, and a detailed discussion of some of the ways in which sentence meaning underdetermines propositional content, see Carston (2002a).

two distinct types of process: on the one hand there is the dumb decoding of encoded linguistic meaning by a dedicated language processing module, while on the other there is higher-level pragmatic inference, which takes as input the output of the language processing module and builds not just implicated meaning but also propositional content. On this view, the output of the language module is standardly (if not invariably) sub-propositional, viewed as something like a proposition schema. This schema will provide just one type of input to the inferential processes involved in forming and confirming hypotheses concerning speaker intentions.

The relevance-theoretic analysis thus recognises a sub-propositional level of representation which equates to the output of the language processing module, a level of representation sometimes referred to as *logical form* within the RT-oriented literature. Beyond logical form, it has become usual within relevance theory to draw a distinction in the class of communicated assumptions between those which are explicitly communicated and those which are implicitly communicated, between *explicatures* and *implicatures* in the RT terminology. This distinction mirrors, although is not exactly coterminous with, the Gricean distinction between what is said and what is merely communicated. In the next section I shall briefly compare these two sets of distinctions in an attempt to spell out some of the points at which the kind of semantic picture which underlies the analysis I shall propose diverges from conceptions of semantics which are traditional in the philosophical literature.

Before coming on to that, however, I want to point towards one further distinction that has become increasingly central to the relevance-theoretic view of semantics. Within the RT framework there has been much work conducted over the last few years into a distinction originally proposed by Blakemore (1987) between two fundamentally different kinds of meaning. Blakemore's original insight was, roughly, this: since utterance interpretation involves two radically different kinds of process, first the retrieval of a logical form and then the manipulation of that logical form through pragmatic inference, we might well expect to find two distinct kinds of meaning corresponding to these processes. And this, for Blakemore, is just what we do find. On the one hand there is *conceptual* meaning, which is straightforwardly representational; on the other there is *procedural* meaning, which lays constraints on the retrieval of implicatures. This distinction has been extensively applied to the analysis of non-truth-conditional meaning within the relevance-theoretic framework.

Thus, for instance, both Blakemore herself (see, for instance, 1989, 2000) and Iten (e.g. 1998, 2000a, 2000b) have employed the conceptual-procedural distinction in an analysis of concessives.

Over the last few years the application of this distinction has been steadily extended within the relevance-theoretic framework. In particular, the assumption that procedural meaning is involved only in the retrieval of implicatures has been challenged. Wilson and Sperber (1993), among others, have suggested that indexicals may encode procedures rather than, or as well as, concepts, procedures that are therefore involved in the retrieval of explicatures. Throughout this thesis I shall return to the question of whether the conceptual-procedural paradigm offers interesting insights into the encoded meaning of different types of singular expression.

## 2.5 Explicature vs. what is said

As discussed above, there is a distinction drawn within the relevance-theoretic framework between logical form and explicature, i.e. between the encoded meaning of a sentence and something which looks very much like the truth-conditional content of a located utterance of that sentence. There are certain key differences, however, between that which relevance theorists take to be an explicature of an utterance and what would be thought of as the proposition expressed by that utterance, or ‘what is said’, within many philosophical frameworks. Sperber and Wilson (e.g. Sperber and Wilson 1986/95, 1998; Wilson and Sperber 1992) and Carston (e.g. 1988, 1996/1997, 2000, 2002a), among others working within the RT framework, have pointed to certain key points at which the philosophical notion of what is said diverges from the level of explicitly communicated meaning relevant to an analysis of utterance interpretation. Irony and metaphor, for instance, are both treated within RT as contributing to the level of explicit content, i.e. as contributing to explicature, while on philosophical analyses these tropes are standardly analysed as part of what is implicated by an utterance. Recently there has also been much debate in the literature on so-called *unarticulated constituents*, elements which, according to certain criteria, appear to form part of the truth-conditional content of an utterance while seemingly unrepresented in the syntax of the sentence uttered. For those of a pragmatic bent, such as Carston (e.g. 2002b) and Recanati (2002), such elements should be viewed as

genuine constituents of propositional content, constituents contributed not by syntax but by pragmatics. For those of a more traditional philosophical bent, such as Stanley (2000), by contrast, if such elements are genuine constituents of propositional content, then they must be represented within the syntax.

What is at stake here? To answer this question we need to look at the theoretical requirements made of the notions of *explicature* and *what is said* by those frameworks in which they are employed. In order to do this I shall have to generalise in a way that may seem unpalatable; there are, after all, wide discrepancies between the takes on the notion of *what is said* or *propositional content* or *the proposition expressed* employed by different frameworks within what might broadly be called the truth-conditional tradition, and there is certainly not the space here to do justice to each of these frameworks separately. In what follows, therefore, my aim is to tread as neutral a path between different takes on the notion of propositional content as possible.

Common to all truth-conditional accounts is the idea that the relatively tractable notion of truth can be used to approach the relatively intractable notion of meaning. To put this idea into practice, a link must be established between linguistic representations and truth. This link is effected via the assumption, which comes in various different guises, that to give the meaning of a sentence is to give the conditions under which that sentence would be true. Now if the truth-conditionalist wants to give the meaning of a sentence via that sentence's truth-conditions, she will need, for any unambiguous sentence, to find a way of giving truth conditions which are constant across contexts of utterance. But this is just what the more pragmatically-minded meaning theorist is claiming not to be possible: if, for instance, metaphor should be viewed as contributing to truth-conditional content, then the same unambiguous sentence can, on different occasions of use, give rise to different truth conditions. To maintain the link between syntactic form and truth required by her framework, therefore, the truth-conditionalist must deny that metaphor does contribute to truth-conditional content. Equally, the hard-line truth-conditionalist is committed to the position that apparently unarticulated constituents must either be genuine constituents of propositional content but articulated at some level of the syntax, or genuinely unarticulated but not constituents of truth-conditional content.



The dyed-in-the-wool truth-conditionalist is thus committed to a fully propositional level of truth-conditional content which is intimately anchored to the sentence uttered. The relevance theorist, by contrast, is, as we have seen above, committed to a sub-propositional level of representation which is intimately anchored to the sentence uttered and one or more propositional levels of representation which, although linked to the sentence uttered, may depart from encoded sentence meaning in fundamental ways. Carston (2002a) argues that the philosophical level of what is said can, with certain caveats, be eliminated from an account of utterance interpretation. One of the themes I wish to explore in what follows is whether such a level, viewed as a propositional level determined by sentence meaning, can also be omitted from a theory of meaning. I clearly do not want to suggest that one could not work with this level of representation if one so wished. What I want to explore, however, is the following possibility: that, at least as far as an analysis of the meaning of singular expressions is concerned, one can get away without an appeal to the traditional philosophical notion of what is said, while one cannot get away without including in one's ontology objects that look very like logical form and explication. To pave the way for the exploration of this idea, I want to say a few words about the place of intuitions in semantic theorising.

## 2.6 Semantic theory and intuition

When we talk about the truth-conditional content of an utterance, what is it that we are talking about? Looking at this question from another angle, are there any objective, theory-independent criteria by which we might say what the truth-conditional content of a particular utterance is, i.e. by which we may determine what is said by that utterance? There have certainly been attempts made to pinpoint such criteria. Within the relevance-theoretic framework it has been suggested that the distinction between explicitly and implicitly communicated assumptions comes down to whether a particular communicated assumption is a development of the logical form of the sentence uttered. Carston (1988) offers a possible cashing out of this suggestion in what she calls the *Functional Independence Principle*. This principle rests on the idea that explicatures and implicatures play independent roles in inferential processes. Thus, speaking very roughly, any communicated assumption

which is a development of logical form should be viewed as an explicature, whereas any communicated assumption which is not a development of logical form should be viewed as an implicature. Recanati (1989b/91), however, shows that Carston's principle throws up some unwanted results; in particular he shows that the principle categorises some clearly implicated communicated assumptions as explicatures.

Recanati himself (e.g. 1989b/91) offers two more promising candidate criteria for discriminating explicitly communicated meaning from implicitly communicated meaning. The first candidate, derived from a test used by Cohen (1971), Recanati terms the *Scope Principle*:

*Scope Principle*: A pragmatically determined aspect of meaning is part of what is said (and therefore not a conversational implicature) if – and, perhaps, only if – it falls within the scope of logical operators such as negation and conditionals.

Recanati (1989b/91, p. 114)

Thus Recanati asks us to consider the following sentences:

- 2) The old king has died of a heart attack and a republic has been declared
- 3) A republic has been declared and the old king has died of a heart attack

It seems that there is a temporal sequence interpretation associated with the use of 'and' in these examples; we take it, in other words, that there is a difference in the ordering of events between (2) and (3). The question we want to address, however, is whether this temporal ordering is part of what is said by utterances of these sentences, or whether it arises at the level of implicature. In order to apply the scope principle we must embed the sentences under a logical operator:

- 4) If the old king has died of a heart attack and a republic has been declared, then Tom will be quite content.
- 5) If a republic has been declared and the old king has died of a heart attack, then Tom will be quite content.

It seems, Recanati points out, that a rational agent could assert (4) at the same time as denying (5), and that (4) and (5) must thus express distinct propositions. This seems to offer us strong evidence that the temporal sequencing associated with 'and' falls within the scope of the conditional and should therefore, according to the scope principle, be treated as an element of what is said<sup>7</sup>.

The second criterion Recanati offers for distinguishing what is said from what is merely implicated is embodied in his *Availability Principle*:

*Availability Principle*: In deciding whether a pragmatically determined aspect of utterance meaning is part of what is said, that is, in making a decision concerning what is said, we should always try to preserve our pre-theoretic intuitions on the matter.

Recanati (1993, p. 248)

The import of this principle is self-evident: if intuitions suggest that a particular utterance has a particular truth-conditional content, then we should take that as strong evidence that that utterance does indeed have that truth-conditional content. Although this principle is the only one to make overt mention of the place of intuition in theorising about meaning, all the principles we have so far looked at depend for their utility on intuitions concerning what is said. This is hardly surprising, since both Carston and Recanati are in the business of providing cognitively realistic accounts.

Of course such principles are at odds with the take on the notion of propositional content employed by those who seek to develop a pure truth-conditional account of meaning by anchoring what is said by an utterance to the syntax of the sentence uttered. What place is there for truth-conditional intuitions on accounts of this type? There seem to be two distinct trends within the recent hard-line truth-conditional literature. On the one hand there are those, such as Stanley (2000), who are prepared to take intuitions as genuine evidence concerning truth-conditional content. For theorists of this stripe, as mentioned above, any element of meaning which intuitively forms part of truth-conditional content must be represented within

---

<sup>7</sup> There are, of course, other stories that may be told about such examples, including that told by Cohen himself. My interest here is not in defending the Scope Principle, although I find the arguments that Recanati provides in support of the principle convincing; rather I am interested in examining the sort of criteria that have been proposed according to which we may distinguish between explicitly and implicitly communicated meaning.

the syntax, albeit covertly. On the other hand there are those, such as, for instance, Borg (forthcoming), who take the view that truth-conditional intuitions are a very poor guide to what is 'literally' said. For Borg, the literal meaning of any sentence can be identified with a fully determinate proposition, although such propositions will by and large be a long way from the assumptions which a speaker wishes to communicate. Even accounts of this sort, however, standardly make appeal to truth-conditional intuitions. Borg, for instance, makes the claim that interlocutors can be relied upon to have intuitions concerning the distinction between literal sentence meaning and speaker meaning; in other words, that they have reliable intuitions concerning what is 'literally' said. She then makes use of such evidence in defence of her claim that supposedly unarticulated constituents should be viewed as non-semantic, i.e. as contributing not to what is said but to what is implicated.

Why this digression through various takes on the notion of what is said and the role of truth-conditional intuition in each? The first lesson I want to draw is that, however much we might want to, we cannot abandon such intuitions altogether; they are, after all, one of the few types of evidence available to us if we want our take on what is said to have an empirical dimension. I thus propose to take intuitions on truth-conditional content seriously throughout this thesis, along much the lines suggested by Recanati's Availability Principle: it seems to me that a theory which respects such intuitions should, *ceteris paribus*, be preferred to one that rides roughshod over them. It is in satisfying the *ceteris paribus* caveat, of course, that difficulties may arise.

Secondly, however, I want to point to an obvious problem in attempting to build a theory of meaning which respects intuitions on truth-conditional content: such intuitions are notoriously variable. To take just one example, the field on the referential-attributive distinction in the interpretation of definite descriptions is divided between those who find it intuitively plausible that a referentially-used definite description may contribute to propositional content an individual other than that which satisfies its descriptive content, and those who find such a notion intuitively incredible. Each camp uses their own intuitions in support of their particular position, often simply refusing to recognise opposing intuitions. I believe that, if we are to take truth-conditional intuitions seriously at all, then we are honour bound to take competing intuitions seriously too; we should, that is to say, expect an

analysis of the semantics and pragmatics of singular expressions to be able to account for contradictory intuitions where they arise.

There are two suggestions I want to make here which may help us when it comes to explaining contradictory intuitions later in the thesis. Firstly, split intuitions may be reconciled by taking them to concern not the content of one and the same utterance but rather the content of two distinct utterances. Clearly the sort of split intuitions we are concerned with here are intuitions about utterances of the same sentence, but it is not so clear that those who claim to have contradictory intuitions are holding the context in which the sentence is located constant. In particular, as I shall discuss at some length in later chapters, I want to explore the possibility that standard philosophical notions of what is said tend to presuppose particular types of discourse contexts, discourse contexts which are idealised in certain significant respects. Such contexts are, I shall suggest, typically idealised with respect to the epistemic states of their speakers and hearers: the intuitions adduced by hard-line truth-conditionalists often presuppose a discourse context in which speaker and hearer have mutual knowledge of all relevant contextual factors. This stands in clear juxtaposition to the relevance-theoretic view outlined above, on which limitations on the speaker's abilities, built in to the definition of optimal relevance as they are, play a key role in determining the explicature of an utterance. I shall return throughout the thesis to the idea that contrary intuitions which seem to bear on the semantics of particular expressions may in fact bear on no more a difference in the contexts in which the utterances are taken to be located.

Finally in this chapter, I want to return to something I flagged up at the end of section 2.3. Sperber and Wilson, as has already been discussed, distinguish between a speaker's informative intention and her communicative intention. I wish to draw a further distinction between types of speaker intention. To say that a speaker has an informative intention is just to say that she has an intention concerning which propositional representation(s) she wishes to communicate to her audience. What I want to suggest is that, separate from the informative intention itself, a speaker will have specific intentions concerning the route via which her audience should reach the content of her informative intention; I shall call this the speaker's *derivational* intention. To give a taste of this distinction, let's look at a brief example. Bert and

Ernie are on a country walk when Bert notices a particularly lovely butterfly hovering near a flower. Pointing in the direction of the butterfly, Bert says:

6) That is very pretty

It seems clear that, speaking in rough terms, the content of Bert's informative intention is the thought that the butterfly he is pointing at is very pretty. But, in addition, Bert has intentions regarding how Ernie should reach this thought, i.e. regarding the strategy which Ernie should follow in deriving the intended interpretation: he intends, for instance, that Ernie should take his ostensive gesture as part of his communicative act, that he should therefore follow the line of Bert's finger to see what Bert is pointing at, that the act of pointing should thus have the effect of making a particular object, the butterfly, highly salient in the context, and that the butterfly should therefore be the most accessible candidate interpretation for Bert's utterance of 'that'. In this way, he maximises the relevance of an interpretation on which the butterfly is taken to be the referent of 'that' by minimising the processing effort to which Ernie must go in order to access such an interpretation<sup>8</sup>. This distinction, between informative and derivational intentions, offers another strategy for reconciling apparently divided intuitions within a single account: in cases where these two intentions are divided, i.e. in which the speaker's derivational intention does not, in fact, lead to the content of her informative intention, intuitions may be weighted towards either of the intentions, thus allowing for the possibility of two distinct intuitions concerning what is said by one and the same utterance; it may, in other words, be that two apparently conflicting intuitions are intuitions about different speaker intentions.

---

<sup>8</sup> You may not be convinced by my analysis of this particular example, although later in the thesis I shall argue that it is along something like the right lines. I hope, however, that it nonetheless serves to illustrate the distinction I am intending to draw between informative and derivational intentions.

## Chapter three

### Proper names

#### 3.0

In this chapter, I want to embark on the project of applying the theoretical machinery presented in the last chapter to an analysis of the semantics and pragmatics of singular expressions. The enquiry will broadly focus on three questions:

Question 1: what contribution do singular expression make to the truth conditions of utterances of sentences in which they appear and how do they make this contribution?

Question 2: what contribution do singular expression make to the cognitive significance of utterances of sentences in which they appear and how do they make this contribution?

Question 3: what are the conditions for understanding singular expressions? What, in other words, has to be the case in order for us to want to say that a hearer has understood an utterance of a singular expression?

I want to start with an examination of the semantics and pragmatics of proper names. The reason is partly historical: as discussed in chapter 1, it was with questions on proper names, and in particular with the questions on names that give rise to Frege's puzzle, that the debate on singular expressions began to take its current shape. But there are good reasons to start with names beyond the merely historical: names, as the Millian intuition attests, have the appearance of being paradigmatically referential; there seems intuitively to be no more to the meaning properties of a proper name than that which it names. Since the position I intend to advocate will entail that there are no *linguistically* referential expressions, i.e. no expressions which are constrained to refer by their linguistic meaning, an analysis of proper names should make clear, at least in outline, the shape that that position will take.

In this chapter I want first to run briefly through the core data to which any theory of names aiming to answer the three questions above must address itself. Once I've done that, I shall lay out the account of the semantics and pragmatics of proper names that I favour, and explore how this account can be put to work in tackling the data already discussed. There is, however, a wide range of positions on names currently on the market. Once I have detailed my own account, I want to examine some of the more plausible of these recent positions, in order to establish how they compare with my own. I hope in doing this to show two things: firstly that many of these positions are, to a large extent, aiming at just the same insights about the workings of proper names as my own account and, secondly, that none of them is able to capture the range of facts about proper names accessible to my account.

Before examining the data, however, I want to say a few words on what proper names are; on how, in other words, we want to individuate the object of our enquiry. I should start with a warning: I do not intend to lay down a set of necessary and sufficient syntactic and/or semantic conditions for name-hood. An intuitive grasp of what makes an expression a proper name, or at least the intuitive ability to discriminate between those expressions which are proper names and those which are not, will have to do for my purposes. It is important, however, to raise two points here, which should act as a reminder of some of the issues raised in chapter 2: firstly, I am intending my claims to apply to an independently identifiable class of expressions; I am not, that is, merely intending to make the general claim that there are natural language expressions that work in such-and-such a semantic and/or pragmatic way. And, deriving from this point, it is important that the enquiry not presuppose a circular criterion of identification for names. The kind of (partial) criterion I have in mind would go something like: given that proper names are referential, only expressions which are referential can be proper names. This kind of criterion will do fine for formal languages, in which a very tight rein can be kept on the syntactic and semantic properties of expression-types, but it will not do for a study of natural languages: if proper names are, in fact, referential, that must only be discoverable *a posteriori*. Any account on which this is an *a priori* truth is not an account of the semantics and pragmatics of a class of natural language expressions of the sort I aim to develop, but is, rather, an account of the contents of certain kinds of thoughts. As I hope is already clear, an account of the contents and, in particular, of



the truth conditions of thoughts is, I believe, a key component of any adequate theory of the semantics and pragmatics of singular terms, but it is not the whole story. This warning against circularity will play a central role not just in the examination of proper names, but across all types of singular expression: I am interested in explaining the behaviour of natural language expression types, in all their complexity, not the behaviour of logical constants and quantifiers.

### 3.1 The central data on proper names

In this section my aim is to lay out in very brief form what have been taken to be the central data for an account of proper names, data which we first saw in chapter 1. There is something not entirely happy about this: after all, once I have presented my own analysis, I shall go on to examine other available accounts and thus also the data that each adduces in its own defence. This division of data is for ease of presentation: my aim is to present a body of central data, examine how my analysis can cope with (or reanalyse) those central data, and, having done that, to ask whether there are any other data which could cause my account difficulties. The division between the two groupings is thus not intended to have any theoretical significance.

So what are the central data on proper names? They group broadly into three classes, corresponding to co-reference, emptiness and rigidity:

#### 3.1.1 Co-reference

Recall that, on the account advocated by Mill (1843), the semantic properties of names are exhausted by their bearers. Distinct names which share a bearer, therefore, should also share semantic properties. Consider, however, the following sentences:

- 1) Satchmo is Louis Armstrong.
- 2) Satchmo is Satchmo.
- 3) Ella believes that Satchmo is a great trumpeter.
- 4) Ella believes that Louis Armstrong is a great trumpeter.

On a Millian account, the move from sentences (1) to sentence (2), and also that from sentence (3) to sentence (4), should not affect the meaning properties of the sentence: in each case we are substituting one name for another co-referential name and since, for Mill, two proper names which have the same reference also have the same meaning, we are substituting one name for another with the same meaning. It seems that there are good reasons to think, however, that the meaning properties of sentence (2) are not identical to those of sentence (1) and, equally, that the meaning properties of sentence (4) are not identical to those of sentence (3). The puzzle as it applies to sentences (1) and (2) concerns cognitive significance: sentence (1) seems to lack the cognitive significance of sentence (2), in that the proposition expressed by sentence (2) seems *a posteriori*, whereas that expressed by sentence (1) seems *a priori*. As Salmon (1986/91) has pointed out, this problem is not particular to the identity predicate; just the same kind of difficulty arises with examples such as (adapting Salmon's own examples):

- 5) Satchmo is a trumpeter if Louis Armstrong is.
- 6) Satchmo is a trumpeter if Satchmo is.

Here again we have two sentences which differ only in the substitution of one name for another co-referential name, yet the substitution seems to have effected a change in cognitive significance: just as for (1) and (2), (6) seems *a priori*, whereas (5) seems *a posteriori*.

For sentences (3) and (4) the puzzle no longer concerns (only) cognitive significance but, rather, truth value: since there appear to be possible circumstances in which (3) might be true and (4) false (any circumstance in which Ella is unaware that Satchmo is Louis Armstrong is such a possible circumstance), it seems that we cannot substitute co-referential proper names, at least in some contexts, *salva veritate*<sup>1</sup>. We are thus pushed towards the conclusion that, contra Mill, names which share a reference do not thereby share identical meaning properties.

---

<sup>1</sup> I shall discuss below the question of whether failure of substitution *salva veritate* is restricted to those contexts which Quine labelled 'opaque', or whether, as Saul (1997) suggests, the phenomenon goes beyond such contexts.

### 3.1.2 Emptiness

The Millian account is also challenged by data on empty proper names, i.e. proper names that fail to refer to any individual. In 1998, a New York publishing company published what purported to be the biography of a forgotten member of the abstract expressionist school. Nat Tate had, supposedly, after years of depression, destroyed most of his works before leaping to his death from the Staten Island Ferry. It eventually transpired, however, that the book was a hoax, devised by the writer William Boyd: Nat Tate had never existed, the reproduced art works that had supposedly survived were, in fact, by Boyd himself, and the photographs of Tate were from Boyd's collection of photographs of unidentified people. In light of this story, consider the following sentences:

- 7) Nat Tate lived and worked in New York.
- 8) Nat Tate does not exist.

On a Millian account, both (7) and (8) should lack meaning: both contain an empty proper name, 'Nat Tate'; the meaning of a proper name is exhausted, on the Millian picture, by its bearer; 'Nat Tate' should therefore lack meaning and both (7) and (8) should fail to be meaningful, since they both apparently predicate a property (in (7) the property of living and working in New York and in (8) the property of not existing) of nothing. But this once more runs counter to some reasonably robust intuitions: it would seem that (7) expresses a complete thought, while (8) not only expresses a complete thought, it expresses a true thought (I am following Fregean tradition in using the terms 'proposition' and 'thought' interchangeably here). It might be thought that the puzzle illustrated in (8), the puzzle that concerns the ability of sentences containing empty proper names to express true propositions, is specific to negative existentials, or more generally to negative sentences, but this is not so, as illustrated by (9) and (10):

- 9) Nat Tate did not live in New York, he didn't live anywhere.
- 10) Nat Tate is a figment of William Boyd's imagination.

In (9), a predicate other than the existence predicate is used, while (10) involves no negation, yet both intuitively express true propositions. The meaningfulness of empty names thus seems to constitute another fundamental challenge to the Millian position.

### 3.1.3 Rigidity

While data on co-reference and emptiness seem to weigh against any account of the semantics of proper names along Millian lines, data on rigidity seem to weigh against any account along non-Millian lines. The intuitions which form the evidence for rigidity concern a comparison between the behaviour of different expression types in modal contexts. There are different ways of getting at these intuitions on modal profile: imagine (if you need to) that in the actual world Louis Armstrong is (leaving tense to one side) the greatest jazz trumpeter. The expressions 'Louis Armstrong' and 'the greatest jazz trumpeter' are thus co-extensive in the actual world. Now consider the following sentences:

- 11) Louis Armstrong performed with Ella Fitzgerald.
- 12) The greatest jazz trumpeter performed with Ella Fitzgerald.

The question we want to ask is: on whom do the truth values of these two sentences depend? In the actual world, given the above supposition, they will both depend on Louis Armstrong; they will in other words both be true if Louis Armstrong performed with Ella Fitzgerald, and both be false if he didn't. Consider, however, on whom the truth values of the propositions expressed by these sentences would depend in a world in which Miles Davis, not Louis Armstrong, is the greatest jazz trumpeter. It seems clear that the proposition expressed by (11) will still depend on Louis Armstrong, whereas that expressed by (12) will depend on Miles Davis. So what does this suggest? It has been taken by many to suggest that proper names pick out the same individual across possible worlds (leaving aside issues concerning worlds in which the bearer of a name does not exist, and questions raised by names with more than one bearer), that they are, in Kripke's terminology, *rigid*, whereas definite descriptions are

*non-rigid* in that they pick out in any possible world the individual that satisfies their descriptive content in that world.

Essentially the same sorts of intuition can be approached from another angle, from the angle of linguistically modal contexts. Thus consider:

13) Louis Armstrong might not have played the trumpet.

14) The greatest jazz trumpeter might not have played the trumpet.

It seems that, whereas (13) has only one reading, in which it is true so long as there is a possible world at which Louis Armstrong doesn't play trumpet, (14) has two readings: one on which it's true in the same circumstances as (13), i.e. on which it is true iff there is a possible world at which whoever is the greatest jazz trumpeter at the actual world, i.e. Louis Armstrong, does not play trumpet, and one on which it is true so long as there is a possible world at which whoever is the greatest jazz trumpeter (at that world) doesn't play trumpet.

What are these intuitions supposed to tell us? The lesson that Kripke (1972) and those who have followed him draw is that, whereas definite descriptions contribute complex descriptive conditions to propositional content, names contribute nothing more than their referents (or, perhaps more accurately, constant functions from Kaplanian circumstances of evaluation to their referents); they see this evidence, in other words, as supporting a return to Millianism. I shall not, for the time being, be too concerned with these conclusions, although I shall have more to say about certain recent versions of them once I have laid out my own analysis of proper names; what I am interested in here is the intuitions themselves, constituting, as they do, a testing ground for analyses of proper names.

Before leaving the intuitions on rigidity, I want to flag up one issue that will play a significant role later on. The same issues as illustrated in sentences (11) to (14) above seem to arise with the following sentences:

15) Louis Armstrong is a jazz musician.

16) The person called 'Louis Armstrong' is a jazz musician.

In (16), the subject position is occupied by what has been called a *nominal description*<sup>2</sup>, i.e. a description which makes explicit mention of the name that its denotation bears. All I want to point out here is that, on the face of it, the use of a nominal description seems to have no effect on the truth-conditional intuitions adverted to above. Just as (11) will be truth-conditionally dependent on the same individual in every possible world, whereas (12) will not, so (15) is anchored across worlds to Louis Armstrong, whereas (16) is not: (16) will be true at a possible world so long as the person who at that world is called 'Louis Armstrong', be that Louis Armstrong or someone else, is a jazz musician.

### 3.2 Proper names and individual concepts

Now that what I have called the *central data* on proper names are in place, I want to outline the kind of analysis of names that I favour. The shape of this analysis will be brought out in three ways: firstly, and probably least helpfully, I will directly lay out the bare skeleton of the analysis, detailing how, as I see it, the theoretical notions explicated in chapter 2 can best be applied to the semantics and pragmatics of proper names; secondly I want to examine how the analysis can be used to address the central data above, and finally, as already mentioned, I want to compare my own analysis with certain of the most plausible accounts currently available. I do not expect the full nature of the analysis I am advocating to be clear until all three of these stages are completed.

#### 3.2.1 The bare bones

The analysis that I favour sees proper names as tools for the communication of individual concepts. As I hope is clear from the discussion in chapter 2, I am approaching singular expressions from a relevance-theoretic perspective on which the central role of language is to constrain inference, specifically to constrain the inference that a hearer undertakes in order to access the content of a speaker's informative intention. What I want to claim, therefore, is that proper names guide a

---

<sup>2</sup> See, for instance, Bach (1987)

hearer's inferential processes in particular ways, ways determined by their lexically encoded meaning.

In what, then, might the constraint laid by proper names on interpretation consist? In broad terms, the answer I want to suggest is that, by dint of their encoded linguistic meaning, names constrain interpretation to certain types of individual concept. As throughout, I am assuming that the process of interpreting an utterance involves building a primary propositional representation, the explicature of that utterance (and may also involve building further, secondary propositional representations, i.e. implicatures). As detailed in chapter 2, there is, on the relevance-theoretic view, a structural homomorphism between sentences and the explicatures they give rise to on occasions of use: explicatures are taken to be developments of a sentence's logical form. Given such a sentence-explicature homomorphism, I am assuming that, for any utterance of a name sentence (i.e. a sentence with a name in subject position), the explicature of that utterance will contain a structural slot corresponding to the syntactic slot which is occupied by a name in the name sentence<sup>3</sup>. To claim that names constrain interpretation to certain types of individual concept is thus to say that this name-slot in the utterance's explicature must be occupied by a certain type of individual concept.

It is, however, going to be one of my central claims that all singular-expression-types are tools for expressing individual concepts. In what way, then, do the constraints encoded by proper names differ from those encoded by other singular-expression-types? Before pointing to a difference, I want to suggest a further similarity: I do not believe that the semantically encoded meaning of any singular-expression-type is sensitive to what in chapter 2 I called the *external dimension* of individual concepts; that is, on the analysis I shall propose, the linguistic meaning of a singular expression is blind to such properties of individual concepts as being *de re* or being descriptive. This of course, given the discussion in chapter 2, has a knock-on effect for which internal properties of individual concepts are exploitable by the linguistically encoded meaning of singular expressions; in particular, that linguistic meaning cannot make use of the REF property, since to be sensitive to REF would be

---

<sup>3</sup> This assumption is, of course, at odds with Russell's position on proper names. However, even those who now advocate a descriptive analysis of proper names along Russellian lines would, I am sure, shy away from accepting Russell's full incomplete-symbol analysis, given its implausibility for a compositional semantics.

the internal corollary of being sensitive to a concept's external dimension. This claim, that the linguistic meaning of singular expressions is blind to the external properties of the individual concepts which, on occasions of use, constitute their interpretations, does of course have wide-reaching implications for the take on the notion of reference embodied in this account. The most obvious of these implications is that reference is not a linguistic notion; that is to say referentiality is not a property of linguistic expressions themselves, but rather of the thought components to which they may, in certain contexts, correspond. Any account of this sort is going to come up against some obvious objections; in particular it seems that one of the three central categories of data for any account of proper names, the data from intuitions on rigidity, may be out of reach. I shall discuss this issue in due course, when I come on to explore more broadly how my account might fare with all three types of central data outlined above.

What I have claimed so far is that the linguistically encoded meaning of proper names is such that the interpretation of names is constrained to certain types of individual concept, but is silent on facts concerning the external dimensions of those concepts<sup>4</sup>. I have not yet addressed the question posed above concerning the difference between the encoded meaning of proper names and that of other types of singular expression, i.e. concerning which properties of the internal dimensions of individual concepts the encoded meaning of proper names is sensitive to. Since we are dealing only with the internal dimension, this reduces to the question of which type or types of descriptive information the semantics of proper names exploit. It is in the answer to this question that the special nature of proper names comes out: essentially proper names are sensitive to themselves. On, as far as I can see, more or less all accounts of their semantics, names are viewed as two-dimensional tools: on the one hand they are used to assign properties to individuals (the property of bearing) and on the other they are used to talk about individuals via the exploitation of this assigned property. As I will discuss in greater detail when I come to compare my own account

---

<sup>4</sup> This is not to say that in order to understand a particular use of a proper name it is not necessary for a hearer to establish whether the individual concept to which the name is intended to correspond is descriptive or *de re* (contains REF or doesn't contain REF). The claim I am here making is simply that the interpretation of proper names does not *always* involve entertaining a *de re* concept (or *always* involve entertaining a descriptive concept). Thus the linguistic meaning of proper names is blind to the *de re*-descriptive distinction. Speaker intention, of course, may be anything but blind to the distinction and it will be my claim, which I shall outline in greater detail below, that inferring whether the speaker's intention involves a *de re* concept or a descriptive concept may be a key part of inferring an overall interpretation.



with others currently on the market, one of the key questions that distinguishes different accounts is the precise nature of the role played by the assigned property in the referential (using the term loosely) exploitation of names: for some, the property of being the bearer of 'PN' contributes directly to the truth conditions of utterances of sentences in which 'PN' appears; for others, the property makes no appearance in propositional content, but is manipulated so as to constrain that to which the name may refer, and on yet others it simply determines reference. The account I advocate can be seen as following a version of the middle path: I believe that the property of being the bearer of 'PN' is used to constrain that to which 'PN' can refer, although not, as I hope will become clear, quite the same constraint as has been championed by others who have been tempted to follow the middle path.

Here I must come clean on just what I am claiming about the semantics of proper names. As are all linguistic expressions, on the framework within which I take myself to be working, proper names are viewed as tools used by speakers to give clues about the content of certain of their mental representations (those representations which they wish to communicate). What, then, does a proper name 'PN', uttered by a speaker S as part of the sentence 'PN  $\phi$ s', tell a hearer H about the mental representation which S wishes to communicate? It tells H two things: firstly, it tells him that S's mental representation contains an individual concept; and secondly it tells him that S's individual concept contains the entry *x is called 'PN'*. My claim is that this information essentially constitutes the linguistic meaning of a proper name. If this were so, what knock-on effect would it have for H's interpretation procedure? H would know, by dint of knowing the meaning of the proper name, that, in order to interpret S's use of 'PN', he must access an individual concept, and, beyond that, that it should be an individual concept which he either already associates, or which he comes to associate, with the information *x is called 'PN'*. That, I claim, is as far as the linguistic meaning of 'PN' will take him. The rest will be down to the application of pragmatic principles in context.

Much of the rest of this chapter will be concerned with addressing potential objections to this sort of account. However, I want to address two of those objections before going any further. Firstly, imagine that S wishes to communicate to H the singular proposition that a is  $\phi$ ; while a is in fact the bearer of 'Bill', S mistakenly believes him to be the bearer of 'Ben'; in order to communicate the proposition that a

is  $\phi$ , therefore, S utters the sentence 'Ben is  $\phi$ '; although H knows that  $a$  is the bearer of 'Bill', he either knows or works out that S mistakenly believes  $a$  to be called 'Ben' and is intending to say something about him; H thus interprets S's utterance as expressing the proposition that  $a$  is  $\phi$ . It thus seems that, on the kind of story that I have laid out above, there is nothing to stop S literally expressing a proposition about Bill by using the name 'Ben', a result which, in one form or another, has seemed unacceptable to many<sup>5</sup>.

My response to this objection is to challenge the intuition underlying it; I am, in other words, prepared to accept, given the approach to propositional content developed in chapter 2, that S has indeed expressed (rather than merely communicated) a proposition about Bill by using 'Ben'. This does not, of course, mean that nothing has gone wrong here; clearly something has: S entertains a false belief, and an attribution of this false belief to S is one of the premises used by H in the inferential process via which he retrieves the proposition expressed by S's utterance. Thus the interpretation of S's utterance has gone wrong, in that, in order to access the intended interpretation, H has had to follow an inferential strategy other than that which S intended him to follow. Putting this in the terms developed in the last chapter, there is a mismatch between S's informative and derivational intentions. From that, however, we don't need to conclude that the interpretation reached by H does not equate to the explicature of S's utterance. Approaching the question from within the framework outlined in chapter 2, there is every reason to suppose that it does: there are, as far as I can see, reasonably strong intuitions that, on the assumption that Bill is indeed  $\phi$ , S has, in uttering 'Ben is  $\phi$ ' got something right and something wrong. We have already seen what she's got *wrong*: she has the false belief that Bill is called 'Ben', a false belief that H must access to retrieve the intended content. It seems natural, then, to say that what she has got *right* is that the proposition expressed by her utterance is true.

Given that my response follows on from the discussion in chapter 2, I am not making quite the substantive claim that I might at first sight seem to be making. Identifying, as I did in that chapter, the object of the truth-conditionalist's intuitions with the explicature of an utterance of a sentence in an idealised discourse context, it

---

<sup>5</sup> See, for instance, Kripke (1972, 1977).

is going to be true that, on this view of propositional content, an utterance of 'Ben  $\phi$ s' will never be able to express the proposition that Bill  $\phi$ s. The reason is this: an omniscient speaker utters 'Ben  $\phi$ s'; by doing so, she indicates, given the meaning for proper names proposed above, that she believes the individual she is talking about to be called 'Ben'; since she is omniscient, she must be right about this – her intention, that is, must be to talk about an individual who is called 'Ben'; so an omniscient speaker cannot use 'Ben' to talk of Bill<sup>6</sup>. Recall the role played by a speaker's abilities and preferences within the communicative principle of relevance discussed in chapter 2. Another way of presenting the point I am making is that it is key to the original example that H should take as a premise in his inferential process that S is fallible, that her abilities may skew what she takes to be the optimally relevant interpretation of her utterance. Without this premise, H would never have been able to reach the intended interpretation. But of course that premise is unavailable in the case of an omniscient speaker, who can thus never use a name to talk about an individual other than the name's bearer. I therefore, as again should be clear from the discussion in chapter 2, have no substantive argument against standard philosophical/semantic claims on this point: it is of course open to the truth-conditional semanticist to posit a level of semantic representation defined according to this idealised discourse context and doing so makes his claim true. My argument, as presented in chapter 2, is simply that this level of representation is dispensable. This discussion should make it clear, however, why the truth-conditionalist's intuitions on these cases are as they are.

The second potential objection to my analysis that I want to raise at this point is the following: if the meaning of a proper name is really as I have suggested, then it seems that a speaker may never use a proper name to refer to an individual whom he does not believe to be the bearer of the name. Yet there are clear cases where this is possible. Consider S and H again: H now wishes to communicate to S the proposition that *a* is  $\psi$ ; he could of course say to S something like: *that person who you think is called 'Ben', well he's actually called 'Bill' and Bill's  $\psi$* . But he could instead exploit his knowledge that S believes *a* to be called 'Ben' and convey the same message with an utterance of 'Ben is  $\psi$ '. It seems I am committed to saying that, in this case, H will succeed in expressing the proposition that *a* is  $\psi$ , and yet he has used a proper name to

---

<sup>6</sup> Or at least she cannot do so if she is communicating with an omniscient hearer.

talk about an individual whom he does not believe to be the bearer of the name. This seems incompatible with the claim that, in using a proper name, a speaker indicates that she believes that which she wishes to talk about to be the bearer of the name. There is, in fact, no incompatibility. In this case, H is manipulating the discourse context and the nature of pragmatic inference. It is of course not the case that he believes Bill to be called 'Ben'. It is nevertheless crucial to the success of his communicative aim that S should take him so to believe. If S does not believe that H believes that a is called 'Ben', then she will never reach H's intended interpretation (unless of course she sees the light, realises that a is, in fact, called 'Bill', that H believes her to have a false belief and that he is trying to exploit that belief in order to communicate the proposition that a is  $\psi$ ). The imagined objection, therefore, simply fails to appreciate the nature of the interaction between linguistic meaning and pragmatic principles, and the uses to which speakers can put this interaction.

### 3.2.2 Names, individual concepts and questions 1 to 3

I have, so far, laid out the bare bones of an analysis of the lexically encoded meaning of proper names, although, as I indicated earlier, I do not expect the full nature of the analysis I support to be clear until I have addressed how this account fares with the central data, and how it compares with other accounts currently available. Before getting on to the central data however, I want to do a bit more spelling out; in particular, I want to examine in broad outline what answers my account suggests to the three questions presented at the start of this chapter, questions concerning the truth-conditional contribution and cognitive significance of names and the conditions under which we would want to say that a speaker has understood an utterance of a name.

What I have so far claimed is that proper names are tools for communicating individual concepts; specifically, a speaker uses a proper name 'PN' where she wishes to communicate a proposition her mental representation of which contains an individual concept which she associates with the information 'x is called 'PN''. What constraints, then, does the use of a proper name, on this account, lay on the truth-conditional content of an utterance containing it? It is a consequence of my claim that the constraints laid on truth conditions by the lexically encoded meaning of proper

names are minimal; in particular, since proper names are silent on the external dimensions of the concepts which constitute their interpretations on occasions of use, they can be used to express either *de re* or descriptive individual concepts according to speaker intention. This is so not just at the level of utterance content that I take to be theoretically significant, but also at the level taken to be significant within the truth-conditional literature, i.e. the level which I have identified with the proposition expressed in an idealised discourse context. On my analysis, even an omniscient speaker communicating with an omniscient hearer could use a name to express either a *de re* or a descriptive concept.

One of the very few points of agreement between opposing camps in the debate on names has been that, whatever story you tell, it has to be a story on which either names *always* express *de re* concepts (i.e. always give rise to singular truth conditions), or on which they *always* express descriptive concepts (give rise to general propositions). It is of course the case that, given the assumptions discussed in chapter 2, any account on which name sentences can be used to express either singular or general truth conditions, is, *ipso facto*, an ambiguity account of proper names, and ambiguity accounts have, quite rightly, seemed to many to be an unappealing last resort. I shall leave a detailed discussion of the question of ambiguity until I have presented my analyses of a range of different singular-expression-types, since the same considerations will apply across the board. My claim, however, will be that, although names can, on my analysis, be used to express concepts with either truth-conditional profile, i.e. with either type of external dimension, this does not make my analysis an ambiguity account, or, at least, it does not make it an ambiguity account in any way that should worry us from a methodological point of view.

What reason might there be to hold that name sentences can express either singular or general truth conditions? I have already outlined, in section 3.1.3 above, the intuitions which seem to support the view that names are rigid designators, i.e. that they give rise to singular truth conditions. This evidence is of course not uncontroversial: it is, for instance, central to Russell's position, discussed in chapter 2(?), that proper names, as we use them in natural language, never give rise to singular truth conditions<sup>7</sup>. The Russellian picture on proper names is not without its

---

<sup>7</sup> I am assuming, I hope uncontroversially, that, from the perspective of modern linguistic theory, we would not be tempted to go along with Russell in claiming that demonstratives are true names.

latter-day supporters, and I shall come on to discuss some of those Russell-flavoured accounts below. For now, however, I want to point to some of the evidence used in support of the view that name sentences (can be used to) express general propositions. There are (at least) two types of example which have been taken to demonstrate non-referential uses of names. Firstly there is the type of descriptive name discussed by Evans (1982) and also by Grice (1969). Evans asks us to consider a naming convention on which whoever invented the zip is to be called 'Julius'. Basing our use on this convention, we can, it seems, use 'Julius' as any other proper name. It is plausible, however, that our use of 'Julius' is not referential (in the terms I favour, the interpretation of our use of 'Julius' does not involve a *de re* individual concept), since we do not know, in the appropriate sense, which individual 'Julius' refers to (or, rather, since the concept we associate with 'Julius' is not linked to an individual in the way required for *de re*-ness).

In an attempt to rein such cases back into the direct reference fold, Recanati (1993) argues that to use a proper name in such a way that it is backed only by a descriptive psychological mode of presentation (individual concept) is to use that name non-literally. For Recanati, Evans-type descriptive names are, essentially, referential expressions in search of a referent:

If we use a descriptive *name* rather than a description, this is precisely because we look forward to a richer state of knowledge in which we will be able to think of the referent non-descriptively. A descriptive name ... is created only in the expectation that more information about the bearer will accumulate, thus eventuating in the possibility of thinking of the latter non-descriptively. This possibility is simply *anticipated* by the use of a descriptive name.

Recanati 1993, p. 180.

I think there is something right about this but also something wrong. As I shall discuss when I come on to examine how my account fares with intuitions on rigidity, it does seem to be the case that names are mostly used as tools of reference, i.e. used to communicate *de re* individual concepts. I shall argue however that this is not a fact concerning the semantics (i.e. encoded meaning) of proper names, but is a pragmatic

fact, deriving from considerations of cognitive utility. Thus, although Recanati is right to claim that there is an intimate connection between proper names and *de re* individual concepts, it is my belief that he's looking for an account of this relation in the wrong place.

Beyond that, however, his account seems to make certain counterintuitive predictions. Imagine that I decide that I am going to call whoever is the male star of the film *Notorious* 'Eric'; now I happen to know that the male star of *Notorious* was Cary Grant; on Recanati's account, I am in a cognitive position to assign a referent to 'Eric', thus my use of 'Eric' should be referentially anchored to Cary Grant. Compare, however, sentences (17) and (18):

17) Cary Grant might not have been Eric.

18) Hesperus might not have been Phosphorus.

Kripke (1972) has convincingly argued that the proposition expressed by (standard uses of) (18) is necessarily false; Hesperus, that is, just could not have failed to be Phosphorus. I have the intuition, by contrast, that there is a reading of (17) on which it is true given that James Stewart (or anyone for that matter) might have been the star of *Notorious*. This intuition seems to weigh against the position that 'Eric' is referential and refers to Cary Grant. I presume that Recanati's response would be to claim that the proposition which I have in mind involves a non-literal use of 'Eric', whereas the proposition *literally* expressed by (17) is just as necessarily false as that expressed by (18). Since, however, I have stipulated that 'Eric' should express a descriptive concept of whoever is the male star of *Notorious* and since that is just the concept that it seems to express in what I take to be the proposition expressed, we should require a powerful argument to convince us that this proposition is non-literal (inasmuch as the idea of literality or non-literality has any role to play within the kind of framework I have adopted). Recanati, it seems to me, has no such argument.

Evans-type descriptive names thus seem to offer evidence that proper names can be used non-referentially. It might be objected, however, that we are, in examining these names, a long way from proper names as they are actually used in natural language; we do not, that is, go around deciding to call the inventor of the zip 'Julius' or the male star of *Notorious* 'Eric'. There are two responses to this: firstly,

even if it were the case that natural language names were never, in fact, assigned via descriptive concepts, what Evans shows is that they could be so-assigned without thereby ceasing to be names, a fact which should be accommodated by an analysis of names. Secondly, however, there are cases in which names are so assigned, the best-known example, discussed by, among others, Kripke (1972), being 'Jack the Ripper'. Our use of the name 'Jack the Ripper' is backed not by acquaintance, in Russellian terms, but by a description something like 'the person who committed such-and-such murders'. And again we can use this name to show that all is not well with any attempt to salvage referentiality for Evans-type names, along the lines followed by Recanati. Imagine we discover that Jack the Ripper was in fact Liverpool businessman James Maybrick. Now consider a possible world in which precisely those murders which Maybrick committed in the actual world were committed by Edward Prince of Wales. It seems clear that, as used in that world, the name 'Jack the Ripper' would designate Edward rather than Maybrick. But beyond that, it seems that this is a world in which *James Maybrick is not Jack the Ripper*. Again, Recanati would have to claim that, here, I am (in the imagined context) using the name 'Jack the Ripper' non-literally, since, once it is established that Jack the Ripper is James Maybrick, 'Jack the Ripper' becomes referentially anchored to Maybrick. I can see no evidence for non-literality, however, in the claim that Maybrick is not Jack in the imagined world.

The behaviour of Evans-type names thus offers strong evidence that (a) names can be (and sometimes are) assigned via descriptive concepts and (b) these descriptive concepts can constitute the truth-conditional contribution of such names. Kent Bach, whose overall account of the semantics of proper names I shall discuss in greater detail later in the chapter, offers another sort of example: firstly, Bach (1987) asks us to consider a discussion on the US presidential electoral process; one contributor suggests that, given recent results, alphabetical order might prove a better criterion for selection; under those conditions, she claims:

19) Aaron Aardvark might have been president.

Bach 1987, p. 143.



The speaker of (19) clearly does not intend to use the name 'Aaron Aardvark' referentially – not only does she have no individual in mind, her utterance does not require that she should believe there to be an individual who the name picks out. It seems, therefore, that 'Aaron Aardvark' is used descriptively and that the truth conditions of an utterance of (19) are general. More specifically it seems plausible that in this case 'Aaron Aardvark' is interpreted as equivalent to the description 'the person called 'Aaron Aardvark'', giving (19) an interpretation on which it is equivalent to (20):

- 20) It might have been the case that the person called Aaron Aardvark was president.

Bach (2002) offers a further example on which a proper name seems to give rise this sort of nominal interpretation:

- 21) If his parents had named him 'Aristocrates', Aristotle would have been Aristocrates instead of Aristotle.

Again, it seems that the property of being called 'Aristotle' is part of the contribution made by (at least one of the occurrences of) 'Aristotle' to truth-conditional content.

I am not sure how convinced a dyed-in-the-wool Millian would (or should) be by Bach's examples, or rather by the conclusions he draws from them. While Bach, advocating as he does an account on which name sentences express general propositions, must allow that they can nevertheless be used to *communicate* singular propositions, it is open to the Millian to claim the reverse: that name sentences express singular propositions but may be used to communicate general propositions. For sentence (19) in particular this kind of story seems to offer a plausible analysis: given what seems to me to be an intuitive oddness about (19), it may well be that an utterance of (19) fails, in the truth-conditionalist's terms, to express a literal proposition, but may nevertheless be used successfully to communicate propositions such as that the president might have borne the name 'Aaron Aardvark'. I shall return to a discussion of Bach's position below, but for the time being I shall put his examples to one side, and rely instead on the evidence from Evans-type names to

support my contention that proper names may be used non-referentially. This claim only makes sense, of course, within a framework on which proper names are not taken to be defined according to the type of their truth-conditional contribution. As discussed in section 3.0 above, however, my claims, as certainly, for instance, those of Bach, should be taken to concern the class comprising those expressions which are intuitively taken to be proper names, not a semantically-individuated sub-class.

In answer to question 1, it is central to the account I propose, then, that proper names can make either singular or general contributions to truth-conditional content, according to speaker intention. The answer my account suggests to question 2 will become clearer when we come on to discuss the data concerning co-reference and emptiness. In outline, however, it should be clear that, on this analysis, the cognitive significance of a proper name is determined by the internal dimension of the individual concept which, on an occasion of use, constitutes its interpretation. It may seem, however, that there is an obvious objection to any account of this sort: the cognitive significance of a singular expression, its sense, in Fregean terms, is not a subjective property that differs from cognitive agent to cognitive agent, but must be constant as between different cognitive agents; how else, the Fregean might ask, could we account for successful communication? On my analysis, however, there is no reason to suppose that mere co-reference would guarantee any overlap between the internal dimensions of the individual concepts of two cognitive agents, leaving successful communication as something of a mystery. This objection, however, has no force, since there is no requirement that the cognitive significance of proper names should be intersubjective in the sense that an utterance of a name sentence should give rise to the same inferential effects for different cognitive agents. Consider sentence (1) again, repeated here as (22):

22) Satchmo is Louis Armstrong.

Sentences such as this have received the attention they have because, although they seem to predicate the identity of an object to itself, they are nevertheless cognitively significant. But the cognitive significance of this sentence is clearly relative: to some it may express a new and surprising piece of information, to others what it expresses may be familiar. The importance of such sentences is not, therefore, that what they

express *must* be cognitively significant, but that it *can* be so. And this fact can be explained, on the picture I am presenting, by the observation that the names such sentences contain *can* be associated by a cognitive agent with different individual concepts (as can any two proper names)<sup>8</sup>. For the time being that is all I shall have to say but, as mentioned above, the details of how my account deals with cognitive significance will be spelled out when we come to discuss co-reference and emptiness.

Finally, what answer does the analysis I have proposed suggest for question 3, the question of what it is to understand a proper name? On the sort of picture I am painting, a picture on which there are two distinct theoretically-significant levels of representation for sentence meanings, this question conceals two sub-questions. Firstly, there is the question of what it is to grasp the linguistic meaning of a proper name. Given the discussion in the previous section, the answer that my analysis offers to this question is quite straightforward: to grasp the linguistic meaning of a proper name 'PN' is to grasp that any speaker uttering 'PN' as part of a name sentence is thereby indicating that a constituent part of the mental representation which forms the content of her informative intention is an individual concept associated with the information 'x is called 'PN''. This is not sufficient, however, to understand the use of a name at the level of propositional content, at the level of explicature. To understand an utterance, it is not enough merely to entertain what one might call *meta-conceptual* representations concerning the speaker's thought; it is also necessary to entertain a propositional representation which stands in certain relations of similarity to the speaker's own representation. To see what this means, consider a complete stranger approaching you and uttering (23):

23) Lars is a liar.

It seems that you are in a position to understand certain things about the mental representation which the speaker is trying to communicate: you can, for instance, infer

---

<sup>8</sup> The names they contain being associated with different individual concepts is not in fact quite enough for these sentences to be cognitively significant; it is, after all, entirely plausible that I have different individual concepts associated with 'Superman' and 'Clark Kent', while the sentence 'Superman is Clark Kent' would still be uninformative for me. It is therefore necessary for cognitive significance not just that the two names should be associated with different individual concepts, but that those concepts should not contain information about mutual identity, i.e. that, in this case, the concept I associate with 'Superman' should not contain the information 'x is the bearer of 'Clark Kent'' and vice versa.

that her mental representation contains an individual concept (that she is talking about an individual) and that her individual concept contains the information ‘x is called ‘Lars’’ (that she intends to talk about someone she thinks of as ‘Lars’). There is what seems to me to be a robust intuition, however, that you are not in a position to grasp the full propositional content of her utterance, just as you would not have been if, instead of (23), she had uttered ‘he is an impostor’<sup>9</sup>. This intuition is not universally held: as I shall discuss later in this chapter, there are accounts of the semantics of proper names which are committed to the claim that you are able to grasp what the speaker has said, in something like the Gricean sense, by an utterance of (23), without any further information<sup>10</sup>. As will become clear, I have no knock-down argument against such positions; they do, however, give rise to theoretically unwelcome results, not the least of which is their violation of the intuition under discussion. Assuming that this intuition is as robust as it seems to me to be, what extra information do we need in order to grasp the propositional content of an utterance of (23)? We need to know *who Lars is*. In the sort of framework I have laid out, this is equivalent to saying that we need to be able to entertain an individual concept of Lars of the appropriate sort. The question thus comes down to how we want to explicate the notion of appropriateness used here. What makes an individual concept entertained by a hearer an appropriate interpretation of a speaker’s use of a proper name?

It would seem to be at least a necessary condition that the hearer’s concept and the speaker’s concept should not differ in their external dimensions, that they should, in other words, be truth-conditionally equivalent. The question then becomes: is this condition sufficient? It seems clear that it is not: if all that were required for understanding were truth-conditional equivalence, then we would be unable to distinguish between the conditions for understanding of utterances of (1) and (2). Yet it seems clear that a hearer who merely grasps a proposition that is true iff Satchmo is Satchmo has not thereby understood an utterance of (1). What more is required for understanding? What I want to suggest is that we should look to the process of interpretation for an answer to this question. In interpreting (1) a hearer will arrive at the individual concept which constitutes his interpretation of ‘Satchmo’ via the property of being the bearer of ‘Satchmo’ just as he will, *mutatis mutandis*, for ‘Louis

---

<sup>9</sup> For more on the parallels between proper names and indexicals/demonstratives, see section 3.2.4.2 below.

<sup>10</sup> Bach’s account, mentioned above, is one such.

Armstrong'. Given this, the hearer will have understood an utterance of (1) iff he entertains two distinct concepts which are truth-conditionally anchored to Louis Armstrong, one of which presents its referent in a Louis-Armstrong kind of way and the other of which presents its referent in a Satchmo kind of way. But what does a Satchmo way of presenting a referent look like? All we can say about this is that an individual concept of Louis Armstrong presents him in a Satchmo way if it is associated with the information *x is called 'Satchmo'*. Although this will standardly require that the concept entertained by the hearer as a constituent of his interpretation will itself contain the information *x is called 'Satchmo'* this need not be the case; there will, after all, be some contexts in which it will be optimally relevant to refer to an individual by using a name of which that individual is not the bearer, a fact of which both speaker and hearer may be aware. Nevertheless, for a name successfully to refer to an individual, it must be the case that that individual is appropriately associated with the property of being the bearer of that name. And it is this fact about the referent that a hearer must grasp for understanding.

### 3.2.3 Individual concepts and the core data

As discussed in section 3.1, data concerning co-reference and emptiness have been taken to provide an acid test for accounts of the semantics of proper names. In this section I want to examine these data in rather greater detail. Once we have a better understanding of their significance, I hope to show that they are naturally captured within the analysis I have proposed.

#### 3.2.3.1 Co-reference

Co-referring names, as I outlined in 3.1, give rise to two distinct problems for analyses of the semantics of proper names: on the one hand they seem to be capable of differing in cognitive significance despite sharing a referent, and, on the other, they seem, in certain contexts, not to be substitutable *salva veritate*; in Fregean terms, they seem to raise problems both at the level of sense and at the level of reference. Taking these problems in order, think back to sentences (1)–(2), repeated here as (24)–(25):

- 24) Satchmo is Louis Armstrong.
- 25) Satchmo is Satchmo.

The difficulty lies in explaining how these two sentences, differing only in the substitution of one proper name for another co-referential name, can differ in cognitive significance. As is well-known, it is in an attempt to address this question that Frege (1952) introduces the distinction between sense and reference. But a closer look at Frege's proposed solution might offer us a clearer insight into what is really going on here. Frege's proposal, as it applies to the examples at hand, would claim that 'Satchmo' and 'Louis Armstrong' differ in cognitive significance because they differ in sense. There is a tension, however, between this claim and another central element of Frege's analysis, the intersubjectivity of sense: it seems clear that I do not need to associate the name 'Satchmo' with precisely the same descriptions with which you associate it in order to understand your utterance of 'Satchmo' and in order for (24) to be cognitively significant to me. It thus seems that, on a Fregean picture, sense is needed to account for the shift in cognitive significance from (24) to (25), yet the senses on which Frege must rely lack one of the key elements required by his non-cognitivist perspective: intersubjectivity. Frege himself recognised this tension. In an oft-quoted passage, he says:

In the case of an actual proper name such as 'Aristotle' opinions as to the sense may differ. It might, for instance, be taken to be the following: the pupil of Plato and teacher of Alexander the Great. Anybody who does this will attach another sense to the sentence 'Aristotle was born in Stagira' than will a man who takes as the sense of the name: the teacher of Alexander the Great who was born in Stagira. So long as the reference remains the same, such variations of sense may be tolerated, although they are to be avoided in the theoretical structure of a demonstrative science and ought not to occur in a perfect language.

Frege (1893, p.159, fn. 4)

It is a commonplace to point out that providing an articulated semantics for natural language expressions was no part of Frege's aim, but it is nevertheless worth noticing that this tension strikes a blow at the applicability of the sense-reference distinction, as formulated by Frege, to an analysis of natural language proper names. For Frege, what it is to understand an expression is to grasp that expression's sense, yet Frege himself is prepared to accept the possibility of a hearer understanding a use of a proper name without associating the same sense with that name as does the speaker<sup>11</sup>. The tension that Frege's account thus faces is illuminating: on the one hand it seems that the difference in cognitive significance between distinct but co-referential names depends on there being more to the *meaning* of a proper name than its referent, while on the other hand it seems that understanding of a proper name is possible where the only common ground between speaker and hearer is at the level of reference, i.e. where all that connects them is that they are *thinking of the same thing*.

Beyond this tension within Frege's analysis, there are also, of course, the problems for descriptive accounts raised by Kripke (1972)<sup>12</sup>. In particular, the sentence:

26) Aristotle was born in Stagira.

will end up, on Frege's analysis, being analytic for anyone who takes the sense of 'Aristotle' to be *the teacher of Alexander the Great who was born in Stagira*. This result, as Kripke has argued, seems plainly wrong: (26) expresses what is inescapably a contingent truth, if truth it be. Yet there is some intuitive basis for the Fregean position on this question: while the proposition expressed by (26) seems clearly contingent, it is nevertheless the case that, for a cognitive agent who associates the description *the teacher of Alexander the Great who was born in Stagira* with the name 'Aristotle', sentence (26) will not be cognitively significant in the way that it is for an agent who does not associate this description with the name. There is thus a relation between the descriptions associated by an agent with a proper name and the cognitive

<sup>11</sup> It is open to the Fregean to claim that in such cases the hearer does not truly understand the speaker's use of the proper name. Such a move, however, takes us a long way from any intuitive notion of understanding.

<sup>12</sup> See chapter 1.

significance that an utterance of the name will have for that agent. It is in the relation he posits between such descriptions and *meaning* that Frege goes wrong.

These apparent problems for a Fregean analysis of proper names fall away when we approach the question of co-reference from within the framework I have outlined. Remember that, on my analysis, to understand a name 'PN' it is necessary to access an individual concept which shares the object of its external dimension with that of the speaker. The only constraint laid on the internal dimension is that it should contain the information *the speaker believes that x is called 'PN'*. How, on this picture, would we want to account for the shift in cognitive significance between, for instance, (24) and (25)? The answer is that such a shift is just what we should expect: the cognitive significance of a proper name is a function of the internal dimension of the individual concept which constitutes its interpretation; it is thus independent of facts about reference and, *a fortiori*, of facts about co-reference. That is just to say that we should expect different proper names to be associated with individual concepts with different internal dimensions and thus with different cognitive significance. When it comes to questions about cognitive significance, therefore, co-reference is simply a red herring.

On the face of it, there are certain similarities between an account along these lines and Frege's account: on both, the reference of a proper name is distinct from the bearer of that name's cognitive significance. So how can this analysis tackle the problems for Frege outlined above? For Frege, a proper name is such that in order to grasp its meaning, it is necessary to grasp a sense, i.e. a particular descriptive condition<sup>13</sup>. This descriptive condition then determines which object it is, if any, that the name refers to. On my account, by contrast, all that is required to understand a proper name is to entertain *some* individual concept which is co-extensive with the speaker's concept and which is appropriately associated with the name itself. The key point, as mentioned above, is that there are no constraints on the relations of similarity that must hold between the internal dimensions of the two concepts. It is thus the case that, on this picture, there is more to the meaning/interpretation of a proper name than its referent: there is the internal dimension of the individual concept which constitutes that interpretation. However, there is no obstacle to speaker and hearer associating

---

<sup>13</sup> I am here discussing what I take to be the views of Frege, not the views of those who, following in his footsteps, have sought to do away with the descriptiveness of Fregean senses. See, for instance, Evans (1982) and McDowell (1977).



different descriptions with the name's referent, since there are no constraints on how similar the internal dimensions must be. To put it another way, the meaning of a proper name is such that it marks the *fact* of associated descriptive information, without determining the *nature* of that information.

This analysis also offers a solution to the second problem, the problem of accounting for cognitive significance without analyticity. Given the anatomy for individual concepts which I have proposed, the cognitive significance of (a referential use of) a name is independent of its truth-conditional contribution; the counterpart of sense in my picture, that is, does not determine reference, for *de re* individual concepts. On such a picture, then, cognitive significance depends on the individual concept which constitutes the interpretation of a proper name, while the descriptive content of that individual concept, its internal dimension, does not determine reference and is not, therefore, analytic. It is, however, clear on this analysis why there should be a difference between the cognitive significance that (26) has for an agent who associates the description *the teacher of Alexander the Great who was born in Stagira* with the name 'Aristotle' and the significance for an agent who does not associate description and name. It is not that for one agent the proposition expressed by (26) is analytic while for the other it is synthetic, it is rather that for one it is informative while for the other it is uninformative; for one, that is, it offers a novel piece of information while for the other it does not.

It is important to note, however, that these points apply only to referential uses of proper names, i.e. to those uses the intended interpretation of which is a *de re* individual concept. It seems clear, for instance, that Kripke's objection concerning analyticity does not apply to descriptive uses of proper names. Consider, for instance:

27) The Whitechapel murders were not committed by Jack the Ripper.

For most (although not all) utterances of (27), the proposition expressed by this sentence *seems contradictory*; it seems, that is, false at all possible worlds<sup>14</sup>. This suggests that, unlike (26), sentence (28):

---

<sup>14</sup> The *on-most-uses* caveat is necessary due to the variability of the external dimension of descriptive concepts discussed in chapter 2. Continuing with the example of 'Jack the Ripper', there are various pieces of information standardly associated with the name, two of which might be *x committed the Whitechapel murders* and *x wrote a series of letters to the police signing himself 'Jack'*. Now since the

28) Jack the Ripper committed the Whitechapel murders.

is (again, on most uses) analytic, or at least expresses a necessary truth. The problem of differing cognitive significance also disappears when we come to look at descriptive names, since there is no more temptation to equate the cognitive significance of co-extensive descriptive names than there is to equate the cognitive significance of co-denoting definite descriptions.

My claim is, therefore, that the account I am proposing offers a natural analysis of those data concerning the cognitive significance of co-referring proper names which have been taken as a key testing ground for theories of the semantics of proper names. What, then, of the data concerning the failure of substitutivity of such names, *salva veritate*, in certain contexts? I want first to raise a caveat: how opaque contexts should best be handled within a general theory of meaning is one of the most fundamental areas of debate in current work in formal and philosophical approaches to meaning, and it goes well beyond the scope of this thesis to provide a fully articulated account of the semantics of such contexts. However I want at least to suggest a direction in which it might be worth looking for such an account, given the kind of framework I have advocated.

There is much evidence to suggest that any analysis on which belief is taken as a direct relation between individuals and propositions faces formidable problems when it comes to the semantic analysis of belief contexts<sup>15</sup>. It seems quite natural, however, on the sort of account I favour, to view belief not as a direct relation between individuals and propositions but as a direct relation between individuals and mental representations, and thus as an indirect relation between individuals and propositions. Again, it is beyond the scope of this thesis to examine in depth precisely what sort of relation this belief relation might be, but it will most likely look something like:

---

concept associated (by most of us) with 'Jack the Ripper' is descriptive, either of these descriptions may be central to the contribution to truth conditions of an utterance of 'Jack the Ripper'. Thus, if a speaker intends to link her use of the name to the first of these pieces of information, then her utterance of (28) will indeed express a necessary truth, although her utterance of 'Jack the Ripper wrote letters to the police' will not, whereas the situation will be reversed if she intends her use to be linked to the second piece of information.

<sup>15</sup> For one particularly well-known presentation of such problems, see Schiffer (1987).

**A believes that p iff A entertains a mental representation M which is a representation of p and A adopts an endorsing attitude to M.**

Clearly this formulation leaves a great many questions open, but it will, I believe, do for our present purposes. What implications would such a construal of the belief relation have for problems concerning the interpretation of proper names in belief contexts? Consider once more (3) and (4), repeated here as (29) and (30):

**29) Ella believes that Satchmo is a great trumpeter.**

**30) Ella believes that Louis Armstrong is a great trumpeter.**

The first things to point out is that each of these sentences seems to have two distinct readings, one reading on which the embedded name is referentially transparent and one on which it is referentially opaque. Since 'Satchmo' and 'Louis Armstrong' are co-referential, the readings on which they are interpreted transparently will be truth-conditionally equivalent, each predicating of Louis Armstrong that he is such that Ella believes him to be a great trumpeter. On the view of belief sketched above, these transparent readings give us only partial information about Ella's beliefs: they give us the external dimension but not the internal dimension; they tell us, in other words, what are the objects of Ella's beliefs, but nothing about how she thinks of those objects.

On the opaque readings, we are not left so much in the dark: with the embedded names interpreted opaquely we know that Ella thinks of the object of her belief in (29) as 'Satchmo' and of the object of her belief in (30) as 'Louis Armstrong'. Put in the cognitive terms I favour, we know, for (29), that the individual concept which forms a constituent of Ella's belief is associated with the information *x is called 'Satchmo'* and, for (30), that Ella's individual concept is associated with the information *x is called 'Louis Armstrong'*. On this analysis of opacity the failure of substitutivity typical of opaque contexts falls out naturally: in essence, what (29) tells us, on an opaque reading, is that Ella stands in the belief relation to a mental representation containing an individual concept associated with the information *x is called 'Satchmo'*, while what (30) tells us is that she stands in this relation to a

representation containing an individual concept associated with the information *x* is called 'Louis Armstrong'. Since 'Satchmo' and 'Louis Armstrong' are different names, we have no reason to suppose that Ella associates these names with the same individual concept (as we know, co-reference is a red herring as far as cognitive significance is concerned) and thus no reason to suppose that, in the move from (29) to (30) we are exchanging like for like. On this analysis, the names 'Satchmo' and 'Louis Armstrong' have just the same encoded meaning as in any other context; what differs is to whom the information associated with the name is attributed, this attribution being a matter of speaker intention.

This is only the briefest outline of a possible approach to the behaviour of proper names in referentially opaque contexts within the framework I am trying to develop. There is, however, a potential challenge to the analysis, even in this vaguest of forms. Saul (1997) points to what she perceives as failures of substitutivity of proper names parallel to those occurring in belief contexts but which occur in unembedded positions. She asks us to consider the following sentences:

- 31) Clark Kent went into the phone booth and Superman came out.
- 32) Clark Kent went into the phone booth and Clark Kent came out.

It seems, at least to Saul, that if (31) is an accurate description of an event (within the fictional world of Metropolis) then (32) is not. That is to say, her claim is that if (31) is true (in the fiction), then (32) is false. This presents an apparent problem for those who hold to a direct reference flavoured account of proper names, just as does failure of substitutivity in opaque contexts: it seems that (31) and (32) differ in truth value while being identical but for the substitution of one name for another co-referential name. But Saul's data also seems to present a problem for those who adopt a more Fregean approach: the Fregean accounts for referential opacity in propositional attitude contexts by appealing to indirect references, i.e. by positing that names (and other singular expressions) refer not to their customary referents when they appear in such contexts, but rather to their customary senses. As Saul points out, this story is simply not open to the Fregean in accounting for failure of substitutivity in unembedded contexts, since it is unclear what it could mean to claim that names do not refer to their customary referents in simple contexts.

It might also seem that Saul's examples pose a threat to the kind of analysis which I have advocated, problems along much the same lines as those faced by a Fregean analysis: although my account can provide what I take to be a plausible analysis of the difference in cognitive significance between (31) and (32), an analysis of the potential difference in truth value seems beyond it. On the account I favour, the truth-conditional contribution of a referential use of a proper name is just that name's referent. Since, in the fiction, 'Clark Kent' and 'Superman' refer to the same individual, the two names should make the same contribution to truth conditions, and (31) and (32) should therefore be truth-conditionally equivalent. The analysis of belief contexts I have sketched above does not extend naturally to such cases, since it depends on two features, neither of which seem necessarily present for Saul's examples: the involvement of a cognitive agent other than the speaker and hearer, and a failure to realise, on the part of that cognitive agent, the identity central to the examples. Taking (29) and (30) again, the potential difference in truth value depends on the involvement of Ella as a cognitive agent, and Ella's failure to realise that Satchmo is Louis Armstrong: it is only in those worlds at which Ella is not aware of this identity that (29) and (30) may differ in truth value. Saul's examples, by contrast, seem to rely on no such lack of enlightenment: if we know anything about the Superman stories, we know that Clark Kent is Superman, yet it still seems that, for us, (31) and (32) may differ in truth value.

How might we tackle these problems within the framework I favour? One obvious strategy might be to claim that there is something special about Saul's examples; they do, after all, involve fictional names, which have often been seen as requiring a separate account from non-fictional names<sup>16</sup>. We can, however, dismiss this idea for the simple reason that the same sorts of failure of substitutivity seem to occur in non-fictional cases. Consider sentences (33) and (34):

33) *Norma Jean Baker wasn't a film star but Marilyn Monroe was.*

34) *Norma Jean Baker wasn't a film star but Norma Jean Baker was.*

---

<sup>16</sup> Although I shall be arguing that fictional and non-fictional names do not in fact need to be treated separately (see section 3.2.3.2 below).

Here we see much the same pattern as in Saul's examples: it seems that (33) may be true whereas (34) is a contradiction, although the sentences differ only in the substitution of one name for another co-referring name.

There are two key questions that arise: firstly, if we cannot look to the fictional nature of these examples for a key to the failure of substitutivity involved, where should we look? And secondly, should we expect whatever story we end up telling about these examples to extend to failure of substitutivity in propositional attitude contexts? If we answer 'yes' to this second question, then it would seem that the sort of account of such contexts which I have sketched above would receive a serious blow, given that, as we have seen, it does not naturally extend to simple contexts. Taking this second question first, there are good reasons to suppose that, whatever the source of the failure of substitutivity in Saul's examples, it is different from the source of failure of substitutivity in propositional attitude contexts. Consider the following example: Bill has read two novels, one a novel in Spanish called 'Rayuela' by a novelist called 'Julio' and another a novel in English called 'Hopscotch' by a novelist called 'Julius'. Being rather unobservant, Bill has failed to notice that Hopscotch is a translation of Rayuela. Now this scenario gives us a more or less standard context for building failure-of-substitutivity-in-propositional-attitude-context cases. Thus, sentences (35) and (36) may have different truth values:

35) Bill believes that Julio is a great writer.

36) Bill believes that Julius is a great writer.

although 'Julio' and 'Julius' are co-extensive. This is not however a context in which Saul-type examples can get off the ground. If, for instance, sentence (37) is true, then so is sentence (38):

37) Julio wrote Rayuela and Julius wrote Hopscotch.

38) Julio wrote Rayuela and Julio wrote Hopscotch.

Examples such as this offer at least *prima facie* evidence that Saul's examples are not of the same kind as examples of failure of substitutivity in propositional attitude contexts.

Predelli (2001a) offers convincing evidence that, whatever the source of the failure of substitutivity in Saul's examples, it is not down to the semantics of proper names (or other singular expressions for that matter). He points out that much the same sort of thing happens where there is no singular expression to give rise to it. He asks us to consider a situation in which the notoriously unsuccessful-with-women Clark Kent is sitting in a conference room with equally timid Art and Bart. In this situation, it seems that an utterance of (39) will be true:

39) Nobody in the conference room is successful with women.

Later that day, however, the notoriously successful-with-women Superman is being interviewed in the conference room by Art and Bart. In this context, it seems that an utterance of (39) will be false. Thus, although, given the identity of Clark and Superman, the set of individuals in the conference room seems unchanged from one context to the next, the truth value of an utterance of (39) has changed.

What might all this suggest? Many of those who have addressed the questions raised by Saul's paper have introduced what we might call *sub-individual* metaphysical entities to do so: Forbes (1997, 1999) makes use of *ways of dressing* and *personae*; Moore (1999) appeals to *aspects*, Pitt (2001) talks of *alter egos* and Saul herself (1997) mentions the possibility of using *temporal phases* to do the job. The evidence presented by Predelli (1999, 2001a) strongly supports an analysis in terms of some such entities: what Predelli's evidence suggests is that, in the move from (31) to (32), we are not swapping like for like, and thus that, as they are used in these two sentences, the names 'Superman' and 'Clark Kent' are not co-referential. How could this idea be developed within a cognitive framework making use of the notion of individual concepts? Ask yourself whether the following sentence is true:

40) Superman works as a reporter on the Daily Planet.

It seems to me that there are conflicting intuitions on this question: on the one hand there seems to be a strong intuition that Superman does not work as a reporter on the Daily Planet, while on the other there is the intuition that, since Superman is Clark Kent and Clark Kent works for the Daily Planet, so does Superman. How might we

explain these conflicting intuitions? On an individual concept analysis there is a straightforward story to tell: we have three separate individual concepts, one of Superman-Clark Kent, one of Superman and one of Clark Kent. In other words, we conceptualise as an individual not just that entity who appears sometimes as a superhero and sometimes as a reporter but also both of the sub-parts of that entity. This analysis receives support from a range of beliefs concerning Superman and Clark Kent: Superman wears a red cape but Clark Kent doesn't; Clark Kent wears glasses but Superman doesn't, and so on. Were we to have just a single Superman-Clark concept, then, given these beliefs, that concept would be riddled with contradiction. It seems, however, that there is, intuitively, no contradiction in these beliefs. What this comes down to is that, as well as conceptualising Superman and Clark Kent as distinct dimensions of the same individual, we also conceptualise them as discrete individuals.

This analysis leads to a natural interpretation of the conflicting intuitions concerning the truth value of (40): we will take the proposition expressed by (40) to be true when we interpret 'Superman' as corresponding to our Superman-Clark Kent concept, whereas we will take it to be false when we interpret 'Superman' as corresponding to our Superman concept. I conclude, therefore, that Saul's examples of failure of substitutivity in non-embedded contexts require separate treatment from failure of substitutivity in propositional attitude contexts, and not only pose no threat to the kind of analysis of proper names that I favour, but can, in fact, be handled naturally within such an analysis.

### 3.2.3.2 Empty names

In the previous section I aimed to show that an account of the semantics and pragmatics of proper names along the lines I favour can offer a plausible solution to the puzzles concerning co-referring names that have so worried philosophers of language. In this section I want to turn to how my account might handle data concerning empty names, i.e. names that fail to name any individual. As with co-referring names, there are, broadly speaking, two distinct problems raised by empty names, one concerning significance and the other concerning truth value. I shall address these problems in order.

Recall sentence (7), repeated here as (41):



41) Nat Tate lived and worked in New York.

The problem raised by examples such as (41) can be illustrated thus: shortly before William Boyd's hoax is revealed to the world, Bill, a leading light on the New York art scene, utters (41) to Ben; a few days later Bill and Ben discover that they have been taken in by a hoax and that Nat Tate is no more than a figment of Boyd's imagination. What will Bill and Ben now think of Bill's utterance of (41)? Intuitively it seems that Bill will now believe that what he said in uttering (41) was not true, given the non-existence of Nat Tate, but he will surely not believe that it was meaningless; equally it will seem to Ben that his lack of awareness of the hoax barred him from grasping the truth value of Bill's utterance, but did not bar him from grasping its *meaning*. These intuitions of course pose something of a problem for accounts of proper names within the direct reference tradition: if, as the classical direct reference position holds, the meaning properties of a proper name are exhausted by that name's referent, then 'Nat Tate', failing as it does to designate any individual, should lack meaning. Yet both Bill and Ben's intuitions seem incompatible with this result: Bill's intuition suggests that his utterance is fully meaningful, while Ben's intuition suggests that understanding an empty proper name is compatible with failing to realise that it is empty.

There have been many attempts, within the direct reference framework, to address these problems. Some have bitten the bullet and accepted the conclusion that Bill's utterance of (41) does fail to express a full proposition, explaining away Bill and Ben's intuitions by appealing to extra implicated descriptive propositions<sup>17</sup>; others have found this result unacceptable and have instead eliminated empty names by positing an ontologically special class of objects for them to refer to<sup>18</sup>. There is not the space here to examine these accounts in detail (although I shall be saying more about accounts such as Salmon's below). The problem for almost all, however, is that they violate one or other half of the key intuition: either, as with bite-the-bullet accounts, they violate the intuition that empty names sentences are fully meaningful,

---

<sup>17</sup> See, for example, Braun (1993), Taylor (2000) and Reimer (2001).

<sup>18</sup> See Salmon (1998) for one of the most recent versions of this type of view.

or, as with Salmon-type accounts, they violate the intuition that empty names are empty.

Accounts of proper names within the direct reference tradition thus seem to face serious difficulties in dealing with the data concerning empty names. That is not to say that these difficulties are fatal. If, however, we could find an account which could respect the intuition that empty name sentences are fully significant as well as the intuition that the names they contain are genuinely empty, then we should, *ceteris paribus*, prefer it. The account of proper names proposed by Frege seems to be able to reconcile just these two desiderata. For Frege, the name 'Nat Tate', as a meaningful linguistic expression, must be associated with a sense. That Nat Tate does not exist simply tells us that whatever the sense of 'Nat Tate' may be, it does not happen to determine any individual. However, since, for Frege, the bearer of a sentence's cognitive significance is identical to the proposition expressed by that sentence is identical to that sentence's sense, the failure of the sense of 'Nat Tate' to pick out an object in the world is neither here nor there for the propositional content or significance of a sentence such as (41). (41) will still express a complete proposition and will thus be fully significant. The problem, of course, lies in the *ceteris paribus* caveat above: as Kripke (1972), Donnellan (1970) and many others have shown, all things are not equal for descriptive theories of names. It thus seems that our best chance of accounting for the significance of empty name sentences has been snatched away.

The account I have outlined, however, is able to reconcile Fregeanism on empty names with Kripke's objections to description theories (in the case of referential uses of proper names, the only type of use which gives rise to the problems under discussion). The key to this reconciliation lies once more in the anatomy of individual concepts which I proposed in chapter 2 for *de re* individual concepts, descriptive content is overridden, as far as truth conditions are concerned, by causal relation; what a *de re* individual concept is a concept *of*, that is to say, is determined relationally, not satisfactionally. One corollary of this mutual independence between internal and external dimensions is that a *de re* individual concept may have a fully articulated internal dimension, without its external dimension happening to pick out anything in the world. Once more we return to the idea, mentioned in the previous section, that as far as the cognitive significance of referentially-used proper names

goes, reference is simply a red herring. There is thus no reason, on the present picture, to suppose that a sentence containing an empty name should be any more or less cognitively significant than an equivalent sentence containing a non-empty name. This analysis of the cognitive significance of referential uses of empty names is not necessarily in conflict with direct reference positions. The key claim of the direct reference position on names is that referential uses of names (which is to say, on the DR position, *all* uses of names) contribute nothing to truth conditional content other than their referents, and this is just what my analysis holds. Given that it does not identify meaning with truth-conditional contribution, however, my analysis makes available an account of the cognitive significance of names based on their linguistic meaning.

It is of course the case that, for non-referential uses of proper names, there is no problem with intuitions on the significance of name sentences. Imagine that it turns out that those women who it is believed were murdered by Jack the Ripper died, in fact, of a peculiar illness which leaves its victims looking as if they had been gruesomely murdered. Under those circumstances a use of 'Jack the Ripper' intended by the speaker to be truth-conditionally anchored to the property of having committed the Whitechapel murders would, presumably, be empty. Given the analysis of descriptive individual concepts in chapter 2, however, there is no temptation to think that a sentence such as (42) should, on such a use, be meaningless:

42) Jack the Ripper came from Liverpool.

for more-or-less Fregean reasons: given the speaker's intentions, there will be an entirely determinate condition for denotation, and thus an equally determinate proposition expressed.

What I hope I have established so far is that the account I favour ties in with intuitions on the significance of sentences containing empty names; it ties in, that is, with the intuition that a sentence containing an empty name can be as *meaningful* as one containing a non-empty name. But what, on the present picture, does this notion of *meaningfulness* amount to? Given the picture of individual concepts with which I am working, it amounts to the idea that, from a solipsistic perspective, there is nothing to distinguish the type of thought expressed by a sentence containing an empty name

from that expressed by a sentence containing a non-empty name. What it does not amount to is a claim that all utterances of sentences containing empty names express determinate sets of truth conditions. In this respect, my account inherits what has been seen by many to be an unacceptable aspect of the Fregean picture: it is committed, as, *pace* the views of Evans (1982), is Frege, to the possibility that a sentence may be fully meaningful while lacking a truth value. I first want to outline why and how my account is committed to this conclusion; once I have done so, I shall examine why this result has seemed so unpalatable and whether we should really be worried by it.

For variety's sake, consider sentence (43):

43) Nat Tate was over six feet tall.

Has a speaker, in uttering (43), said something true or false? Intuitions on this sort of question are notoriously variable, but there is surely one point on which all sides can agree: an utterance of (43) does not express a true proposition<sup>19</sup>. Beyond this it seems that there is no consensus: for the direct reference theorist the failure of sentences such as (43) to express any proposition and thus to bear a truth value is in line with the referential view of proper names, while for the descriptivist, (43) expresses a complete and false proposition, thus supporting a descriptive view of names. Of course, as with most of the extensive use made of truth-value intuitions in the literature, opposing intuitions are either ignored or dismissed out of hand<sup>20</sup>. My intuition is that the imagined utterance of (43) fails to have any truth value. But what is my intuition an intuition of? And why (other than for reasons of theoretical expediency) should others have opposing intuitions? An account able to capture such apparently opposing intuitions should, I suggest, be preferred to one that embraces one intuition at the expense of another.

How should we analyse what a speaker of (43) has said on the sort of analysis I have proposed? There is no one answer to this question: how we analyse the content of (43) will depend on the external dimension of the individual concept corresponding to the speaker's use of 'Nat Tate'. My intuition, on which (43) has no truth-value,

---

<sup>19</sup> Things are not, I believe, in fact as simple as this. The judgment that (43) is not true depends, I shall claim below, on our taking it as a statement about the actual world. Taken as such, however, this intuition seems reasonably robust.

<sup>20</sup> Notable exceptions to this generalisation can be found in the growing literature in support of multi-proposition semantics. See, for instance, Bach (1999) and Neale (1999).

might naturally be seen to correspond to an utterance of (43) on which the speaker takes her individual concept to be TC-anchored relationally or to be TC-anchored to the actual satisfier of some descriptive condition. In other words, my intuition concerns referential uses of proper names: on a referential use, a speaker cedes control over the truth-conditional content of her use of a name, with the resulting possibility that the name may fail to have any truth-conditional content. It is, of course, one of my central claims that these are not the only truth-conditional possibilities for individual concepts: it is open to a speaker to TC-anchor her use of a name directly to a descriptive condition. It seems plausible to suppose that it is this sort of use which the descriptivist has in mind when considering these examples, and which underlie his intuitions on truth value. There is, after all, quite an impetus in this direction: as we have seen, an utterance of (43) seems to be meaningful although there is no individual designated by 'Nat Tate', thus 'Nat Tate' must be semantically linked to some descriptive condition. It is only in the light of a two-component analysis of individual concepts that this line of reasoning loses its force. Given that it makes available more than one type of proposition expressed, truth-conditionally individuated, the two-component picture is thus able to accommodate both the referentialist's intuition that (43) lacks a truth value and the descriptivist's intuition that (43) expresses a false proposition.

As I have already mentioned, however, it accommodates these intuitions at the cost of allowing for the possibility that an utterance may be, in some sense, fully meaningful, while failing to bear a truth value, a result that many have seen as unappealing. So why this distaste? The account which standardly bears the brunt of this criticism is Frege's, and I suspect that we can trace the criticism back not to the result itself, but to Frege's version of it, or rather to how it fits in to Frege's overall account. As mentioned above, Frege identifies the proposition expressed by a sentence with the cognitive significance of that sentence and with that which determines the sentence's truth value, all roles played by the sentence's sense. For Frege, the claim that a sentence containing an empty name can be fully significant is the claim that it has a complete sense. However, to claim that a sentence has a complete sense is, on Frege's picture, also to claim that it expresses a complete proposition, and this is where I think the problem lies: it has seemed to many that on Frege's (or any other) account, it is simply incoherent to claim that a sentence can

express a complete proposition yet lack a truth value. But it is important to realise that this is not the claim I am making: on my analysis, an utterance containing a referential use of an empty name fails to express a complete proposition, i.e. fails to express a determinate truth condition, for the very reason that the name is empty and thus fails to contribute anything to truth-conditional content. This does not, however, rob the utterance of meaning since, from an internal perspective, its interpretation will be type-indistinguishable from that of an utterance containing a non-empty name. This is simply the result of the abdication of responsibility for truth-conditional content which characterises referential uses of names (and other singular expressions). It seems to me, therefore, that we should not be too worried about allowing in meaningful utterances without truth values. I shall discuss in chapter 7, however, the implications of such a conclusion for truth-conditional approaches to natural language meaning.

On my analysis the truth-conditions of a name sentence, and thus the possible truth values of sentences containing empty names, are not predetermined by linguistic meaning: some utterances containing empty names may fail to express any determinate set of truth conditions, while other may be false. There is, however, one class of empty name sentences which seems to resist such an analysis, and which has been the subject of much interest in the literature: the class of negative existentials. Recall sentence (8), repeated here as (44):

44) Nat Tate does not exist<sup>21</sup>.

The problem here is that, not only is (44) clearly significant, it also has every appearance of being true. This presents problems for all sides: for the referentialist, (44) should fail to express a determinate set of truth conditions, since 'Nat Tate', being empty, fails to contribute anything to propositional content. And, while the possibility that a sentence might express a complete proposition yet lack a truth value has seemed unappealing, the possibility of a sentence failing to express a complete proposition yet having a truth value seems simply incoherent. For the Fregean descriptivist there are problems too: as a meaningful name 'Nat Tate' surely has a sense; it is thus to be expected that (44) should express a full proposition and be fully

---

<sup>21</sup> I am once more abstracting away, I believe harmlessly, from issues of tense.

cognitively significant. However, since, on Frege's picture, the truth value of a sentence is a function not of the senses of the sentence's component parts but of their references, the failure of 'Nat Tate' to refer should rob (44) of a truth value. Yet there is what is at the least a strong intuition that (44) is true. Negative existentials are, in fact, true in precisely those cases where the name they contain does fail to refer. Drawing back a bit, we can see that the problems for both sides are of the same kind: on both referentialist and Fregean positions, a name sentence will be true iff the object designated by the name satisfies the property expressed by the predicate; in other words, a name sentence is viewed, on both these pictures, as a tool used to say of some object that it has some property. But on such a view it is entirely mysterious that negative existentials should be capable of being true.

The version of descriptivism espoused by Russell seems to offer a way out of these problems<sup>22</sup>. On Russell's analysis, discussed briefly in chapter 1, a natural language proper name such as 'Nat Tate' must be reanalysed as a disguised description<sup>23</sup>. Thus sentence (45) is semantically equivalent to sentence (46):

45) Nat Tate exists.

46) The  $\phi$  exists.

where  $\phi$  is to stand for some property which Nat Tate uniquely satisfies. This, in turn, given the Theory of Descriptions, has the logical form:

47)  $\exists x (\phi x \wedge \forall y (\phi y \rightarrow y = x) \wedge Ex)$

Now, since, on Russell's picture, the locutions 'there is a such-and-such' and 'a such-and-such exists' are taken to be equivalent, the existence predicate in (47) is otiose, and we can thus reduce (47) to (48):

48)  $\exists x (\phi x \wedge \forall y (\phi y \rightarrow y = x))$

---

<sup>22</sup> See, for instance, Russell (1905).

<sup>23</sup> *Pace* Sainsbury (1995)'s interpretation of Russell's views.

given which, the logical form of (44) comes out as:

$$49) \neg \exists x (\phi x \wedge \forall y (\phi y \rightarrow y = x))$$

This analysis seems to be in better shape than the alternatives canvassed so far: it has a story to tell not just about why negative existentials appear to be true when the name they contain fails to designate an individual, but why they actually *are* true in those circumstances. Given the logical form in (49), a negative existential will be true so long as there is no individual which satisfies the description for which the name is shorthand, i.e. so long as the name fails to denote, in Russell's terms. Russell's analysis however will not do: not only does it suffer from the problems faced by any descriptive general account of names, it also violates grammatical appearance by treating 'exists' not as a genuine first-order predicate. It seems therefore that we must look elsewhere for a solution to the puzzle set by negative existentials.

The proper treatment of existence and of negative existentials is a complex question and one which we shall only have the space to skate over here. I do however want to suggest the direction in which I would be inclined to look for a solution, given the sort of analysis of the cognitive underpinnings of reference outlined in chapter 2. The first point to notice is that, whatever, the right solution may be, it is not to do with the specifics of the semantics of proper names. Just the same phenomenon arises with all types of singular expression. Consider Bill and Ben again, who are playing a virtual reality game in which slavering monsters appear to be rushing towards them with unfriendly intentions. Bill, concerned that one particularly fearsome monster might seem too lifelike for Ben's comfort, may utter any of the following:

50) That doesn't exist.

51) He doesn't exist.

52) That monster with green eyes and red fur doesn't exist.

53) The monster with green eyes and red fur doesn't exist.

Any solution that we may want to propose to the problem raised by (44) will have to apply equally to (50)–(53). So where might we look for such a solution? I want to start not with the sort of empty names that we have been concentrating on so far, but with



fictional names, and with what I take to be a key intuition about such names. Consider the following sentence:

54) Sherlock Holmes plays the violin.

Is the proposition expressed by (54) true or false? The intuition that I want to focus on is that, although in the actual world (54) does not express a true proposition, it is *true in the fiction*. How might we read this intuition? One natural way might be to say that what truth value we will be inclined to assign to (54) will depend on the context in which it is used, and, in particular, on what, given the context, we take (54) to be a statement about. If we take (54) to be a statement about how things are in the actual world then we would not be inclined to judge it as true. If, however, we take it as a statement about how things are, not in the actual world, but in the non-actual world of Conan Doyle's fictions, then it seems entirely natural to judge (54) as true. What might this suggest? On the face of it, it seems to suggest that we might try and approach fictional names via a notion of parallel non-actual worlds inhabited by non-actual individuals; via, in other words, a version of Lewisian realism on possible worlds<sup>24</sup>.

How might such an approach help us with an analysis of (54)? What we would have to say is something like this: although 'Sherlock Holmes' doesn't have a referent in the actual world, it does have a referent in Conan Doyle's non-actual world. So, when we take (54) to be a statement about that non-actual world, there is no problem with it expressing a full set of truth conditions, and no problem with it being true. What if we take it as a statement about the actual world? Here we return to a case similar to that discussed in relation to (43), although concerning an overtly fictional name, rather than a supposedly non-fictional name. The hearer who takes (54) to be a statement about the actual world must choose whether the name 'Sherlock Holmes' is intended to refer in the actual world or not. This is clearly a problematic question: if he decides the former, then, since Sherlock Holmes does not exist in the actual world, (54) will fail to express a complete proposition; on the other hand, if he decides the latter, then, since individuals who do not exist in the actual world cannot hold

---

<sup>24</sup> That is not to say, as will become clear below, that the conception of non-actual worlds with which I want to work is Lewis's own (see, for instance, Lewis (1986)), or that I ultimately want to commit myself to any version of modal realism construed in the strong sense used by Lewis himself.

properties in the actual world, (54) will, although expressing a complete proposition, express a proposition that must be false.

This analysis glosses over two interesting questions: firstly, if we can entertain both *de re* and descriptive concepts of individuals existing in the actual world, can we entertain both kinds of concepts of individuals existing in non-actual worlds? As far as (54) is concerned, not much hangs on this, but there are, it seems to me, examples which suggest an answer. Consider (55) and (56):

55) Sherlock Holmes might not have been a detective.

56) Sherlock Holmes might not have been called 'Sherlock Holmes'.

It seems to me that both of these sentences are straightforwardly true in Conan Doyle's fictional world. This suggests that fictional names used in statements about fictional worlds can be used every bit as referentially as non-fictional names in statements about the non-fictional world. But via what mechanism might this work? I do not believe it is plausible that we stand in a causal relation appropriate for reference with non-actual individuals. The only other option, then, is that this referentiality is effected via what in chapter 2 I called 'stipulative *de re*-ness'; in other words, the referentiality of 'Sherlock Holmes' as used in statements about the world of Conan Doyle's fictions is based on our ability to anchor individual concepts truth-conditionally to worlds.

The second question concerns the sort of abstract objects posited by Salmon (and by ersatzists more generally). For Salmon, (54) will be true in the fiction not because 'Sherlock Holmes' refers to an individual who has the property of playing the violin in any world, but because 'Sherlock Holmes' refers to an abstract object existing in the actual world, and, as part of his fictions, Conan Doyle pretends that that abstract object plays the violin. Should we accept such abstract objects into our ontology, and, if so, what role should they be assigned? I believe we should accept them, but that they should not be taken to be the only referents of fictional names. Consider the following two sentences:

57) Sherlock Holmes was so-named by Conan Doyle.

58) Sherlock Holmes was so-named by his parents.

(57) seems to be true if taken as a statement about the actual world, although false if taken as a statement about the fictional world created by Conan Doyle, since Conan Doyle himself does not exist in that fictional world. (58), by contrast, seems (likely to be) true in the fiction, although false in the actual world (or failing to express a complete proposition). This seems to suggest that ‘Sherlock Holmes’ *can* be used to pick out some actual entity, and the abstract fictional characters of Salmon’s ontology will do as well as anything else for this rôle. If, however, we take these abstract entities to be the only referents of fictional names (if we take, for instance, ‘Sherlock Holmes’ always to refer to the real-world fictional character created by Conan Doyle), then we will be forced to provide an analysis of the truth conditions of (58) of an entirely different kind from the analysis we provide of the truth conditions of (57), and this is just what Salmon does. This seems at least a methodologically undesirable result; an account which could assign truth conditions to (57) and (58) via the same mechanism would presumably be preferable to one that required an appeal to entirely separate mechanisms. Furthermore, as Phillips (2001) points out, if Salmon-type abstract object are the *only* entities to which empty names can refer, then there is no principled way of preventing non-empty names from also referring to abstract objects. If, however, we allow for the possibility that ‘Sherlock Holmes’, for instance, refers not only to an actual abstract object but to a non-actual concrete object, then we can begin to see a way out of this difficulty. An account along these lines would also be likely to fare better in its approach to negative existentials, to which we now return.

How might taking non-actual, fictional worlds seriously help us to address the problems raised by negative existentials? We first need to say something about existence. When we say that something exists or doesn’t exist, we are usually taken as meaning that it exists or doesn’t exist in the actual world. But, just as other predications can be taken to be true or false relative to non-actual worlds, so can existence predications. In Hitchcock’s *Psycho*, Norman Bates spends much time talking off-screen with his mother. As the film progresses, however, it becomes clear that he’s actually talking to himself and that there is no mother. Now someone watching the film could, it seems to me, gasp half way through and, having picked up on the way things are going, utter (59):

59) Oh no, Mrs. Bates doesn't exist.

and I think we should be inclined to say that she's right, that Mrs. Bates doesn't exist *in the world of the fiction*. Of course Mrs. Bates doesn't exist in the actual world either, but it is surely not this that makes what the speaker of (59) has said true: (59) could, as far as the speaker knows, turn out not to be true, i.e. it might turn out that Mrs. Bates does exist in the world of the fiction, although she still doesn't exist in the real world. So it seems that, as with any other predication, we need to know which world an existence statement is intended to be a statement about in order to establish what proposition has been expressed and thus whether it's true or false.

It seems to me, however, that it is in respect of just this question, the question of which world a statement is intended to be interpreted as being about, that the special nature of existence statements comes out. Consider (60):

60) Sherlock Holmes doesn't exist.

How will a hearer interpret an utterance of (60)? The first question he will have to answer is which world it's intended to be taken as a statement about, a question that he should resolve on the basis of pragmatic principles and context. Let's take it that it's intended as a statement about the actual world (this is of course not necessary – Sherlock Holmes after all doesn't exist in War and Peace). In that case, how should we interpret the name 'Sherlock Holmes'? As with (54), we have two choices: either 'Sherlock Holmes' is intended to pick out an individual in the actual world, or it is intended to pick out an individual in a non-actual world. The former possibility should be eliminated on grounds of relevance: if 'Sherlock Holmes' were intended to pick out an individual existing in the actual world, then, given its predicate, (60) would necessarily express a false proposition. We are left, then, with an interpretation on which 'Sherlock Holmes' is taken to pick out an individual in a non-actual world, i.e. on which it picks out the individual in Conan Doyle's fictions. On this interpretation, (60) will express a proposition which looks something like (61):

61) Sherlock Holmes [the detective in Conan Doyle's stories] does not exist in the actual world.

Putting this in the cognitive terms developed in chapter 2, the interpretation of 'Sherlock Holmes' should consist in an individual concept which is truth-conditionally anchored not to the actual world but to the non-actual world of Conan Doyle's fictions. What it would take for this Sherlock Holmes, or any other fictional character, to exist in the real world, is an open question. It may very well be that Sherlock Holmes necessarily does not exist in the actual world. But this isn't a question we need to worry about here – we can, it seems to me, assume that, since sentence (61) is meaningful we do have some intuitive grasp of what conditions would have to be fulfilled in order for Sherlock Holmes to exist in the real world, although those conditions may in fact be unsatisfiable.

So far we have restricted ourselves to a discussion of fictional names. How might we extend this account to non-fictional empty names? There seems to be a very natural extension to the example we began with, the example in (44), repeated here as (62):

62) Nat Tate does not exist.

It seems that what we need to say is that, just as Sherlock Holmes inhabits the world of Conan Doyle's fiction, so Nat Tate inhabits the world of William Boyd's hoax biography. It seems, in fact, that Nat Tate is really a fictional character, although this is hidden from us by the fact that at one time some people believed he was not, i.e. believed that the person who inhabited the book existed in the actual world. So maybe, to see how the story goes, we should set ourselves a slightly harder task. Consider sentence (63), the example which is universally used in the literature on empty names:

63) Vulcan doesn't exist.

Nineteenth century astronomers looked into the skies and noticed an anomaly in the orbit of Mercury. They posited that this anomaly could only be caused by the presence of another planet in an orbit between Mercury and the Sun and they named this planet 'Vulcan'. Of course it actually turned out that there was no such planet,

and that the oddities were simply a result of general relativity, thus 'Vulcan' turned out to be an empty name. How, on the sketch above, should we analyse this story? The natural analysis would look something like this: when astronomers posited the existence of Vulcan, they thereby posited a world that was, in fact, different from the actual world; they posited, in other words, a non-actual world. And it is this world that Vulcan inhabits. Given this analysis, sentence (64):

64) Vulcan is a planet affecting the orbit of Mercury.

should be taken as true if intended as a statement about the relevant theoretical constructs, although of course it will not be true if taken as a statement about the actual world. So the negative existential in (63) will fall out in much the same way as the others we've seen: on the intended interpretation, the individual concept which constitutes the interpretation of 'Vulcan' will be truth-conditionally anchored not to the actual world but to the non-actual world of astronomical theory, and (63) will express the proposition that the object to which it is truth-conditionally anchored does not exist in the actual world. It looks, therefore, as if our story may be applicable to non-fictional as well as fictional negative existentials. The analysis it offers comes down to this: what negative existentials express is not the problematic proposition that some actual object does not exist, but rather a proposition concerning the failure of an individual existing at one world to exist at another.

But, given how unappealing a metaphysical commitment to possible worlds has seemed to so many, how should we take this talk of non-actual worlds? Should we accept that, given the explanatory advantages to be bought by an appeal to non-actual worlds, we must, in the spirit of Lewis, bite the bullet and accept the reality of non-actual worlds? Or should we, as many analyses of modality have done, attempt to have our cake and eat it; attempt, in other words, to retain the explanatory advantages of possible worlds without the unappetising metaphysical commitment? I want to try to steer a middle course. Again, it is not my aim here to provide a fully articulated analysis of the metaphysical status of non-actual worlds. My aim is the much more modest one of sketching the outline of what might prove a profitable line of enquiry.

On the cognitive picture I have so far elaborated, what is the relation between individual concepts and the actual external world? We might look at it in these terms:

the world is an *information base* for individual concepts; that is to say, our cognitive aim is to maximise the correlation between the information we store in our individual concepts and the way things are with the objects our concepts are concepts of. If, for instance, I have a concept of Tipper Gore, I am only justified in adding the piece of information *x is married to Al Gore* to that concept if the world is such that Tipper Gore is indeed married to Al Gore. Now what I want to suggest is that we should, in providing an analysis of empty names, commit ourselves metaphysically to non-actual worlds only so far as is necessary in order to accept that the actual world is one of many information bases. This sort of analysis is not entirely dissimilar to some varieties of what Lewis (1986) calls *linguistic ersatzism*; for Adams (1974), for instance, possible worlds are to be identified with *maximally consistent sets of propositions*. However, I believe that, restricting ourselves to an application of this notion to empty names, rather than to the broader question of modality, and viewing it from a cognitive perspective, we need not fall prey to the problems facing accounts of a linguistically ersatz stripe.

What I want to suggest, then, is this: we treat information-bases other than the actual world conceptually in just the same way that we treat the actual world. We take such information bases to support the formation of individual concepts, to verify or falsify predications, to support inferences and so on. And, just as with concepts that correspond to the actual world, we can perform particular operations on concepts linked to non-actual information-bases. We can, for instance, actualise them to the particular information base to which they relate; we can, in other words, attach a REF feature to them which anchors them truth-conditionally to that particular information base. This is just what we saw above in relation to examples (55) and (56). On this analysis, (54), for instance, will be true in the fiction, because the predication it expresses is verified by the fictional information base. What, then, is it for a particular predication to be verified by an information base? This is where, I suspect, the actual information base (i.e. the actual world) and non-actual information bases part company: whether a predication is verified by the actual world will depend, assuming that some correspondence theory of truth is right, on whether it corresponds to how non-mental things are arranged in the world; whether a predication is verified by a non-actual information base, by contrast, will presumably depend on whether it corresponds to how its creator has decided things are with the information base. I do

not, however, wish to suggest that information bases, as I conceive of them, are simply Adams' sets of maximally consistent propositions in disguise; they are in some ways narrower and in others broader. On the one hand, there are many predications which any particular non-actual information base will neither verify nor falsify. So, for instance, the information base that Conan Doyle created in writing the Sherlock Holmes stories has nothing to say on whether (65) is true or not:

65) Sherlock Holmes met Florence Nightingale.

On the other hand, the predications supported by an information base go well beyond the propositions explicitly expressed in the process of its creation. Thus the Sherlock Holmes information base, would, I imagine, falsify (66), although of course no claim one way or the other is explicitly made in the stories themselves:

66) Sherlock Holmes met Philip Marlowe.

This conception of non-actual worlds as information bases is clearly a long way from Lewis's notion of possible worlds. In particular it is incompatible with the causal isolation that Lewis takes as a fundamental property of possible worlds. On the conception of non-actual worlds which I want to make use of, there are causal relations running in both directions from actual to non-actual worlds: things are as they are in non-actual worlds because of facts about the actual world, facts concerning the intentions, beliefs and so on of human agents; Sherlock Holmes plays the violin, for instance, because of an intention on the part of Conan Doyle. At the same time, we may *believe* that Sherlock Holmes plays the violin because of the way things are with the information base created by Conan Doyle in writing the Sherlock Holmes stories.

This is of course only the barest outline of a possible strategy for tackling empty names within a cognition-based framework. The idea that I want to focus on, however, is that the puzzles raised by empty names and by negative existentials in particular are a product of our conceptual abilities, specifically of our abilities both to create individual concepts which fail to pick out anything in the actual world, and to treat information bases which are not intended to correspond to the actual world conceptually in the same way as we treat the actual world. My claim is that the



combination of these two factors provide the key to unlocking the puzzles posed by negative existentials.

### 3.2.3.3 Individual concepts and rigidity

So far I have attempted to show how my account might be applied to the sort of data that have traditionally been taken to weigh on the descriptive side of the debate on names. Now I come to the key data adduced by the referential side, data concerning the apparent rigidity of proper names. As discussed in 3.1.3, it has been taken as the fundamental intuition underlying direct-reference accounts that name sentences (and sentences containing certain other singular expressions) depend truth-conditionally on the same individual across possible worlds. Consider once more (11) and (12), repeated here as (67) and (68):

67) Louis Armstrong performed with Ella Fitzgerald.

68) The greatest jazz trumpeter performed with Ella Fitzgerald.

The intuition we are focusing on is that, whereas (67) will depend truth-conditionally on the same individual, Louis Armstrong, across possible worlds, i.e. the proposition expressed by (67) will be true at a possible world iff Louis Armstrong performed with Ella Fitzgerald at that world, (68) will depend truth-conditionally on different individuals at different worlds; specifically it will depend on whoever is the greatest jazz trumpeter at that world.

Such intuitions constitute one of the most serious challenges to descriptive accounts of proper names, and much ingenuity has been expended by those who favour such accounts in attempting to explain them away. The problems for my account, as a mixed referential-descriptive account, are rather different, however. I have no reason to deny the referentialist's intuition; it is, after all, central to my account that proper names, as used in certain contexts, can and do contribute nothing but their referents to truth-conditional content. I am thus happy to accept that an utterance of (67) may be truth-conditionally anchored to the same individual across possible worlds. At the same time, I have claimed that proper names, as used in other contexts, contribute descriptive conditions to truth-conditional content. Thus I argued

that, even were it to be established that Jack the Ripper were actually James Maybrick, one could still consistently conceptualise a possible world at which Jack is not Maybrick, a world at which someone other than Maybrick committed the murders attributed to Jack. In order to conceptualise such a world, we must take 'Jack the Ripper' not to be truth-conditionally anchored to any individual in the actual world, but rather to be anchored to a descriptive condition which will be satisfied by different individuals at different worlds. If this is right, that proper names can give rise to descriptive truth conditions, then the challenge posed by intuitions of rigidity to my account is not to explain why utterances with descriptive truth conditions appear to have singular truth conditions, but is rather to explain why, when both types of truth-condition are available, referential uses of proper names so heavily outweigh descriptive uses. If you already accept the conclusion that names can give rise to general truth conditions, then it may be that not a great deal hangs on this. You may, however, feel that, in the absence of an explanation of this datum, you are inclined to reserve judgment on the descriptive use of names. In which case, what I am trying to persuade you of is this: there are good reasons why names should predominantly be used referentially which have nothing to do with the constraints laid on interpretation by their encoded meaning; we can, that is, get predominant referentiality without encoding universal referentiality.

The idea I want to explore is that the predominant referentiality of uses of proper names derives (at least partly) from facts not about the semantics of names but rather about naming practices, and that these facts about naming practices derive, in turn, from general facts about human cognition. Just as most uses of proper names are referential, so most acts of naming are demonstratively *de re*; they are constituted, that is, by utterances such as 'I name this child...', 'I shall call this place...', where the demonstratives involved are truth-conditionally anchored to objects of immediate perception. Thus acts of naming which are backed by non-perceptual descriptive conditions are clearly the exception, although not, of course, unheard of, as examples such as 'Jack the Ripper' and 'Deep Throat' attest. Why should it be, then, that most acts of naming are demonstrative rather than descriptive? The answer I want to suggest is that we name in this way because of facts about our cognitive architecture; in particular, we are cognitively designed primarily to track physical objects rather than their properties. Things could of course have been otherwise. As Bach (2002)

points out, we could have been cognitively geared towards properties, such as name-bearing, rather than to the objects that held those properties:

Suppose we cared about the proper names people had regardless of whose names they were. An employer might want to hire someone because his name was 'Cedric Scampini', a tourist might visit a city because its name was 'Cincinnati', and a diner might be tempted to try a restaurant called 'Colestra'. However frivolous such sentiments might be, people could attach great importance to names and come to regard bearing a certain name as a noteworthy property, regardless of who or what the name belongs to.

Bach (2002)

What I want to suggest is that Bach's idea can be extended to properties other than name-bearing. Although we are cognitively designed so as primarily to discriminate individuals on a physical basis, we do sometimes discriminate according to properties, both where we have to, for lack of the appropriate relation with a physical object, and where the property concerned is, for whatever reason, highly significant in its own right. The examples we have seen so far, 'Jack the Ripper' and 'Deep Throat' seem to combine both of these features: it is presumably the case that both names were originally assigned descriptively because those doing the naming did not in some sense know *who* they were naming. But why should we bother to assign a name descriptively at all? Because the description via which we assign the name is significant enough that we want a way to talk about *whoever* satisfies it. It seems, in fact, that these two features standardly come together for descriptive naming: if the property in question is not significant in its own right, then there will be no motivation to coin a name at all, while if we know from the start who the description picks out and which name that individual bears, then we will usually use this pre-existing name. As we have seen earlier in this chapter, however, this does not entail that descriptive names are names in search of a referent: once they have been coined, descriptive names can continue to be used to express descriptive individual concepts, whether or not a *de re* concept of the description's satisfier is available.

On this picture, on which names are descriptive because of facts about the circumstances of their coining, it seems that context can play only a limited role in determining whether a use of a name is descriptive or referential: it cannot, for instance, allow a name which has been assigned demonstratively to be used descriptively. But is this right? Is referentiality a once-and-for-all thing for proper names? There is at least some evidence to suggest it may not be. Returning to Superman and Clark Kent, and taking all statements to be statements about the fictional world created by Siegel and Shuster, we would presumably want to say that the name 'Superman' was demonstratively assigned; it was assigned, that is, while either explicitly or implicitly demonstrating Superman himself, i.e. while demonstrating the Superman-Clark Kent individual in his Superman guise. Now imagine that the Superman stories had taken a rather different turn; that when Lois finally uncovers the man of steel's true identity, it is Lex Luthor she finds beneath the cape. In that turn of events, sentence (69) is presumably true:

69) Superman is Lex Luthor.

in which case, in the actual world of the stories (actual relative to the stories, that is), (70) will be true (on one reading):

70) Superman might have been Lex Luthor.

If this is right, then 'Superman', although demonstratively assigned to Clark Kent, seems not merely to contribute Clark Kent (or, given the discussion in 3.2.3.1, a sub-part of the Clark-Superman individual) on all uses, but may instead contribute some descriptive condition, presumably a condition something like *x flies around Metropolis in a cape and tights saving people*.

This raises a further question however: if names that are demonstratively assigned *can* be used descriptively, why do we come across so few clear-cut examples of this? One answer might be that it is simply a matter of salience. What we are concerned with here is individual concepts for which there is a possible *de re* anchor but which are nevertheless capable of being used descriptively. The claim might go: such concepts are typically so rich in descriptive content that no one condition is

sufficiently salient to be retrievable consistent with the presumption of optimal relevance. I doubt whether this claim will stand up as it is, however. Consider the name 'George W. Bush'. It seems that, for most of us, there will be one descriptive condition associated with this name which has much greater salience than any other, the condition *x is the President of the United States*. Yet it does not seem that we can use the name 'George W. Bush' in such a way that it is truth-conditionally equivalent to the description 'the President of the United States'. To see this consider (71) and (72):

- 71) The President of the United States might have been a woman (if Hilary Clinton had stood instead of Al Gore).
- 72) George W. Bush might have been a woman (if Hilary Clinton had stood instead of Al Gore).

Such examples seem to weigh heavily against any straightforward salience story. So why can we get a descriptive reading of (70), but not a descriptive reading of (72)? What I want to suggest is that the salience story is right inasmuch as it locates the problem in how rich in information certain individual concepts are, but that it is taking things from the wrong end. The problem with retrieving a descriptive interpretation of (72) is not that there is no one salient descriptive condition, but that there are too many non-salient conditions. In order to interpret 'George W. Bush' descriptively, we would have first to access an individual concept which is rich with descriptive conditions, conditions such as *x is the son of former president George Bush*, *x is the former Governor of Texas*, *x is the brother of the Governor of Florida* and of course *x is the President of the United States*. Having accessed this concept, we would then be forced to ditch all but one of these conditions in order to interpret the name descriptively. Compare this with a descriptive interpretation of (70). The individual concept associated with 'Superman' will presumably contain conditions such as *x can fly*, *x wears a cape*, *x saves people in Metropolis*, and so on. In order to retrieve the imagined descriptive interpretation of (70), we need jettison none of this information; more or less the only information we will have to jettison, in fact, is the condition *x is Clark Kent*. It seems plausible that this distinction between the descriptive interpretations of (70) and (72) should have a knock-on effect for

relevance: while (72), on the intended interpretation, will put a hearer to gratuitous processing effort, gratuitous since there is a more relevant utterance available, i.e. (71), and will thus fail to be optimally relevant, (70) will not put the hearer to gratuitous effort and should thus pass the relevance test.

In this section, then, we have discussed why, although descriptive interpretations of proper names are available on the story I want to tell, they are so markedly less common than referential interpretations. The answer I have given is essentially that there are two causes: firstly, descriptive interpretations of demonstratively-assigned names are hard to come by because the jettisoning of descriptive information standardly robs descriptive interpretations of optimal relevance, and secondly descriptively-assigned names, although standardly open to descriptive interpretation, are rare due to underlying facts about human cognitive architecture. The rarity of descriptive uses of proper names does not, however, render such uses insignificant: it is, so I want to claim, central to an understanding of proper names as linguistic expressions, rather than an understanding of the individual concepts that standardly constitute the interpretations of proper names, that we should recognise the availability of non-referential interpretations.

### 3.3 Names, individual concepts and alternative accounts

So far in this chapter I have laid out the bare bones of the account of the semantics and pragmatics of proper names which I favour and I have examined how this account might offer solutions to some of the more pressing problems for analyses of names. Now I want to move on to look at some other accounts of proper names within the philosophical literature, and in particular at how these accounts compare with my own. In doing so, we will come across questions on proper names that we have not yet addressed, in particular how best to analyse the quantified use of proper names and what story to tell about names with more than one bearer.

As discussed in chapter 1, truth-conditional accounts of singular expressions work on the assumption that tokens of a singular expression-type must make the same type of contribution to the truth conditions of any utterance of a sentence in which they appear. The main positions on proper names can be divided, therefore, according to the type of contribution they take proper names to make to truth conditions. The

current dominant positions might best be viewed as occupying a spectrum as regards truth-conditional contribution. At one end of the spectrum are those who take the meaning properties of a proper name to be exhausted by the name's referent, while at the other end there are those who believe that its referent constitutes no part of the meaning of a name, with names contributing descriptive conditions to truth-conditional content. Between these two extremes are those who claim that the meaning properties of a proper name do go beyond that name's referent but that, nevertheless, its referent is all that a name contributes to propositional content. Let me say a bit more about each of these positions in order to draw out the parallels and differences between each and my own account.

### 3.3.1 Neo-descriptivism on names

Recent champions of extreme descriptivism, i.e. the thesis that proper names are semantically equivalent to definite descriptions, are hard to come by, given the attack on descriptivism launched by, among others, Kripke (e.g. 1972), Donnellan (e.g. 1970) and Kaplan (e.g. 1989a). I want here, however, to examine two such accounts, that of Bach (1987, 2002) and that of Geurts (1997). These accounts have broadly the same shape: for both, a proper name is semantically equivalent to a nominal description, i.e. a description which contains the name itself. For Bach 'PN' is semantically equivalent to 'the bearer of 'PN'', while for Geurts it is equivalent to 'the individual named 'PN''. As far as I can see nothing hangs on this minor difference in formulation. Both also accept that, although the proposition literally expressed by a name sentence (i.e. the proposition determined by linguistic meaning) is descriptive, name sentences may be and often are used to convey singular propositions: in this regard, Bach draws a distinction between what is literally said and what is communicated, while Geurts distinguished between the proposition expressed and the asserted content of an utterance. If I understand these distinctions correctly, a little more may hang on which of these formulations one adopts. Beyond these general similarities, Bach and Geurts diverge in the overall framework within which they place their accounts of names: Bach's account is placed within a largely traditional Russellian framework, while Geurts adopts a DRT-flavoured framework,

within which all definites are taken to be presuppositional. Given this divergence I shall briefly examine these two accounts separately.

Bach's account is chiefly motivated by two types of evidence: what he takes to be descriptive uses of names constituting complete NPs and quantified uses of proper names. We have already seen examples of the former kind in sentences (19) and (21), repeated here as (73) and (74):

73) Aaron Aardvark might have been president.

74) If his parents had named him 'Aristocrates', Aristotle would have been Aristocrates instead of Aristotle.

Sentences (75) to (77), taken from Bach (2002), offer examples of the latter kind:

75) There are seven David Smiths in the APA.

76) There are other intelligent David Kaplans.

77) Only one state has a Salem that is its capital.

Bach takes these examples to show firstly that proper names when they appear as full NPs are semantically equivalent to nominal descriptions and secondly that names can appear as predicates within larger NPs. What should we say about such examples? As I have already indicated, I am not sure how convinced a committed Millian would or should be by the examples in (73) and (74). There is, it seems to me, no pressing obstacle to prevent the Millian claiming that these uses are non-literal. Bach's own account depends on the underlying idea that speakers rarely if ever intend to communicate precisely what is said by the sentences they utter. He accepts, therefore, that names are predominantly used to communicate singular propositions, although, on his account, the propositions expressed by name sentences are general. Given this, it seems that he is, at least in principle, open to the reverse position, i.e. the position that name sentences express singular propositions but can be used to communicate general propositions.

Beyond this potential counter-argument there is a more fundamental problem facing Bach's account in light of the sort of account which I have proposed above. As Bach himself admits, the level of propositional content with which he is working,



what he takes to be essentially Grice's *what is said*, is not accessible to speaker intuitions. Thus it seems to intuition as if the proposition expressed by (78), for instance:

78) Louis Armstrong was born in Storyville.

will be true at any possible world iff Louis Armstrong was born in Storyville at that world. On Bach's story, by contrast, what is literally said by (78) will be true iff there is a unique bearer of 'Louis Armstrong' at that world and he was born in Storyville (or rather at the unique bearer of 'Storyville' at that world). Equally, if someone utters (79) to you:

79) Dave has left town.

two things seem intuitively to be the case: firstly, you will not have fully understood the proposition expressed by the utterance if you do not *know*, in some sense, who Dave is, but secondly that, if you did know who Dave was, you might find that the proposition expressed by (79) was true. On Bach's picture, neither of these hold at the level of what is said. For Bach the literal propositional content of (79) is the general proposition in (80):

80) The bearer of 'Dave' has left town.

the grasping of which requires no knowledge of any particular individual, while, since there are in reality many Daves, the proposition in (80), construed Russellianly, is doomed to be false. Of course Bach accepts that the proposition literally expressed is not standardly an object of speaker intuitions, and that such intuitions usually concern some extra pragmatically communicated proposition. He nevertheless holds that proper names *must* literally express general propositions since, for him, examples such as (73) and (74) show that they *can* express general propositions<sup>25</sup>. This

---

<sup>25</sup> I do not want to suggest that this is the only basis on which Bach makes this claim. My point is that, in the absence of speaker intuitions concerning what Bach takes to be the proposition literally expressed, any argument for the literal propositional content of name sentences which Bach proposes must depend on broad theoretical considerations, i.e. on how things are with other parts of his theory.

argument rests on the truth-conditional assumption that whatever type of contribution we can show proper names to make to the literal content of some sentences, they must make to the literal content of all sentences. This form of argument is, however, undermined by the availability of the type of account which I have outlined. If proper names may, without being ambiguous in any theoretically interesting way<sup>26</sup>, be used to express either singular or general propositions, then we may accept that certain uses of proper names do give rise to general truth conditions without thereby being committed to the claim that names always give rise to general truth conditions. This move allows us to accommodate the sort of intuitions on descriptiveness that Bach takes to be central to an account of proper names, without riding roughshod over the very robust intuitions of rigidity that standard uses of proper names give rise to.

What of the examples in (75) to (77)? As Bach points out, such predicative uses of proper names have been largely ignored in the philosophical literature on names<sup>27</sup>. The reason they have been so ignored is, I imagine, that they simply fall outside the concerns of many of those who have addressed the question of the semantics of proper names. For many, the interest of proper names lies in the light they shed on the nature of reference and thus uses of proper names which do not involve reference are simply beside the point. I, however, committed as I am to providing an account of proper names as linguistic expressions, must take such uses rather more seriously. How might predicative uses of proper names fit in to the picture I have so far been painting? On my account, as on many others, there is a strict division between the property of bearing a proper name and the linguistic use which we may make of that property. The uses of proper names which I have been looking at so far, and which have been the main focus of attention in the philosophical literature on names, use the name-bearing property in order to pick out particular individuals. Thus when I utter (78), for instance, I use the name 'Louis Armstrong' in order to pick out a particular individual, the property of being the bearer of 'Louis Armstrong' playing an intimate role in determining which individual it is that I am talking about. What examples such as (75) to (77) seem to show, however, is that this is not the only linguistic use we may make of the name bearing property. We may also, it seems, use proper names as predicates, with the extension of 'PN' being

---

<sup>26</sup> We shall come on to the discussion of ambiguity in chapter 7.

<sup>27</sup> The most notable exception is Burge (1973), of which more below.

bearers of 'PN'. This much I hope is reasonably uncontroversial. But what does this tell us about the semantics of proper names? Does it, for instance, tell us that proper names just are predicates applying to their bearers? Although I have a certain sympathy with this view, I do not believe it can be quite right.

There are certain clear distinctions between proper names and standard one-place predicates. In particular, standard one-place predicates standing alone cannot have the individual-concept uses typical of proper names. Thus if I call the desk I am working at 'Florence' I can use (81) to express the proposition that my desk has six drawers, but I cannot use (82):

81) Florence has six drawers.

82) ? Desk has six drawers.

One can get round this by claiming that when proper names appear as complete NPs, as in (81), there is a covert determiner in the logical form of the name sentence. You can then, with Bach, semantically identify this covert determiner with 'the', or you can, with Burge (1973), identify it with 'that'. Either way, however, you need to be able to tell some story about why proper names license such covert determiners, whereas standard one-place predicates do not.

This and other differences between the syntactic behaviour of names and that of one-place predicates<sup>28</sup> seem at least to cast doubt on the idea that proper names, *wherever they appear*, are predicates. In the light of this, I would like to suggest an alternative analysis of the predicative use of names. I have, in fact, already more or less pointed to the kind of account I am drawn to: predicative uses of names are just a different way in which we can exploit the name-bearer relation linguistically. I suspect, however, that this use is in some sense derivative of the individual concept use of proper names. What I mean to suggest by this is that proper names are bestowed on individuals primarily for the purpose of allowing us to think and talk about those individuals efficiently. However, the referential exploitation of proper names works as it does because, in bestowing names, we assign properties to individuals. Although the property thus assigned is assigned primarily for referential exploitation, it is nevertheless a property in its own right and, as any other property,

---

<sup>28</sup> For an interesting discussion of such differences, see Segal (2001).

can be exploited for a variety of linguistic purposes. On this picture, then, a name is bestowed for the purpose of allowing us linguistically to express and to entertain an individual concept of the bearer and it is this use to which names are most often put. A by-product, however, of this process is that we may make use of the nominal property assigned to the bearer in other linguistic ways.

Such an analysis gives rise to two questions. Firstly, does this amount to positing an ambiguity in proper names? Yes and no. Proper names are, on this picture, ambiguous in that they can make fundamentally different kinds of contribution to truth conditions: they can either contribute individual concepts, loosely speaking, or they can contribute properties. But it's hard to see how we might avoid this conclusion: even those such as Bach and Burge who take the referring use of proper names to be derivative of the predicative use must accept that the same expression can make these two different types of contribution to truth conditions. Bach, for instance, accepts that, when appearing as a full NP, a proper name is semantically equivalent to a definite description, whereas when appearing as a part of a larger NP, proper names simply express nominal properties. There thus seems to be no escape from truth-conditional ambiguity. On the other hand, the predicative use of proper names is, on the picture I favour, a by-product of the individual-concept use: without the one, the other cannot get off the ground. The ambiguity, such as it is, is therefore of a very different kind from standard cases of lexical ambiguity.

The second question concerns how, on this picture, predicative and individual-concept uses of proper names relate truth-conditionally. Will the individual concept use of proper names end up, as Bach and Burge would have it, being semantically equivalent to a determiner + predicative use? Again the answer is yes and no. There is, of course, nothing to prevent us from defining a determiner, call it *Dname*, such that when it is concatenated with a predicative use of a name it delivers the appropriate truth-conditions (or, better, delivers the appropriate individual concept). I do not believe, however, that English has a determiner which is equivalent to *Dname*. The only two obvious candidates would, I presume, be 'that' and 'the'. Segal (2001), following Higginbotham (1988), runs a convincing argument to show that, if there is an implicit determiner in individual-concept uses of proper names, then it isn't 'that'. He asks us to consider Fred Schmidt stumbling out of a pub. I, who mistake him for Fred Bloggs, might say either (83) or (84):

- 83) Fred has had a little drink.  
84) That Fred has had a little drink.

I would not, however, be making the same claim in both cases. My utterance of (83) seems to make a possibly false statement about Fred Bloggs, whereas my utterance of (84) seems to be a probably true statement about Fred Schmidt. It seems, therefore, that 'Fred' and 'that Fred' are not equivalent. But nor are 'Fred' and 'the Fred' (or perhaps better 'the person called 'Fred'', given the ungrammaticality of 'the Fred'). Imagine that we have been invited to a party by our good friend Fred. Turning up we find the room is full of people who are all called 'Fred'. You could in that context, utter either (85) or (86):

- 85) Fred is surrounded by Freds.  
86) The person called 'Fred' is surrounded by Freds.

There seems, however, to be a clear distinction between these two utterances in the context imagined: the first appears to express a true proposition about our friend Fred, whereas the second does not. The problem here, as I shall argue when I come on to discuss definite descriptions, is that the use of a definite description requires that its nominal element should be uniquely denoting in an available context, whereas the use of a proper name involves no such requirement. Given evidence such as the above, it seems that 'the' is out of the running too. Since, therefore, the property of being the bearer of a proper name, the very property expressed by predicative uses, plays a role in the interpretation of individual concept uses of proper names, such individual concept uses are potentially susceptible to a determiner + predicate analysis. No English determiner, however, will do the job for us.

Overall, then, the evidence Bach adduces does not, it seems to me, force us towards a nominal description account of proper names. Given how robust intuitions of the referentiality of standard uses of proper names are, an account which must posit that all such uses are non-literal should be rejected, *ceteris paribus*, for one that can allow for such uses literally to express singular propositions. Within the *ceteris paribus* proviso, however, is the requirement that the account we adopt should allow

for names in argument positions to make descriptive contributions to truth-conditional content on occasion, given what I take to be the equally strong intuitions in support of that claim. The account I have proposed stands as a possible way out of this theoretical deadlock between referentialists and descriptivists.

The descriptivist account of names offered by Geurts (1997) is of a rather different kind from Bach's. In rough outline, Geurts' central claim is that all definites are presuppositional expression. In the DRT-flavoured approach to presupposition with which Geurts is working, this amounts to the claim not just that definites encode an existential presupposition, but that the presence of this presupposition has certain implications for the interpretation procedure followed by a hearer. For Geurts, presupposition and anaphora are closely linked. A hearer interpreting a definite will first search for an antecedent for that definite in his current DRS. If he fails to find such an antecedent then he will accommodate the definite's presupposition while attempting to link it to an element of his pre-existing representation of the world. Given this background story on definites, Geurts' claim is that an utterance of a sentence containing a proper name 'PN' will give rise to the same proposition expressed and the same interpretation procedure as an utterance of an equivalent sentence in which the definite description 'the individual named 'PN'' is substituted for 'PN'. Consider (87):

87) John is a stockbroker.

This will, on Geurts' picture, give rise to a DRS which will look something like:

88) [x: individual called 'John' x, stockbroker x]

This, as it stands, equates to a descriptive proposition, given a standard analysis of DRSs and reference markers. However, the underlining indicates that there is a presupposition which requires resolution. The hearer is thus semantically mandated to search for an antecedent for the reference marker within his DRS. In this case the search will fail, although Geurts offers an example using a proper name in which the search will succeed:

89) I have three friends: John, Jack and Joe. My best friend is John.

Once the search has failed, the hearer of (87) will do two things. First he will accommodate 'John', leading to a representation which will look something like (90):

90) [x: individual called 'John' x, speaker believes: [x: individual called 'John'  
x, stockbroker x]]

and secondly, he will attempt to link the discourse referent with some pre-existing representation. If, for instance, he knows of an individual called 'John', he may link the reference marker in his DRS with his mental representation of this individual.

The picture Geurts paints thus makes use of a variety of distinct levels of interpretation. Firstly there is the initial DRS constructed by a hearer, a DRS which is entirely determined by linguistic meaning and is therefore context-independent. Secondly there is the DRS which results from following the instructions encoded by presuppositional expressions, i.e. a DRS in which the discourse referent introduced by the proper name is either bound or accommodated. And finally there is the representation which results from linking this DRS to pre-existing conceptual material. In parallel with these distinct levels of interpretation, Geurts draws a distinction between the proposition expressed by an utterance and the asserted content of that utterance. Given this distinction, he runs a story on rigidity which has much in common with that told by Bach: in very rough terms, most utterances containing proper names appear to express singular propositions because their asserted content is demonstratively linked to objects of perception, although the propositions they literally express are in reality descriptive.

Geurts bases his claims on a range of evidence which he takes to show that names pattern with definite descriptions and should thus be taken to fall within the same semantic class. Among these pieces of evidence is, for instance, that names are subject to definiteness effects in just the way that definite descriptions are. Thus (91) is grammatical whereas (92) and (93) are not:

91) There is a philosopher available.

92) \*There is the philosopher available.

93) \*There is John available.

I find much to agree with both in the weight Geurts attaches to such evidence and in the account he proposes. As will become clear once I have come on to discuss the semantics and pragmatics of definite descriptions, I believe Geurts is right to hold that there is much semantic similarity between names and descriptions: both, in the terms I favour, are tools for expressing individual concepts, although, as argued above, I do not believe that proper names are semantically equivalent to nominal descriptions. To say, however, that there is much semantic similarity between names and descriptions is not, on the picture I am painting, equivalent to the claim that proper names can only literally give rise to general truth conditions. Given the framework within which he builds his account, Geurts does seem committed to this claim, a claim that I have already argued against above. In his distinction between proposition expressed and asserted content, however, Geurts appears to be pointing towards the sort of account that I have proposed. One could look at his story on names as follows: the context-independent meaning of a proper name consists in a descriptive condition combined with an instruction constraining the use to which a hearer may put this descriptive condition; given this context-independent meaning, a hearer will build an interpretation which may be anchored to certain of his pre-existing conceptual representations. Expressed in these terms, there is much common ground between my account and Geurts'. There are, however, certain key differences.

For Geurts, the context-independent meaning of a name sentence determines a set of (general) truth conditions, whereas I take there to be no such complete context-independent truth-conditional content. How might we decide between these two positions? Given that, for Geurts as for Bach, this literal content is standardly not accessible to intuition, it may be difficult to find conclusive evidence. There do, however, seem to be at least some considerations that weigh in favour of my account: as argued in chapter 2, there are many sentences for which no context-independent proposition, however minimal, is retrievable. If, then, we must accept that the linguistic meaning of these sentences cannot be identified with propositional content, then much of the impetus behind maintaining this identification for other sentences is lost: if we can't analyse linguistic meaning via propositional content for all sentences, then why would we want to do so for any sentence? Secondly, as argued above, the



descriptive propositions on which nominal-descriptivist accounts of names depend not only are not accessible to intuition but run directly counter to some fairly robust intuitions. Thus it should be the case that we can fully understand the literal propositional content of (79) despite having no idea who Dave is. It certainly seems that, in such a case, we can grasp the linguistic meaning of the utterance, since that is just to say we can grasp its context-independent meaning, and, on the strength of this, we may infer certain entailments of the proposition expressed, such as that someone left town, but there seems a strong intuition that we do not know under what conditions the utterance itself is true. Equally it seems at least counterintuitive that if I utter:

94) Noam Chomsky is a Gemini.

the meaning of my sentence determines a proposition which is true iff someone called 'Noam Chomsky' is a Gemini, regardless of whether Noam Chomsky is a Gemini. All this is, however, inconclusive: given his willingness to accept that literal propositional content is inaccessible to intuition, Geurts may maintain that the counterintuitiveness of these results is neither here nor there, being a result not of the truth-conditions of the linguistically determined DRS, but of the asserted content of the utterance. It is also, perhaps, to do Geurts an injustice. He is, after all, happy to accept that the eventual interpretation of an utterance of, for instance, (94) will be anchored to a mental representation of Noam Chomsky, and that this anchoring may take any of a variety of different forms, some of which lead to apparent rigidity. The key distinction between my account and that of Geurts may thus come down to whether one takes the propositions expressed by referentially-used name sentences actually to be rigid, or whether one takes them to be descriptive, but so linked to *de re* mental representations as to appear rigid. It seems to me that, whether the possibility of denying the significance of intuitions is available to us or not, we should be inclined to prefer an account that does not do so if one is available. What I have attempted to show in this chapter is that such an account is available. There does not, in fact, appear to be any very pressing reason why Geurts should not accept such a conclusion, other than the standard truth-conditional assumption that one type of singular expression should make one type of contribution to truth-conditional content. He could, after all, claim

that the reference marker in the DRS which constitutes the eventual interpretation of a name sentence may be anchored to an individual or not according to context.

How overall then does the account I have proposed compare to the two neo-descriptivist accounts assessed? On the one hand I am sympathetic to such accounts: given the analysis of names that I have outlined and the analysis of definite descriptions that I shall outline in chapter 5, I go along with the neo-descriptivists in seeing a high degree of similarity between the encoded meaning of proper names and that of definite descriptions. On the other hand, there seems to be overwhelming intuitive evidence, which we should respect if at all possible, that names do not standardly give rise to general truth conditions. The account I have proposed aims to respect this intuitive evidence at the same time as allowing for the possibility that names can give rise to general truth conditions given an appropriate context.

### 3.3.2 Proper names as indexicals

For the descriptivist names are univocal and make the same contribution to the truth conditions of all sentences in which they appear. Thus, taking sentence (95):

95) Louis is a trumpeter.

the descriptivist is committed to the following claim: although there are many bearers of 'Louis' and thus many statements that (95) can be used to make, this fact is semantically irrelevant. On all uses, (95) will literally express something like the proposition in (96):

96) The person called 'Louis' is a trumpeter.

It is of course the case, for the descriptivist, that (95) may be used to convey a proposition that is truth-conditionally anchored to, say, Louis Armstrong, but that is a pragmatic fact, not a semantic one.

The second position I want to examine goes along with descriptivism in holding names to be univocal, but does not hold that they make the same contribution to propositional content on all uses. This is essentially to claim that names are

context-dependent expressions, i.e. that they should be viewed on the model of indexicals (or perhaps demonstratives) rather than on the model of definite descriptions<sup>29</sup>. This claim has taken various forms. For Burge (1973), as has already been mentioned, proper names express properties; specifically 'PN' expresses the property of being a PN, where an individual is a PN iff it bears 'PN'. Burge then posits that, when appearing in argument positions, names take covert demonstrative determiners. Thus the 'Louis' in (95), for instance, will be semantically equivalent to 'that Louis'. I have already argued, following Segal (2001) and Higginbotham (1988), that if there is a covert determiner it cannot be semantically equivalent to 'that'. Should we therefore abandon a covert determiner analysis altogether? Segal (2001) offers some interesting, if merely suggestive, evidence in support of the position that we should. He points, in particular, to deficit and developmental data both of which appear to suggest that proper names and predicates are fundamentally different kinds of linguistic objects. As he himself accepts, however, his evidence is not conclusive. There is also a certain amount of syntactic evidence, marshalled by Longobardi (1994), which weighs against a determiner + predicate analysis of proper names. Overall, however, as Segal (2001) argues, the jury is still out on whether we may analyse proper names along these lines. The question I shall address below is in what respects such an analysis differs from others in the same general logical vicinity.

More recent versions of the names-as-indexicals story have been told by Recanati (1993), Pelczar and Rainsbury (1998) and Pelczar (2001) among others. I shall here focus on just one of these accounts, that of Recanati, in assessing how my own account compares with indexical accounts of names<sup>30</sup>. Recanati's analysis of proper names is quite straightforward: making use of the Kaplanian distinction between character and content, Recanati's central claim is that proper names, like indexicals but unlike definite descriptions, have different character and content. On his analysis, the character of a proper name 'PN' is the nominal description 'the bearer of 'PN'' while the content of 'PN' is simply PN; the name 'PN' is, in other words, a referring expression which is constrained by its linguistic meaning to refer to the bearer of 'PN'.

<sup>29</sup> Although this defines a clear distinction on most accounts, there are some, such as that of Geurts discussed above, on which indexicals and definite descriptions are viewed along similar lines. As I indicated above, I have a certain sympathy with Geurts' position.

<sup>30</sup> I will, however, briefly return to Pelczar and Pelczar and Rainsbury below.

Recanati's position, then, differs from a descriptivist position such as those discussed above not in the linguistic meaning it assigns to proper names, since on both views the linguistic meaning of 'PN' is simply a nominal description. Rather the difference lies in the step from linguistic meaning to propositional content: for the descriptivist the propositional contribution of a proper name is simply its linguistic meaning; for Recanati, by contrast, the descriptive meaning of a proper name drops out of the picture before the level of propositional content, with proper names contributing nothing but their referents to truth conditions. This middle position thus agrees with descriptivism on meaning but not on propositional content. The situation is entirely reversed when we come to compare the indexical account of names with the neo-Millian account which I shall come on to discuss in more detail below: whereas the indexical story agrees with the neo-Millian on propositional content, the accounts disagree on linguistic meaning. The essential distinction between Recanati's account and the neo-Millian story comes down to whether one takes the conventions linking names and their bearers to be linguistic or not. For the Millian, these conventions are linguistic and therefore the *meaning* of a name is just its bearer. For the indexicalist like Recanati, by contrast, the name-bearer convention is non-linguistic, and names therefore depend for their content on non-linguistic contextual features, i.e. on the name-bearer relations that obtain in the context of utterance.

Two questions immediately arise from Recanati's account: what evidence might there be that name-bearer conventions are non-linguistic and, if it turns out that he is right and that the conventions are non-linguistic, does this entail that names are indexical? There is some evidence that seems to suggest that name-bearer relations are non-linguistic: both Recanati (1993) and Bach (2002) argue that, whereas we *would* be inclined to say that someone who did not know what, for instance, the word 'tiger' designated was thereby a less competent speaker of English, we would *not* be inclined to level the same charge at someone who did not know who, for instance, the name 'Noam Chomsky' designated. Equally, any view on which the name-bearer relation is taken to be linguistic, i.e. on which the meaning of a name is taken to be just its bearer, will have to posit massive ambiguity for names such as 'John' and 'Jane' which are shared by large numbers of bearers. While such massive ambiguity is probably not fatal to accounts on which name-bearer relations are linguistic, it is at

best counterintuitive: 'Jane' does not, on the face of it, seem to be ambiguous in the way that, for instance, 'bank' is.

It seems plausible, then, that name-bearer relations are non-linguistic. On such a view, being the bearer of a name subserves reference just as might any other property. Thus, wanting to talk about someone in a crowd, I may, if I know it, use his name. Failing this, however, I may pick on his particularly garish green jacket as a means to facilitate reference; I may, for instance, say:

97) Green jacket looks like he's having a good time.

On an indexical view, the relation between a name and its bearer is much the same as the relation between the green jacket and its wearer. There is, of course, the obvious difference, that names are designed for the purpose of facilitating reference, whereas green jackets are not. Does the view that name-bearer relations are non-linguistic, however, entail an indexical view of names? On the account I have developed above, name-bearer relations are indeed taken to be non-linguistic; my account is not, however, an indexical account in quite the sense that Recanati's is. To see how the two accounts differ in this respect, it may be useful to examine a recent objection to indexical accounts of proper names of the sort championed by Recanati raised by Predelli (2001b).

Predelli's objection concerns our intuitions on the truth conditions of certain sentences as assessed against particular contexts. He asks us to consider a context, call it *c*, in which there is one and only one bearer of the name 'London' which is not the city that we call 'London' but rather the city that we call 'Boston'. Against such a context, on what will the truth value of (98) depend?:

98) London is in England

Note that we are not concerned with the truth value of an utterance of (98) in *c* – it is clear that an utterance of (98) in *c* would depend for its truth value on how things are with Boston in *c*. Rather what we are interested in is the truth conditions of (98), used in just the way we use it, with the reference of any indexical expressions assigned against *c*. Predelli's intuition is that against *c*, (98) says something that is true iff

London is in England, regardless of how things are with Boston; that, in other words, 'London' picks out the actual bearer of 'London' rather than the bearer of 'London' in c. Compare this to sentence (99):

99) I am in England

It seems to Predelli that, as assessed against any context  $c^*$ , (99) will be true iff the agent of (99) is in England in  $c^*$ . If this intuition is right, it seems to offer us a clear distinction between proper names such as 'London' and indexicals such as 'I': whereas 'London' seems to pick out London across contexts, regardless of whether London is the bearer of 'London' in those contexts, 'I' seems to pick out for any context the agent in that context. It thus seems that, whereas 'I' encodes a character which looks something like 'the agent of this context', 'London' does not encode the character 'the bearer of 'London''. This, in other words, appears to be good evidence against Recanati's position<sup>31</sup>.

What should we say about Predelli's claims? Firstly, they depend on intuitions not about the truth conditions of utterances, but about the truth conditions of sentences in contexts. It seems less than clear that we should have reliable intuitions on such objects. In particular Predelli requires that we should, in accessing the appropriate intuitions, hold constant meaning properties while shifting both the context against which indexical reference is resolved and the circumstance against which the sentence is assessed for truth or falsity. It would hardly be surprising if a certain amount of slippage were to occur in the attempt to access such intuitions. Nonetheless it does seem that there is some intuitive difference between the examples in (98) and (99), a difference which would appear to weigh against a straightforwardly indexical account on which 'PN' simply refers to whatever is the bearer of 'PN' in the appropriate context. It seems to me, however, that this does not force us towards Millianism on names.

What game is Predelli asking us to play? Although, for reasons that need not concern us here, he makes his object of study sentences in contexts rather than utterances, we can, I believe, translate Predelli's claims into the language of utterances. When Predelli asks us to consider (98) in c, what he is essentially asking is

---

<sup>31</sup> Of course this evidence, if right, also weighs heavily against descriptivist stories on proper names.

that we should consider what truth-conditions an utterance of (98) would have were it to be uttered in *c* by an agent speaking the language that we speak; he is asking us in other words to consider what proposition we would express by an utterance of (98) in *c*. Viewed from this perspective, it does seem intuitively right that if I utter (98) to you in *c* I thereby express a proposition not about Boston but about London, and if I utter (99) to you in *c* I thereby express a proposition about myself. This, however, is just what the account of proper names which I have outlined above predicts. In interpreting my utterance of (98) in *c*, you know, given the encoded meaning of proper names, that the mental representation I wish to explicate contains an individual concept which I associate with the information *x is called 'London'*. This is, of course, a very different thing from knowing that I am entertaining a concept of whatever is the bearer of London in the context of utterance. You should, on this analysis, accept London, rather than Boston, as the optimally relevant interpretation of my utterance of 'London' given what you know about the naming conventions which I associate with the name 'London'. Put another way, when I use the name 'London', it refers to London, regardless of the naming conventions in force in the context of utterance, because of the naming conventions which I and the rest of my linguistic community associate with the name 'London'. These conventions, although non-linguistic, get carried over to the counterfactual contexts against which Predelli asks us to assess (98) and (99). There are, by contrast, no parallel context-independent conventions governing the use of 'I' which can be carried over to counterfactual circumstances in assessing the truth conditions of (99). Whereas the name-bearer relation which links 'London' to London is available for linguistic exploitation across a wide range of different contexts, 'I' can only refer to me if I am the speaker.

It seems to me, then, that Predelli's examples do pose a threat to indexical-flavoured accounts of the type supported by Recanati (1993), Pelczar and Rainsbury (1998) and Pelczar (2001) among others. They do not, however, pose a threat to the kind of account which I favour, on which 'PN' is not constrained to refer to the actual bearer of 'PN' in the context of utterance, but refers (speaking loosely) rather to the optimally relevant individual appropriately linked to the information *x is the bearer of 'PN'*.

### 3.3.3 Neo-Millianism on names

Finally in this chapter I want to say a couple of words about the neo-Millian position on proper names, a position that has for some time been dominant in the literature. I don't want to go into any great detail, since both the position and its motivation are clear and have already been discussed at some length in this chapter. For the Millian, following Kripke (1972 etc), Kaplan (1989a etc) and Donnellan (1970), the semantic properties of a proper name are exhausted by that name's referent. In Kaplan's framework this comes out as the claim that the character of a proper name is a constant function from contexts of utterance to a content which is a constant function from circumstances of evaluation a referent. The motivation for this kind of story comes largely from the intuitions on rigidity which I discussed in section 3.1.3 above: proper names seem to pick out the same individual across possible worlds. The position is also motivated negatively by the battery of arguments which Kripke (1972) levels at descriptivism, arguments that are sufficiently well-known not to need repetition here.

As far as referential uses of proper names go, my account has much in common with certain flavours of Millianism. In particular, I agree with the Millian that a referentially-used proper name contributes nothing to truth-conditional content but its referent. I am also, so far as referential uses are concerned, broadly in agreement with the Millian, such as Wettstein (1986), who holds that facts concerning cognitive significance should be treated separately from facts concerning truth conditions. Given the anatomy for *de re* individual concepts outlined in the last chapter, on which the internal dimension of a *de re* individual concept does not determine the referent of that concept, cognitive significance, which is the domain of the internal dimension, is separate from truth-conditional contribution, the domain of the external dimension. I of course part company with the Millian on the question of how a use of a proper name is anchored to its referent. On this question, my analysis of referential uses of proper names comes closer to the indexically-flavoured account of Recanati (1993), although, as discussed in the previous section, there are significant differences between Recanati's account and my own.

In its analysis of referential uses of proper names, then, my account has a certain amount in common with analyses of a neo-Millian stripe. It is in the treatment of non-referential uses of proper names that the two accounts come apart most clearly.



For the neo-Millian, descriptive uses of names simply fall outside his area of enquiry, whereas I have taken such uses seriously. Whereas my account can handle both referential and non-referential uses with the same theoretical machinery, the neo-Millian must account for descriptive uses of names with a machinery other than that which he uses to account for referential uses, if he is to account for them at all.

## Chapter 4

### Indexicals

#### 4.0

In this chapter I want to turn away from proper names and towards another class of singular expressions: the class of indexicals. I shall, for the time being, be restricting my discussion to syntactically simple indexicals, i.e. the so-called *pure* indexicals (expressions such as 'I', 'today', 'here') and simple demonstratives ('this', 'that' etc), leaving a discussion of complex demonstratives ('this cat', 'that moose') to a later chapter<sup>1</sup>. I also want to stress once more a caveat that I raised earlier: I shall not be directly concerned with the anaphoric uses of indexicals and demonstratives (although they will be mentioned in passing). This omission is for reasons of space: as the vast literature on the subject attests, the complexities of anaphora are worthy of many theses in their own right. I make this omission, however, in the belief that a proper theory of indexicals should be able to provide a unified account of anaphoric and non-anaphoric uses, and that, suitably developed, the story I shall go on to outline below should provide the basis of just such a theory of indexicals.

Both the key questions which I want to address in this chapter and the structure within which I shall address them will be more or less familiar from the last chapter. Again I shall centre my discussion on questions concerning linguistic meaning, truth-conditional contribution, cognitive significance and conditions for understanding. I shall do so by once more laying out a body of core data, some of which will be closely related to the core data of the previous chapter. Once I have done this I shall present the account of the semantics and pragmatics of simple indexicals which I favour, and show how this account can deal with the core data.

Before launching into the meat of the chapter, however, I want, as I did for names, to say a few words about how we would want to individuate the class of indexicals. In general terms, the answer must be that the class is individuated

<sup>1</sup> My use of the term 'indexical' to cover both pure indexicals and demonstratives may strike some as less than entirely felicitous. There are two things to be said about this: firstly, so long as it is borne in mind that my intention is to use the term in this way it shouldn't prove too pernicious, and secondly, as Schiffer (1981) points out, it is a usage which has been reasonably widely adopted.

according to certain facts concerning context-dependence. Whereas it has usually been held that proper names make the same contribution to propositional content regardless of the context in which they are uttered (although I have, of course, taken issue with this claim in the preceding chapter), there seem to be other expressions which make different propositional contributions according to context of utterance. Thus, if I say 'Elvis Presley loved peanut butter' and you say 'Elvis Presley loved peanut butter', then it seems intuitively clear that we have said the same thing. If, on the other hand, I say 'I love peanut butter' and you say 'I love peanut butter' then, although we have again both uttered the same sentence, it seems we have said different things: I have said something which depends for its truth value on my predilection for peanut butter whereas you have said something that depends on your predilection for peanut butter. Equally, if you say 'he loved peanut butter' and I say 'he loved peanut butter', you may succeed in saying something about Elvis, while I may succeed in saying something about my dog Rex. At the heart of the notion of indexicality, therefore, lies context-dependence, the ability of certain expressions to give rise to different interpretations in different contexts. But is there more to indexicality than that? When we talk about indexicals, are we talking about a subset of context-dependent expressions? It seems clear that the answer to this question is 'yes', but quite how we would want to demarcate the class of indexicals from other context-dependent expressions is less clear.

For Perry (1997a, 1997b, 2001 etc.), following Reichenbach (1947), the key to indexicality is token reflexivity: an expression is indexical iff the referent of its utterance in a context is determined by certain facts holding between the referent and the utterance itself. Thus  $x$  is the referent of an utterance  $u$  of 'I' iff  $x$  is the speaker of  $u$ ;  $y$  is the referent of an utterance  $u_1$  of 'here' iff  $y$  is the place in which  $u_1$  is produced and  $z$  is the referent of an utterance  $u_n$  of 'now' iff  $z$  is the time at which  $u_n$  is produced. Even if we accept this picture, and I shall argue later in this chapter that we should certainly think twice before doing so, the token-reflexivity criterion only applies (relatively) unproblematically to pure indexicals. It is far from clear how we might extend such a picture to expressions such as 'that', 'he', 'it' and so on. In what relation must an object stand to an utterance of 'that' in order to be the referent of that utterance? As we shall see, the answer given by some is that it must be the object of an accompanying demonstration. When we start to pick at the notion of demonstration

presupposed by such accounts, however, it becomes apparent that, if it is to do the job we want of it, then to say that something is the object of an associated demonstration amounts to no more than saying that it is the object referred to.

How then should we discriminate the class of indexicals from other context-dependent expressions? As was the case for proper names, I'm not sure that there is any clear criterion which we can point to in advance of developing an account of the semantics and pragmatics of indexicals: it may, for instance, turn out that proper names are context-sensitive in much the same way as many of those expressions which have traditionally been thought of as indexicals<sup>2</sup>. We shall, for the time being, therefore, have to rely on what amounts to little more than a list: I shall take the class of indexicals to include personal pronouns, 'that', 'this', 'now', 'here', 'today', *and other similar expressions*. We must, however, bear in mind the same caveat as I raised in the last chapter: the semantic properties of this group of expressions are an empirical matter, and cannot be viewed as stipulated in advance. Looking at this from another direction, what I am interested in exploring are the meaning properties of a group of linguistic expressions, not the semantic properties of particular kinds of use to which these expressions can be put. I shall return to this point towards the end of the chapter, when I come on to discuss how the account I favour relates to other accounts of the semantics of indexicals, since it will turn out that my account may not be quite as antithetical to certain others as it might at first seem.

#### 4.1 The central data on indexicals.

What are the core data to which any account of the semantics of indexicals must address itself? There is much overlap here with the core data laid out for accounts of proper names in the last chapter: just as with proper names, our analysis must be able to account for the behaviour of co-referring and empty indexicals, and must also be able to address their apparent rigidity.

---

<sup>2</sup> As discussed in the previous chapter, the claim that proper names are best viewed as a variety of indexical is not without its supporters.

#### 4.1.1 Indexicals and co-reference

In the last chapter I discussed at some length what has come to be known as Frege's puzzle, the puzzle of how identity statements can be both true and informative. So, although 'Satchmo' and 'Louis Armstrong' are co-referential, sentence (1) seems *a priori*, whereas sentence (2) seems *a posteriori*:

- 1) Satchmo is Satchmo.
- 2) Satchmo is Louis Armstrong.

Equally, although 'Satchmo' and 'Louis Armstrong' are co-referential, a rational agent may assent to (3) while dissenting from (4):

- 3) Satchmo is a great trumpeter.
- 4) Louis Armstrong is a great trumpeter.

so long as she is unaware of the truth of (2).

Puzzles concerning co-reference and cognitive significance arise for indexicals just as for proper names, although, given the particular nature of indexicals, they raise somewhat different questions. Whereas the proper-name version of Frege's puzzle arises when we substitute one name for another co-referential name, two co-referential uses of one and the same indexical can give rise to much the same effect. To see that this is so, consider the following scenario: Bert and Ernie are standing in Bert's garden; in front of him, Ernie can see what he takes to be a particularly hideous piece of garden statuary, let's call it *g*; at the same time he can also see what he takes to be a rather lovely ornament through Bert's French windows; unknown to Ernie however, what he has taken to be the lovely ornament inside is no more than a reflection of the back of *s*. Wishing to apprise Bert of his views, Ernie says 'I like that (pointing at the reflection of *g*) but I don't like that (pointing directly at *g*)'. This immediately gives us a corollary of the second part of the puzzle above: in this context, Ernie will be disposed both to assent to and dissent from utterances of (5), associated with distinct demonstrations, although the tokens of 'that' refer to one and the same object in both utterances:

5) That is lovely.

We need take the scenario just one step further to produce a corollary of the first part of the puzzle above. Bert, realising Ernie's mistake, says to Ernie:

6) But that is that.

again pointing first at the reflection of  $\underline{s}$  in the window and then directly at  $\underline{s}$ . Here, despite the fact that Bert's identity statement involves two uses of the same expression with the same referent, it is nonetheless informative: it should, for instance, cause Ernie to revise his willingness to both assent to and dissent from utterances of (5). The puzzle then is this: on the face of it, demonstratives seem to be about as referential as any expression can be; given that they appear to encode no particular property, what else but their referent might they introduce into propositional content<sup>3</sup>? It also seems reasonable to hold that a hearer who understands an utterance must grasp the proposition expressed by that utterance. Yet what the above evidence seems to show is that these two premises are incompatible: if demonstratives introduce their referents to propositional content, and if grasping propositional content is a necessary condition for understanding, then so long as Ernie is rational and understands the two utterances of (5), he should not be able to assent to one while dissenting from the other. Equally, if he has understood Bert's utterance of (6), then what he has grasped in so doing should be uninformative to him, since, if 'that' is referential, (6) simply predicates of  $\underline{s}$  its identity with itself. It seems that one of our two premises must go.

We have seen, then, that co-referential demonstratives give rise to puzzles concerning cognitive significance in much the same way that co-referential proper names do. What of co-referential pure indexicals? This is a rather more fraught question. Corazza and Dokic (1992) attribute to Wettstein (1986) the view that pure indexicals, or at least 'here', may give rise to similar puzzles<sup>4</sup>. Thus they ask us to consider Pierre who, falsely believing that he has moved between  $t_1$  and  $t_2$ , is inclined

---

<sup>3</sup> As I shall discuss below, there are a variety of possible answers to this question.

<sup>4</sup> It seems less than entirely clear whether Wettstein does indeed hold this view.

to assent to 'here is fine' at  $t_1$  but inclined to dissent from 'here is fine' at  $t_2$ . Whether such examples can be taken as genuine instances of the puzzle we are discussing is unclear: the time lapse between  $t_1$  and  $t_2$  may, as Recanati (1993) points out, influence the cognitive difference in these cases. There are, however, much clearer examples in which pure indexicals, or at least the expression types which are standardly identified as pure indexicals, can give rise to the puzzle. Bert and Ernie are discussing Bert's recent holiday; in front of them on a table are two maps; Bert points to a spot on one map and says:

7) I stayed here.

He then turns to a point on the other map and says:

8) I didn't stay here.

Of course, unbeknownst to Bert, the two map co-ordinates he has pointed at represent the same spot. So again we have our puzzle: Bert is simultaneously disposed to assent to one utterance of (7) (accompanied by a demonstration of a particular spot on the first map) and inclined to dissent from another utterance of (7) (accompanied by a demonstration of a particular spot on the second map), although the two occurrences of 'here' refer to the same place<sup>5</sup>.

It seems, then, that 'here' can give rise to the same sort of puzzle as can 'that'. How about other pure indexicals<sup>6</sup>? Examples with 'now' are rather harder to come by, but there are some plausible candidates. Bert and Ernie are playing a musical game: they have separate recordings of all the instrumental parts of a particular orchestral work; at one point in the work there is a particularly loud cymbal crash; what they have to do is listen to one of the instrumental parts and say exactly when they think

---

<sup>5</sup> I have no doubt that some will find themselves instinctively resisting this and other subsequent examples. I am not at this point, however, making any theoretical claims about their proper analysis; I am merely attempting to establish that there are at least some uses of some expression types standardly identified as pure indexicals which can be co-referential without being cognitively equivalent. I shall come on to discuss the particulars of such examples below, where I shall be defending the position that they can be accounted for with the minimal univocal semantics that I shall outline, along with an adequate pragmatic machinery. For the time being, however, I am not intending to presuppose any particular analysis of the examples I present.

<sup>6</sup> I am here using 'pure indexical' to denote a set of expression types, not a subset of the uses of these expression types. I shall come on to discuss this distinction below.

the cymbal crash happens. Bert first listens to the piccolo part and at a particular point says:

9) The cymbal crash happens now.

He then listens to the oboe part and, just as they are coming to the same point in the work, Ernie asks Bert whether he thinks this is the point at which the cymbal crash happens. Bert replies:

10) No, the cymbal crash doesn't happen now.

Again we seem to have same expression, same referent, but different cognitive significance.

Finally, can we find the same puzzle with uses of 'I'? While such cases are inevitably bizarre, there may again be some candidates. Bert and Ernie are on a mountaineering holiday. At the top of a mountain, Bert shouts as loud as he can:

11) I am Bert.

Such is the acoustic of the mountain that thirty seconds later an echo of Bert's shout comes back to him. His voice has been so distorted by its journey, however, that Bert does not recognise this as an echo of his own utterance, and believes instead that someone else must be repeating his message. Thus, although he would be inclined to assent to his original utterance of (11) he might well be inclined to dissent from the echo of his utterance. So again we appear to have a situation in which a rational agent can both assent to and dissent from utterances of (or, indeed, one and the same utterance of) the same indexical sentence in which the indexical refers to the same object.

What we have seen is that co-referring indexicals give rise to puzzles similar to those raised by co-referring proper names. Just as with names, distinct co-referring indexicals appear capable of differing in cognitive significance. Moreover, two distinct but co-referential utterances of the *same* indexical expression also seem



capable of bearing different cognitive values<sup>7</sup>. There does, however, seem to be a disparity between the ease with which the puzzle can be set up using demonstratives and the relative difficulty of finding examples using pure indexicals. I shall come on to discuss why this might be so when I examine how my account of the semantics and pragmatics of indexicals can be used to address the core data.

#### 4.1.2 Indexicals and emptiness

In the last chapter I discussed how the Millian position on proper names appears to be threatened by uses of empty (i.e. non-referring) proper names. Thus, although 'Nat Tate' does not in fact refer, sentence (12) appears to express a complete thought:

12) Nat Tate lived and worked in New York.

Furthermore, certain sentences containing 'Nat Tate' appear capable not just of being fully meaningful but of being true:

13) Nat Tate does not exist.

Although this has sometimes been overlooked, parallel examples threaten a referential semantics for indexicals. In the last chapter I asked you to imagine Bill and Ben, who are playing a virtual reality game populated by monsters. During the game, Bill, seeing one particularly fierce monster, may utter either (14) or (15):

14) That is horrible.

---

<sup>7</sup> Looking back at proper names, we can see that they too can give rise to this version of Frege's puzzle. Salmon (1986/91) asks us to consider Elmer, a bounty hunter, who is on the trail of notorious mobster Bugsy Wabbit. Hearing that Elmer is after him, Bugsy changes his appearance so radically that he feels safe introducing himself to Elmer still using the name 'Bugsy Wabbit'. Just as Bugsy intends, Elmer fails to realise that his new friend Bugsy is the same Bugsy he is chasing. In this context, Elmer may be inclined to assent to an utterance of 'Bugsy Wabbit is a dangerous criminal' produced while, for instance, pointing at a picture of Bugsy before the change of appearance, while at the same time being inclined to dissent from an utterance of 'Bugsy Wabbit is a dangerous criminal' produced while pointing at his new friend. It thus seems that two co-referential utterances of the same proper name can have differing cognitive significance. I do not, of course, intend to offer this as evidence in favour of my position: if you remain to be convinced by my account of proper name I'm sure you will find nothing to convince you in my analysis of this example. The point I wish to make is merely that, if you are inclined to accept my account of proper names, then such cases provide the proper-name counterpart of the indexical examples discussed above.

15) That doesn't exist.

It seems that (14) is fully meaningful, just as is (12), and (15) is true, just as is (13). So once more we have the following problem: demonstratives seem paradigmatically referential; yet the meaningfulness of (14) and the truth of (15) appear to demonstrate that the meaning properties of demonstratives must go beyond their referents. If all there was to the meaning of a demonstrative was its referent then, lacking a referent, 'that' in (14) and (15) should be meaningless.

So demonstratives have meaningful empty uses just as proper names do. What of pure indexicals? Again it seems to be a considerably trickier matter setting up empty uses of pure indexicals, but once more there are some candidates. Detective Bert comes upon the body of Smith who has been foully murdered. Shocked by the killing, and horrified at the brutality of whoever has done it, Bert shouts:

16) You are insane.

at the top of his lungs. Of course, unbeknown to Bert, Smith has just met with a particularly unfortunate accident. It seems, once more, that Bert has expressed a complete thought, a thought that would be true iff whoever murdered Smith was insane. Yet, since there is no murderer of Smith, his use of 'you' is empty. 'I' too seems capable of meaningful empty uses. A poster appears on Bert and Ernie's street; on it there is a photograph of a man, beneath which is written the legend:

17) I am your Labour candidate for the forthcoming local elections.

However, the photo is, in fact, a mock-up produced on some prankster's computer; it does not depict any actual person. Nevertheless it seems plausible that the poster legend is fully meaningful in just the same way as (12), (14) or (16); what it says could, after all, be closely paraphrased by pointing at the picture while uttering 'that man is your Labour candidate for the forthcoming local elections'. If this is right, then 'I' appears capable of being meaningful even when empty.

It seems therefore that emptiness threatens a referential semantics for indexicals as it does a referential semantics for proper names: the claim that the

propositional contribution of an indexical is exhausted by its referent appears to be incompatible with the meaningfulness of sentences containing non-referring indexicals. There again, however, seems to be a disparity between the meaningfulness of empty demonstratives and of empty pure indexicals, a disparity which should be accounted for by our theory of the semantics of indexicals.

#### 4.1.3 Indexicals and rigidity

As I discussed in the previous chapter, while data on co-reference and emptiness appear to threaten a Millian analysis of the semantics of proper names, data on rigidity threaten any non-Millian analysis. The same is true of indexicals, although, given their context-sensitivity, issues surrounding rigidity are a degree more complex. In the last chapter I asked you to compare the following sentences:

- 18) Louis Armstrong performed with Ella Fitzgerald.
- 19) The greatest jazz trumpeter performed with Ella Fitzgerald.

In particular I asked you to consider upon what the truth values of these sentences would depend at a world at which Louis Armstrong was the greatest jazz trumpeter and upon what they would depend at a world in which he was not the greatest jazz trumpeter. The key intuition for Kripkean notions of rigidity is that sentences such as (18) will depend for their truth values on how things are with the same individuals across possible worlds, whereas sentences such as (19) will not; they will depend instead upon how things are with whichever individual satisfies their descriptions at a possible world. Consider, then, the following:

- 20) I performed with Ella Fitzgerald (as uttered by Louis Armstrong).
- 21) You performed with Ella Fitzgerald (uttered while addressing Louis Armstrong).
- 22) He performed with Ella Fitzgerald (uttered while pointing at Louis Armstrong).

Again we need to ask ourselves the following question: upon what will the truth values of utterances of these sentences depend in counterfactual circumstances? We need to be slightly careful what we are asking here. What we are *not* asking is: if these sentences are uttered in counterfactual circumstances (or, rather, in circumstances other than those supposed above) what will their truth values depend upon? As uttered in other circumstances their truth values could, after all, depend on more or less anything. What we *are* asking is: supposing that they are uttered in the imagined circumstances, upon what will the truth values of the propositions thus expressed depend in counterfactual circumstances? To clarify matters, it might be worth reminding ourselves of the distinction drawn by Kaplan (e.g. 1989a and 1989b) between *contexts of utterance* and *circumstances of evaluation*. A context of utterance is taken to be that against which the values of indexicals are fixed, a circumstance of evaluation is taken to be that against which the proposition expressed by an utterance is evaluated for truth or falsity. Thus, if Bert says:

23) I live in Sesame Street.

then the value of 'I', in the particular context of utterance, is Bert. Given this, Bert's utterance will express the proposition that Bert lives in Sesame Street. Now once we have this proposition, we can hold it up to any number of possible circumstances of evaluation and see whether it is true or false in each. It will of course be true in all those circumstances in which Bert lives in Sesame Street, and false in all those in which he doesn't. With this terminology in place, what we are asking is not what propositions (20) to (22) might express in different contexts, but rather, taking (20) to (22) to be uttered in the imagined contexts, upon what will their truth values depend in different circumstances.

Intuitively it seems that indexicals pattern with proper names rather than with definite descriptions; it seems, in other words, that each of the utterances of (20) to (22) will be true, in a counterfactual circumstance, iff Louis Armstrong performed with Ella Fitzgerald in that circumstance. The apparent rigidity of indexicals is, then, something which our semantic theory must be able to account for.

#### 4.1.4 Indexicals, answer machines and post-it notes

We have seen, then, that much of the data which motivated the theory of proper names laid out in chapter 3 carries over to indexicals. There is one other type of evidence which I wish to discuss before presenting the kind of story on indexicals which I favour, a type of evidence which has received much attention in the recent literature, and which we have already touched on above. As discussed in the introduction to this chapter, Perry takes the view that the key to indexicality is token-reflexivity. This analysis is very closely allied to that developed by Kaplan in his series of seminal papers (Kaplan 1978, 1989a and 1989b among others). For Kaplan, as mentioned in chapter 1, the phenomenon of indexicality forces upon us a distinction between two types of meaning, labelled by Kaplan *character* and *content*. The content of an expression is that which, in a particular context, the expression contributes to the proposition expressed; the character of an expression is that which determines the expression's content for any context. Thus for Kaplan the character of 'I' can be identified with the description 'the person who utters this token', that of 'now' with 'the time at which this token was uttered' and that of 'here' with 'the location at which this token was uttered'. Given the distinction between character and content this is not to say that 'I' has the same meaning as 'the person who utters this token' but rather that, in any particular context, 'I' will refer to whoever satisfies this description.

Recently there has been much discussion of a class of uses of indexicals which seem on the face of it to pose some serious problems for any analysis along these lines<sup>8</sup>. One such problem is that raised by answering machine messages: we may assume, as a starting point, that whoever produces an utterance *u*, and wherever and whenever she produces it, it must be the case that the person who utters *u* is at the place at which *u* is uttered at the time at which *u* is uttered. Now, given the semantics for 'I', 'here' and 'now' proposed by Kaplan, this appears to force us to the conclusion that (24) is valid, i.e. that no utterance of (24) could be false as evaluated against its context of utterance<sup>9</sup>.

---

<sup>8</sup> See, for instance, Sidelle (1991), Predelli (1998), Corazza et al (2002).

<sup>9</sup> This is not, of course, to say that any utterance of (24) is necessarily true, since it may well be false in circumstances of evaluation other than the context of utterance.

**24) I am here now.**

In order for an utterance of (24) to be false, after all, it would need to be the case that the speaker was not at the place of utterance at the time of utterance. And, if (24) cannot be false in any context of utterance, then the negation of (24), i.e. (25), cannot be true in any context:

**25) I am not here now.**

But this is where the problem lies: there seem to be cases in which (24) is truly uttered. Bert has an answering machine on which he has recorded (25) as his outgoing message. While Bert is out one day, Ernie phones him and listens to this answering machine message. Given that Bert is in fact out at the time Ernie calls, it seems that the message Ernie has listened to is true. Thus it seems that we have a true utterance of (25), contrary to the predictions forced on us by a Kaplanian semantics for pure indexicals.

While the key to these answer machine puzzles seems to lie in the interpretation of 'now', there are other much-discussed examples which hinge on the interpretation of 'I'. Bert and Ernie have spent a hard morning working in the office; Bert decides to take his lunch break, but Ernie decides to slave on; as he works, Ernie notices a series of people approaching Bert's desk and looking around for him; in order to save them trouble, he writes the message in (26) on a post-it note and sticks it on Bert's desk:

**26) I am not here now, I've gone to lunch, but I'll be back at 2.**

The first thing to notice is that we have here another example of a true tokening of (25). But this example also seems to show something else: although it is Ernie who is the producer of this tokening, i.e. although it is Ernie who has written the message, 'I' in (26) seems to refer not to Ernie but to Bert. Again this appears, on the face of it, to be beyond the reach of a Kaplanian semantics for 'I' given that, for Kaplan, 'I' is semantically constrained to refer to the agent of its tokening.

## 4.2 Indexicals and individual concepts

I have now laid out what I take to be the core data for an account of the semantics and pragmatics of indexicals. Any such account must have a story to tell about the cognitive significance of co-referring and empty indexicals<sup>10</sup>, the apparent rigidity of indexicals and also about a range of uses that do not seem to fit neatly within the currently dominant theories. Following the pattern of the previous chapter, I now want to lay out the bare bones of the account of indexicals which I favour. Once I have done that shall I come on to examine what analyses this account offers of the core data.

### 4.2.1 The bare bones

The account of indexicals I favour has much in common with the account of proper names laid out in the last chapter. There are, however, certain key respects in which the two accounts diverge. I shall start by examining the common ground.

Firstly it will be my claim that indexicals, just as proper names, are tools for explicitly communicating individual concepts. As discussed in the last chapter, this is to say that a speaker uttering a sentence containing an indexical in an argument position thereby indicates that the propositional mental representation which constitutes the content of her informative intention contains an individual concept. As with proper names (and all other individual concept communicators) the linguistic meaning of an indexical is silent as to whether the corresponding individual concept is *de re* or descriptive, although again this does not equate to the claim that establishing whether the appropriate individual concept is *de re* or descriptive is not a necessary condition on understanding. It may well be the case, and indeed I shall argue that it is the case, that, just as with proper names, a hearer will not have understood an indexical utterance unless he entertains an individual concept of the appropriate truth-conditional type. The point here is that indexicals are not semantically constrained,

---

<sup>10</sup> I want to stress that I am not at this stage intending to presuppose any particular kind of analysis of these data: it may turn out, once we have examined the facts concerning cognitive significance in rather more detail, that our semantics for indexicals need not be answerable to such facts (from a purely truth-conditional angle I have much sympathy with this view). My point is merely that, if our interest lies in an analysis of the meaning properties of natural language expressions, then we must have some story to tell about the data concerning co-reference and emptiness.

i.e. constrained by their encoded linguistic meaning, either to express *de re* concepts or to express descriptive concepts: they can do either. Once more, then, the account I favour runs foul of what appear to be robust intuitions concerning rigidity: it will be my claim that indexicals are not rigid designators by dint of their linguistic meaning, although they can be, and often are, used rigidly. I shall, therefore, need to address two questions below: firstly, what evidence is there that indexicals can be used to express descriptive individual concepts, and secondly, if they can be so used, why are intuitions on their rigidity so strong?

So far, then, my analyses of the semantics of indexicals and proper names run along much the same lines. To see where the accounts diverge we need to ask ourselves what are the specific constraints that indexicals lay, by virtue of their linguistic meaning, on the kinds of individual concepts that may constitute their interpretations on occasions of use. The first point to observe is that, on more or less any account of their semantics, members of the class of indexicals are taken to be alike with regard to some of their meaning properties, but different with regard to others; that is just to say that there are some semantic features which are common to all indexicals, although all indexicals do not have the same meaning. From Kaplan's perspective, all indexicals are alike in that their characters are not constant functions; despite having a constant linguistic meaning, they can, that is, refer to different objects in different contexts. They differ, on the other hand, in the specific characters which they encode: as we have seen 'I' is taken to have as its character a function from utterances to their utterers; 'here', a function from utterances to their locations and 'now', a function from utterances to the times of their occurrence. Pure indexicals thus seem to encode descriptive meaning, albeit descriptive meaning that, on Kaplan's analysis, drops out of the picture before the level of propositional content. And it is hardly controversial to suggest that it is not just pure indexicals which encode descriptive meaning. There in fact seems to be a continuum: at one end stand the pure indexicals, which encode, at least according to Kaplan, Perry et al., descriptive meaning sufficiently rich to determine a content for any context; at the other end stand what we might call the pure demonstratives, 'this' and 'that', which seem to encode at most a minimal descriptive meaning: if 'I' has descriptive meaning 'the utterer', then we might identify the descriptive meaning of 'that' as 'some object' plus a distance feature to distinguish 'that' from 'this'. Between these two extremes we find a range



of indexicals which, while encoding richer descriptive meaning than the pure demonstratives, still do not encode a sufficiently rich meaning to determine content as (supposedly) do the pure indexicals. The most obvious members of this category are 'he' and 'she'.

The question then arises, what role do these descriptive features play in the interpretation of indexicals? There is a range of possibilities: Bach (1987) takes the view that they contribute to propositional content, or, more specifically, to *what is said*. Thus, for Bach, what is said by an utterance of, for instance, (27):

27) He is a fireman.

is the proposition that a certain male is a fireman. Given the framework within which Bach is working, however, this does not amount to the claim that the descriptive features of indexicals appear in asserted content, and thus quite how substantive the disagreement between Bach and others really is seems unclear. The second option is that adopted by Kaplan and others of a direct reference stripe: descriptive meaning goes to determine the referent of an indexical in a context, but makes no appearance in propositional content. Thus for Kaplan, 'he' in (27) will constrain reference to a male entity, but will not appear in the proposition expressed. Finally there is a third possibility: descriptive meaning might be taken to provide only pragmatic guidance to the interpretation of indexicals. This amounts to the claim that such features serve only to facilitate interpretation, and play no semantic role whatsoever. As far as non-pure indexicals go, this position has been championed by, for instance, Larson and Segal (1995). They ask us to consider a speaker who, walking past King's College London, points towards it and says:

28) She is going to be closed over the Christmas period.

While there is clearly something odd about this utterance, they argue, we would not want to say that 'she' cannot refer to King's, and that this utterance therefore cannot express a proposition that is true iff King's is going to be closed over Christmas.

How convinced should we be by this claim, and therefore what place should the descriptive meaning of indexicals hold in our theory? I suspect that dyed-in-the-

wool Kaplanians will not find much to convince them here. What is to stop them, after all, from rolling out the familiar distinction between semantic reference and speaker reference to account for intuitions on this example? It's certainly true, they might say, that an utterance of (28) could succeed in *communicating* a proposition about King's College London, but that is very far from showing that 'she', as used in (28) has King's as its semantic reference. Given the intuitive strangeness of (28) we would do well to locate reference to King's at the level of speaker rather than semantic reference. It seems to me, however, that we should not let such a response sway us. Firstly, there are, I shall claim, good reasons why (28) might seem anomalous even if, as Larson and Segal claim, the gender feature of 'she' plays no more than a pragmatic role; and secondly, there are examples in which 'he' and 'she' can be used to refer to individuals who do not satisfy their gender features and which have none of the intuitive oddness of (28). Consider the story of Dr James Barry: Barry was a prominent British doctor of the nineteenth century, one of the first to perform a caesarean section in which both mother and baby survived. He appears to have had many eccentricities, among them that he would never allow anyone to see him undress. After his death the reason for this particular eccentricity became clear: despite having lived his entire life as a man, it transpired, as he was being laid out for burial, that Barry was a woman. Consider, then, a conversation between Bert and Ernie about James Barry: Bert knows the truth about Barry's gender, but Ernie thinks of Barry merely as an interesting (male) figure of nineteenth century medical history. Given this, Bert says:

29) You know, when he was laid out for his funeral it turned out he was really a woman.

Now there seems to me to be a robust intuition that Bert's utterance of (29) *literally* expresses a proposition that is true iff when Barry was laid out for his funeral, it turned out that Barry was really a woman. Indeed (29) seems the natural way of expressing such a proposition. Consider the alternative (30) as uttered in the imagined context:

30) You know, when she was laid out for her funeral it turned out she was really a woman.

An utterance of (30) seems to me a markedly less natural way to express the intended proposition than (29). Yet if the gender features of 'he' and 'she' did indeed play a semantic role, things should be just the other way around. Since 'he' would thus be constrained to refer to a male, we would have to say that (29) expressed no proposition at all, although it might of course succeed in communicating a proposition. Indeed we would have to say that a vast number of utterances that purported to be about Barry during his lifetime were, in fact, about no-one. This result seems at best highly counterintuitive<sup>11</sup>.

What should we conclude from this example? I think we should conclude that Larson and Segal are right, that the gender features of 'he' and 'she' are merely pragmatic. Indeed it is not just the gender features that are pragmatic; the animacy feature is also pragmatic: it might, after all, have turned out at his/her death that Barry was a remote-controlled robot, but we could still express this by saying 'he was a robot'. But if this is so, why does the example in (28) seem so odd? I think that from a relevance-theoretic perspective there is a straightforward answer to this question: since the speaker of (28) has chosen to use 'she' to refer to King's College London, and given the proposed semantics for 'she', the speaker has indicated that the property of being female will assist her hearer in his interpretation process. Yet to reach the intended interpretation, the hearer of (28) in fact has to bypass this indication. In chapter 2 I drew a distinction between two kinds of speaker intention: on the one hand the speaker has an intention concerning which interpretation her hearer should arrive at, while on the other she has an intention concerning the route via which he should arrive at it. In these terms, the oddness of (28) becomes clear: the speaker's two sets of intentions seem incompatible. On the one hand she intends her hearer to arrive at an interpretation which will be true iff King's is closed over Christmas, while on the

---

<sup>11</sup> You may be unconvinced by the example in (29) on the grounds that the imagined use of 'he' is anaphoric on an antecedent use of 'James Barry'. There are two things to be said about this: firstly, since a proper theory of indexicals should, in my view, be able to account for both anaphoric and non-anaphoric uses within a single framework, this does not render (29) uninteresting; but, more significantly, even if we do take this example as involving anaphora, it is entirely inessential to the example: we could, for instance, imagine a discourse-initial utterance of (29) produced while pointing at a picture of Barry.

other she has indicated, via her choice of expression, that she intends her hearer to use the property of being female as an aid to his interpretation strategy. It is from the mismatch of these two intentions that the intuitive strangeness of (28) arises.

From a relevance-theoretic perspective, there is also a good story to be told about why, despite the apparent tension between referring to an individual as 'he' and predicating femaleness of that individual, an utterance of (29) is not only an entirely legitimate way to express the intended proposition, but seems to be a more natural way to do so than an utterance of (30). How might Ernie's interpretation of (29) proceed on a relevance-theoretic analysis? Ernie is presented, as a constituent of the utterance he must process, with a use of the indexical 'he'. Given the semantics of 'he', he will be able to conclude two things: firstly, that Bert's mental representation contains an individual concept, and secondly, that Bert intends him to use the property of maleness in his search for an appropriate concept. Ernie will, presumably, have in his conceptual repertoire a large number of individual concepts associated with the information *x is male*, and the natural first strategy, faced with an utterance of 'he', is to narrow the search to this class of individual concepts. This first strategy will not always pay dividends: it may be that Bert is using 'he' to refer to an individual who, although salient in the mutual cognitive environment, has not yet come to Ernie's attention, as when Bert points to a picture of Barry of whom Ernie has not previously heard. In this case, trawling through pre-existing x-is-male concepts won't be of much use. Either way, Ernie is justified in restricting his search to individual concepts, be they pre-existing or formed on the hoof, which contain the information *x is male*. Why is he justified in doing this? For the following reason: if Bert intended Ernie to access an individual concept which it is mutually manifest to both does not contain this information, then to use 'he', encoding as it does the property of maleness, would be to put Ernie to gratuitous processing effort. This is just what we saw with the example in (28). Given the guarantee of optimal relevance communicated by all utterances, Ernie is thus justified in restricting his search to, essentially, what he takes to be concepts of males. One of these concepts will, *ex hypothesi*, be a concept of James Barry, since, at the stage of the discourse at which he interprets 'he', Ernie believes Barry to have been male. Moreover, given the context (previous discourse about Barry, a demonstration of Barry's picture etc.), Ernie's concept of Barry should be the most accessible x-is-male concept in his repertoire. He should therefore try out an

interpretation on which 'he' refers to Barry, to see whether it satisfies his expectations of relevance. If it does, and in the imagined context it should, he should accept it as the intended interpretation.

We can see, therefore, how, on a relevance-theoretic picture, the utterance in (29) can express a proposition which is about James Barry. But why should (29) be a more natural way of expressing this proposition than (30)? It seems at least strange that there should be a preference in this case for talking about Barry by using an expression whose descriptive meaning he does not satisfy, rather than by using an expression whose descriptive meaning he does satisfy. But in the above analysis of (29) we can see why (30) should be a less preferred option. A speaker who uses 'she' offers up the property of femaleness as a guide to interpretation. Her hearer is thus justified in trawling through concepts containing the information *x is female*. But in the imagined context, Ernie's concept of the intended referent, far from containing the information *x is female*, contains the information *x is male*. This trawl through *x-is-female* concepts will thus be a red herring, putting Ernie to gratuitous processing effort once more. An utterance of (30) should not, therefore, be optimally relevant on the intended interpretation.

There thus seems to be good reason to suppose that the descriptive components of the meanings of indexicals such as 'he' and 'she' play a purely pragmatic role in interpretation. What of pure indexicals? In order to show that the descriptive components of pure indexicals are pragmatic just as those of non-pure indexicals, we would minimally require cases in which occurrences of pure indexicals take as their interpretation something other than the value of their Kaplanian character (or something similar) in context. Smith (1989) points to a range of cases in which 'I' can be used to talk about things other than the speaker. Consider a race-goer who, having bet on a particular horse, finds that it is performing disappointingly. She may utter (31):

31) I am in last place.

in order to express the proposition that the horse she has backed is in last place. Equally a baseball manager, asked where the team he manages is based, might say:

32) I'm in Chicago.

thus expressing the proposition that his team is based in Chicago. Finally, one may utter (33) or (34):

33) I'm out of gas.

34) I'm out of ammunition.

and thereby express propositions that will be true iff your car is out of gas and your gun is out of ammunition. Smith's examples would seem, on the face of it, to suggest that being the agent of the utterance, i.e. satisfying what, with Kaplan, we may take to be the descriptive meaning of 'I', is not a necessary condition for being the referent of 'I'. This in turn might suggest that the descriptive meaning of 'I', just as the descriptive meaning of 'he' and 'she', plays a merely pragmatic role. But might there be reason to doubt this conclusion? Nunberg (1993, 1995) offers an alternative analysis of examples such as (33) and (34) in terms of what he calls *predicate transfer*. The broad idea is that the property of satisfying a particular predicate can be transferred from one object to another appropriately related object. In terms of the examples above, the property of being out of gas or out of ammunition might be seen to transfer from a car or a gun to its owner. If Nunberg is right, then 'I' in these examples may still refer, as it standardly does, to its tokener, since this tokener will have inherited the properties of being out of gas or ammunition from his car or gun. Nunberg's analysis, if correct, seems to extend straightforwardly to Smith's other examples, (31) and (32) above.

Nunberg (1990, 1993) offers some other examples in which 'I' (or 'me') supposedly receives an interpretation other than the speaker. He asks us to consider (35) as spoken by the condemned man on death row:

35) I am traditionally allowed to order whatever I like for my last meal.

(36) as spoken, by, for instance, George W. Bush:

36) The founders invested me with sole responsibility for appointing Supreme Court justices.

and (37) uttered in response to a confidence from an unknown interlocutor:

37) You should be more careful. I might have been a communist.

Nunberg's claim is that, in each of these utterances, 'I' must receive a descriptive interpretation, and thus cannot merely be interpreted as referring to the utterer. There are two things to be said about this claim: firstly, as Recanati (1993) has shown, not all of Nunberg's examples necessarily demonstrate quite what he takes them to demonstrate. For Recanati, we may analyse 'I' in (37) as referring to the speaker, just as a Kaplanian semantics demands, if we take (37) to be a statement not about metaphysically possible worlds, but rather about epistemically possible worlds. The proposition expressed by the imagined utterance of (37), therefore, will be true, on Recanati's analysis, iff there is a world which is epistemically possible for the addressee at which the speaker himself is a communist; it will be true, in other words, iff *for all the addressee knows* the speaker himself might have been a communist. Nunberg's example in (35) also seems at least questionable. Consider the following paraphrase of (35):

38) By tradition, I'm allowed to order whatever I like for my last meal.

For (38) there seems no reason to suppose that 'I' is interpreted in anything other than the standard Kaplanian way, i.e. as referring to the speaker. Given this, how can we be sure that *traditionally* in (35) is not interpreted as equivalent to *by tradition* in (38)? Nunberg himself seems to have doubts about the interpretation of *traditionally* in this example: in a discussion of (38) he, apparently inadvertently, slips from using *traditionally* to using *usually*, the semantics of which seem much clearer (see Nunberg (1993, p 32)). It's interesting to note, however, that the reading he claims for (35) becomes markedly less accessible when we replace *traditionally* with *usually*:

39) I am usually allowed to order whatever I like for my last meal.

The second thing to note about Nunberg's examples is that, even if they do give rise to descriptive interpretations, we cannot necessarily take this to show that the descriptive dimension of the encoded linguistic meanings of pure indexicals plays a merely pragmatic role, and indeed this is not what Nunberg himself takes them to show. For Nunberg, the encoded meaning of a pure indexical such as 'I' is indicative in the sense that it simply takes one to a particular object; his claim is, however, that identifying this object is not the end of the matter as far as interpretation is concerned, that interpretation may, in other words, go beyond the object which satisfies the descriptive meaning of the indexical used. Thus, even on Nunberg's analysis, the descriptive meaning of an indexical plays a semantic role, in that it determines which object in the context of utterance is to be the *index*, in Nunberg's terms, although, since it does not determine the particular relation that holds in context between index and interpretation, it will not determine the indexical's interpretation.

Might there be other examples for which the interpretation of a pure indexical does not pass through, so to speak, the object which constitutes the value in context of the relevant Kaplanian character? Smith (1989) shows that 'now' and 'here' can be used to refer to entities other than the time and place of their tokening. So, for instance, a lecturer talking about Napoleon's march on Moscow might say:

40) Winter was now closing in and, with his lines of supply cut, Napoleon had no choice but to retreat.

The token of 'now' in (40) refers not to the time of its own tokening, but to a particular time in the past. Moreover, the interpretation of 'now' does not pass through its Kaplanian referent: it is not by dint of some relation in which it stands to the time of utterance that the referent of 'now' in (40) is secured. The same lecturer, moving on to a discussion of Napoleon's arrival on St Helena, may then say:

41) Here he would die, some say at the hands of his British captors.

Again, the token of 'here' in (41) refers not to the place of its tokening, but to St Helena, and again, not through any relation in which St Helena stands to the place of



utterance. Later I shall come on to discuss what status, given the account I shall by then have defended, we should accord Kaplan's distinction between pure indexicals and demonstratives. I'd like, however, to flag up here a point that will be significant in that discussion. It is worth noticing the semantic parallels between these syntactically simple pure indexicals and certain complex demonstratives<sup>12</sup>. Compare, for instance, uses of 'here' and 'this place'. It seems that, in just the same way as we have seen with 'here', 'this place' can be used either to pick out the place of its own tokening, or to pick out another contextually salient place. Thus in (42) and (43) 'here' and 'this place' have token reflexive uses, while in (44) and (45) they both have non-token reflexive uses:

42) I have lived here for twenty years.

43) I have lived in this place for twenty years.

44) Eventually the train reaches Vladivostok. Here the line ends.

45) Eventually the train reaches Vladivostok. At this place the line ends.

Certainly (45) seems less colloquial than (44), but it is nevertheless grammatical and seems, on the face of it, to be truth-conditionally equivalent. I shall argue below that evidence such as this should lead us to reconsider Kaplan's distinction.

We have seen, then, that many pure indexicals seem capable of giving rise to interpretations that do not pass through what we may call their Kaplanian referents<sup>13</sup>. To put this another way, they can receive interpretations which neither are their Kaplanian referents, nor are secured via some relation in which they stand to those referents. What lesson should we learn from this? So far I have been suggesting that we should learn the following lesson: that, for these pure indexicals, the descriptive material which Kaplan identifies as character plays a merely pragmatic role in interpretation. But this is not the only lesson available to us here. There is a more radical conclusion we could come to: that these indexicals simply do not encode what Kaplan identifies as their character. Such a conclusion seems supported by (42)–(45):

---

<sup>12</sup> I shall be saying a great deal more about complex demonstratives in chapter 6.

<sup>13</sup> Kaplan does not intend his account to cover such uses of pure indexicals, uses of which he is fully aware. If, however, we are interested in providing a semantics for natural language expressions, then the broader the range of uses of an expression a particular account can satisfactorily account for, the more well-disposed we should feel towards that account. I shall discuss this point in greater detail below.

as these examples show, we can achieve those interpretations which are at the core of Kaplan's account, those interpretations which are token reflexive, using expressions which surely do not encode such token reflexivity, i.e. using complex demonstratives. Given this, what is to stop us identifying the encoded descriptive dimension of 'here' simply with the property of being a place, and that of 'now' with the property of being a time? Once I have outlined the account of complex demonstratives which I favour, a task which I shall undertake in chapter 6, it will become clear that, as far as I am concerned, the idea under discussion just is the idea that 'here' and 'now' are semantically equivalent to 'this place' and 'this time'<sup>14</sup>. This is not, of course, to say that they are syntactically equivalent: the former are clearly syntactically simple, whereas the latter are syntactically complex. It seems plausible that the differences in distribution demonstrated by, for instance, 'here' and 'this place' may be a result of the combination of their syntactic difference and the constraints laid on interpretation by the nature of pragmatic inference. I shall come back to this once I have laid out my account of complex demonstratives. For the time being, I shall simply say that I believe this idea to be basically right: the descriptive dimension of the encoded meaning of 'here' goes no further than the property of being a place and that of 'now' no further than the property of being a time.

This still leaves us, however, with the question of what roles these properties play in interpretation. Do they directly constrain interpretation, or is the role they play pragmatic? Smith (1989) claims that 'here' can refer to non-spatial entities, and 'now' to non-temporal entities. Thus he asks us to consider a lecturer who, reaching a natural pause, says:

46) I will stop here.

and the author of an academic paper who, at a certain point in that paper, writes:

47) Now I am going to prove the bundle theory of objects is false.

---

<sup>14</sup> I am setting to one side the fact that utterances of 'here' and 'now' can carry a range of implicit prepositions; that is to say, 'here' can, on this story, equate semantically to 'this place' (as in an utterance of, for instance, 'here looks like a good place for the picnic'), but it can also equate to 'at this place' ('I've bought my newspaper here for forty years'), 'to this place' ('come here') and no doubt others.

I'm not sure that these examples are overly convincing as evidence for the pragmatic status of the place and time features of 'here' and 'now': there may well be plausible analyses on which 'here' in (46) receives as its interpretation an individual concept of a place, and 'now' in (47) an individual concept of a time. There are, however, other examples which seem more promising. Bert has discovered the secret of time travel. Travelling back to the 1970s, he finds that everyone's wearing crocheted tank-tops and flares. He contacts Ernie and says:

48) Here in 1976, no-one has any fashion sense.

There is what seems to me to be a robust intuition that Bert's utterance of 'here' refers not to the place he is in but rather to the time he is in. Convincing examples with 'now' seem a little harder to come by, but again there are candidates. Bert is giving Ernie directions to his new house. At one point he says:

49) Now the road takes a sharp turn to the left.

Firstly it seems clear that the interpretation of the token of 'now' in (49) is entirely independent of the time of it tokening. But secondly it seems at least plausible that 'now' receives no time as an interpretation, but rather a place: Bert is not saying, on this analysis, that the road turns at a particular time (it presumably takes just this turn at all times), but rather that it turns at a particular place. We seem thus to have evidence that the place and time features encoded by 'here' and 'now' cannot be semantic; it cannot be the case, that is, that 'here' and 'now' are constrained, by dint of their linguistic meaning, to receive as their interpretations places and times respectively.

Before we return to the case of 'I', let me summarise the conclusions we have reached so far. The encoded meaning of indexicals can be analysed along two dimensions: on the one hand they encode their status as individual concept communicators, on the other they encode a more or less specific property. What are these properties? I have identified them as follows: 'he' encodes the property of being male, 'she' the property of being female, 'here' the property of being a proximal place and 'now' the property of being a proximal time. On the analysis I wish to defend,

these properties play a purely pragmatic role in interpretation: they do not constrain reference other than in tandem with our pragmatic machinery.

On this picture, as should by now be clear, the Kaplanian distinction between pure indexicals and demonstratives becomes somewhat blurred: all the indexicals so far discussed encode an element of descriptive meaning, although none encode uniquely identifying descriptions. Such a blurring seems desirable independently of the arguments I have given: in particular, as has been widely appreciated, the meaning of most supposedly-pure indexicals underdetermines reference in any context even on token-reflexive uses. Thus, on a standard direct-reference view, 'here' and 'now' are seen as constraining reference to the place of utterance and the time of utterance, but are silent as to the extent of this place or time. If Bert says:

50) I live here.

while standing by Oscar's garbage can on Sesame Street, he can have different intentions concerning the interpretation of his utterance of 'here': he may, for instance, intend 'here' to be interpreted as referring narrowly to this particular spot next to Oscar's garbage can, in which case what he has said will be false, or he may intend 'here' to refer broadly to Sesame Street, in which case what he has said will be true. Equally, discussing his favourite baseball team, Bert may say:

51) The Giants are now playing at Pacific Bell Park.

Again he may have different intentions concerning the interpretation of his utterance of 'now'; he may, for instance, intend it to refer broadly to the 2002 season, or narrowly to the very moment of his utterance.

What status does the Kaplanian distinction between pure and non-pure indexicals have on my analysis? The answer seems to be that it becomes a distinction of use rather than of meaning. Recall that, for Perry (1997a, 1997b etc) the key to indexicality is token-reflexivity: while there may be many expressions which are broadly speaking context-sensitive, what is special about indexicals is that they are sensitive to certain features of their own utterance. Now on the above analysis, as we have seen, token-reflexivity is a matter not of encoded linguistic meaning, but of use:

pure indexicals and certain complex demonstratives are both open to token-reflexive use and are also both open to non-token-reflexive use. In chapter 2 I discussed the distinction between two different types of *de re* individual concept, types that Recanati (1993) labels *encyclopedic* and *egocentric* concepts. I have as yet not asked which of these types of concept the indexicals discussed correspond to. I shall therefore now suggest an answer: these expressions do not correspond to one or other of these types of concept as a matter of meaning. What we have seen in (42) to (45) is that they may be used to express either encyclopedic or egocentric concepts. What I am claiming, therefore, is that token-reflexive uses of these indexicals correspond to egocentric concepts, while non-token-reflexive uses correspond to encyclopedic concepts. This in fact seems to me to be at the heart of the notion of token-reflexivity. Consider a speaker whose informative intention contains a mental representation a constituent of which is an egocentric concept of the place where she is, what Recanati calls her *hic* concept. In order to communicate this constituent of her mental representation she utters 'here'. As part of his interpretive task, her hearer must infer, on the basis of pragmatics and context, whether the speaker's concept is egocentric or encyclopedic. Let's assume that he infers it to be egocentric; he must then establish which place it is a concept of. How will he undertake this task? He will do so via the following inferential route: the speaker intends to communicate a propositional representation containing her *hic* concept; one's *hic* concept is an egocentric concept of the place where one is; the speaker's *hic* concept must therefore be an egocentric concept of *the place where the speaker is*; she must be talking about the place where she is<sup>15</sup>. We can see here that what Kaplan identifies as the character of 'here' can in fact be derived as an element of interpretive strategy via an inferential chain starting from the premise that the speaker wishes to communicate an egocentric concept. It thus seems that, whereas non-token-reflexive uses are out of reach for the Kaplan/Perry picture (they must be somehow hived off from token-reflexive uses), the account I have proposed not only can handle such uses, but can show why the token-reflexive uses on which Kaplan and Perry focus receive, on their accounts, the analyses they do.

Finally we come on to 'I'. What meaning does 'I' encode and what status does this meaning have? Recall that, in examples (31) to (37), repeated here as (52) to (58):

---

<sup>15</sup> I am here, as should be clear, abstracting away from the pragmatic status of 'here''s place feature.

- 52) I am in last place.
- 53) I'm in Chicago.
- 54) I'm out of gas.
- 55) I'm out of ammunition.
- 56) I am traditionally allowed to order whatever I like for my last meal.
- 57) The founders invested me with sole responsibility for appointing Supreme Court justices.
- 58) You should be more careful. I might have been a communist.

we saw that there are cases in which 'I' and 'me' appear to receive an interpretation other than the speaker. In the discussion above I pointed out, however, that, even if we take these examples at face value, the interpretations in (52) to (58) may pass through the Kaplanian referent, i.e. to pass through the actual speaker. As far as I can see, there are no entirely impure uses of 'I', i.e. uses which equate to, for instance, the use of 'now' in (40) and the use of 'here' in (41). On many accounts (Peacocke 1983, Wettstein 1984 etc), the meaning of 'I' is taken to be exhausted by its use as a tool for self-reference; all one can say about the meaning of 'I', in other words, is that it is an expression used by speakers to refer to themselves. I think there is much that is right in this sort of account: whereas it seems plausible, given the evidence adduced above, that indexicals such as 'he' and 'here' encode properties as part of their linguistic meanings, there seems to be good reason to suppose that 'I' encodes no such property. Consider mental tokenings of indexicals, for instance: if I think to myself 'she is very tall' it seems plausible that I am thinking of whichever individual my thought is about as female (although I may of course be mistaken); it seems far less plausible, however, that when I think to myself 'I would like another cup of coffee' I am thinking of myself as the agent of my own thought. This mental tokening of 'I' is surely just a direct way of thinking about myself; it directly equates, in other words, to my Ego concept. Given this picture we can derive a Kaplanian semantics for 'I' just as we did for 'here' above: a speaker who uses 'I' is thereby indicating that the mental representation which constitutes the content of her informative intention contains her Ego concept; one's Ego concept is a concept of oneself; therefore the speaker's Ego

concept is a concept of the speaker; thus in order to understand the speaker's utterance, the hearer must entertain a concept of the speaker.

This analysis sets 'I' apart from the indexicals discussed above, in that 'I' is taken to be marked as token-reflexive by dint of its linguistic meaning. Given this, we are justified, if we so wish, in identifying the linguistic meaning of 'I' with something like its Kaplanian character, since the inferential chain leading from 'I' to *speaker of this utterance* will apply on all uses of 'I', although we should bear in mind, I believe, that this meaning is derived from more basic facts about the use of 'I'. If we do forge such an identification, however, what status should we accord the descriptive content thus attributed to 'I'? Should we, in other words, take it that 'I' is semantically constrained to refer to its actual speaker or to something which stands in a particular relation to that speaker, or should we rather take the view that the property of being the tokener plays the same role in the interpretation of a tokening of 'I' as do the properties encoded by other indexicals? I believe we should take the latter view: that this derived property of being the tokener plays an essentially pragmatic role in interpretation. I shall, however, delay further discussion of this point until we come to examine the data from answering machine messages and post-it notes later in this chapter.

I have now laid out the bare bones of the account of indexicals which I wish to defend: indexicals are marked as individual-concept communicators by their linguistic meaning. Beyond this they encode properties, although these properties play only a pragmatic role in interpretation. Most indexicals are not semantically marked as token-reflexive, although 'I' is an exception. It is now time to see how this account can be put to use in tackling the data laid out in section 4. 1.

#### 4.2.2 Indexicals, individual concepts and the core data

Recall from section 4.1 that I laid out four types of data which I intend to use as a litmus test for the account of indexicals which I have proposed: data concerning co-

reference, emptiness, rigidity and what we might broadly call deferred uses<sup>16</sup>. I shall address these types of data in turn.

#### 4.2.2.1 Indexicals, individual concepts and co-reference

As I showed in 4.1, the semantic puzzles surrounding co-referential indexicals are of a slightly different sort from those surrounding co-referential proper names. Whereas the classic formulation of Frege's puzzle using proper names requires two distinct names, similar puzzles can arise with two distinct co-referential uses of one and the same indexical expression. Thus, when Bert utters (6), repeated here as (59):

59) But that is that.

pointing once directly at his garden ornament and once at the reflection of the ornament in his French windows, he has said something that is informative to Ernie. The puzzle we need to address is how this can be so, given that all the meaning properties of his two utterances of 'that' appear to be the same: they must, presumably, have the same linguistic meaning, since they are tokens of the same expression, and, *ex hypothesi*, they share the same referent.

What conclusions should we draw from the sort of puzzle raised by (59)? The most obvious conclusion seems to be that, contrary to appearances, the two tokens of 'that' in (59) do not in fact share all their meaning properties. But how can this be so, given that we have same expression-same reference? According to one response, that broadly advocated by, for instance, Kaplan (1989a) and Braun (1996), we are missing a very obvious fact about the two tokenings of 'that' in (59). Look at the example a bit more carefully and you see that the two 'that's differ in one key respect: they are associated with distinct demonstrations. If we take this feature away, then the puzzle vanishes: if Bert keeps pointing in just the same direction during both his utterances of 'that', then his utterance will be uninformative, and necessarily so, according to the Kaplan/Braun position.

---

<sup>16</sup> This is perhaps not an entirely felicitous use of terminology, since I do not intend to talk about just those uses which Nunberg (e.g. 1993) labels 'deferred'. As long as this is borne in mind, however, I do not think this terminology should cause too many difficulties.



How might this sort of story be spelled out? For Braun (1996) (as for many others) the cognitive significance of an utterance of an indexical depends not on that indexical's content in the context of utterance, but on its character. Thus the cognitive significance of an utterance of 'here', for instance, derives not from the place it refers to, but from thinking of that place under the character encoded by 'here', i.e. thinking of it as the location of this utterance of 'here'. Character, as we have already discussed, is that dimension of meaning which determines content in context. While 'I' and 'here' seem to encode such a meaning dimension<sup>17</sup>, 'that' and 'this' do not. For Kaplan (1989b), as I mentioned at the end of the previous chapter, the proper object of semantic evaluation is not an utterance but a sentence in context. To say that pure demonstratives seem not to encode a (complete) character is just to say that placing them in a Kaplanian context without any more ado fails to deliver a content. It seems, however, that utterances of pure demonstratives do, in fact, receive referents; Bert's first utterance of 'that' in (59), for instance, receives his garden ornament as its referent (as does his second 'that'). Since it is character that determines content, Braun takes this to indicate that, while a demonstrative by itself lacks a complete character, a demonstrative in a context does not. Where, then, does this complete character come from? It is here that demonstrations enter Braun's picture: on his analysis, it is the demonstration which accompanies the utterance of a demonstrative which supplies that demonstrative with a complete character. Thus, for Braun, the meaning properties of (an utterance of) a demonstrative expression must be analysed along three dimensions: linguistic meaning, i.e. that dimension which remains constant across contexts; character, which is determined by an associated demonstration, and content, which is determined in context by the combination of the demonstrative expression and its associated demonstration.

Any account along these lines is, however, going to face two major difficulties: firstly, there are uses of demonstrative expressions which, although intuitively not associated with any demonstration, nevertheless succeed in securing a referent; and secondly such an account depends on the claim that, where an utterance of a demonstrative is associated with a demonstration, it is that demonstration which

---

<sup>17</sup> Although I am of course defending a position on which this is not so, i.e. on which the encoded meaning of no indexical expression determines content in context.

determines its reference, a claim which, once we examine it in any detail, begins to look less than robust. I shall address these problems in order.

Consider two contexts in which Bert might utter the following sentence:

60) That is my favourite painting.

In the first context, Bert and Ernie are standing in a room which is entirely empty except for the Mona Lisa hanging on one wall. In the second context they are in a room full of paintings, one of which is the Mona Lisa, and Bert utters (60) while pointing at the Mona Lisa. In both contexts it seems that Bert's tokening of 'that' refers to the Mona Lisa. But why does Bert use a demonstration in the second context and not in the first? On an intuitive level the answer would seem to be that, while the first context is such that Bert's utterance of 'that' secures a reference without extra help, this is not the case in the second context: in this context, since there is more than one potential candidate for 'that' to refer to, Bert has to provide Ernie with extra clues. Thus, in the first context it seems that an utterance of 'that' has secured a reference without the aid of any associated demonstration. What might Braun say about such a use? Given that, in line with Kaplan, he has identified the character of an expression with that which determines reference, and that he would presumably accept that the utterance of 'that' in context one does secure a referent, he can say two things: either he can say that 'that', as uttered in context one, has a character without having an associated demonstration, or he can say that, contrary to pre-theoretic intuition, there is an associated demonstration in context two. If he adopts the former strategy, he is accepting that such uses of 'that' are beyond the scope of his theory. This seems a highly undesirable result: firstly, a recourse to intuition shows up no significant semantic difference between the utterances of (60) in the two imagined contexts, but perhaps more significantly, if he has no story to tell about the interpretation of (60) in context one, he has no way of showing that the particular route via which (60) is interpreted in context two is not simply an instantiation of a more general schema for the interpretation of demonstratives. If this is so, then Braun's analysis is missing a significant generalisation. Either way, we should prefer an analysis of the semantics and pragmatics of demonstratives which can account for

the interpretations of (60) in both contexts to one that accounts for one at the expense of the other. Braun would be well advised, therefore, not to adopt this first strategy.

What of the second strategy, to claim that the reference of 'that' is secured, in the first context, by an associated demonstration? This seems closer to Braun's intention: in his terminology, 'demonstration' is defined as 'the "extra" something beyond utterance that a demonstrative needs in order to secure a reference' (Braun 1996, p. 146). Braun fails, however, to tell us what this extra something is. Wettstein (1984) offers a suggestion: what determines reference is a set of contextual cues, those cues which the speaker relies on to convey his intended reference. I believe that Wettstein's account contains much that is right and indeed shares much with the analysis I shall outline below: he accepts that we must adopt a very broad notion of context if we are to account for the reference of demonstratives, and that the cues on which reference may depend are not limited by type. There is, however, a key problem with this account, a problem that goes to the heart of demonstrative reference: contextual cues, if we take them to be objective physical features of the context of utterance, will not, by themselves, do the job we want of them. It is possible for just the same utterance in just the same context, physically defined, to express different propositions. Let's imagine a third physical context in which Bert might utter (60): this time there are two paintings on the wall, one by Monet and one by Dalí, and Bert utters (60) without any accompanying demonstration. Now let's place this utterance within two different backgrounds: in the first, it is mutually highly manifest to Bert and Ernie, given the knowledge about each other which they bring to the discourse context, that Bert hates impressionism; in the second it is mutually highly manifest that Bert hates surrealism. It seems intuitively clear that, set against the first background, Bert's utterance will be true iff the Dalí is his favourite painting, whereas set against the second it will be true iff the Monet is his favourite painting. Thus the utterance may express two different propositions despite the fact that there are, *ex hypothesi*, no objective physical differences between the two contexts of utterance. How might Wettstein respond to this objection? So long as he accepts that different propositions are expressed in the example above, there seems only one way for him to go: he must accept that the cues which determine reference may include mental facts about speaker and hearer which are not objectively observable in the context. This sort of account is thus not amenable to incorporation into a theory such

as Braun's: for Braun, whatever a demonstration is, it must be an independently identifiable feature of context which, along with the demonstrative expression itself, determines a sense. Were Braun to identify his demonstrations with Wettstein's contextual cues, then a demonstration would simply be whatever determines reference; in which case Braun's theory will be reduced to the claim that whatever determines the reference of a demonstrative determines its reference.

What of the second difficulty facing a Braun-type account? If, as I have argued above, we cannot define the notion of demonstration so broadly that it takes in whatever determines demonstrative reference in context on pain of rendering our analysis in terms of demonstrations vacuous, then we need to have some independent criterion via which to identify what counts as a demonstration. On the view adopted by many (see, for instance, Berckmans (1990), Reimer (1991, 1992), Bach (1992) etc.), demonstrations are seen as intentional ostensive acts. If we take this view of demonstrations, then it seems we can hope to give some substance to claims such as Braun's: on this view, Braun's claim would be that what gives uses of demonstrative expressions complete character, i.e. what enables them to secure a referent, is an associated ostensive act<sup>18</sup>. Such a development of Braun's theory would, of course, fall foul of the objection above: uses such as the first utterance of (60) would be beyond its grasp. But there is another problem: this version of a Braun-type view depends on the claim that, where utterances of demonstrative expressions are associated with demonstrations, it is those demonstrations that are criterial to determining reference. This view has met with resistance from a number of different quarters: for Kaplan (1989b) and Bach (1992), it is speaker intention rather than demonstration that determines demonstrative reference; for McGinn (1981) it is spatio-temporal location. As Reimer (1991, 1992) shows, however, there are some apparently robust intuitions underlying the claim that demonstrations are semantically criterial. Consider a well known example used (in a slightly different form) by Kaplan (1978): the speaker, sitting at a desk, believes that on the wall behind him there is a picture of Rudolf Carnap; he wishes to say of this picture that it is a picture of one of the greatest philosophers of the twentieth century; thus, without turning round, he points to the place he believes Carnap's picture to be and says:

---

<sup>18</sup> In light of the discussion above it should be clear that I am not intending to attribute this view to Braun; my aim is to explore how Braun might escape the vacuity to which his account is condemned by his failure to provide an independent criterion of identification for demonstrations.

61) That is a picture of one of the greatest philosophers of the twentieth century.

Unbeknownst to him, however, the picture of Carnap has been replaced with one of Spiro Agnew. There seems a clear intuition that, contrary to his intentions, the speaker in this case has said something false of the picture of Agnew, rather than something true of the picture of Carnap. Given this, it seems that the speaker's ostensive act has trumped his referential intention: he intended to refer to the picture of Carnap, his ostensive act picked out the picture of Agnew, the picture of Agnew is secured as the referent of 'that'<sup>19</sup>. Reimer (1992) offers a similar example: the speaker has left her keys on her desk; next to the speaker's keys are her colleague's keys; the speaker goes to grab her own keys, while uttering:

62) These are mine.

Unfortunately, however, she misses her own keys and grabs her colleague's keys instead. Once more it seems that, despite her intention to refer to her own keys, and say truly of them that they are hers, she has in fact referred to her colleague's keys and said falsely of them that they are hers. Again, demonstration seems to trump intention.

What reason might there be to question this analysis? As Reimer herself concedes (Reimer 1992), intentions do play at least some role in determining reference. If I point vaguely in the direction of a terrace of houses, for instance, and say:

63) That is my house.

it seems plausible to hold that I have succeeded in saying something about the house I intended to refer to. Yet, given the vagueness of my demonstration, it cannot be this that is entirely responsible for securing reference; rather, it seems, it is my intention to refer to one of the things within the range of my demonstration that secures this as the

---

<sup>19</sup> McGinn's account is still in the running here. The arguments that I shall go on to outline, however, should apply to this account just as to the criterial-demonstration account.

referent of my use of 'that'. Indeed, as Kaplan (1978) points out, all demonstrations are to a lesser or greater degree vague: while pointing at the object I intend to refer to, I am inevitably pointing at a range of other things as well. Thus when I point at my house while uttering (63), I also point at, say, the front of my house, some of the bricks in the front of my house etc. Yet it is my intention to say something about my house, not about some of the bricks which go to make up its façade; and it seems that this intention wins out: sure enough, it is the house itself rather than some of its bricks to which my use of 'that' refers. This kind of observation offers us a *prima facie* reason to question the kind of analysis Reimer favours, on which demonstrations are criterial to demonstrative reference: we have shown that speaker intention plays at least some role in determining reference; on Reimer's picture, therefore, reference is determined partly by demonstration and partly by intention; given this, an account on which the whole job is done by speaker intention should be preferred to Reimer's on grounds of parsimony. If, moreover, such an account could encompass those uses of demonstratives which secure reference without the aid of an accompanying demonstration, uses which, as I have already mentioned, seem beyond the grasp of Reimer's account, then this should be taken as further reason to prefer such an account over Reimer's.

Let's take a step back and see where all this leaves us. Presented with the data on co-reference and cognitive significance, there seems reason to believe that the meaning properties of uses of indexicals are not exhausted by the combination of their linguistic meaning and their referents. In response to this, theorists such as Braun have defended a position on which the extra meaning properties that indexicals receive in context derive from associated demonstrations. We have seen, however, that such an analysis faces two problems, or maybe one problem which presents itself in two guises: on the one hand, if we draw our criteria for demonstration-hood sufficiently widely then we can hope to account for all uses of demonstratives, but only at the cost of rendering our account vacuous, since it will end up explaining reference in terms of reference; on the other hand, if we set our criteria for demonstration-hood sufficiently narrowly to render our claims substantive, then we put a large class of uses of demonstratives beyond our reach.

How might we find our way out of this dilemma? I believe the account of the semantics of indexicals in terms of individual concepts laid out above offers just such a way out. Consider once more Bert's utterance in (6), repeated here as (64):

64) But that is that.

How is it, on an individual-concept analysis, that, despite Bert's two utterances of 'that' sharing both linguistic meaning and referent, his utterance is nonetheless informative to Ernie? The key to the answer I wish to give is that it is informative for this simple reason: any two utterances of 'that' may potentially correspond to distinct individual concepts, albeit that these distinct concepts may be concepts of the same entity. How might Ernie's interpretation of (64) proceed? As a part of his interpretive process, he must interpret Bert's first utterance of 'that'. Given the encoded meaning of 'that' he will know that to interpret this token he must access an individual concept. This will essentially be all the linguistic meaning of 'that' tells him; 'that' can, after all, be used, given an appropriate context, to talk about anything of which we can have an individual concept. Ernie's next task, then, is to work out which individual concept corresponds to Bert's first utterance of 'that'. He has a certain amount of information to go on: in particular he knows that Bert is performing an intentional ostensive act. Given this ostensive act (and I shall have more to say about what role demonstrations play in interpretation below), the most accessible candidate-concept should be a concept of Bert's garden ornament. Not just any concept of the garden ornament will do however. Given the previous discourse, Bert knows something about Ernie's conceptual repertoire: he knows that Ernie has one concept of the garden ornament presented in one way and a distinct concept of the garden ornament presented in another way. His ostensive act thus has the following effect: it manipulates the accessibility hierarchy of Ernie's concepts in such a way that one of these two concepts becomes the most accessible candidate for the interpretation of the first 'that'. Much the same can be said about Ernie's interpretation of the second 'that': he must, again, search for an individual concept, and, given Bert's ostensive gesture on this occasion, the most accessible concept should be his other concept of the garden ornament, i.e. the concept which contains information like *x is inside Bert's house, x is lovely* etc. He will thus build a mental representation which contains

both of these concepts, flanking whatever the conceptual counterpart of the identity predicate may be. Given that, in the context, the two individual concepts which appear in this representation will be *de re*, Ernie's interpretation of Bert's utterance will be necessarily true. Since, however, on the conceptual picture I painted in chapter 2, the internal and external dimensions of *de re* concepts are independent in that neither determines the other, this does not have any effect on the cognitive significance of the utterance. From the mental representation which he has built as an interpretation of Bert's utterance, Ernie will be able to glean all sorts of new information: he will, for instance, be able to establish that the garden statue he dislikes is the ornament he likes, that the ornament which appears to be inside is the statue which appears to be outside and so on. Thus, as soon as we appreciate that grasping a proposition containing an indexical expression (or any other singular expression) involves accessing an individual concept which corresponds, on the occasion of use, to that expression, the puzzles raised by co-reference fall away.

What role, on this picture, do demonstrations play in interpretation? Recall that, for Kaplan (1989a) and for Reimer (1991, 1992) demonstrations are semantically significant in that they directly constrain reference. I suggested above, however, that we have good reasons for looking for an alternative account: any account on which we define demonstrations narrowly enough to render our claims substantive will be unable to account for a wide range of uses of indexicals, those uses on which an indexical secures reference without the aid of an accompanying demonstration. The approach of those, such as Braun, who wish to see demonstrations as bridging the gap between the linguistic meaning of demonstratives and whatever they refer to on an occasion of use, has been, as we have seen, to assimilate those uses on which reference is secured without a demonstration to those on which it appears to be secured via a demonstration. Given the difficulties facing such an account, however, I suggest that we should adopt precisely the opposite strategy: we should assimilate those uses with demonstrations to those uses without. What might such a strategy amount to? Let us return to the example in (60), repeated here as (65):

65) That is my favourite painting.



Recall that we considered two contexts in which (65) might be uttered, one in which it secured reference without a demonstration and one in which it required a demonstration to secure reference. How is reference secured in the first context? Intuitively there is a straightforward answer: the Mona Lisa is secured as the referent of 'that' because it is the only potential candidate in the context of utterance. Things are of course not quite as simple as this. In any real world context, there will be a multitude of candidate referents: Bert might, for instance, use 'that' to refer to the wall upon which the Mona Lisa is hanging, the frame in which the painting is mounted, or the top right hand corner of the painting itself. So what is really going on in this case? Looking at it from a relevance-theoretic perspective, we might say something along these lines: the combination of the context of utterance and the utterance itself are such, in this case, that one candidate referent for 'that' is significantly more salient than any other: Ernie is after all in a room with one and only one painting in it, trying to interpret an utterance in which Bert has indicated that the referent of his use of 'that' is a painting. In this case, therefore, Bert can rely on the context of utterance, as it is, plus the utterance itself, to get Ernie to his intended referent. This is not so, however, in the second case: here the context, as it is presented prior to Bert's utterance, is not such that Ernie can work out which object Bert is talking about. What can Bert do about this? He can manipulate the context so that it does lead Ernie to the intended referent, i.e. so that an interpretation on which the utterance of 'that' refers to the Mona Lisa will be optimally relevant to Ernie. What effect must the appropriate manipulation have? It must rearrange the saliency of the objects in the context of utterance so that the most accessible interpretation of Bert's utterance of 'that' which satisfies Ernie's expectations of relevance is the intended interpretation. It is only under these circumstances that Bert's utterance will be optimally relevant to Ernie on the intended interpretation. It is to effect this contextual manipulation that Bert produces a demonstration. As an intentional act, a demonstration has just this effect: it raises the saliency of one particular object in the context of utterance, thus making a concept of this object highly accessible to the hearer.

It is worth noting that this sort of manipulation of relative contextual saliency can be put to a number of different uses. It is not, for instance, the case that a speaker will always make a particular object more salient in order to talk about that object

itself; in some circumstances, raising the saliency of one object will so rearrange the context that, along with the utterance itself, another object becomes the optimally relevant interpretation of an indexical (or any other singular expression). Bert and Ernie are again in a room surrounded by paintings, one of which is the Mona Lisa, but, instead of uttering (65), Bert points towards the Mona Lisa and says:

66) He is my favourite painter.

Bert, it would seem, has succeeded in saying something which is true iff Leonardo is Bert's favourite painter. He has done so not by demonstrating the object he wishes to talk about, i.e. not by demonstrating Leonardo himself, but by demonstrating something which is sufficiently closely associated with Leonardo that it raises the salience of Leonardo in the context of utterance. Along with the gender and animacy features of 'he', discussed above, and the particular predication of (66), this demonstration of the Mona Lisa should have the effect of making a concept of Leonardo the optimally relevant interpretation of Bert's utterance of 'he'.

On my analysis, then, demonstrations serve to manipulate context: they make particular features of the context more salient and thus particular individual concepts more accessible. This story has much in common with that proposed by Bach (1992):

The fact that a certain dog is being demonstrated is no different in kind from the fact that a certain dog is barking raucously, is much larger than the others, or was just mentioned.

Bach (1992, p. 144)

As I have already mentioned, however, there are examples which seem to suggest that demonstrations play a greater role than this: that they determine reference. Recall examples (61) and (62), repeated here as (67) and (68):

67) That is a picture of one of the greatest philosophers of the twentieth century.

68) These are mine.

We will have to rethink the claim that demonstrations determine reference in light of examples such as (66), which show that what is demonstrated is not always what is referred to. Examples such as (67) and (68) may still, however, be taken to show that demonstrations are semantically significant in that they either determine reference or they determine an object via which reference is secured, an index in the terms used by Nunberg (1993). Bach, however, as the above passage indicates, wants to claim that demonstrations do not even have this degree of significance; for him, as for others including Kaplan (1989b), what is criterial in determining demonstrative reference is not an associated demonstration but the speaker's intention to refer. But how, in that case, can he account for cases such as (67) and (68) in which it seems that the speaker intends to refer to one thing but actually refers to another? His answer is that when we take the view that the speaker of, for instance, (67) refers to an object other than that which she intends to refer to, we are looking at the wrong speaker intention. It is certainly the case that she has an intention to refer to the picture of Carnap, an intention which goes unfulfilled, but this is not the semantically significant intention. Her semantically significant intention, what Bach calls the 'specifically referential' intention, is to refer to the picture on the wall behind her, and in this she succeeds. I believe that Bach's account is on very much the right lines. I think however he is wrong to identify the key intention as specifically referential. Given that, as illustrated by examples such as (66), demonstrations can be used as part of the fulfillment of an intention to refer to something other than the demonstratum, I think we would do better to think of the intention which Bach focuses on as a demonstrative intention.

Consider how we might analyse Reimer's central examples, (67) and (68), on the kind of picture I have drawn. The speaker of (68), for instance, produces her utterance with the intention that her hearer should come to entertain a mental representation which is true iff her keys are hers. Prior to her utterance, however, she is presented with a context in which an utterance of (68) will not be optimally relevant on the intended interpretation, since there are two potential candidates for 'these' to refer to. She thus needs to do something to manipulate the contextual salience hierarchy. With this in mind, she forms the demonstrative intention to produce an ostensive gesture which will make one of the potential candidates more salient than the other, and thus make a concept of this candidate more accessible for her hearer. She does this because, as discussed in chapter 2, she has an intention not just

concerning which propositional representation her hearer ends up entertaining but also concerning which route he should follow in order to reach this propositional representation; she has, that is, both informative and derivational intentions. In this case, therefore, she has the intention that her hearer should recognise her ostensive gesture as ostensive, that, on the strength of that recognition, he should access a concept of the keys which her ostensive act is aimed at, and that he should try this concept out for relevance as an interpretation of her utterance of 'these'. What has happened in this case, however, as in the other cases Reimer focuses on, is that the path via which the speaker intends her hearer's interpretation to proceed does not lead to the interpretation she intends him to access. That is to say that in this example she has so manipulated the context of utterance that the optimally relevant interpretation of her utterance will not be the intended interpretation. On this analysis the speaker need have no intention, as she must on Bach's account, to say something which is true iff the keys in her hand are hers; the ostensive gesture of grabbing the keys is simply a strategy for making her own keys more salient, a strategy that, as it turns out, misfires.

Let me briefly summarise, then, the solution to the puzzles posed by co-referring uses of indexical expression suggested by the account I have proposed. Given that indexicals are individual concept communicators, any use of an indexical will correspond to an individual concept. Since there is no reason to suppose that, merely because two utterances of the same indexical expression pick out the same object, they must therefore correspond to the same individual concept (we are, after all, entirely capable, as all the puzzles concerning co-reference attest, of failing to realise the identity between something presented to us in one way and the same thing presented to us in another way) there is thus no reason to suppose that two co-referring utterances of the same indexical expression need have the same cognitive significance. Given this, we need posit no extra level of meaning which is constant as between speaker and hearer: it is the *fact* of there being discrete concepts, rather than any particular feature of their content, which lies at the heart of this puzzle.

#### 4.2.2.2 Indexicals, individual concepts and emptiness

Having said a great deal about co-referring uses of indexicals, there will not be a great deal to say about empty uses of indexicals. Recall that we're thinking of examples such as (14) and (15), repeated here as (69) and (70):

69) That is horrible.

70) That doesn't exist.

uttered by Bert to talk about the slaving monster that appears to be charging towards him and Ernie as part of the virtual reality game they are playing. I propose to say very little more about such examples than that I believe the analysis of empty uses of proper names outlined in the previous chapter can naturally be extended to account for empty uses of indexicals. This is to say the following: utterances such as (69) are meaningful, in the sense that they are every bit as cognitively significant as utterances containing denoting indexicals, for the reason that, as the above discussion should have made clear, the cognitive significance of indexicals is a function of the internal dimensions of individual concepts. Given this, and given our cognitive ability to gather together information into individual concepts in the absence of any *actual* individuals for those concepts to pick out, the meaningfulness of (69) will be unaffected by the failure of its token of 'that' to refer. The proposed analysis in terms of individual concepts thus offers a straightforward account of how utterances containing uses of empty indexicals can be cognitively significant.

It seems, however, that (70) (and maybe (69)) is not only cognitively significant, but is also true. In the last chapter I proposed an analysis of empty uses of proper names in terms of intentions to talk about entities in and makes statements about information sources other than the actual world. Again I believe this analysis carries over to empty uses of indexicals. Thus, if we are inclined to say that (69) is *in some sense* true, that will be, on my account, because (69) is taken to be a statement not about how things are in the actual world, but rather about how things are in the non-actual world of the virtual reality game, and the game constitutes an information source which verifies (69). (70) will be true, by contrast, not because it is taken to be a statement about a non-actual world, but because it is, in essence, a trans-world

statement: (70) will be true iff the slaving monster of the virtual reality game has no existence in the actual world; iff, in other words, the x-is-a-monster concept deriving from the game information source picks out nothing in the actual world.

Such an analysis can also be applied to examples (16) and (17), repeated here as (71) and (72), in which we find empty uses of Kaplanian pure indexicals:

71) You are insane.

72) I am your Labour candidate for the forthcoming local elections.

Recall that (71) is detective Bert's response to finding the body of Smith, who he wrongly believes to have been the victim of a brutal murder, and (72) is the legend found beneath the mock-up photograph of a non-existent candidate. Again it seems that the cognitive significance of these utterances is unaffected by the emptiness of their indexicals. And again this is down to the relation between the cognitive significance of an utterance of an indexical and the internal dimension of the individual concept which constitutes its interpretation on that utterance. Thus we can see that, for both demonstratives and pure indexicals, cognitive significance, determined as it is by internal rather than external aspects of individual concepts, is unaffected by emptiness.

#### 4.2.2.3 Indexicals, individual concepts and rigidity

In section 4.1.3 I showed that indexicals appear to pattern with proper names in terms of rigidity. That is to say, there is a robust intuition that an utterance with an indexical in subject position will be true in any counterfactual circumstance depending on how things are in that circumstance with the *actual* referent of the indexical. In that section we considered examples (20) to (22), repeated here as (73) to (75):

73) I performed with Ella Fitzgerald (as uttered by Louis Armstrong).

74) You performed with Ella Fitzgerald (uttered while addressing Louis Armstrong).

75) He performed with Ella Fitzgerald (uttered while pointing at Louis Armstrong).

What we found was that, from an intuitive standpoint, it seems as if each of (73) to (75) will be true against a circumstance of evaluation iff Louis Armstrong performed with Ella Fitzgerald in that circumstance. Now it seems on the face of it as if the account of indexicals I have outlined above may be unable to accommodate this intuition. On my account, after all, indexicals may be used referentially, i.e. they may be used to express *de re* individual concepts, but they are not marked so to do by their linguistic meaning: they may also be used to express descriptive concepts. In order to defend my position against this apparent objection I want to show two things: firstly, I want to show that indexicals genuinely are capable of giving rise to descriptive interpretations and secondly I want to show that, even if I am right, there are good reasons why our intuitions concerning the rigidity of indexicals are so robust.

What evidence is there, then, that indexicals can give rise to descriptive interpretations? Nunberg (1993) discusses a wide range of uses in which this appears to be the case, uses which involve what he calls *deferral*. I have already introduced some of Nunberg's examples, examples such as (35) to (37), repeated here as (76) to (78):

76) I am traditionally allowed to order whatever I like for my last meal.

77) The founders invested me with sole responsibility for appointing Supreme Court justices.

78) You should be more careful. I might have been a communist.

As discussed above, however, I am not sure we should read too much into these examples: there seem to me to be analyses for each on which they may be viewed as referring to the speaker. Nunberg does, however, offer some other examples which seem more promising. Consider (79), which appeared in the University of Arizona newspaper the Friday before classes began:

79) Tomorrow is always the biggest party night of the year.

It seems that 'tomorrow' in (79) is not interpreted as referring to the day after the tokening of (79), as it should on a Kaplanian semantics, but is rather interpreted as

equivalent to 'the Saturday before classes begin'. There are two things to be said about this example however: firstly, I am not sure that one couldn't run an analysis on which tomorrow is taken to be acting as a bound variable, in which case the interpretation will not be descriptive as Nunberg claims; but secondly, even if we accept that Nunberg is right and that the interpretation of 'tomorrow' genuinely is descriptive, this does not necessarily undermine the claim that indexicals are semantically rigid designators. We have come across examples such as this before, on which interpretation passes via an indexical's Kaplanian referent to something which is appropriately related to that referent. Recanati (1993) runs a defence of the direct reference picture against such uses along the following lines: what Nunberg's examples, such as (79) show is that the propositions expressed by indexical sentences are not always singular. They do not, however, threaten the claim which is at the heart of direct reference, that there is a fundamental semantic difference between descriptions and directly referential expressions such as indexicals. To maintain this claim we need to distinguish between two distinct levels of content, the proposition ultimately expressed, and what Recanati terms the *basic level of interpretation*. The direct referentialist's claim then becomes that at the basic level indexicals and other directly referential expressions are semantically marked to give rise to singular interpretations. The fact that they can be interpreted descriptively at the level of ultimate propositional content presents no problems for this claim. Nunberg attempts to undermine this argument on the grounds that, were the interpretation of indexical sentences to have to pass through a fully propositional basic level, that basic level would end up in many cases being incoherent. Thus the basic level interpretation of (79) would be true iff September 3<sup>rd</sup> 1987, for instance, were always the biggest party night of the year. Given that it only comes around once, Nunberg would claim, September 3<sup>rd</sup> 1987 can't *always* be anything. While this objection may stand up against the letter of Recanati's analysis, it does not, as far as I can see, have much impact on the spirit of that analysis. It seems clear that in cases such as (79), interpretation of 'tomorrow' does proceed via its Kaplanian referent, i.e. via September 3<sup>rd</sup> 1987. It is because it is instantiated by this referent that the property of being the Saturday before classes begin ends up as the interpretation of 'tomorrow' in (79). So long as Recanati does not require that a global basic level interpretation be entertained as part of the interpretation process, but rather restricts himself to the



claim that the local interpretation of 'tomorrow' passes through a basic level, Nunberg's objection would seem to have no force.

My interest, however, is in showing that Recanati's fundamental claim, that there is a semantic asymmetry between one class of singular expressions which are marked as descriptive and another class which are marked as referential, is wrong. To do so, I will need to show that there are cases in which indexicals can give rise to descriptive interpretations without passing through any basic level object along the interpretive path. In the case of demonstratives there seem to be some plausible candidates. Schiffer (1981) asks us to consider Ralph, who, coming across what looks like a big footprint in the sand, says:

80) He must be a giant.

Since Ralph has no idea whose footprint it is, his utterance cannot correspond to a *de re* individual concept of his own. Intuitively, moreover, Ralph's addressee can fully grasp the meaning of Ralph's utterance without herself knowing who produced the footprint. It seems therefore that Ralph's utterance will be true iff whoever produced the footprint must be a giant; it seems, in other words, as if the interpretation of Ralph's utterance of 'he' will be descriptive. There is also, in this case, no temptation to believe that the interpretation of 'he' passes through any intermediate referential stage: there is, after all, no candidate for this role unless one wishes to make the highly dubious claim that the interpretation of 'he' passes through the footprint.

To see another example of the descriptive use of demonstratives, we can return to the case which was central to my claims about the descriptive use of proper names. Bert and Ernie are at Scotland Yard's Black Museum when they come across a painting of a cloaked figure prowling the streets of London, a painting which is supposed to be of Jack the Ripper. Bert points at the picture and says:

81) He was a Liverpool businessman.

Here it seems that 'he' will receive much the same interpretation as would 'Jack the Ripper' in the same context. If this is the case, and given the arguments I deployed to show that 'Jack the Ripper' can be used descriptively in the last chapter, we should be

inclined to accept that 'he' in (81) receives a descriptive interpretation. Assume that Bert is right, that Jack the Ripper was indeed Liverpool businessman James Maybrick. On a standard direct-reference picture, 'he' should therefore receive (a non-truth conditional mode of presentation of) Maybrick as its interpretation, albeit via a rigidified description that Maybrick satisfies. It should thus be the case that, as evaluated against any counterfactual circumstance, the proposition expressed by (81) will be true iff Maybrick is a Liverpool businessman in that circumstance. But this does not seem to tally with intuition. Imagine a counterfactual circumstance in which Maybrick was indeed a Liverpool businessman, but in which he murdered no-one; now imagine that the same murders which Maybrick committed in the actual circumstance, or so we are assuming, were committed in this circumstance by London painter Walter Sickert. There is a robust intuition that in this circumstance Sickert is Jack the Ripper. Furthermore, it seems reasonably clear that, as evaluated against this circumstance, Bert's utterance in (81) will be false, since in that circumstance Jack the Ripper is not a Liverpool businessman but a London painter. If this is right, then it forces us to the conclusion that 'he' in (81) receives a descriptive interpretation; its interpretation, that is, corresponds to an individual concept the external dimension of which is descriptive. And again there is surely no temptation to think that interpretation passes via a particular individual as Recanati would have it: Bert's demonstration of the picture does not make Maybrick salient, after all, it makes Jack the Ripper salient; and, although Maybrick is Jack the Ripper in the actual world, there are countless possible worlds in which he is not.

Thus there seems good reason to believe that demonstratives can receive descriptive interpretations. What of pure indexicals? Given the analysis presented above, the interesting cases are going to be 'I' and possibly 'you'. Cases on which these indexicals are used descriptively are inevitably going to be hard to come by. The semantic properties of 'I' derive, I claimed above, from its use as a tool of self-reference. To find a descriptive use, it seems, we will thus need to find a case in which the speaker is thinking of herself descriptively, hardly a common state of affairs. I claimed above, however, that the property of being the speaker plays only a pragmatic role in the interpretation of 'I'. We might thus do better, in the hunt for descriptive uses of 'I', to look for a case in which 'I' does not pick out the actual agent of the utterance. These cases will be discussed at greater length in the next

section. I shall, however, present one example here. Back in Victorian London, the painting of Jack the Ripper which Bert and Ernie are to see in the Black Museum over a century later has become the best-known image of the Ripper, adorning the cover of every newspaper. Oscar, who has been following the events in Whitechapel, has become convinced that James Maybrick is responsible. He does not, however, want to get involved. He thus tears the front cover, bearing the picture of the Ripper, from his daily paper, scrawls the words:

82) I am James Maybrick.

across the picture, and, in the dead of night, attaches it to the gates of Scotland Yard. What proposition is expressed by Oscar's message? It seems to be a proposition which is true iff Jack the Ripper is James Maybrick. Much the same analysis can be given, then, for the use of 'I' in (82) as for the use of 'he' in (81): the proposition expressed by (82) will be true in all circumstances in which Jack is Maybrick and false in all circumstances in which he is not. And once more, there is no temptation to claim that interpretation passes through the Kaplanian referent: we would presumably want to say that, coming to work the next morning, the officers of the Yard will understand the scrawled message, yet they will have no idea who has left it there. As mentioned above, a more detailed analysis of such examples must wait until the next section. Utterances such as (82) do, however, seem to suggest that even the purest of pure indexicals can receive descriptive interpretations, even at Recanati's basic level of interpretation. This result is to be expected given the analysis of such expressions outlined above. What I have suggested is that the encoded meaning of 'I' is token reflexive in a way that the encoded meanings of (most) other indexicals are not. On the face of it this analysis is closely related to Kaplan's semantics for 'I' which can be viewed as claiming that 'I' encodes the property of being the speaker of this token. If Kaplan were right, however, and the linguistic meaning of 'I' did so constrain the interpretation of 'I' that tokens could only pick out their own agents, we should not expect to find descriptive uses, even if 'I' were not semantically marked as referential: as mentioned above, it is hard to imagine cases in which speakers think of themselves descriptively and thus on which understanding a token of 'I' would require entertaining a descriptive individual concept. It is because this token-reflexive

property, the property of being the speaker, plays a pragmatic role in interpretation that we do come across cases such as (82).

I hope that I have met the first of my two aims: to demonstrate that there are descriptive uses of indexicals. What of the second aim, to explain why, if such uses are available, intuitions concerning the rigidity of indexicals seem so robust? This is not of course a semantic issue, and is thus not essential to the current discussion: so long as I have convinced you that there are genuinely descriptive uses of both demonstratives and pure indexicals, it is in some sense neither here nor there why speaker intuitions on this matter are as they are. The question seems, nonetheless, to throw interesting light on facts about our use of indexicals and thus I want to pursue it briefly here. The first point to note is that, even though there are, as I have shown above, genuinely descriptive uses of indexicals, they are heavily outnumbered by referential uses. Why should this be? I discussed similar issues concerning proper names in the last chapter. The points I made then, however, do not straightforwardly carry over to indexicals. What I suggested was that the heavy preponderance of referential uses of proper names may have something to do with the nature of naming practices, essentially that we tend to name things demonstratively rather than descriptively. It is essential to indexicals, however, that the relation between an expression and its interpretation on an occasion of use is not a once-and-for-all matter as it is with proper names: whereas a specific convention links me to 'George' and thus allows uses of 'George' to refer to me (roughly speaking), there is no convention linking me to 'he' which allows uses of 'he' to refer to me. Since, therefore, the link between indexicals and their referents is so context-specific, facts about naming practices are not going to provide us with an explanation for the predominance of referential uses of indexicals. Underlying my account of the bias towards referential uses of names, however, were some foundational claims about the nature of cognition which may be carried over to the case of indexicals. In the last chapter I suggested that we are cognitively designed primarily to track physical objects rather than properties. There are, however, exceptions to this general rule, where a property is significant enough in its own right that we have an interest in thinking about whoever bears that property. Thus we may have an interest in talking about whoever committed the Whitechapel murders, whoever left a particular footprint in the sand and so on. It seems that this underlying fact about human cognition, the bias towards

discriminating individuals on a physical basis rather than according to their properties, may offer us an explanation for the predominance of referential uses of indexicals. Why, then, are intuitions on the rigidity of indexicals so strong? Because, given the facts about human cognition, most uses of indexicals are referential. This does not entail, however, that all uses of indexicals are referential. If the story I have told is along more or less the right lines, then we can account for the preference for referential uses of indexicals without having to locate any specifically referential features in the linguistic meaning of indexicals. Putting the same point another way, we should expect to see a bias towards referential uses of indexicals even if they are not semantically referential, i.e. even if they are not marked as referential by their encoded meaning.

#### 4.2.2.4 Indexicals, individual concepts, answer machines and post-it notes

In the last three sections I have shown how the account I favour can handle data concerning the cognitive significance of indexicals, as well as data concerning their apparent rigidity. I want now to come on to the last type of data discussed above. Kaplan (1989a) draws attention to a particular type of use of indexicals:

There are certain uses of pure indexicals that might be called 'messages recorded for later broadcast', which exhibit a special uncertainty as to the referent of 'here' and 'now'.

Kaplan (1989a, fn.12)

Kaplan makes it clear that he does not believe such uses pose any threat to his semantic system. There is however, as discussed in section 4.1.4, a *prima facie* problem. Consider again sentence (25), repeated here as (83):

83) I am not here now.

It seems that Bert can record (83) as his outgoing answer-machine message, and thereby use 'now' to refer to a time other than that at which he produces his utterance of (83), i.e. to refer to the time at which the message is played back, rather than the

time at which it is recorded. There has been a series of attempts in the literature to reconcile this kind of data with a Kaplanian semantics for pure indexicals. Sidelle (1991) draws a distinction between the context of encoding and the context of decoding. At each of these times, for Sidelle, there will be a distinct utterance with a distinct set of parameters. When Bert records his answer-machine message he is, on Sidelle's analysis, setting things up in such a way that he can produce an utterance at a later time. The idea is, then, that utterances can, in certain circumstances, be deferred. If this is right, then we can still hold that 'now' refers to the time of utterance, since there is an utterance at the time at which the message is played back and it is to the time of this utterance that 'now' refers in (83). Predelli (1998), however, raises an objection to Sidelle's account. He asks us to consider Jones who, expecting his wife to return home at 6pm and expecting not to be at home himself at that time, writes a note at 4pm, just before walking out the door, which says 'I am not here now'. As things turn out, however, Mrs Jones does not get home and therefore does not read the message until 10pm. On Sidelle's account there are two possible time parameters which could determine the referent of Jones' use of 'now': the time of encoding, i.e. 4pm, or the time of decoding, i.e. 10pm. However, neither of these seems, to Predelli, to give the right result: regardless of when Mrs Jones actually returned home, Jones intended the message to be read at 6pm and it is thus 6pm to which his use of 'now' refers. If Predelli is right, then Sidelle cannot be. I'm not sure, however, that Predelli's claims concerning his example are entirely convincing: imagine that Jones has returned home before his wife, say at 8pm, but has forgotten to retrieve his message from the prominent position where he left it; Mrs Jones then returns home at 10, reads the message and is thus surprised to find her husband already in bed; it seems at least plausible that, given these circumstances, she would be right to think that the message was not true. If this is so, however, Predelli's example does not stand as a counterexample to Sidelle's claims. Corazza et al. (2002) offer a more convincing counterexample however. They ask us to consider Ben and Joe who work together in an academic department. One day Joe is off sick and Ben notices a number of students approaching the door of his office and looking bemused. To help them out he writes:

84) I am not here today.

on a post-it note which he sticks to the door of Joe's office. It seems that the message will be true iff Joe is not here today, yet, on Sidelle's analysis, it is hard to see how this can be so. Sidelle would, presumably, have to say that Joe was the agent of the utterance which takes place when his students read the message, yet Joe is not even aware that the message exists. An account which requires that someone can be the agent of an utterance he has no part in producing and is not aware of seems questionable. It seems then that, whether or not the example he uses is a helpful one, Predelli is right to challenge Sidelle's position. Having done so, he offers an alternative analysis of such examples. What we need to distinguish, for Predelli, is not the context of encoding and the context of decoding, but the actual context of utterance and the intended context of utterance. So what is going on in examples such as (83) and (84) is that the context against which the speaker intends his utterance to be interpreted differs from the actual context in respect of the values assigned to some of the indexical parameters. Thus when Bert leaves (83) as his answer message he intends it to be interpreted against a context in which the time parameter is not that of the context of utterance, but that of the context in which it is played back; and when Ben leaves the message in (84) on Joe's door, he intends it to be interpreted against a context in which the value of the agent parameter is not Ben himself, as it is in the actual context, but is Joe. There are, however, two key problems with Predelli's account as it stands. Firstly, as Corazza et al. point out, it massively overgenerates: since Predelli has nothing to say on the question of what constraints there might be on which contexts a speaker may legitimately intend her utterance to be interpreted against, there is no clear constraint on what she may use an indexical to refer to. Consider (83) again: on Predelli's story there is nothing to stop Bert leaving (83) on the answering machine at his home in Sesame Street and thereby saying something that is true iff he is not in Las Vegas; all that is required for him to do so is that he should intend his utterance to be interpreted against a context in which the value of the location parameter is Las Vegas. Yet this seems highly counterintuitive. The second problem with Predelli's account is related to the first: since he provides no independent criterion for what can count as an intended context, his attempt to save a Kaplanian semantics renders that semantics essentially vacuous. Predelli's strategy for saving the Kaplanian picture is to say that 'I', for instance, does indeed refer to the

agent but that the context against which the values of indexicals are fixed may not be the actual context of utterance. But how are we to determine which context an utterance is to be interpreted against? Essentially the only answer that Predelli gives is that we should interpret an utterance against that context for which the location parameter is wherever the speaker intends to refer to as 'here', the time parameter is whenever the speaker intends to refer to as 'now' and the agent parameter is whoever the agent intends to refer to as 'I'. But then what does his analysis amount to? No more than the claim that we can save the Kaplanian semantics for 'here', for instance, by the simple move of saying that wherever a use of 'here' refers to is the place of utterance (in the intended context). This looks to be a defence of Kaplan by sleight of pen.

In response to the problems faced by Predelli's account, Corazza et al. offer an alternative based not on intentional contexts but on conventional contextual parameters. Their idea is that, for any potential context, there are conventions governing which entities serve as values of the contextual parameters:

For any use of the personal indexical, the contextual parameter of the agent is *conventionally* given – given by the *social or conventional setting* in which the utterance takes place. For instance, with 'now', the setting or context in which it is used changes the time that the term refers to: if 'now' is heard on an answering machine, we take the relevant time to be the time at which it is heard ... In contrast, if we read 'now' on a postcard ... we take the relevant time to be the time at which the words were written.

Corazza et al. (2002, p. 11)

Thus we can again salvage a Kaplanian semantics, but this time by invoking conventions rather than intentions: 'now' in Bert's answer machine message refers to the time of playback because convention determines that 'now' in an answer-machine message refers to the time of playback. Again there are two key problems with this analysis. Firstly, the sort of individuation of contexts upon which this account relies will not do the work required of it. Imagine that you phone a friend and hear the following answer machine message:



85) Hi, I'm not here now (pause with knocking sound in the background) OK, someone's at the door now, I've got to go.

Here it seems that the two tokens of 'now' will refer to different times: the first token will refer to the time of playback, the second to the time of utterance. If, however, as Corazza et al. would have it, the convention governing what is taken to be the time parameter in context is determined by the social or conventional setting, then the two utterances of 'now' should not be able to differ in interpretation in this way. I suspect this is not a knock-down argument against the position of Corazza et al. More worrying for them is the interpretability of utterances of indexicals in contexts in which there is no pre-existing convention. Bert and Ernie are examining Bert's collection of fossils; Bert points to one of the fossils and says:

86) There were so many fossils here that we couldn't fit them all in the car.

Now it seems plausible that Bert has said something that is true iff there were so many fossils at the place where he found that fossil that he couldn't fit them all in the car. However there is surely no temptation to think that there is a convention to the effect that in contexts in which someone is pointing at a fossil, the location parameter is the place where the fossil was found.

What might we learn from this objection to the account of Corazza et al.? The lesson I believe we should learn is that we can achieve with pragmatics what they want to do with conventions. The only sacrifice we have to make, and I imagine that all those discussed above might regard it as quite a large sacrifice, is to accept that the role of the encoded meaning of indexicals is pragmatic. How might this help us? The key to the idea I want to explore is that 'I', for instance, is not semantically constrained to refer to the utterer but is pragmatically constrained to pick out something which it is *relevant* to think of as the utterer. Consider Ben's note in (84). In the scenario envisaged, Ben has done two things: firstly he has written 'I am not in today' on a post-it note and secondly he has attached that post-it note to Joe's door. Coming across the note attached to Joe's door, how is one of Joe's students, let's call him Elmo, going to interpret it? As part of his interpretation process, Elmo must ask

himself which concept he is intended to entertain as his interpretation of 'I'; since this token is, presumably, intended to be referential, this comes down to the question of who 'I' refers to. Given that 'I', by dint of its linguistic meaning, raises the contextual salience of the property of being the agent, and that contexts in which it is relevant to think of someone other than the actual agent as the agent will be rare, Elmo is justified in starting his search for a referent with the actual agent. Now Elmo may or may not know who the actual agent is, but, whoever he is, Elmo knows that he has attached the message to Joe's door. The effect of this, as an intentional element of the act of communication, is to raise the salience of Joe in the context. Now the agent may or may not be Joe himself. If Joe is the agent, and he wishes 'I' to refer to himself, then raising his own salience and thus raising the accessibility of a concept of himself for his addressee, will increase the relevance of his utterance by minimising the processing effort required to interpret it. If, on the other hand, the agent is not Joe and he wished 'I' to refer to himself, then the reverse is the case: by raising Joe to salience he will have reduced the relevance of his own utterance by putting his addressee to gratuitous processing effort. It is thus the case that an agent other than Joe who wished the token of 'I' in (84) to refer to himself would be reneging on the presumption of optimal relevance communicated by his own utterance by placing the message on Joe's door. Given that, as part of his interpretation strategy, Elmo is justified in trusting that the message will be optimally relevant on the intended interpretation, he can therefore exclude the possibility that an agent other than Joe could intend 'I' in (84) to refer to himself. In this case, i.e. if the agent is not Joe, then who might that agent intend to refer to with his use of 'I'? In other words, and given the encoded meaning of 'I', who might it be relevant to think of as the agent of this message? The answer would seem to be that, since the message has been left on his door, it might be relevant to think of Joe as the agent. Elmo should thus try out a concept of Joe for relevance and, if it satisfies his expectation of relevance as it should, he must accept it. Although I do not propose to go through them one by one, I believe that the answering machine examples are susceptible to much the same sort of pragmatic analysis.

What might it mean to say, as I have above, that it may be relevant to think of someone other than the agent as the agent? I suspect that it comes down to something like asking: who might the actual agent be pretending is the agent for communicative

purposes? If this is so, then my account looks like it has much in common with Predelli's: we could, after all, look at Predelli's intended non-actual contexts as pretend contexts. However I believe that my account faces neither of the problems to trouble Predelli's: firstly it does not overgenerate, since, as illustrated by the example above, the ability of an agent to use 'I' to refer to an individual other than himself is tightly constrained by context and pragmatic inference; and secondly it does not salvage a Kaplanian account at the cost of making that account vacuous, since it does not attempt to salvage a Kaplanian account at all: on my story, indexicals are not semantically constrained by their linguistic meaning.

### 4.3 Conclusion

I have now laid out the account of the semantics and pragmatics of indexicals which I favour, an account on which indexicals are marked as individual concept communicators by their linguistic meaning, and on which the elements of descriptive meaning which they encode play a pragmatic role in interpretation. I have also shown how this account handles what I have taken to be the core data on indexicals. In conclusion I want to say a couple of words about how this account compares with the other main positions currently on the market. In particular I want to stress a couple of observations made in the course of this chapter concerning comparisons between my account and the mainstream direct-reference flavoured accounts of Kaplan, Perry etc.

What is the direct-reference story on indexicals, as originally formulated by Kaplan (1989a, 1989b in particular), a story of? Kaplan is essentially concerned with logical properties; he is in other words concerned with truth. Given this, there is good reason to think that the account he proposes should not be answerable to data concerning cognitive significance<sup>20</sup>. Moreover he is not, apparently, concerned with the logical properties of linguistic expressions, but rather with the logical properties of certain uses of those expressions:

What is common to the words or usages in which I am interested is that the referent is dependent on the context of use and that the meaning of

---

<sup>20</sup> For a particularly influential defence of the claim that it should not, see Wettstein (1986).

the word provides a rule which determines the referent in terms of certain aspects of the context.

Kaplan (1989a)

As Kaplan himself accepts, this narrows his area of interest to classes of use rather than classes of expression. Given these facts, Kaplan's position and the position I have defended above may not be quite as antithetical as they might at first seem. In particular, given the limits of his interest, Kaplan need not deny that two co-referring uses of an indexical expression may differ in cognitive significance, nor that the very expressions in which he is interested may have uses other than those he is concerned with, uses on which they give rise to descriptive rather than singular truth conditions. In doing so, however, he must accept that his account is not an account of the uses of natural language indexical expressions, but rather an account of the logical properties of some uses of those expressions. His theory and my own are, therefore, theories of different things, since what I am interested in is the semantics and interpretation of a class of linguistic expressions.

Not all direct-referentialists are willing, however, to swallow this pill; to accept that a full account of the meaning properties of natural language indexicals lies beyond the scope of their theories. For Perry (e.g. 1997a, 1997b), for instance, facts concerning cognitive significance can be accommodated within a direct-reference framework by the simple move of identifying cognitive significance with referent-role pairings (or, in Kaplan's terms, with character-content pairings). The idea, very briefly, is that the cognitive significance of, for instance, my own mental tokenings of 'I' derives not from the fact that they refer to me, nor from the fact that I am entertaining the character of 'I', but from the fact that I am thinking of me under the character of 'I'. Such accounts are more clearly in conflict with my own and, although I do not intend to go into any detail on the question here, they seem to come off worse: as, for instance, Recanati (1993) has shown, there are good reasons why we would not want to identify the cognitive significance of an indexical with its linguistic meaning, or indeed with the pairing of its referent and linguistic meaning.

What does all this show us then? What I believe it shows us is that the direct reference picture of indexicals may well offer a satisfactory account of the logical properties of some types of use to which indexicals may be put. What it cannot, by

itself, achieve, is an account of the meaning properties of a class of natural language expressions. If natural language is what we are interested in, therefore, we should prefer the sort of account which I have defended above.

## Chapter five

### Definite descriptions

#### 5.0

In this chapter I want to turn my attention towards definite descriptions, and, in particular, to the debate over descriptions initiated by Donnellan (1966). Donnellan points out that there are two apparently different ways in which definite descriptions (expressions typically of the form 'the F') can be used. On the one hand they can be used to talk about whoever or whatever satisfies their descriptive content, and on the other they can be used to pick out a particular individual. Take, for instance, the sentence in (1):

- 1) The Ferrari driver has an unfair advantage

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

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

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

### 3) Jones has an unfair advantage

If, in uttering (1), the speaker means (2), then she is using the description 'the Ferrari driver' *attributively*, in Donnellan's terminology; if she means (3) she is using the description *referentially*. This, then, in pre-theoretical terms, is Donnellan's referential-attributive distinction. In this chapter I want to reverse the order of proceedings of the last two chapters; I shall, that is, start by examining some of the dominant positions on descriptions to be found in the literature, and only once an outline of those positions is in place shall I develop the account of descriptions which I favour.

#### 5.1 The debate on descriptions

There are certain key questions that have guided the debate on referential and attributive uses of definite descriptions. The central question, as commonly formulated, is whether the referential-attributive distinction is semantically significant. But what does this mean? There are two ways that it is commonly cashed out in the literature on descriptions:

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

Given the type of theoretical presuppositions discussed in chapter 2, many contributions to the referential-attributive debate have assumed a particular relation between these two formulations to hold: they have assumed, in essence, that if you answer 'yes' to (a), you are thereby committed to answering 'yes' to (b), and if you answer 'no' to (a) you are committed to answering 'no' to (b). It should be clear, however, that such an assumption is unwarranted from the perspective of the framework within which I have been working, a framework based on the theoretical underpinnings developed in chapter 2. Once it's

appreciated that this assumption is unwarranted, I shall argue, a great number of the apparently thorny issues which have surrounded descriptions dissolve.

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

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

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

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

#### 5.1.1 Donnellan's position

What is Donnellan's answer to question (a)? It seems clear to me (although not to Kripke (1977)) that for Donnellan there is a truth-conditional difference in the proposition expressed between referential and attributive uses. An utterance of a sentence containing an attributively-used definite description in subject position expresses an object-independent proposition: in Donnellan's terms the description occurs *essentially*, in that its descriptive content appears in a specification of the truth conditions of this proposition. An utterance of a sentence containing a referentially-used definite description in subject position, by contrast, expresses an object-dependent



proposition: the descriptive content of the description plays no part in a specification of truth conditions, serving merely to guide the hearer towards a particular individual, and then falling out of the picture.

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

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

In an attempt to demonstrate that propositions with different truth conditions really are expressed on referential and attributive uses, Donnellan turns to contexts in which truth values for the different uses do not coincide; in other words, he turns to cases of misdescription. To see how this works, consider (1) again but, this time, imagine that the speaker and hearer have made a mistake: Jones is not, in fact, the Ferrari driver

---

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

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

This reliance on cases of misdescription has opened Donnellan to some of the more forceful criticism of his account, in particular that of Kripke (1977). The central problem is that this phenomenon seems not to be specific to definite descriptions. It is, for instance, possible, in an appropriate context, to achieve reference to Jones by using the proper name 'Smith'. This should not lead us, so the argument goes, to posit an ambiguity in proper names. Many of those who argue in favour of a Donnellan-type position now accept that, at the very least, the referential/attribution distinction is not best illustrated by recourse to misdescription (see, for instance, Recanati (1993) and Wettstein (1981)).

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

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

Donnellan (1966 p. 59)

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

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

### 5.1.2 Yes-theory vs no-theory

On, then, to the two main strands in the debate following from Donnellan's paper. I will start with a few general introductory comments. It seems to me that these strands are based on two good insights and one bad assumption: the one bad assumption, as I have already mentioned, is that answering 'yes' to question (a) above commits you to answering 'yes' to question (b); the two good insights are that, on the one hand, descriptions are univocal and, on the other, that the referential-attributive distinction corresponds to a difference in truth conditions. Anyone who relies on the bad assumption, however, cannot consistently maintain both the good insights: if you answer 'yes' to both questions (a) and (b), then you are denying the univocality of descriptions,

if you answer 'no' to both (a) and (b) then you are denying that the referential-attributive distinction is truth-conditional. And these are just the positions that the two main strands take.

The second of these positions, that the answer to both (a) and (b) is 'no', that the referential-attributive distinction is not truth-conditional and that descriptions are not ambiguous, is adopted by a group of theorists including Grice (e.g. (1969)), Neale (e.g. (1990)), Kripke and others who I shall call *no-theorists*. For the no-theorist the picture looks something like this<sup>2</sup>: descriptions encode a Russellian semantics; on the basis of the subsidiary bad assumption that, with slight caveats, the same answer must be given to question (c) as to question (d), the no-theorist therefore holds that the proposition expressed by any utterance containing a description must have Russellian truth conditions; in other words, for the no-theorist, even where 'the F' is used referentially, the proposition expressed by 'the F is G' will be true iff there is exactly one F and that one is G.

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

---

<sup>2</sup> Clearly there are significant differences between the accounts that I am grouping together here; I believe, however, that there are sufficient similarities for it to be valid for my purposes to discuss them together.

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

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

A: I am out of petrol

B: There is a garage round the corner

Grice (1967, p. 32)

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

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

What evidence is there for this claim, beyond the Wiggins-type intuition on cases of misdescription? The most heavily-used weapon in the no-theorist's armoury is

methodological rather than empirical: the no-theorist argues that his account is more parsimonious than the alternative; he need only posit one semantically encoded meaning for 'the' and the rest of the work is done by independently motivated pragmatic machinery. Of course any account which argues for a truth-conditional difference between referential and attributive uses must, so long as it's based on the bad assumption, posit two semantically encoded meanings for 'the' if it's going to maintain that the referential-attributive distinction corresponds to a difference in the truth conditions of the proposition expressed. And that, so the no-theorist argues, violates Modified Occam's Razor, Grice's principle that 'senses are not to be multiplied beyond necessity' (Grice (1967)). As Recanati (1993) points out, however, this argument only carries any weight so long as the bad assumption is maintained; so long, in other words, as you equate pragmatically derived meaning and implicated meaning. Given a view of meaning along relevance-theoretic lines, however, there is no temptation to equate these two. On such a view, therefore, the no-theorist's methodological argument has no force<sup>3</sup>.

The no-theorist is left, then, with the Wiggins-type intuition as the bedrock of his claim that referential interpretations are communicated as implicatures; this intuition is the heart of his answer to (a) and, as I shall argue later, it seems to me that this intuition is not the evidence for the no-theorist's claims that it has been taken to be. The no-theorist has rather more arguments at his disposal when it comes to answering (b), the question of whether 'the' is lexically ambiguous or not (and remember that, for the no-theorist who holds to the bad assumption, a 'no' answer to (b) is every bit as good as a 'no' answer to (a), since the one is supposed to follow from the other). Kripke (1977 fn 28), for instance, shows that in some perfectly intelligible cases it is hard to say whether a description is being used referentially or attributively, an unlikely state of affairs for a true lexical ambiguity.

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

---

<sup>3</sup> I shall have more to say about ambiguity when I come on to discuss possible objections to my overall view of singular expressions in chapter 7.

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

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

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

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

A speaker uttering (7) could certainly intend to convey an object-independent quantified proposition. But she could also intend to convey an object-dependent proposition about the particular students in her office. Given that referential and attributive uses can be shown to arise with quantified expressions generally, it becomes a very unattractive option, so the no-theorist argues, to claim that each of these expressions is lexically ambiguous. And for the no-theorist, holding as he does to the bad assumption, that's

strong evidence in favour of the claim that the referential-attributive distinction is not truth-conditional.

The no-theorist, then, believes the following: the referential-attributive distinction is non-truth-conditional, and all utterances with a definite description in subject position, however that description is used, have Russellian truth conditions; the referential use of descriptions arises not at the level of what is said but at the level of what is communicated, i.e. it arises as an implicature. And I have pointed towards a few of the stronger arguments that the no-theorist relies on in support of his claim. I next want to turn to the other major strand in the debate over definite descriptions: the yes-theoretic position<sup>4</sup>. At the heart of this position are the claims that the referential-attributive distinction does correspond to a difference in truth conditions and, therefore, on the basis of the bad assumption, that 'the' is lexically ambiguous<sup>5</sup>.

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

8) The policeman is pulling cars over

In the appropriate context, for instance, if speaker and hearer are driving along and see a policeman ahead, an utterance of (8) may be used to convey a determinate proposition. It is, however, so the yes-theorist argues, ludicrous to believe that that proposition has Russellian truth conditions; if that were the case it would be true iff there were exactly one policeman and that one was pulling cars over. Since there exists more than one policeman, such a proposition would be false, whereas we may easily imagine a situation in which what is conveyed by (8) is true. There have been various attempts by neo-

---

<sup>4</sup> Versions of this position are taken by Wettstein (1981) and Peacocke (1975), among others.

<sup>5</sup> I suspect that I am doing many of those who might think of themselves as yes-theorists a disservice: their thesis is, in most cases, based exclusively on the first of these two claims, with little or no mention of whether 'the' is ambiguous or not.



Russellian no-theorists to defend their position against this objection. Generally speaking, there are two possible approaches: on one approach, the incomplete descriptive content of a definite description is supplemented at some syntactic level with contextually available information so as to make it uniquely denoting; on the other, the domain of the definite description is constrained, on the basis of context, so as to allow only one interpretation<sup>6</sup>. However, both of these approaches face apparent difficulties. The key problem for the first is that there is no principled way to decide which uniquely denoting description appears in the proposition expressed, given that there will standardly be more than one available. The second solution is challenged by cases in which there is no satisfactory way to restrict the domain so as to give rise to the right interpretation (for a discussion of the shortcomings of these approaches see Breheny (1999)<sup>7</sup>). I shall however have more to say about this second strategy below.

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

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

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

---

<sup>6</sup> Whether these two approaches really are distinct is not entirely clear. Neale (1990) suggests that 'when all is said and done, the explicit [contextual supplementation] and implicit [domain restriction] methods might turn out to be notational variants of one another' (Neale (1990, p115, fn 48).

<sup>7</sup> Breheny (1999) proposes an interesting alternative to these approaches to incompleteness, an approach on which the determiner itself is marked as requiring contextual supplementation in line with the descriptive content of the description.

11) That table is broken

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

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

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

in which the indexical 'she' is not syntactically bound by a quantified antecedent; if 'she' is not acting as a bound variable then it must be acting as a genuine referring expression; 'she' receives its interpretation from the definite description 'the girl in the pink suit'; therefore 'the girl in the pink suit' must be a genuine referring expression; so definite descriptions must be ambiguous as between quantified and referential senses. However,

as Neale (1990) points out, free indexicals can receive their interpretation from definite descriptions which are clearly being used attributively. Consider (13), taken from Neale (1990, p.175):

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

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

The two main positions on definite descriptions should now be clear: yes-theorists believe that the referential-attributive distinction is truth-conditional (and maybe that 'the' is lexically ambiguous); no-theorists believe that 'the' is univocal and that the referential-attributive distinction is non-truth-conditional. I hope, however, that in outlining some of the arguments used on both sides of this debate, I have suggested that each camp has better arguments in favour of one of its two claims than it does for the other. The yes-theorist seems to have good arguments in favour of the claim that the referential-attributive distinction is truth-conditional; he seems rather less well equipped with arguments to support the view that 'the' is lexically ambiguous to the extent that many yes-theorists simply do not address the issue. The no-theorist, on the other hand, has some very good arguments to show that 'the' is univocal; whereas his arguments for referential uses arising as implicatures are more or less entirely theory driven: for the no-theorist the proposition expressed by an utterance is essentially defined as the product of semantic decoding (with a certain amount of give and take to allow indexicals and ambiguity into the picture); pragmatic inference occurs, for the no-theorist, as indeed for the yes-theorist, only in the retrieval of implicated meaning. Since the only way to get from a univocal semantics to two distinct interpretations is via pragmatic inference, one of the interpretations of descriptions must, for the no-theorist, arise as an implicature<sup>8</sup>. Neale (1990), in particular, seems firmly committed to this assumption as it applies to

definite descriptions. In a discussion of the account of Recanati (1989), on which a univocal semantics leads in some contexts to an object-dependent proposition and in other contexts to an object-independent proposition, Neale remarks:

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

Neale (1990, p. 112, fn 36)

Neale seems to be rejecting out of hand the possibility that the gap between semantically encoded meaning and propositional content can, here, be filled by pragmatic inference. This is particularly puzzling in the light of his willingness to accept the proposals of Carston (1988), who argues that the temporal and causal implications associated with certain uses of 'and' are, while not the product of a lexical ambiguity, nevertheless part of the meaning explicitly communicated by an utterance (see Neale (1990, p 108, fn 27)). Why Neale should be prepared to accept the one proposal while rejecting the other is unclear<sup>9</sup>.

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

## 5.2 The semantics and pragmatics of definite descriptions

Before getting into the details of my account I want to draw out the general considerations which will guide it. As I have already made clear, it seems to me that both of the strands of thought above capture certain key facts about the interpretation of

---

<sup>9</sup> This is really no more than a reformulation of the bad assumption.

definite descriptions. In particular, the no-theorist is right to claim that definite descriptions are univocal and the yes-theorist is right to claim that the referential attributive distinction is truth-conditional (I am, of course, not the first to make this claim; see, for instance, Recanati (1993), Rouchota (1992, 1994), Bezuidenhout (1997)). By taking this position, I am committing myself to a further claim: that the bad assumption is wrong; that it is consistent to give different answers to questions (a) and (b) (and also to questions (c) and (d)). Once we accept this claim, the need to give the same answers to (a) and (b) (or (c) and (d)) evaporates. It is now possible to answer 'yes' to (a) and 'no' to (b): referentially and attributively used definite descriptions may well make different contributions to propositional content, but that is no longer any reason to suppose that 'the' is ambiguous; the truth-conditional difference may be the result not of semantic encoding, but of pragmatic processes operating in the retrieval of the proposition expressed. And it is, essentially, a version of this view that I shall advocate.

Probably the clearest way to present my account is by answering the four questions set out earlier. I therefore intend to answer first questions (b) and (d) together and then questions (a) and (c) together. The answer I propose to (b) and (d) is this: definite descriptions are univocal and encode particular constraints on the hearer's interpretation, constraints that may be viewed as procedural in the terms of Blakemore (1987). Looking at this from the other end, a definite description 'the F' provides the hearer with certain information about the propositional representation which constitutes the content of the speaker's informative intention. It tells the hearer firstly that the proposition the speaker wishes to explicate contains an individual concept, be it *de re* or descriptive, and secondly that this individual concept is associated by the speaker with the information *x is F*. This stands in contrast to indefinite descriptions which, speaking very roughly, do not indicate that the speaker is entertaining an individual concept.

Is this, however, enough to do the job we want? Is it sufficient for 'the F' to guide the hearer to an individual concept which the speaker associates with the

---

<sup>9</sup> More recently, Neale does seem to have accepted a truth-conditional distinction between referential and attributive uses, although not the distinction proposed in this chapter (see Neale (1999)).

information *x is F*? It seems that it is not. Consider Bert and Ernie who are standing in a furniture warehouse full of tables. A particularly lovely table catches Bert's eye, and, wanting to share this aesthetic pleasure, Bert says to Ernie:

14) The table is very beautiful

There seems to be something at least infelicitous about Bert's utterance in (14): putting things in pre-theoretic terms, there seems to be something unhappy about using a definite description in a context in which there is more than one entity satisfying its nominal. But maybe this is not a semantic fact, i.e. a fact concerning encoded meaning, but a pragmatic fact; maybe the infelicity of (14) derives from Ernie's inability, given the context of utterance, to reach Bert's intended interpretation. There seems to be good evidence to suggest that this cannot be right. Compare (14) with (15):

15) That table is very beautiful

as uttered in the same context. While it still seems that, without any further contextual cues, Ernie will be unable to establish which table Bert is intending to talk about, and thus to reach Bert's intended interpretation, just as with (14), (15) nevertheless seems to lack the infelicity of (14). What does this tell us about the semantics of definite descriptions? It seems to tell us that their nominal elements must be uniquely denoting, just as per Russell's semantics. As discussed above, however, many uses of definite descriptions are incomplete; their nominal elements are anything but uniquely denoting. How might these uses fit in to the picture so far sketched? I believe that a version of the contextual restriction story mentioned above is likely to do the job for us. That is to say, the nominal elements of definite descriptions are not uniquely denoting simpliciter, but against a pragmatically restricted context. The example in (14) is infelicitous, on this sort of story, due to the lack of any available context in which the nominal 'table' is uniquely denoting. But there are well-known objections to stories along these lines. Although I do not want to go into a detailed defence here, I do want to suggest the outline of a

response. The standard objections involve utterances for which the use of a definite description seems felicitous despite the apparent unavailability of a context in which the description's nominal is uniquely denoting. Thus, for instance, I may have two workmen working on my house, both of whom are qualified electricians, but only one of whom is actually at work on my electrics (say the other is helping him out by doing some plumbing). In this context it seems that I can felicitously utter (16):

16) The electrician arrived late this morning

although there seems to be no available context which includes the electrician I intend to talk about but excludes the other electrician who is currently at work on the plumbing. It seems to me, however, that such examples cast light not on the semantics of definite descriptions but on another dimension of the functioning of lexical semantics.

Within the relevance-theoretic framework there has been much research conducted recently into concept narrowing and loosening (see, for instance, Carston (1996/97, 2002a etc)). The idea is this: the concept expressed by a predicate on an occasion of use may not be identical with the full concept that that predicate lexicalises; the expressed concept may be logically stronger or weaker than the lexicalised predicate. Thus, to take an example much discussed in the literature, the lexical item 'bachelor' presumably encodes a concept which looks something like *unmarried male human*<sup>10</sup>. Consider, however, an utterance of (17):

17) Mary is bored of being single and wants to meet a bachelor

The relevance-theoretic claim is that 'bachelor', as used in (17), expresses not the full BACHELOR concept, but something logically stronger, something which narrows the concept's denotation down to bachelors who are young, straight etc. What I want to suggest is that examples such as (16) hinge not on the semantics of definite descriptions

---

<sup>10</sup> I don't wish to presuppose any particular take on the anatomy of concepts such as BACHELOR. For the time being, however, this approximation will, I hope, suffice.

but on concept narrowing. In an utterance of (16), 'electrician' will express not a concept which denotes anyone qualified as an electrician, but rather a concept which denotes only individuals who perform the functions of an electrician. If this suggestion is on the right lines, then the apparent problem which examples such as (16) pose for a semantics on which the nominal of a definite description expresses a uniquely denoting property relative to a pragmatically-retrieved context evaporates. 'Electrician', as it appears in the nominal of 'the electrician' in (16) does, on this construal, express a uniquely-identifying property, since there is only one individual in the context who is performing the functions of an electrician.

My answers to questions (b) and (d), then, are that definite descriptions are univocal, with 'the F' providing the hearer with the information that the proposition the speaker wishes to explicate contains an individual concept associated with the information *is a unique F (in a pragmatically-retrievable context)*. How about questions (a) and (c)? As I hope I have made clear, it seems to me that the arguments put forward by the ambiguity theorist in favour of his contention that the referential-attributive distinction corresponds to a difference in truth conditions are very convincing. There is, however, still the issue of the Wiggins-type intuition to deal with, the intuition that in cases of misdescription, even if the predicate is true of the intended referent, so long it is false of the actual denotation of the description used, the proposition expressed by an utterance containing a referentially-used definite description will be false. I intend to address this issue once I have laid out my proposal a bit more fully: it's my belief that a reasonable explanation for this intuition should fall naturally from the account. My answer to (a), then, is that there is a truth-conditional distinction between referential and attributive uses.

How about question (c), the question of what contribution definite descriptions make to the propositions expressed by utterances in which they appear? I believe, along with Donnellan and Recanati, among others, that the proposition expressed by an utterance containing an attributively-used definite description is object-independent, whereas the proposition expressed by an utterance containing a referentially-used definite description is object-dependent. How, then, do these different interpretations arise? I



want first to give the outline of an answer before working briefly through a couple of examples to show how the whole thing might go in practice. In outline the answer looks like this: as I've already suggested, what is semantically encoded by definite descriptions does no more than constrain the hearer's search for an interpretation to an individual concept containing the information *x uniquely satisfies the description's nominal property*. The gap between this encoded meaning and the propositional constituent corresponding to the description is filled by inferential pragmatic processes guided by considerations of relevance. The difference between referential and attributive interpretations of definite descriptions amounts to no more than this: in some contexts the optimally relevant interpretation of a description will be a *de re* individual concept, in other contexts it will be a descriptive individual concept.

As I suggested, this might all be a bit clearer if I work briefly through a couple of examples. Let me remind you of sentence (1) and its two proposed interpretations, repeated here as (18), (19) and (20):

- 18) The Ferrari driver has an unfair advantage
- 19) The Ferrari driver, whoever he is, has an unfair advantage
- 20) Jones has an unfair advantage

Let's deal with the referential interpretation, (20), first. Remember that, in the imagined context for (20), speaker and hearer are standing in the pit lane looking at a driver, Jones, who they believe to be the Ferrari driver and who is scheming with the race organiser. In that context it is mutually manifest to speaker and hearer that both have (or are capable of acquiring) a *de re* concept of Jones, since, however we define the epistemic relationship requisite for *de re* thought, we're going to want to say that the speaker and hearer in this case both have sufficiently close epistemic contact with Jones to be able to entertain a *de re* concept of him. The speaker then utters (18) and sets the hearer on a hunt for an individual concept which includes the information *x is a unique Ferrari driver*. There seem to be two obvious candidates, one a descriptive concept something which denotes whoever is the Ferrari driver in this race and the other a *de re*

concept of Jones, standing in front of him. The question is, which one does he accept? And the answer to this, on a relevance theoretic picture, is going to depend on which concept offers him more contextual effects. It's not hard to see, then, which this will be in the proposed context: since the hearer believes Jones to be the Ferrari driver, any contextual effects carried by the descriptive interpretation come along for free with the *de re* interpretation. However, the *de re* interpretation also has additional contextual effects: it might, for instance, contain the information that Jones is wearing a red helmet, in which case the hearer could infer that the driver with a red helmet has an unfair advantage and so on. In this context, then, the *de re* interpretation is going to be optimally relevant since it offers the hearer more contextual effects for no extra processing effort; the description will, therefore, be interpreted referentially.

Now let's consider (19), the attributive interpretation. Remember that in the imagined context for (19), speaker and hearer are discussing a race; they do not know who the Ferrari driver is, they merely know that whoever he is, he has an unfair advantage since he will be allowed to start the race in front of the other cars. Again the speaker utters (18) and the hearer starts his hunt for an individual concept. What options are open to him this time? Notice that the hearer, in this context, has no *de re* concept of the Ferrari driver in question: that is simply what is meant by saying that he doesn't know *who the Ferrari driver is*. So the first individual concept he will access will be the descriptive concept of the Ferrari driver in the race under discussion, i.e. a satisfactoral concept containing information such as 'is a racing driver', 'drives a Ferrari', 'is driving a Ferrari in this race' and so on. He will then, as ever, assess this interpretation for optimal relevance and, given that it will have enough contextual effects - the race is unfair, the Ferrari driver is likely to be the winner, the other drivers may complain - for no undue processing effort, he should accept it as the intended interpretation.

Does this entail that wherever a referential interpretation is available, a hearer should accept it as the intended interpretation? Although this will standardly be the case, for reasons along the lines presented above, there are certain situations in which, although speaker and hearer are both capable of entertaining *de re* representations of the individual which uniquely falls within the denotation of a description, nevertheless an

attributive interpretation of that description will be optimally relevant. The clearest examples are those in which, by using a definite description, a speaker who intended her utterance to express a singular proposition would be putting her hearer to gratuitous processing effort. Consider the following situation: Wimbledon fortnight has just ended and Lleyton Hewitt has won the men's singles title, a fact that is mutually manifest to A and B. A then utters (21):

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

How is B going to interpret the definite description 'the winner of this year's Wimbledon men's singles title'? As with all definite descriptions, the hearer's task is to find the optimally relevant individual concept containing the descriptive content of the description. But will this be a *de re* concept or a descriptive concept in the given context? The *de re* interpretation is eliminated on the grounds that, had that been A's intended interpretation, there is at least one alternative utterance available to her which would have put her hearer to less effort for the same contextual effects, namely (22):

- 22) I bet Lleyton Hewitt also wins the US open

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

This, then, is the outline of my proposal: definite descriptions encode procedures instructing the hearer to access an individual concept associated with the information *x uniquely satisfies the description's nominal*. The rest is left to pragmatic inference, guided by considerations of relevance. The distinction between referential and attributive

interpretations arises, on this picture, not as a product of the encoded meaning of descriptions, nor, indeed, in some sense, as a product of the interpretation *process*: in all cases the interpretation of definite descriptions proceeds along the same lines, leading the hearer from encoded meaning to the optimally relevant individual concept in the context. The referential-attributive distinction lies, rather, at the level of the mind-world relationship: although all definite descriptions lead to individual concepts, some of those concepts, those which are *de re*, are anchored in a particular way to objects in the world, and thus give rise to singular truth conditions, whereas others, those which are descriptive, are not so anchored and give rise to general truth conditions.

As I have already mentioned, however, I am far from being the first to argue that the referential-attributive distinction is truth-conditional at the same time as advocating a univocal semantics for definite descriptions. How, then, does the account proposed above differ from others within this 'yes-no' tradition, and why? What we might call the 'standard' account within this tradition takes the attributive reading as prior, i.e. as the lexically encoded meaning of descriptions. The referential reading is then accessed in the appropriate context via an extra interpretive step. Recanati (1993), for example, proposes a synecdochic step from the descriptive concept encoded by the description to a *de re* concept of which the encoded concept is a part. Within the framework I have adopted, however, there is a problem for any account of this type. If we are to take the central rôle played by relevance seriously, then we must accept, with Sperber and Wilson (1986/95) that each utterance communicates a guarantee of its own optimal relevance. One aspect of this guarantee is that there should be no alternative utterance (compatible with the speaker's abilities and preferences) which would have allowed the hearer to access the same interpretation for less processing effort. On any 'attributive-encoding' account, however, there will always be an utterance available to the speaker which would satisfy precisely this condition: the equivalent complex demonstrative. In other words, if attributive encoding is right, a speaker who intends her hearer to access a referential interpretation of 'the F' could have achieved the same result while putting her hearer to less effort by uttering 'that F'. There is, of course, no current consensus on the precise encoding and interpretation of complex demonstratives and I shall come on to discuss

this question at some length in the next chapter. On most analyses, however, the proposition expressed by an utterance containing a complex demonstrative in subject position is singular; for many this is the key distinction between definite and complex demonstratives. If this is the case, then it seems clear that the one-step interpretive procedure from demonstrative utterances to singular propositions will put the hearer to less processing effort than the two-step procedure from 'the F' to singular proposition proposed on the 'attributive-encoding' account. This leaves us with the worrying result that referential interpretations of definite description sentences should never achieve optimal relevance and should, therefore, never be accepted.

The analysis I have proposed above, however, avoids this unwanted conclusion. On this account, the encoded meaning and interpretation of definite descriptions is essentially blind to the referential-attributive distinction: the optimally relevant interpretation of a particular utterance of a description will simply be the most accessible concept containing the information *x is uniquely F* to offer enough contextual effects for no undue processing effort, whether that concept is descriptive or *de re*. In contrast to the 'attributive-encoding'-type account, the interpretative path from definite description to referential reading involves just one step. It should, thus, put the hearer to no more processing effort than the equivalent complex demonstrative and so can achieve optimal relevance.

A further advantage of the account proposed is that it can deal with the data upon which the main alternative to a Russellian semantics relies. While Russell's account focuses on the uniqueness implication apparently carried by definite descriptions, a parallel tradition, advocated, in one form or another, by, for instance, Heim (1988) and Kempson (1986), focuses on the familiarity or accessibility associated with definite descriptions. Consider the following sentence:

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

For Heim (1988), such examples suggest that what definite descriptions encode is not uniqueness but familiarity. Building this into her overall *file change* semantics

programme, definite descriptions instruct a hearer to update an existing file, whereas indefinite descriptions instruct the hearer to initiate a new file. This view has come under attack on the basis that, although definite descriptions are standardly used to talk about the familiar, they can be used to introduce new referents. Consider, for instance, (24):

24) The oldest panda in China is coming to London Zoo

An utterance of (24) can clearly be understood by a hearer who had no pre-existing file on the oldest panda in China. Heim has mounted a defence against this objection on which there are mechanisms which link such descriptions to existing files. It seems to me, however, that a more natural explanation of this data falls from the account proposed above, along with standard communicative principles. There seems to me to be no need to propose that definite descriptions encode anything like familiarity. Remember that, on the account I propose, what a definite description does encode is an instruction to access an individual concept. What, then, if the hearer does not have, or is not able to acquire, the intended individual concept? In this case the speaker would be putting her hearer to gratuitous (indeed fruitless) processing effort, thus guaranteeing that her utterance will fail to be optimally relevant. In order to comply with the presumption of optimal relevance, a speaker who uses a definite description must have reason to believe that the individual concept which is the intended interpretation of the description is either already in the hearer's conceptual repertoire (i.e. is familiar), or is accessible to the hearer. The familiarity or accessibility standardly associated with definite descriptions therefore fall out naturally from the semantics proposed, in tandem with the presumption of optimal relevance.

At this stage I should tidy up a couple of loose ends. I suggested earlier that my account would be able to handle the Wiggins-type intuition that, in cases of misdescription, the truth value of an utterance will depend on whoever the description denotes, not on the individual the speaker is intending to refer to. There are two things I would like to say on this point. Firstly, I agree with Reimer (1998), who argues that the supposedly pre-theoretical intuitions of the opposing camps in this debate are, in reality,

anything but pre-theoretical. They are, rather, heavily influenced by the theoretical preconceptions that each camp brings to the debate. But secondly, it does not seem to me to be necessary to oppose Wiggins' intuition in order to maintain a truth-conditionally significant referential/attributional distinction. Recall the distinction that I drew in chapter 2 between two types of speaker intention, informative and derivational. What seems to be going on in cases of misdescription is that these intentions are in some sense conflicting. Of course it may be that speaker and hearer are entirely unaware of any conflict, since they are both labouring under the same misapprehensions. Nevertheless, in the sort of cases under discussion it is part of the speaker's derivational intention to use her intended referent's unique F-ness to guide her hearer towards a particular interpretation. From the omniscient perspective adopted by the meaning theorist, there are therefore two clear and conflicting intentions, one to refer to a particular object and the other to guide the hearer to that object via its unique F-ness. As discussed in chapter 2, intuitions may be weighted towards either of these intentions, thus giving rise to the split intuitions so commonly found in discussions of the semantics of definite descriptions.

Finally, I want briefly to examine a couple of other uses of definite descriptions mentioned earlier: generic and functional uses. On generic uses I have no more to say than that, so long as one is prepared to accept types as well as tokens into one's ontology, generic uses of definite descriptions are only to be expected on the above account: generic interpretations merely involve retrieving individual concepts which represent type-level individuals. Functional uses, such as (5), repeated here as (25):

25) The president of the United States changes every four years

raise rather more complex issues and have recently received much attention in the literature on descriptions. Recanati (1993), for instance, abandons an earlier account (see Recanati (1989)) partly due to what he perceives as its inability to handle such uses. I would like tentatively to suggest that functional uses may, in fact, be something of a red herring. There are, after all, certain features which distinguish functional from other uses

of definite descriptions: most notably, unless one wishes to introduce individuals comprising sequences of individuals into one's ontology, functional uses do not denote individuals. It seems to me that this might support an account of functional uses along somewhat different lines from those proposed for, among others, referential and attributive uses. One attractive possibility is that functional uses may be metalinguistic; in concrete terms, an utterance of (25) may express a proposition to the effect that the denotation of the descriptive individual concept corresponding to 'the President of the United States' changes every four years. But how might a hearer reach this interpretation? As ever, his first step must be to search for an optimally relevant individual concept. So what candidate concepts could there be? There might be a candidate *de re* concept of George W. Bush, giving the overall utterance interpretation in (26):

26) George W. Bush changes every four years

But this interpretation will, under most circumstances, fail to achieve optimal relevance: finding a context in which (26) offers sufficient contextual effects to be relevant will put the hearer to a great deal of processing effort<sup>11</sup>. This interpretation will, therefore be rejected and another accessed, maybe a descriptive concept of whoever happens to be the president. But again, under most circumstances, such an interpretation will fail to achieve optimal relevance for much the same reasons as above. Having failed to find an optimally relevant individual concept, the hearer will, on the strength of the communicated presumption of optimal relevance, be justified in searching for an interpretation beyond those immediately made available by the semantic encoding of the definite description. A metalinguistic interpretation will yield sufficient contextual effects to be relevant - such as, for instance, that Presidential elections must take place every four years, the current President has nearly four years of his term left – with minimal processing effort and should thus be accepted as optimally relevant. I admit that this is a



speculative story on functional uses, but it seems to me that it opens up a possibility that might well bear further investigation. The key point is that, if something along these lines is right, functional uses should receive a different kind of analysis from those given for referential and attributive (etc.) uses.

### 5.3 Conclusion

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

The account I have proposed posits a univocal semantics for definite descriptions which is indeterminate between referential and attributive readings (among others). One upshot of the particular account I have proposed is that the distinction between referential and attributive interpretations of descriptions is the product not of the relationship between language and mind, but of the relationship between mind and world: it arises because of a distinction between the kind of link that exists between *de re* concepts and the world and the kind of link that exists between descriptive concepts and the world. This seems to me to be yet more evidence, if evidence is needed, that the real

---

<sup>11</sup> Of course, that's not to say that there aren't some contexts in which this interpretation would achieve optimal relevance and would thus be accepted – imagine, for instance, (26) as a contribution to a

domain of truth-conditional semantics is not sentences of natural language but propositions, i.e. sentences in the language of thought.

### Complex demonstratives

#### 6.0

Over the last few years, the philosophical debate on the referring-denoting distinction has shifted its focus. Donnellan (1966) placed definite descriptions at the heart of the debate, and for three decades or so, they stayed there. Recently, however, definite descriptions have been usurped by what have variously been called *demonstrative descriptions*, *complex 'that'-phrases* and *complex demonstratives* (I shall stick to this last throughout). Complex demonstratives, which at a first pass we can think of as any expression of the form 'that F', pose a particular problem for the truth-conditional semanticist toiling on the border between referring and denoting: on the one hand it would seem that they are semantically closely linked to simple demonstratives, standardly taken to be the prototypical referring expressions, but on the other they display the kind of syntactic complexity associated with definite descriptions and, more broadly, with the class of quantifiers. For the standard truth-conditionalist, who either overtly or tacitly accepts the hypothesis that all noun phrases must either be referential or quantificational<sup>1</sup>, this is an alarming state of affairs, and much ink has been spilled in the effort to show that these apparently anomalous expressions do, after all, behave in familiar ways. None of the stories currently on offer, however, embodies the truth, the whole truth and nothing but the truth: some, while offering nuggets of truth, are empirically inadequate; others offer much of the truth as far as they go, but seek to delimit their field of inquiry in ways that, under scrutiny, reveal themselves as unprincipled; and yet others offer what seem to me to be profound insights into the workings of complex demonstratives, but fall at least at the hurdle of theoretical parsimony.

I shall structure this chapter as I did the last: first I shall examine a cross section of accounts from along the theoretical spectrum. My aim will be to unearth the nuggets hidden within these analyses at the same time as demonstrating the shortcomings of each. In the second part of the chapter, I want to take the insights

---

<sup>1</sup>For overt formulations of this assumption, see, for instance, Neale (1993) and Dever (2001).

offered by these accounts and, with the help of a type of evidence used by none of them, approach complex demonstratives from within the relevance-theoretic framework.

## 6.1 Referentialism

It is more or less a given of philosophical semantics that if any expressions are referential then simple demonstratives are. They seem to display more clearly than any other expression-type what have been taken to be the two key features of referentiality: they track individuals across possible worlds and their empty uses lead to propositional failure. If I utter (1):

### 1) That is filthy

then the truth of my utterance will depend, in any possible world, on the object I have referred to in the actual world; if, say, I have pointed at a whiteboard while uttering (1), then my utterance will be true at a possible world iff that very whiteboard is filthy at that world<sup>2</sup>. If, on the other hand, I utter (1) without intending to refer to anything by my use of 'that', then intuitively it seems that my utterance fails to express any proposition; what would have to be the case, after all, for my utterance to be true?

If one accepts that simple demonstratives are referential, there would seem to be at least a strong pull towards extending this story to complex demonstratives<sup>3</sup>. Simple demonstratives, after all, form a syntactic part of their complex brethren and furthermore their uses seem semantically very closely allied. If I utter (1) while pointing at a particular whiteboard, then it would seem that what I have said is much the same as what I would have said by pointing at the same whiteboard while uttering (2):

### 2) That whiteboard is filthy

---

<sup>2</sup> I don't intend to suggest here that the demonstration is itself semantically significant, but rather that it gives the hearer a way in to the speaker's intentions, which are semantically significant.

<sup>3</sup> I hope it goes without saying that I am not myself supporting the view that simple demonstratives are referential; I am merely trying to represent a particular view of demonstratives, a view which differs fundamentally from my own.

It also seems that our use of a complex demonstrative in (2) passes the tests for referentiality mentioned above: the truth of the proposition expressed by an utterance of (2) will depend on the very same object across possible worlds and if, in uttering (2), I fail to pick out a particular object by my use of 'that whiteboard', then again it seems that my utterance fails to express any proposition. Given this evidence, an entirely straightforward account of complex demonstratives seems to present itself: they are referential, just as are simple demonstratives; their contribution to propositional content goes no further than their referent.

But what role will be played on such a story by the nominal element of the complex demonstrative (the 'F' in 'that F')? Maybe, on the model of simple demonstratives, they play no more role than the contextual cues that guide us toward the object a speaker intends to refer to as 'that'; they may, in other words, merely play a pragmatic role in helping the hearer assign reference. This proposal, advocated by, for instance, Larson and Segal (1995), has at least one immediate implication that has seemed unpalatable to many: the nominal element of a complex demonstrative comprises a meaningful NP, so why should the meaning of this NP simply disappear from utterance content? For the compositional semanticist, and for some who are not so wedded to compositionality, disappearing meanings have been seen as very unwelcome theoretical artefacts. But there are other problems beyond this for the simple referential thesis. Utterances with complex demonstratives in subject position (henceforth *complex demonstrative utterances*) seem to give rise to entailments which are incompatible with simple referentiality. Consider (3) and (4):

- 3) That green car is very old
- 4) Some green car is very old

There does on the face of it seem to be an entailment from (3) to (4), yet the simple referentialist has no story to tell about this: for her, sentence (3) expresses the same proposition as would (5), uttered while pointing at the appropriate green car:

- 5) That is very old

and there is surely no temptation to think that (5) entails (4) (there may, of course, be an entailment once we add the extra premise that that is a green car, but this is not what's at issue here). How might the referentialist try to rescue her story, which after all seems to be firmly rooted in intuition, from this challenge? One strategy, adopted in differing forms by Kaplan (1978, 1989a, 1989b), Braun (1994) and Borg (2000), is to locate the contribution of the nominal not in utterance content but in Kaplanian character. On this kind of story, the nominal 'green car' in (4) would constrain the reference of 'that green car' to something that is, in fact, a green car, and then drop out of the picture. How does this help with the difficulties faced by our initial story? Firstly, on this type of account the nominal does make a semantic contribution, albeit not a contribution to propositional content. Secondly, although it still sees no actual entailment between, for instance, the propositions expressed by utterances of (3) and (4), it does at least have a story to tell about why there should *appear* to be an entailment: since the complex demonstrative 'that green car' can only be used to refer to something that is, in fact, a green car, an utterance of (3) cannot express a true proposition unless some green car is very old. In other words, although the proposition expressed by (3) does not entail the proposition expressed by (4), (6) is valid in the sense of Kaplan (e.g. 1989b):

- 6) If that green car is very old, then some green car is very old.

which is, for the character-theorist, close enough to explain the illusion of entailment.

For Dever (2001), however, these solutions fail to do the work we want of them: firstly, what we're trying to explain about (3) and (4) is not the illusion of an entailment but an actual entailment, and secondly, a story on which the nominal element of a complex demonstrative contributes merely to character and not to content violates semantic innocence, the principle that the same expression should display the same semantic behaviour regardless of the linguistic context in which it appears. Dever himself proposes a story which, although closely allied to Braun's character-theoretic analysis, can cope with these supposed problems. His story falls within what seems to be a growing trend in philosophical semantics for multiple-proposition analyses, i.e. analyses on which single utterances standardly express more than one proposition. Abstracting away from the syntactic minutiae, Dever's story

amounts to the proposal that ‘that F is G’ expresses a sequence of two propositions: the primary proposition that that is G, i.e. the referential proposition expressed on all the accounts we’ve looked at so far, and the secondary proposition that that is F.

I am unsure as to how impressed we should be by the purported problems with the character-theoretic account which motivate Dever’s story. Firstly, as Dever himself points out, there’s little evidence to support the view that there is an actual, rather than merely apparent, entailment between, for instance, (3) and (4); and secondly, as I have argued in Powell (2002), there are good reasons why we should not expect semantic innocence to hold at the level of propositional content, but rather should relegate it to the lower level of linguistic meaning. Whether or not you are inclined to agree with me on this, there is a rather more fundamental problem facing Dever’s account, and indeed facing all of the referentialist accounts we have so far looked at: as King (1999, 2001) shows, complex demonstratives seem to have non-referential uses. In particular, King points to three different types of uses of complex demonstratives which, he argues, present the referentialist with difficulties. Firstly there are what he terms *no demonstration no speaker reference (NDNS)* uses. He asks us to consider Scott who, in the course of a lecture on hominid history, utters (7):

- 7) That hominid who discovered how to start fires was a genius.

Here Scott has no *particular* individual in mind to whom he intends to refer by his use of ‘that hominid’ and equally, it seems, his hearer need not think of any *particular* individual as ‘that hominid’ in order to understand Scott’s utterance.

Secondly King points to *quantifying in (QI)* uses, uses on which a quantifier outside the complex demonstrative binds an anaphoric element inside:

- 8) Every father dreads that moment when his eldest child leaves home.

And finally there are *narrow scope (NS)* uses, on which complex demonstratives seem to occur within the scope of other quantifiers:

- 9) That professor who brought in the biggest grant in each division will be honored

There seem to be two possible interpretations of (9), one on which one and the same professor brought in the biggest grant in each division and ‘that professor’ refers to her, and one on which ‘that professor’ occurs within the scope of ‘each division’ and thus picks out, for each division, the unique professor who brought in the biggest grant in that division. On the latter interpretation, the complex demonstrative has an NS occurrence.

On none of these types of use does the complex demonstrative appear to be functioning referentially: NDNS uses pattern with Donnellan (1966)’s attributive uses of definite descriptions in allowing ‘whoever he is/was’ to be inserted after the NP:

10) That hominid who discovered how to start fires, whoever he was, was a genius.

which at least strongly suggests that they should receive a non-referential analysis. As for QI and NS uses, there seems little temptation to analyse them referentially: what could possibly count as the referent of ‘that moment when his eldest child leaves home’ in (8) or ‘that professor who brought in the biggest grant in each division’ in (9)?

On the face of it, the existence of such non-referential uses of complex demonstratives seems to rule out any straightforwardly referential account. Many referentialists, however, have recognised such uses (in particular QI uses) without apparently feeling the need to abandon their accounts and start again. The approach of Lepore and Ludwig (2000) is typical:

Sometimes, of course, “that” is pressed into service as a variant of “the”, and one could imagine someone uttering (21) [‘Every man loves that woman who is his mother’] with that in mind. We are not concerned with such uses of “that”, but rather with demonstrative uses.

Lepore and Ludwig (200, p.219 fn.28)

We can take this exclusion in two ways: either Lepore and Ludwig (and others who support this approach to non-referential uses) are not in the business of studying



language itself, but are, rather, concerned with pinning down the truth-conditions of particular uses of language, or there is an implicit assumption in the passage above that there are two different ‘that’s in play, one referential and one non-referential. If we are to take the first reading, then Lepore and Ludwig are simply not engaged in the same enquiry as, for instance, King, Dever and myself. If, however, we are to take the second reading, explicitly advocated by Dever, then the referentialist may still have a claim to provide a complete account of the expression-type he has in his sights, i.e. referential complex demonstratives.

The burden of proof is, here, on the referentialist. All things being equal, and with our lesson on parsimony learnt from Grice and modified Occam’s razor, we should be ready to reject a story which posits an ambiguity in ‘that’ for one that posits none. Presuming we can tell a story about ‘that’ which doesn’t require us to posit an ambiguity (and I aim to do just this in the latter part of this chapter), the referentialist must therefore convince us that there really is an ambiguity if he is to overcome the methodological bias against his story. Dever attempts to give at least some evidence to show that ‘that’ is ambiguous (apart from his observation that the Oxford English Dictionary supports this view, which I don’t intend to worry about too much). His key evidence is that, in examples such as (8), substituting ‘this’ for ‘that’ leads to infelicity:

11) \*Every father dreads this moment when his eldest child leaves home.

Given that, for referential uses of ‘that’, this substitution is unproblematic, Dever concludes that the ‘that’ in, for instance, (8) is simply a different lexeme from the ‘that’ in, for instance, (2). There is, however, an undefended assumption in Dever’s argument that seems to be at least questionable: that because ‘this’ and ‘that’ are substitutable in referential contexts, their semantics must be such that they are substitutable in all contexts. But there is a clear and familiar difference between the semantics of ‘this’ and the semantics of ‘that’: ‘this’ is used to pick out objects which are by some standard near to the speaker, while ‘that’ is used to pick out objects which are not near the speaker. There is also some evidence to suggest that ‘that’ is, in some sense, a default, whereas ‘this’ is marked: consider, for instance, that, however close something is to the speaker, so long as there is not another further object which

she might be taken to be referring to, she may refer to it as ‘that’; whereas a speaker who refers to something distant as ‘this’, whether or not there is a candidate nearer object, is at least speaking infelicitously. It seems to me that it may be just this dimension of the semantics of ‘this’ and ‘that’ that is responsible for the infelicity of (11): on the one hand the speaker’s use of the complex demonstrative is descriptive, thus picking out no particular individual, but on the other she is indicating, by her use of ‘this’ rather than ‘that’, that what she is talking about is relatively near to her. It may well be the incompatibility between these two that leads to infelicity.

Whether or not you accept this argument, the point I want to stress is that the infelicity of utterances like (11) may well be explicable in terms other than the ambiguity of ‘that’. Until we have stronger empirical evidence to support the view that ‘that’ is ambiguous, therefore, we should, on methodological grounds, prefer a univocal semantics for ‘that’. In the next section I shall look at some accounts of this type.

## 6.2 Quantificationalism

The point we seem to have reached is this: given the lack of evidence for an ambiguity, we are on the hunt for a univocal semantics for complex demonstratives, but one that allows for both referential and non-referential uses. Where might this hunt lead us? On the assumption that all NPs are either referential or quantificational, there seems only one direction open to us: complex demonstratives are quantificational. But what might a quantificational semantics for complex demonstratives look like? Neale (1993) outlines a story on which ‘that F is G’ is equivalent to ‘the actual F I am demonstrating is G’ and on which complex demonstratives take mandatory wide scope. The advantage of this account, for Neale, is that it gives the content of the nominal a role to play in utterance content, thus fitting with the intuition mentioned earlier that meanings shouldn’t simply vanish. It is, however, as Lepore and Ludwig (2000) point out, not only ad hoc in that there is no clear reason why the complex demonstrative should always take wide scope, but also empirically inadequate: it has no story to tell about why the failure of a complex demonstrative to designate should result in the failure of the utterance in which it appears to express a proposition (this is not, after all, a property of definite

descriptions); and it predicts, wrongly, that an utterance by *s* of ‘that *F* is *G*’ entails that *s* exists. We thus need something rather more subtle if we are to place complex demonstratives within the class of quantifiers.

Two recent accounts, those of King (2001) and Lepore and Ludwig (2000), attempt to provide just that something. Lepore and Ludwig’s account is based on the premise that, contrary to superficial appearance, complex demonstratives are not syntactically formed by the concatenation of a quantifier ‘that’ and a nominal. Instead they are restricted quantifiers with the quantifier element suppressed and both ‘that’ and the nominal contributing to the restrictor. In essence, as Lepore and Ludwig make clear, their story treats ‘that *F* is *G*’ as equivalent to [The *x*: *x*=that and *x* is *F*](*x* is *G*)<sup>4</sup>. I suspect that, as far as it goes, this story may have a certain amount to recommend it. But, as mentioned earlier, Lepore and Ludwig are only in the business of accounting for referential uses of complex demonstratives. If, therefore, we accept, as I have argued we should, that there are truly non-referential uses of complex demonstratives, then our hunt for a satisfactory account continues.

King (2001) offers a story which is specifically tailored to account for both referential and non-referential uses of complex demonstratives. On his analysis, complex demonstratives are profoundly context-sensitive quantifiers, depending on context not just for which properties are to be taken as determining the intended referent, but also for whether that referent is to be tracked across possible worlds or not. For King, the lexical meaning of ‘that’ (as it appears in complex demonstrative constructions; he only tentatively suggests that the outline of his account might be extended to simple demonstratives) is a four-place relation:

\_\_\_\_\_ and \_\_\_\_\_ are uniquely \_\_\_\_\_ in an object *x* and *x* is \_\_\_\_\_  
 King (2000, p.43)

The first argument place, on this story, is to be filled by the property expressed by the nominal, the second by a property determined by speaker intention, the third either by *jointly instantiated* or *jointly instantiated in*  $\langle w, t \rangle$  and the fourth by whatever property is expressed by the predicate. There are for King two fundamentally different kinds of intention that can determine the property to saturate the second argument

---

<sup>4</sup> Lepore and Ludwig (2000, p.215)

place: either the speaker may have a perceptual intention, which may include both intentions towards objects of immediate perception and towards remembered objects of past perception, or she may have a descriptive intention, i.e. an intention towards whatever satisfies some descriptive condition. Which of these two kinds of intention she has will make a difference not only to the first-order property that saturates the second argument position but also to the second-order property that saturates the third argument position. If the speaker has a perceptual intention, towards the object a say, then the second argument position will be saturated by the property of being identical to a and the third by the property of being jointly instantiated in  $\langle w, t \rangle$ , where  $w$  and  $t$  are to be taken rigidly to designate the world and time of utterance. If, however, the speaker's intention is descriptive, then the second argument position will be saturated by whatever property her intention determines, and the third by the property of being jointly instantiated.

Let me work through a couple of examples to illustrate how this story is supposed to go. Consider first a classic referential use of a complex demonstrative: Janet and John are at a party and, seeing a man across the room, Janet says:

12) That man wearing a kipper tie has no dress sense.

What proposition will Janet's utterance have expressed on King's story? The first parameter will be saturated with the property of being a man wearing a kipper tie, the second, since Janet has a perceptual intention concerning the man in the kipper tie, with the property of being (identical to) a, where a is the very individual in question, and the third, again because Janet's intention is perceptual, with *are uniquely jointly instantiated in*  $\langle w, t \rangle$ . Finally, the fourth parameter will be saturated with the property of having no dress sense. What, then, will this amount to? It amounts to the claim that Janet's utterance will be true in any circumstance of evaluation iff there is, in the context of Janet's utterance, one and only one thing which is both a man wearing a kipper tie and also a, and that thing has no dress sense in the circumstance.

Next let's take a non-referential use, say an utterance of the sentence in (7). What truth conditions will King's account predict for this? Here the first parameter will be filled with the property of being a hominid who discovered fire, the second with the same property, since the speaker's intention is simply to talk about whoever

satisfies the nominal, and the third, since the speaker's intention is descriptive, with *are uniquely jointly instantiated*. Finally again the fourth parameter will be filled by the property of being a genius. King's prediction here, then, is that an utterance of (7) will be true in a circumstance iff there is one and only one hominid who invented fire in that circumstance and he is a genius in that circumstance<sup>5</sup>.

What advantages does this complex semantics have over the simpler story told by Neale (1993)? One clear advantage is that it falls at neither of the two hurdles mentioned above: firstly, the proposition expressed by an utterance of 'that F is G' will not include any mode of presentation of the speaker, and thus will not entail that the speaker exists. Secondly, where there is nothing which the speaker refers to with a perceptual use of 'that F' (the use that referentialists are worried about), there is no property of being identical to the object of the speaker's perceptual intention to saturate the second argument position, and thus no complete proposition expressed. It also satisfies the referentialist's other intuition about referential uses: that in such uses the referent of the complex demonstrative is tracked across possible worlds. On King's story, the truth value of a complex demonstrative utterance backed by a perceptual intention will, in any possible world, depend on the object of the speaker's intention in the context of utterance.

That it can handle non-referential uses at the same time as accommodating the fundamental intuitions upon which the referentialist bases his story seems to be strong evidence in favour of King's account. Unfortunately, not only is King's account a paradigmatic example of Ptolemaic semantics, it is also empirically inadequate. Any story on which complex demonstratives are quantificational is going to have to show that they enter into scopal interactions with other quantifiers, modal operators etc. Views in the literature are radically divided on this, with the dividing line not surprisingly running between referentialists and quantificationalists. For Dever (2001), for instance, complex demonstratives cannot enter into scopal relations, whereas for King (2001) they can. Dever of course excludes King's QI and NS uses from the class of complex demonstratives, and with NS uses in particular, we've already seen what seems to be a scope ambiguity involving a complex demonstrative: (9) has two readings, one on which the complex demonstrative falls within the scope of the quantifier 'each division', and one on which the scopal relation is reversed. But

---

<sup>5</sup> I'm ignoring the past tense here simply for ease of presentation.

what of those uses of complex demonstratives which are supported by perceptual intentions, i.e. those uses at the heart of the referentialist's story? As Dever points out, cases in which complex demonstratives occur in sentences with uncontroversial quantifiers are not going to be of much help here. To see why, consider (13):

13) Every woman in the room likes that man over there.

On King's story, this sentence should equate to two sets of truth conditions which will look roughly like (14) and (15):

14) For all  $x$  such that  $x$  is a woman in the room the properties of being a man over there and being identical to  $a$  are uniquely jointly instantiated in  $\langle w, t \rangle$  by  $y$  and  $x$  likes  $y$ .

15) The properties of being a man over there and being identical to  $a$  are uniquely jointly instantiated in  $\langle w, t \rangle$  by  $y$  and for all  $x$  such that  $x$  is a woman in the room  $x$  likes  $y$ .

But, although these sets of truth conditions may differ in their scope relations, their truth values will covary across contexts. There is thus no way to establish whether (13) really does equate to (14) and (15). Sentences in which complex demonstratives co-occur with verbs of propositional attitude, however, are a different matter. Dever asks us to consider the example in (16):

16) Albert believes that upright citizen is a spy.

For King there should be a reading of (16) on which Albert holds a contradictory belief, i.e. the belief that the property of being an upright citizen and being  $a$  are uniquely jointly instantiated in  $\langle w, t \rangle$  by  $x$  and  $x$  is a spy. But there just doesn't seem to be any such reading.

For both King and Lepore and Ludwig it is important that such readings should be available, and both offer examples in which such narrow scope readings of perceptually-used complex demonstratives relative to verbs of propositional attitude are supposedly available. Both, however, seem to me to be examples of something

else. King asks us to imagine a party at which Alan has just been named CEO of Chanticleer. Sherry, a Chanticleer employee who believes Alan hates her, arrives at the party to hear the bad news. Two other party-goers are in conversation, when one looks over to Sherry and, seeing that she is looking very glum, asks why. The other, pointing at Alan, replies:

17) Sherry believes that guy who was just named CEO of Chanticleer hates her.

King then mounts an argument that runs like this: the utterance of (17) is an explanation of Sherry's behaviour; if the complex demonstrative in (17) were interpreted referentially (or as taking wide scope) (17) would not be an explanation of Sherry's behaviour, since she has long believed that Alan hates her without wasting any time moping about it; it must therefore be the case that the complex demonstrative is interpreted as taking narrow scope with respect to the belief operator, thus placing the property of being just named CEO of Chanticleer within Sherry's belief.

But this cannot be the reason that (17) is taken as an explanation of Sherry's behaviour. Consider a slight variant on the context above: imagine that both speaker and hearer know that Alan has *not* been appointed CEO of Chanticleer, but that Sherry misguidedly believes he has. In this situation the reading that King takes to be our natural interpretation of (17) will be plainly and straightforwardly true, since Sherry does hold the belief that the properties of being just named CEO of Chanticleer and being Alan are *uniquely instantiated* in *x* and *x* hates her. Yet an utterance of (17) in such a context seems at best highly infelicitous. Given this, we would do well to look elsewhere for a story on why an utterance of (17) acts as an explanation of Sherry's behaviour, and we don't have far to look. Let's assume that, on the evidence above, (17) expresses a singular proposition about Alan. Since it appears at a point in the conversation at which the hearer is expecting an answer to his question, he will interpret the utterance of (17) in that light; in other words, he will ask himself why Sherry's belief that Alan hates her is making her glum. Now there is one property of Alan that the speaker has made contextually highly salient by his choice of singular expression: the property of being the newly-appointed CEO of Chanticleer. Putting together the premise that Sherry believes that Alan hates her and the premise that

Alan is the newly-appointed CEO of Chanticleer, along with various other premises such as that Sherry works for Chanticleer, it shouldn't take the hearer long to reach a conclusion about the causes of Sherry's glumness. (17) can thus act as an explanation of Sherry's behaviour without any need for a narrow scope interpretation.

Lepore and Ludwig try another tack. They ask us to consider Tom, Mary and Mary's companion in a restaurant. Tom leans over to Mary and, pointing at a waiter in white sneakers, says 'that man wearing white sneakers is a good waiter'. Mary, who has failed to hear a part of Tom's utterance but who sees that the waiter Tom is talking about is wearing Nike sneakers, turns to her companion and says 'Tom believes that man wearing Nike sneakers is a good waiter', to which her companion, who has better hearing, replies:

18) No, he thinks that man wearing *white* sneakers is a good waiter.

Lepore and Ludwig's argument then runs like this: if the complex demonstrative in (18) is analysed referentially (or as taking wide scope), Mary and her companion are both attributing the same belief to Tom, i.e. the belief of the man in question that he is a good waiter; since they are in disagreement, this cannot be the intended interpretation; the complex demonstrative in (18) must therefore be interpreted as taking narrow scope in relation to the belief operator, thus placing the property of being a man wearing white sneakers within Tom's belief. But this argument is based on a false assumption: that, if there is a disagreement between Mary and her companion, it must be a disagreement on the content of Tom's belief. Consider the following example, adapted from Carston and Noh (1995):

19) A: We went to the zoo and saw the hippopotamuses.

B: No, we went to the zoo and saw the hippopotami.

In (19) B is disagreeing with A, but her disagreement is not with the content of A's utterance, it is, rather, with the linguistic form A has used to express that content: the negation in (19) is metalinguistic. This type of analysis would seem to extend very naturally to (18): what Mary is taking issue with is not the content of her companion's utterance, which is simply the attribution of a singular belief to Tom, but the linguistic



form her companion has used to express that content, which inaccurately mirrors the linguistic form used by Tom to express his own belief. Given the availability of such an analysis, this example seems to offer no firm evidence of scope interaction between a complex demonstrative and a verb of propositional attitude.

Even clearer than propositional attitude contexts are sentences in which complex demonstratives, used perceptually, co-occur with negation. Consider (20):

20) That policeman is not John's brother.

uttered while pointing, say, at a particular policeman. On any quantificational account, (20) should admit of two interpretations. Sticking with King's story, those two interpretations will look something like:

21) The properties of being a and being a policeman are uniquely jointly instantiated in  $\langle w, t \rangle$  by  $x$ , and it is not the case that  $x$  is John's brother.

22) It is not the case that the properties of being a and being a policeman are uniquely jointly instantiated in  $\langle w, t \rangle$  by  $x$ , and  $x$  is John's brother.

In other words, there should be a reading of (20) on which it will be true just in case there is nothing which is a unique policeman who is a here and now and which is John's brother. But there doesn't seem to be any such reading: this is, after all, a reading which would be true in the circumstance that the individual who is the object of the speaker's intention *is* John's brother, just so long as he is not a policeman. King mounts a defence against this kind of argument by invoking pragmatic processes: in perception-based cases, speakers are interested in getting their hearers to pick out the objects of their perceptual intentions; to this end, they will pick, as the nominal, a predicate which they believe their intended referent to satisfy, and which will thus help their hearers towards identifying the intended referent; given this, they are not going to want their utterances to be interpreted in such a way that they are true merely by virtue of their referents not satisfying the nominal predicate. The key point of this argument is that, with perceptual uses, narrow scope readings of complex demonstratives relative to negation are pragmatically blocked. It seems to me, however, that such readings are not so much blocked as absent. Consider (23):

23) There is a flag hanging out of every window.

The sentence in (23) clearly has two scopal readings, one of which, the reading on which 'a flag' takes scope over 'every window', is pragmatically blocked in most contexts. But there are two things to note about (23): firstly, although this reading is blocked, we can nevertheless make it out, and secondly, it is possible to manipulate the context so that this reading becomes the natural one: suppose we are discussing a factory across the road which prints huge flags and dries them after printing by hanging them out of the window; the flags are often so large that, when drying, they hang out of four of five windows, but this time they've excelled themselves; wanting you to come and have a look, I utter (23). In such a context, the wide-scope existential reading will be at least much less strongly suppressed than in a neutral context. These two properties, the properties of being discernible and de-suppressible by contextual manipulation, are typical of pragmatically suppressed scopal readings. Yet the supposed suppressed reading of (20) has neither of these properties. Given this, the only reason to suppose that there is such a reading is if one's theory requires it.

So where does our discussion of the referentialist and quantificationalist strategies leave us? The referentialist seems to have a lot that is right to say about referential uses of complex demonstratives; in particular, as the discussion of scope above demonstrates, she is right to claim that, on referential uses, complex demonstratives contribute nothing but their referents to propositional content. However, she has nothing to say on non-referential uses which, as I hope I have convinced you above, are genuine uses of complex demonstratives. The quantificationalist, on the other hand, has a lot that is right to say about non-referential uses (I shall expand on just how much of the quantificationalist's story is right later in this chapter), but, again, given the scopal data above, has serious difficulties when trying to account for referential uses.

### 6.3 Hybrid accounts

If referential uses are truth-conditionally singular, non-referential uses are truth-conditionally general, and complex demonstratives are not ambiguous, what kind of

story is there left to tell? Maybe the way forward is an account on which complex demonstratives are hybrids of referring terms and quantifiers. Versions of this sort of account have been advocated by Richard (1993) and Neale (1999)<sup>6</sup>. For Richard, complex demonstratives are what he calls *articulated terms*: they introduce not only their referent to propositional content, but also some quasi-Fregean way of thinking of that referent, as expressed by the nominal. Quite what the truth conditions of a complex demonstrative utterance will be on this story is slightly obscure: Richard suggests that an utterance of 'that F is G' will be truth-conditionally equivalent to an appropriately related utterance of 'that is F and that is G', in which case his story is not really hybrid in the sense we are interested in, and will fall at the same hurdles as, for instance, Dever's account (plus some others). The syntactic structures he assigns to articulated terms, however, suggest that maybe his intention is to claim that complex demonstrative utterances give rise to two parallel sets of truth conditions, one singular and one general.

Whether this is Richard's view or not, it is clearly the view of Neale (1999). For Neale, the lexical semantics of a complex demonstrative consist in a series of instructions: initially, the hearer of 'that F is G' is instructed to build the descriptive proposition in (24):

24) [the x: s is indicating x & Fx] Gx<sup>7</sup>

He is then instructed to find whichever object is the x such that s is indicating x and x is F in the context, and build a second, singular proposition containing that object:

25) Gα<sup>8</sup>

The thrust of Neale's story is thus that complex demonstratives are not referential *or* quantificational, they are referential *and* quantificational. In tandem with this two-proposition account, Neale tells a story about contextual weighting: in most contexts

---

<sup>6</sup> Of course one could, as, for instance, Dever does, see Lepore and Ludwig's account as hybrid, and even perhaps King's, since they posit both quantificational and referential aspects to the semantics of complex demonstratives. On both these stories, however, the referential aspects are embedded within a quantificational structure. They thus seem to be fundamentally quantificationalist.

<sup>7</sup> Neale (1999, p.67)

<sup>8</sup> Neale (1999, p.67)

it will be the proposition in (25) which will carry what he calls the contextual weight, but in some 'exceptional circumstances'<sup>9</sup> it will be the proposition in (24).

This story will not, however, do the work we need. And, to be fair to Neale, it is not intended to do so, since, even in those contexts in which it is the descriptive proposition that carries the contextual weight, Neale is still envisaging that there should be a perceived object towards which the speaker has some referential intention; he is, in essence, giving an account on which the kinds of uses of demonstratives which lie at the heart of the referentialist story can be interpreted (predominantly) descriptively, rather than one that can account for the purely descriptive uses raised by King. On King's descriptive uses, there is no object towards which the speaker has any referential intention; when Scott uses the expression 'that hominid who discovered how to start fires' he has no individual in mind about whom he wishes to be understood to be talking, rather he wants to talk about whoever is the hominid who discovered how to start fires. In this kind of situation, Neale's first proposition will be false, and his second proposition will simply not get off the ground: since there is no  $x$  such that  $s$  is indicating  $x$ , it is, a fortiori, not the case that ((there is a unique  $x$  such that  $s$  is indicating  $x$  and  $x$  is  $F$ ) and  $x$  is  $G$ ); and, since nothing satisfies the descriptive material in the general proposition, there is no object for the singular proposition to be about.

We must therefore look for another kind of hybrid story, and it is the business of the rest of this chapter to outline and defend an account of this sort.

#### 6.4 Complex demonstratives and communication

All the accounts we have looked at so far agree, either tacitly or overtly, on one underlying principle: that, whatever the meaning of complex demonstratives may be, it must be analysed along more than one dimension. For character-theorists like Braun (1994) and Borg (2000), a satisfactory account must recognise the different contributions complex demonstratives make to character and content; for quantificationalists like King (2001), complex demonstratives are semantically mandated to introduce into propositional content not just the property expressed by their nominal, but also a further property determined by context; and for hybridists

---

<sup>9</sup> Neale (1999, p.68)

like Richard (1993) and Neale (1999), complex demonstrative contribute both the property expressed by their nominal and their referent to propositional content, albeit, on Neale's account at least, to the content of different propositions. All of these different two-tier stories on complex demonstratives are expressions of the same underlying intuition: that understanding a complex demonstrative involves both the linguistically-given nominal and some extra non-linguistically-given (i.e. contextually determined) element; when we understand 'that F', we understand firstly that what is being talked about is an F, and secondly that it is a particular F being talked about, and, in order to establish *which* F, we need to make use of something other than the nominal. I have intentionally attempted to leave this intuition very vague: it can, as witnessed by the variety of accounts discussed above, be developed in widely differing ways. But it nonetheless expresses what is possibly the only common ground between all the accounts currently available. Given this, I suggest we should take it seriously.

I propose, then, that we start from the following working hypothesis: 'that F' is a tool that speakers use to talk about particular Fs. The tool works by indicating two things: (i) that the hearer is intending to talk about a particular F; and (ii) that being an F, in the context of utterance, is not the only way in which the speaker is thinking of her intended F. On the face of it, this kind of view seems to have much in common with the quantificational account of King (2001). Indeed, although the account I shall advocate diverges from King's in at least one fundamental respect, there is a fair amount of common ground. Given this, I would like, before progressing to a more articulated formulation of my own account, to highlight a point, in addition to those already discussed above, at which King's story seems to get things wrong. An examination of this shortcoming will, I hope, point us in the direction of a more complete analysis.

For King, it is typical of uses of complex demonstratives backed by descriptive intentions that the property which saturates the second argument position, the argument position reserved for first-order properties determined by speaker intention, is redundant. He asks us to imagine Danielle who knows, on purely general grounds, that there is currently one and only one person swimming across Lake Tahoe. Intending to talk about whoever it is that is currently swimming across Lake Tahoe, Danielle utters:

26) That person swimming across Lake Tahoe now must be cold.

On King's analysis, the first argument position is, as ever, saturated by the property expressed by the nominal, i.e. the property of being a person swimming across Lake Tahoe now. But, for King, that is also the content of the descriptive intention which backs up Danielle's use of the complex demonstrative: her intention is to talk about whoever has the property expressed by the nominal. It is therefore the case that the second argument position is saturated with just the same property as saturates the first argument position. The speaker's intention is thus redundant<sup>10</sup>.

On the rough sketch of the communicative role of complex demonstratives given above, this is at least an undesirable result: what I have suggested is that complex demonstratives are communicative tools designed for a particular purpose, that purpose being to talk about particular members of the nominal class. If the speaker only has one way of thinking of whatever she wants to talk about, then it would seem that, on the picture above, a complex demonstrative is the wrong tool for the job. I accept, however, that if this were the only reason to resist the redundancy in King's account, it would not be reason enough. But there is some empirical evidence to suggest we might do well to think again. Consider an example that seems to be parallel to King's example in (27):

27) That oldest man in the world must be worried about mortality.

Here, as in King's example, the speaker can, self-evidently, believe on purely general grounds that there is one and only one oldest man in the world, given any story on the semantics of superlatives<sup>11</sup>. And here, again, as in King's example, the speaker's intention seems redundant, since, ex hypothesi, she wants to talk about whoever is the oldest man in the world. And yet (27) seems infelicitous in a way that (26) isn't. Why might that be? What distinction can we draw between (26) and (27) that might explain why one is felicitous and the other not?

---

<sup>10</sup> That is not, of course, to say that the speaker's intention must always be redundant on such uses, just that this can be, and standardly is, the case.

<sup>11</sup> If you feel that the points I make about this example are in some way special to superlatives, consider 'that winner of the race...', 'that first day of next year ...' etc.

Another type of evidence seems to point in the direction of a plausible answer. Take the sentence in (28):

28) That dog with three legs is called 'Lucky'.

I want to outline two possible contexts in which (28) might be uttered: in the first context, Janet and John are standing in a room full of dogs, all but one of which have the standard canine allocation of legs. Janet turns to John and utters (28). In the second context Janet and John are again in a room full of dogs, but now all of the dogs are three-legged. Pointing at a particular dog across the room Janet utters (28). There seems to be an intuitive difference between the work done by the nominal 'dog with three legs' in these two contexts: in the first context, Janet is using the nominal not only to pick out which kind of thing she's talking about (she's talking about a dog), but also which individual within that kind she's talking about (the one with three legs). In the second context, by contrast, she is simply using the nominal to indicate which kind of thing she's talking about (she's talking about a dog with three legs), and John must turn to other, non-linguistic factors to establish which particular individual within this kind Janet is talking about (Janet's demonstration is going to give him a big clue).

What might the distinction between Janet and John's two contexts tell us about complex demonstratives? Let me suggest an answer: saturating King's second parameter doesn't always have to be a purely non-linguistic matter; sometimes the speaker can give the hearer information via her utterance that allows him either fully or partially to saturate this parameter<sup>12</sup>. So, returning to (28), in the first context, the property expressed by 'dog', not 'dog with three legs', will saturate the first parameter (the parameter we've so far been thinking of as the nominal parameter), and the property of having three legs will go some way towards saturating the second parameter. In a context such as this, however, where speaker and hearer are in an immediate perceptual relation with the object the speaker is intending to talk about, perceptual/causal properties of the referent will no doubt also enter into the saturation

---

<sup>12</sup> My talk of saturating parameters is not intended to indicate that I am committed to any of the details of King's story; I am simply here using this terminology for ease of presentation. I shall make clear how these conclusions fit within my own account later in the chapter.

of the second parameter. Compare an utterance of (28) in the first context to an utterance of (9), repeated as (29):

29) That professor who brought in the biggest grant in each division will be honored.

for which linguistically given material ('who brought in the biggest grant in each division') will entirely saturate the second argument position.

In the second context above, the nominal of (28), as we've already seen, only goes to identify the type of thing that Janet is intending to talk about; it is left to contextual indicators to determine which individual of that type Janet wants to talk about. Translating this into King's terms, the property expressed by 'dog with three legs' saturates just the first parameter, with the second parameter left entirely to context, albeit a context manipulated by Janet's demonstration.

How does this story help us with the data we started out with, the infelicity of (27) in comparison to (26)? If we reanalyse (26) along the lines just sketched, it would seem natural to say that the kind of thing Danielle is intending to talk about is a person. And which person? Well, the one swimming across Lake Tahoe now. In other words, the first parameter is saturated by the property of being a person, while the second is saturated by the property of being swimming across Lake Tahoe now. Both parameters are thus saturated. But what of (27)? In (27) there are no two properties that are separable in this way: the property of being the oldest man in the world is not the property of being both the oldest man and in the world. It is this, I suggest, that makes (27) infelicitous. Why? Because, contrary to King's account, the parameter which corresponds, on his story, to speaker intention, cannot be redundant: complex demonstratives are tools for picking things out via two or more distinct routes, and uses of complex demonstratives for which, as in (27), there is only one way of thinking of the speaker's intended referent, are thereby rendered infelicitous.

There are two extra pieces of evidence to suggest that this sort of story is along the right lines. Firstly, consider (30):

30) That person who is the oldest man in the world must be worried about mortality.



To my ear, (30) is less infelicitous than (27). On the accounts we have looked at so far in this chapter, all of which treat the nominal as a whole unit, there is no obvious explanation of this difference. On the kind of story I am sketching, however, this contrast is to be expected: in (30), unlike in (27), there are two separable properties (the property of being a person and the property of being the oldest man in the world) which can be prised apart to saturate both of King's first two parameters.

Secondly, imagine that on Tuesday Janet is reading the newspaper and sees an article about the oldest man in the world which she discusses with John. On Wednesday, having mulled things over in the last twenty-four hours, John utters (27) to Janet. Again this seems less infelicitous than does (27) uttered in a context without such a background. Why? Because in this context there is another 'route in' to the person John is intending to talk about: he's the person they were talking about yesterday. The second parameter thus doesn't go unfilled. There still does seem to me to be a certain oddness to an utterance of (27) even in this context, albeit less than before, and this too can be explained in the terms laid out above: complex demonstratives are designed in such a way that the hearer interpreting a complex demonstrative first goes to a class to which the thing talked about belongs, and then goes beyond that to discover which member of the class is being talked about. But of course for any utterance of (27), the nominal property will determine a unique individual, thus leaving the second parameter, although now filled, not doing any semantic work over and above the work already done by the nominal. Again this seems to go against the grain of complex demonstratives, although not as seriously as where there is no property to saturate the second parameter at all. Given this kind of data, it would seem helpful to introduce terminology to distinguish that part of the nominal which goes to saturate King's first parameter, the part of the nominal which makes a complex demonstrative a complex demonstrative, and that optional part of the nominal which can go towards saturating King's second parameter. I propose to call the former a *nominal sortal* and the latter a *nominal individuator*.

Let me then briefly summarise the kind of picture that the evidence above seems to point towards: complex demonstratives are communicative tools of a particular kind, they are tools for talking about individuals; given their meaning, they indicate that the individual in question is being thought of in at least two ways, firstly

as a member of a particular class and then via some other route or routes that distinguish(es) them from all other members of that class. This is, of course, no more than the vaguest formulation of the outline of an analysis of complex demonstratives. It is the purpose of the next section of the chapter to firm up the proposed analysis.

## 6.5 A semantics for complex demonstratives

The analysis I want to propose for the semantics of complex demonstratives is not itself complex, nor should it be surprising, given the analyses of other singular expressions which I have put forward in the last three chapters. The central claim I want to make is, of course, that complex demonstratives are marked by their encoded meaning as individual concept communicators. In what respect, however, does the encoded meaning of complex demonstratives differ from that of other singular expressions? Perhaps the clearest way to answer this question is by comparing the account of complex demonstratives which I favour with that of definite descriptions presented in the last chapter. In that chapter I argued that a definite description 'the F' encodes, over and above its status as an individual concept communicator, the property of being a unique F in some pragmatically-determined context. The nominal property encoded by a definite description thus radically constrains a hearer's search for an appropriate individual concept: since the property is uniquely instantiated (in the appropriate context) there will be only a very limited set of individual concepts available for consideration by the hearer. The meaning encoded by a complex demonstrative, by contrast, leaves the field wide open: a speaker who uses a complex demonstrative 'that F', where F is taken to be a nominal sortal, indicates that the proposition which constitutes the content of her informative intention contains an individual concept associated with the information *x is F*; beyond that, the determination of which is the appropriate individual concept is left entirely up to pragmatic inference. In the last chapter I pointed out that there is a contrast between the sentences in (31) and (32):

31) The table is very beautiful

32) That table is very beautiful

Whereas (31), uttered in a context in which there are many tables, seems at least infelicitous, (32), uttered in the same context, does not. This falls out naturally from the above analysis: the property expressed by the nominal of a definite description must be uniquely instantiated in the appropriate context, thus using ‘the table’ in a context in which there are many tables leads to infelicity. The property expressed by the nominal of a complex demonstrative, by contrast, does not have to be uniquely denoting, thus an utterance of (32) in the same context is entirely felicitous, although the hearer of (32) may require further information before he is able to reach the intended interpretation.

It seems, therefore, that the nominal sortal of a complex demonstrative need not be uniquely denoting. Indeed we can go further: given the sort of examples we saw in the previous section, it seems that the nominal sortal of a complex demonstrative *cannot* be uniquely denoting. Thus the meaning encoded by a complex demonstrative is in some sense the mirror image of that encoded by a definite description: whereas the nominal of a definite description must be uniquely denoting, that of a complex demonstrative must *not* be uniquely denoting. Putting this in the language of individual concepts, whereas a definite description ‘the F’ exploits, by dint of its linguistic meaning, the property of being a unique F in order to guide a hearer to the intended interpretation, a complex demonstrative ‘that F’ exploits the property of being a non-unique F to guide the hearer to an individual concept.

How does this idea tie in with the talk of parameters and saturation which, following King, we have so far been employing? The individual concept model offers a natural interpretation of such talk. Given the semantics outlined above, a speaker who uses a complex demonstrative ‘that F’ indicates firstly that the proposition she wishes to explicate contains an individual concept, and secondly that that individual concept is associated with the information *x is F* (non-uniquely). Now, given that the property of being F is not, in the appropriate context, uniquely-identifying, the hearer knows one more thing: that as an F is not the only way in which the speaker is thinking of that which she wishes to talk about; that, in other words, the individual concept which forms a constituent of the representation which the speaker wishes to explicate contains more information than merely *x is F*<sup>13</sup>. Were it not to contain further information, then it would fail to be an individual concept, since the property

---

<sup>13</sup> Of course this extra information may be of different types, perceptual, propositional etc.

of being an F does not single out an individual. Here we find a corollary of King's parameters: the first parameter represents the property expressed by the nominal sortal, while the second parameter is simply a reflection of the fact that the speaker who uses a complex demonstrative is thereby indicating that she has other ways of thinking of that which she wishes to talk about than merely as an F. Such an analysis also offers an explanation of why King's second parameter, contra King himself, cannot be redundant: for the second parameter to be redundant would, on this analysis, be for the speaker to have no way of thinking of that which she wishes to talk about other than as a satisfier of the nominal sortal. Since the nominal sortal is non-uniquely-identifying, this would be for the speaker to lack an individual concept of her intended designatum.

Where, then, does the nominal individuator fit into this picture? Again the answer is straightforward: beyond the non-uniquely-identifying nominal sortal, a speaker may use any number of different strategies to facilitate her hearer's interpretational task. She may, for instance, use an ostensive gesture or she may provide an extra linguistic clue; she may, that is, use a nominal individuator. The nominal individuator thus serves to bring into the context a further property which the speaker exploits to further her communicative ends, a further property which she indicates is associated with the individual concept she wishes to communicate.

In chapter four I discussed at some length whether the properties encoded by indexicals and simple demonstratives serve a semantic or pragmatic purpose; whether, in other words, they constrain reference or merely guide a hearer's interpretation. I concluded, given certain examples, that they serve a merely pragmatic purpose, that an indexical or simple demonstrative can refer to individuals which do not satisfy their encoded properties. What of the nominal properties of complex demonstratives? It seems that just the same arguments as were used in chapter four carry over to the case of complex demonstratives. In that chapter I argued that (33) was an entirely natural (and literal) vehicle for the proposition that James Barry was really a woman.

33) You know, when he was laid out for his funeral it turned out he was really a woman.

Consider, then, Bert and Ernie who are standing in a gallery admiring a painting of Barry dressed in his usual male attire. Knowing Barry's true history, Bert points to the portrait and says:

34) You know, when that man was laid out for his funeral, it turned out he was really a woman

Again it seems that (34) expresses the same proposition as (33). Indeed it seems that the alternative utterance, on which we replace 'man' with 'woman', is at best anomalous:

35) You know, when that woman was laid out for her funeral, it turned out she was really a woman

It thus seems that, just as with indexicals and simple demonstratives, the property expressed by the nominal of a complex demonstrative does not constrain reference, but merely serves to guide a hearer to the appropriate individual concept.

What, then, are the truth-conditional predictions of my account of complex demonstratives? They will not come as a surprise: on the account I favour, complex demonstratives may give rise to either object-dependent or object-independent truth conditions according to speaker intention. The properties expressed by their nominals serve to guide hearers to the appropriate individual concept and need make no appearance in truth-conditional content. This is not, of course, to say that these properties *cannot* appear in truth-conditional content. It may well be, as it will be, for instance, in the envisaged utterances of (7), (8) and (9), that the intended interpretation of a complex demonstrative is a descriptive concept the external dimension of which is simply the property expressed by the conjunction of the nominal sortal and nominal individuator.

## 6.6 Defending the analysis

The analysis I am advocating is consistent with the data, particularly the scopal data, laid out above. However, I would like, in the last part of this chapter, to raise a couple

of potential objections to the account and to explain why I do not think these objections pose any genuine threat.

I am thinking, in particular, of objections might be levelled at the predictions my account makes on referential uses of complex demonstratives<sup>14</sup>. There are, first of all, the kind of entailment data exemplified by (3) and (4), repeated here as (36) and (37):

36) That green car is very old.

37) Some green car is very old.

On the kind of account I have sketched there is no actual entailment between these sentences, in that (36) could be used to explicate a proposition which did not entail (37): say, for instance, I know you are red-green colour blind and for some reason I do not want to use the word 'red'; I could then utter (36) and thereby explicate a proposition which will be true iff a particular red car is very old, a proposition which would not entail (37). Nevertheless there are two further points to be made: firstly, in chapter 2 I suggested that philosophical intuitions are largely biased towards particular idealised discourse contexts. In such a context, in which all relevant contextual facts are mutually manifest to speaker and hearer, it is hard to see how (36) could express a proposition which did not entail (37). And secondly, such a context is clearly far from unusual: in the vast majority of contexts, the combination of contextual and pragmatic factors will conspire to render explicatures on which (36) entails (37) the optimally relevant interpretations. Thus, although there is, on my account no actual entailment here, there is every reason why there should appear to be such an entailment.

Secondly, Lepore and Ludwig ask us to consider various situations in which the nominals of referentially-used complex demonstratives appear to be playing some role in truth-conditional content. Consider, for instance, (38) uttered while indicating a particular individual:

---

<sup>14</sup> The predictions I intend to discuss are also made by any other account on which referentially-used complex demonstratives contribute nothing but their referents to truth-conditional content.

38) Each woman in this room admires that man whom she sees at the podium<sup>15</sup>.

Their claim concerning this example is that, since there is an element within the complex demonstrative that is anaphoric on an element outside, the content of the nominal cannot drop out of propositional content. This would certainly weigh heavily against accounts, such as mine, on which referentially-used complex demonstratives contribute nothing but their referents to propositional content. I do not believe, however, that there is any problem with examples such as these. Firstly, it seems to me that the referential reading that Lepore and Ludwig are after here, where a particular individual is being picked out as ‘that man who she sees at the podium’, is markedly less natural than a descriptive reading, and of course there is no problem, on my account, with anaphoric relations being set up between a quantifier outside a non-referentially used complex demonstrative and a pronominal element within. There is, however, one way that we can get the anaphoric relation with a referential use, but without having to posit a quantificational semantics. Compare (38) with (39):

39) Each woman in this room admires John whom she sees at the podium.

To my ear (39) is just as (in)felicitous as (38), yet I presume that it would not force us into a quantificational account of proper names. Rather, we should analyse ‘whom she sees at the podium’ as a non-restrictive relative. The same account seems to carry naturally over to (38); the proposition expressed by (38) looks not like (40), but like (41):

40)  $[\forall(x): \text{Woman-in-this-room } (x)][\exists(y): \text{That } (y) \ \& \ \text{Man } (y) \ \& \ \text{Sees-at-the-podium } (x, y)] (\text{Admires } (x, y))$

41)  $[\forall(x): \text{Woman-in-this-room } (x)](\text{Admires } (x, \underline{a}) \ \& \ \text{Sees-at-the-podium } (x, \underline{a}))^{16}$

Let me just mention, before closing, one final and related type of example that might seem to threaten my account. Again, Lepore and Ludwig ask us to consider:

---

<sup>15</sup> Lepore and Ludwig (2000, p.204).

<sup>16</sup> These logical forms are purely for illustrative purposes; I’m not intending to push any particular story on non-restrictive relatives.

42) That shark that took a swimmer off Flager beach last summer attacked him inside the sandbar.

Again there seems, *prima facie*, to be an LF interaction between an element within the complex demonstrative and an element without, once more suggesting that the properties expressed by the nominal appear at LF. Although there is not the space to work through the details of such examples, I see no reason to accept this conclusion<sup>17</sup>. Imagine that I'm right, that the interpretation of the complex demonstrative is simply the particular shark in question. The hearer of (42) is going to face a problem: how should he interpret 'him'? Of course there might be someone who is highly salient in the physical context (the speaker may, for instance be pointing at someone), in which case the swimmer the shark attacked last summer may simply fall out of the picture. But, in the absence of a physically salient individual, the utterance itself has made a particular individual contextually salient: the individual who the shark in question took off Flager beach last summer. It is entirely to be expected that the hearer should thus interpret 'him' as corresponding to this individual. The key point is that, just because the property of being a swimmer who this shark took off Flager beach last summer is contextually salient enough to provide an interpretation for 'him', does not entail that it appears anywhere else in propositional content.

---

<sup>17</sup> Of course there is a very natural descriptive interpretation available here, which Lepore and Ludwig are not interested in but which might well enter into an LF interaction.



## Chapter seven

### Conclusion

#### 7.0

Now that the analyses of the semantics and pragmatics of proper names, indexicals, demonstratives simple and complex and definite descriptions which I favour are in place, I want make a few very brief concluding remarks: firstly I want to take a step back and examine some of the overall claims about singular expressions which are inherent in this account; secondly I shall explore where my account fits in to the range of stories on singular expressions currently on offer; and finally I want to make a few tentative suggestions about the possible implications of the sort of account I favour for theories of natural language meaning.

#### 7.1 The class of singular expressions

What, then, are the overall claims about the class of singular expressions which follow from the specific analyses of chapters three to six? Perhaps the central claim is that the meaning of singular expressions should not be analysed in truth-conditional terms. How might one analyse the meaning of a singular expression in truth-conditional terms? There are two distinct possibilities: firstly one might take the meaning of a singular expression to be identical to that which the expression contributes to the truth-conditional content of an utterance in which it appears, a stance shared by, for instance, the Millian view on proper names and the Russellian view on definite descriptions. Secondly, however, one might take the meaning of a singular expression not to be identical to truth-conditional contribution, but rather to lay a constraint on the type of contribution, truth-conditionally individuated (at least in part), which that expression may make; one might, in other words, view the semantics of a particular singular expression along the lines Kaplan endorses for indexicals. It has been my aim to deny both of these theses: on the analyses I have put forward, the meaning, i.e. the encoded semantics, of a singular expression should be identified with a constraint on truth-conditional contribution, but a constraint which makes no mention, so to speak, of the truth-conditional type of that contribution. The

reason I favour this view essentially comes down to the following argument: there are, intuitively, both referential and descriptive uses of all types of singular expression; we should, following something like Recanati's Availability Principle, respect intuitions on truth-conditional content wherever possible; given the pragmatic machinery made available by Sperber and Wilson's Relevance Theory, it is possible to respect such intuitions in an analysis of singular expressions, i.e. it is possible to develop an account on which singular expressions make distinct types of contribution to the truth conditions of different utterances in which they appear. We should, therefore, prefer such an account.

There is, however, an obvious objection to any such story, one that has, I believe, led many to the view that no account along these lines can be right<sup>1</sup>. According to this objection, what I have proposed for singular expressions is essentially an ambiguity account. On my analysis, one and the same sentence containing a singular expression can, in different contexts, have either genuinely singular or genuinely quantificational truth conditions; what, so the argument might go, could be more ambiguous than that? And we should, of course, reject an ambiguity account wherever a viable alternative is available.

There is certainly something right in this objection: I *am* proposing an account on which singular expressions can give rise to very different kinds of truth conditions, so, if all that is being claimed is that, on my account, singular expressions are *truth-conditionally* ambiguous, then I'm going to have to put up my hands to that. I believe, however, that, on my story, singular expressions are not ambiguous in any way that the meaning-theorist should worry about. Let me briefly rehearse the familiar methodological claims about ambiguity to demonstrate why. According to Grice's Modified Occam's Razor, 'senses are not to be multiplied beyond necessity' (Grice (1967, p.47)). Given two competing theories about the interpretation of a linguistic expression, one of which posits an ambiguity and the other of which posits none, we should be ready to accept the latter theory. How might a theory without an ambiguity do the same work as one with an ambiguity? By handing everything beyond the univocal semantics over to pragmatics. The essence of Modified Occam's Razor, then, is that an expression should be viewed as ambiguous iff one cannot account for its semantic behaviour via a univocal semantics plus pragmatics. But this is just what

---

<sup>1</sup> Among those who clearly express such a view are Schiffer (1981) and Neale (1990).

my account of singular expressions does: it takes a single semantics for each singular expression type, a lexically encoded meaning which constrains the mapping from linguistic to conceptual representation, and leaves the rest to pragmatics. In any theoretically significant sense, therefore, the account I have proposed is not an ambiguity account<sup>2</sup>.

It is thus central to the account I have laid out that the encoded meaning of singular expressions should not be analysed in truth-conditional terms. How, then, should it be analysed? What I have suggested is that we should view singular expressions as encoding constraints on the mapping from language to thought. We can distinguish two dimensions of these constraints: that dimension shared by all singular expressions, i.e. the dimension via which we may individuate the class of singular expressions, and a further dimension, different for each type of singular expression. On the analysis I have proposed, the former dimension comes down to this: that singular expressions are tools for communicating individual concepts. Given the anatomy for individual concepts laid out in chapter two, this equates to the claim that the interpretation of singular expressions is constrained to concepts which are taken to denote individuals<sup>3</sup>. The latter dimension of the encoded meaning of singular expressions introduces a (more or less specific) property, a property which will serve to constrain the inferential process via which a hearer may reach the explicature of an utterance and will thus, indirectly, constrain the explicature itself. Which type of property this is will differ from one type of singular expression to another.

How best should we think of this kind of encoded constraint? I have, at various points throughout the thesis, explored the possibility that the notion of procedural meaning, due originally to Blakemore (1987) and now extensively used within the relevance-theoretic literature, may offer us an interesting way of approaching the encoded meaning of singular expressions. In the final analysis, I am not sure that the use of this label buys us a great deal. On the one hand, it is surely the case that, on the kind of account I advocate, singular expressions encode a fundamentally different kind of meaning from, for instance, predicates: while the former place constraints on the mapping from word to concept, the latter serve, it would appear, simply to point towards a particular concept. However, it is unclear to

---

<sup>2</sup> This argument is essentially a broadening of an argument presented by Recanati (1993) in defence of his account of the referential-attributive distinction in the use of definite descriptions.

<sup>3</sup> I am here leaving to one side the fictional use of singular expressions.

me that we would want to identify the kind of encoded meaning thus proposed for singular expressions with the sort of meaning encoded by non-truth-conditional expressions, upon the analysis of which Blakemore originally based her conceptual-procedural distinction. Whether there are any interesting parallels between the kind of meaning encoded by these two types of expression remains an open question, and, in the absence of any clear answer, one should perhaps be chary of using the procedural label for the constraints encoded by singular expressions.

## 7.2 What kind of account is this?

In this section I want to address two distinct questions: firstly, to what extent is the account of the semantics and pragmatics of singular expressions which I have proposed in this thesis an account of the same thing as mainstream philosophical and formal stories on referring and denoting? And secondly, to the extent that they are accounts of the same thing, where does my story fit in with dominant trends in the literature on reference?

I have so far talked of the analyses laid out in chapters three to six as analyses of the semantics and pragmatics of singular expressions. I have also argued at some length that the accounts I have outlined are, in certain key respects, to be preferred to other theories of the semantics of singular expressions within the philosophical and formal semantic literatures. Might it not be claimed, however, that my analyses are simply not the same sort of thing as the theories I have been comparing them with? While mainstream semantic theories are concerned with delivering truth conditions for natural language sentences and, via a specification of truth conditions, with accounting for logical relations between sentences, I have been concerned with the cognitive processes underlying natural language interpretation, what the inputs to those processes are and how their outputs are reached. These, so the claim might go, are simply different enterprises. Equally, those truth-conditional theories which operate within a Chomskian framework take their task to be modelling native speaker semantic competence; I, however, in focusing on interpretation, seem to have located my interest squarely within the area of performance.

Now there is something right about all this: my primary concerns have not been to pair natural language sentences with their truth conditions, nor to develop an

account of logical relations between natural language sentences. Rather I have indeed focused on the interpretation of utterances containing singular expressions. I do not believe, however, that these considerations support the claim that the account I have presented is simply a different sort of thing from mainstream truth-conditional accounts of reference and meaning. While my account is not directly concerned with pairing sentences with truth conditions and singular expressions with referents, it is nonetheless concerned with accounting for truth-conditional intuitions just as much as are truth-conditional accounts; indeed it takes such intuitions rather more seriously than do many truth-conditional accounts. Thus my account is an attempt to explain the same sort of data as directly truth-conditional accounts.

Secondly, my overall account does, as I shall discuss in the following section, have clear things to say about truth conditions and about the sort of logical relations which have been taken to hold between sentences. The derivation of truth-conditions will, of course, look very different on the kind of story I have told; in particular, the mapping from expression to truth-conditional contribution will pass through cognition. Nevertheless, the account I have proposed should provide the basis for a pairing of located utterances of sentences and truth conditions.

Finally, although I have focused on the interpretation of utterances, I have done so in an attempt to get at what is in essence a matter of competence. In chapter 2, I placed the cognitive dimension of my account within a modular view of the mind, a view on which language processing is taken to be a modular process. The encoded semantics which I have proposed and defended for singular expressions should, within this framework, be viewed as elements of the knowledge base of native speakers, i.e. as what is stored about the meanings of these expressions in a language module. As such, these encoded meanings comprise elements of native speaker competence, although they will, of course, contribute to on-line utterance interpretation. Just as for those semanticists who take themselves to be working within a Chomskian framework, therefore, my interest has been in capturing facts about the specifically semantic knowledge which underlies native speakers' linguistic abilities.

It seems to me, therefore, that, although the particular angle of my approach may differ from that of most within the broadly truth-conditional tradition, I have nonetheless offered an account which addresses those questions which are at the heart

of the truth-conditionalist's endeavour, questions concerning the semantic behaviour of particular natural language expressions. In the answers it gives to these questions, however, where does my account fit in with others currently on the market? There is an obvious sense in which it does not fit in with the dominant schools of thought: since it takes unambiguous expressions to be capable of giving rise to either object-dependent or object-independent truth-conditions in different contexts, it violates one of the fundamental tenets of philosophical/formal semantics, that a single expression-type must always correspond to the same type of propositional constituent. There are, nevertheless, significant points of contact between my own account and the broadly Millian and Fregean traditions. Firstly, the account I have proposed takes seriously the insight fundamental to Fregean epistemology that it is not possible to think of an individual other than under some mode of presentation. It is possible to view individual concepts, then, as a cognitive corollary of Fregean senses. But this equivalence is strictly limited. Firstly, the descriptive content of an individual concept does not, on the picture I have painted, invariably determine which individual the concept is a concept of, whereas it is fundamental to the Fregean picture that sense determines reference. Secondly, for Frege, senses not only serve as modes of presentation, they also serve as the meanings of linguistic expressions. On my account this link is severed: while singular expressions do contribute individual concepts to explicatures, there is no one-to-one pairing between linguistic expressions and individual concepts; we cannot, in other words, specify the meaning of a particular singular expression in terms of a particular individual concept. My account thus preserves one of the fundamental insights of Frege's epistemology, while rejecting his equivalence of modes of presentation with linguistic meanings.

As has been mentioned at other points in the thesis, the Fregean insight that it is not possible to think of an individual other than under a mode of presentation is not necessarily at odds with the Millian view: it is open to those Millians who take truth-conditional content not to be answerable to facts concerning cognitive significance to tell some extra story about cognitive modes of presentation. The account I have proposed, so far as referential uses of singular expressions are concerned, might be viewed as one possible cashing out of such a story. On my account, referentially-used singular expressions contribute nothing but their referents to truth conditions, given the truth-conditional profile of *de re* individual concepts. We thus end up, for these

uses, with the same truth-conditional story as the Millian champions for all uses of supposedly directly-referential expressions. My account fundamentally parts company with the Millian, of course, in the analysis of non-referential uses of certain types of singular expression<sup>4</sup>. It seems to me, therefore, that the account I have defended offers something of a middle way between Millianism and Fregeanism: while those uses which the Millian takes to be rigid do indeed come out as rigid on my account, nevertheless those features of the Fregean account which allow for solutions to the cognitive puzzles raised by co-reference and emptiness are, in essence, retained.

### 7.3 Singular expressions, individual concepts and semantic theory

I want finally to say a few words about the implications of the sort of account which I have proposed for semantic theory generally. In this thesis it has been my aim to explore what sort of account of singular expressions emerges if one takes the relevance-theoretic view of natural language meaning and interpretation seriously; if, that is, one allows for a radical pragmatic contribution to the level of what is said by an utterance. The account that has emerged addresses, as discussed above, those questions which are at the heart of philosophical semantics, but in doing so it offers some answers that are of a fundamentally different form from those standardly offered by philosophical semanticists.

Firstly, on the account I have proposed, the lexical meanings of singular expressions are analysed not in terms of the mapping from language to world, but in terms of the mapping from language to mind. The chain from language to world thus involves two links: that from language to mind and that from mind to world. My account also suggests that those distinctions between mental entities that are of significance to the encoded meaning of singular expressions may not be those distinctions which are significant to the truth conditions of mental representations. I have claimed that, whereas lexical meaning is sensitive to the distinction between individual and non-individual concepts, the truth-conditions of thought are sensitive to distinctions within the class of individual concepts, distinctions to which linguistic meaning is blind.

---

<sup>4</sup> Or at least it fundamentally parts company with the Millian who takes his theory to be a theory of the meanings of certain natural language expressions, rather than of the propositional content of certain types of uses of those expressions.

The account I have proposed thus delivers truth conditions for natural language sentences, although it does so via a two-step process. What of the other central concern of truth-conditional semantics, accounting for the logical relations between natural language sentences. On the relevance-theoretic view of semantics, and on my account of the semantics and pragmatics of singular expressions in particular, the data on logical relations comes out looking very different. Given the account I have proposed, logical relations, as least so far as they are determined by the meanings of singular expressions, are not the sort of thing that hold between sentences; rather they hold between the propositions which, on occasions of use, sentences may be used to explicate. Thus, on my analysis, there is no entailment between (1) and (2) and no contradiction between (3) and (4):

- 1) John went to market
- 2) Someone called 'John' went to market
- 3) She lives in Lake Tahoe
- 4) No female lives in Lake Tahoe

Why, then, do these logical relations appear to hold? Because, for the vast majority of contexts of use, the explication of (1) will entail the explication of (2) and the explication of (3) will contradict the explication of (4). These, however, are facts about the propositions which, on occasions of use, these sentences may be used to express, not about the linguistic meaning of the sentences themselves.

I started this thesis by suggesting that Frege's *Über sinn und bedeutung* had had an effect on twentieth century semantics which paralleled that of Ptolemy's *Almagest* on medieval astronomy: placing truth at the centre of meaning has led to the creation of any number of semantic eccentrics, epicycles and equants. I hope that in this thesis I have at least suggested a strategy for excising these pieces of Ptolemaic machinery by deposing truth and placing mind instead at the centre of meaning.



## **BIBLIOGRAPHY**

Adams, R. 1974. Theories of actuality. *Noûs* 8, 211-231.

Bach, K. 1987. *Thought and Reference*. Oxford: Clarendon Press.

Bach, K. 1992. Intentions and demonstrations. *Analysis* 52, 140-146.

Bach, K. 1999. The myth of the conventional implicature. *Linguistics and Philosophy* 22:4, 327-366.

Bach, K. 2002. Giorgione was so-called because of his name. *Philosophical Perspectives* 16, 73-103.

Berckmans, P. 1990. Demonstrative Utterances. *Philosophical Studies* 60, 281-295

Bezuidenhout, A. 1997. Pragmatics, semantic underdetermination and the referential/attributive distinction. *Mind* 106, 375-409.

Blakemore, D. 1987. *Semantic Constraints on Relevance*. Oxford: Blackwell.

Blakemore, D. 1989. Denial and contrast: a relevance theoretic analysis of *but*. *Linguistics and Philosophy* 12, 15-37.

Blakemore, D. 2000. Indicators and procedures: *nevertheless* and *but*. *Journal of Linguistics* 36, 463-86.

Block, N.. 1986. Advertisement for a semantics for psychology. In P. French et al., (eds.), *Midwest Studies in Philosophy* 10. Minneapolis: University of Minnesota Press.

Borg, E. 2000. Complex Demonstratives. *Philosophical Studies* 97, 225-44.

- Borg, E. Forthcoming. How to Say What You Mean: Unarticulated constituents and communication. Forthcoming in R. Elugardo & R. Stainton (eds) *Ellipsis and Non-sentential Speech*. Dordrecht: Kluwer
- Braun, D. 1993. Empty names. *Noûs* 27:4, 449-469.
- Braun, D. 1994. Structured Characters and Complex Demonstratives. *Philosophical Studies* 74, 193-219.
- Braun, D. 1996. Demonstratives and their linguistic meaning. *Noûs* 30:2, 145-173.
- Breheny, R. 1999. *Context Dependence and Procedural Meaning: The Semantics of Definites*. University of London Ph.D. thesis.
- Burge, T. 1973. Reference and proper names. *Journal of Philosophy* 70, 425-439.
- Carston, R. 1988. Implicature, explicature and truth-theoretic semantics. In Kempson (1988), 155-181.
- Carston, R. 1996/1997. Enrichment and loosening: complementary processes in deriving the proposition expressed? *UCL Working Papers in Linguistics* 8, 205-32; reprinted, 1997, in *Linguistische Berichte* 8, Special Issue on Pragmatics, 103-27.
- Carston, R. 1998. *Pragmatics and the explicit-implicit distinction*. University of London Ph.D. thesis.
- Carston, R. 2000. Explicature and semantics. *UCL Working Papers in Linguistics* 12, 1-44. Revised version to appear in S. Davis and B. Gillon (eds) *Semantics: a reader*. Oxford: OUP.

- Carston, R. 2002a. *Thoughts and utterances: the pragmatics of explicit communication*. Oxford: Blackwell.
- Carston, R. 2002b. Linguistic meaning, communicated meaning and cognitive pragmatics. *Mind and Language, Special Issue on Pragmatics & Cognitive Science* 17 (1/2), 127-48.
- Carston, R and E-J Noh. 1995. A truth-functional account of metalinguistic negation, with Evidence from Korean. *UCL Working Papers in Linguistics* 7, 1-26. Reprinted 1996 in *Language Sciences* 18, 485-504.
- Cohen, L. J. 1971. Some remarks on Grice's views about the logical particles of natural language. In Y. Bar-Hillel (ed.) *Pragmatics of Natural Language*, 50-68. Dordrecht: Reidel.
- Corazza, E. and J. Dokic. 1992. On the cognitive significance of indexicals. *Philosophical Studies* 66, 183-196.
- Corazza, E., W. Fish and J. Gorvett. 2002. Who is I? *Philosophical Studies* 107, 1-21.
- Dever, J. 2001. Complex Demonstratives. *Linguistics and Philosophy* 24, 271-330.
- Donnellan, K. 1966. Reference and definite descriptions. *The Philosophical Review* 75, 281-304.
- Donnellan, K. 1970. Proper names and identifying descriptions. *Synthese* 21, 335-358.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.
- Fodor, J. 1975. *The Language of Thought*. New York: Crowell.

- Fodor, J. 1983. *The Modularity of Mind*. Cambridge: MIT Press.
- Forbes, G. 1997. How much substitutivity? *Analysis* 57.2, 109-113.
- Forbes, G. 1999. Enlightened semantics for simple sentences. *Analysis* 59.2, 86-91.
- Frege, G. 1891. Function and concept. Translated in Geach and Black 1960, 21-41.
- Frege, G. 1893. Über sinn und bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik* 100, 25-50. Translated as 'On sense and reference' in Geach and Black 1960, 56-78. Reprinted in R. Harnish (ed.), *Basic Topics in the Philosophy of Language*, 142-160. New York: Harvester Wheatsheaf. Page references are to this edition.
- Geach, P. and M. Black. 1960. *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell.
- Geurts, B. 1997. Good news about the description theory of names. *Journal of Semantics* 14, 319-348.
- Grice, H. P. 1967. Logic and conversation. In *Studies in the way of words*. Cambridge: Harvard University Press, 22-40.
- Grice, H. P. 1987. Further Notes on Logic and Conversation. In *Studies in the way of words*. Cambridge: Harvard University Press, 41-57.
- Grice, P. 1969. Vacuous names. In D. Davidson and J. Hintikka (eds.), *Words and Objections*, 118-145. Dordrecht: Reidel.
- Heim, I. 1988. *The Semantics of Definite and Indefinite Noun Phrases*. New York: Garland.

- Higginbotham, J. 1988. Contexts, models and meanings: a note on the data of semantics. In Kempson (1988), 29-48.
- Iten, C. 1998. The meaning of *although*: a relevance theoretic account. *UCL Working Papers in Linguistics* 10, 81-108.
- Iten, C. 2000a. '*Non-truth-conditional*' meaning, relevance and concessives. PhD thesis, University of London.
- Iten, C. 2000b. *Although* revisited. *UCL Working Papers in Linguistics* 12, 65-96.
- Kaplan, D. 1978. Dthat. *Syntax and Semantics* 9, 221-243.
- Kaplan, D. 1989a. Demonstratives. In J. Almog et al. (eds.) *Themes from Kaplan*, 4 81-563. Oxford: OUP.
- Kaplan, D. 1989b. Afterthoughts. In J. Almog, H. Wettstein & J. Perry (eds.), *Themes from Kaplan*, 565-614. New York: OUP.
- Kempson, R. 1986. Definite NPs and context-dependence: a unified theory of anaphora. In T. Myers et al. (eds.) *Reasoning and Discourse Processes*. New York: Academic Press, 209-239.
- Kempson, R. 1988. *Mental Representations: The interface between language and reality*. Cambridge: Cambridge University Press.
- King, J. 1999. Are Complex 'That' Phrases Devices of Direct Reference? *Noûs* 33, 155-182.
- King, J. 2001. *Complex Demonstratives: a Quantificational Account*. Cambridge: MIT Press.

- Kripke, S. 1972. Naming and necessity. In D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, 253-355. Dordrecht: Reidel.
- Kripke, S. 1977. Speaker reference and semantic reference. In P. French, T. Uehling and H. Wettstein (eds.), *Contemporary Perspectives in the Philosophy of Language*, 6-27. Minneapolis: University of Minnesota Press.
- Larson, R. and G. Segal. 1995. *Knowledge of Meaning*. Cambridge: MIT Press.
- Lepore, E and K. Ludwig. 2000. The Semantics and Pragmatics of Complex Demonstratives. *Mind* 109, 199-240.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Loar, B. 1988. Social content and psychological content. In R. Grimm and P. Merrill (eds.), *Contents of Thoughts*. Tucson: University of Arizona Press.
- Longobardi, G. 1994. Reference and proper names: a theory of N-movement in syntax and logical form. *Linguistic Inquiry* 25, 609-665.
- Ludlow, P. and S. Neale. 1991. Indefinite descriptions: in defense of Russell. *Linguistics and Philosophy* 14, 171-202.
- McDowell, J. 1977. On the sense and reference of a proper name. *Mind* 86, 159-185.
- McGinn, C. 1981. The mechanism of reference. *Synthese* 49, 157-86.
- Mill, J. S. 1843. *A System of Logic*. New York: Harper Brothers.
- Moore, G. 1999. Saving substitutivity in simple sentences. *Analysis* 59.2, 91-105.

- Neale, S. 1990. *Descriptions*. Cambridge: MIT Press.
- Neale, S. 1993. Term Limits. In James E. Tomberlin (ed.) *Philosophical Perspectives* 7, *Language and Logic*, 89-123. Atascadero: Ridgeview.
- Neale, S. 1999. Coloring and composition. In K. Murasugi and R. Stainton (eds.), *Philosophy and Linguistics*, 35-82. Boulder: Westview.
- Nunberg, G. 1990. Indexicality in Contexts. Paper delivered at the conference on Philosophy and Cognitive Science, Cerisy-la-Salle, France, June 1990.
- Nunberg, G. 1993. Indexicality and Deixis. *Linguistics and Philosophy* 16, 1-43.
- Nunberg, G. 1995. Transfers of Meaning. *Journal of Semantics* 12, 109-132.
- Peacocke, C. 1975. Proper names, reference and rigid designation. In S. Blackburn (ed.) *Meaning, Reference, Necessity*. Cambridge: Cambridge University Press, 109-32.
- Peacocke, C. 1983. *Sense and Context: Experience, Thought and their Relations*. Oxford: Clarendon Press.
- Pelczar, M. and J. Rainsbury. 1998. The indexical character of names. *Synthese* 114, 293-317.
- Pelczar, M. 2001. Names as tokens and names as tools. *Synthese* 128, 133-155.
- Perry, J. 1997a. Reflexivity, indexicality and names. In W. Kunne, M. Anduschus & A. Newen (eds.) *Direct Reference, Indexicality and Propositional Attitudes*, 3-19. Stanford: CSLI-Cambridge University Press.

- Perry, J. 1997b. Indexicals and demonstratives. In B. Hale and C. Wright (eds.) *A Companion to the Philosophy of Language*, 586-612. Oxford: Blackwell.
- Perry, J. 2001. *Reference and Reflexivity*. Stanford: CSLI.
- Pitt, D. 2001. Alter egos and their names. Paper delivered at the 75<sup>th</sup> annual meeting of the Pacific Division of the American Philosophical Association.
- Powell, G. 1998. The deferred interpretation of indexicals and proper names. *UCL Working Papers in Linguistics* 10, 143-172.
- Powell, G. 2000. Compositionality, Innocence and the Interpretation of NPs. *UCL Working Papers in Linguistics* 12, 123-144.
- Powell, G. 2001. The Referential-Attributive Distinction: a Cognitive Account. *Pragmatics and Cognition* 9:1, 69-98.
- Powell, G. 2002. Underdetermination and the Principles of Semantic Theory. *Proceedings of the Aristotelian Society* 102:3, 271-278.
- Predelli, S. 1998. I am not here now. *Analysis* 58:2, 107-115.
- Predelli, S. 1999. Saul, Salmon, and Superman. *Analysis* 59.2, 113-116.
- Predelli, S. 2001a. Art, Bart and Superman. *Analysis* 61.4, 310-313.
- Predelli, S. 2001b. Names and character. *Philosophical Studies* 103, 145-163.
- Ramachandran, M. 1996. "The ambiguity thesis versus Kripke's defence of Russell". *Mind and Language* 11, 371-387.



- Recanati, F. 1989a. "Referential/attributive: a contextualist proposal". *Philosophical Studies* 56, 217-249.
- Recanati, F. 1989b/91. The pragmatics of what is said. *Mind and Language* 4, 295-329.  
Reprinted, 1991, in S. Davis (ed.), *Pragmatics: A Reader*. Oxford: OUP.
- Recanati, F. 1993. *Direct Reference: from language to thought*. Oxford: Blackwell.
- Recanati, F. 2002. Unarticulated Constituents. *Linguistics and Philosophy* 25, 299-345.
- Reichenbach, H. 1947. *Elements of Symbolic Logic*. New York: Free Press.
- Reimer, M. 1991. Do demonstrations have semantic significance? *Analysis* 51.4, 177-183.
- Reimer, M. 1992. Three views of demonstrative reference. *Synthese* 93, 373-402.
- Reimer, M. 1998. The Wettstein/Salmon debate: critique and resolution. *Pacific Philosophical Quarterly* 79, 130-151.
- Reimer, M. 2001. The problem of empty names. *Australasian Journal of Philosophy* 79.4, 491-506.
- Richard, M. 1993. Articulated Terms. In James E. Tomberlin (ed.) *Philosophical Perspectives* 7, *Language and Logic*, 207-230. Atascadero: Ridgeview.
- Rouchota, V. 1992. On the referential/attributive distinction. *Lingua* 87, 137-167.
- Rouchota, V. 1994. *The Semantics and Pragmatics of the Subjunctive in Modern Greek - a Relevance-Theoretic Approach*. University of London Ph.D. thesis.

- Russell, B. 1905. On denoting. *Mind* 14, 479-493.
- Russell, B. 1911. Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society* XI, 108-128.
- Sainsbury, M. 1995. Philosophical Logic. In A.C. Grayling (ed.), *Philosophy: a guide through the subject*, 61-122. Oxford: OUP.
- Salmon, N. 1986/91. *Frege's Puzzle*. Cambridge: The MIT Press.
- Salmon, N. 1998. Nonexistence. *Noûs* 32, 277-319.
- Saul, J. 1997. Substitution and simple sentences. *Analysis* 57.2, 102-108.
- Schiffer, S. 1981. Indexicals and the Theory of Reference. *Synthese* 49, 43-100.
- Schiffer, S. 1987. *Remnants of Meaning*. Cambridge: MIT Press.
- Segal, G. 1996. The modularity of theory of mind. In P. Carruthers and P. Smith (eds.) *Theories of theories of mind*. Cambridge: Cambridge University Press.
- Segal, G. 2001. Two theories of names. *Mind & Language* 16:5, 547-563.
- Sidelle, A. 1991. The answering machine paradox. *Canadian Journal of Philosophy* 21, 525-39.
- Smith, Q. 1989. The multiple use of indexicals. *Synthese* 78, 167-191.
- Sperber, D and D. Wilson. 1986/95. *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge: Harvard University Press. Second edition 1995.

- Sperber, D and D Wilson. 1998. Irony and relevance: a reply to Seto, Hamamoto and Yamanashi. In R. Carston and S Uchida (eds.), *Relevance Theory: Applications and Implications*. Amsterdam: John Benjamins.
- Stalnaker, R. 1976. Possible Worlds. *Noûs* 10, 65-75. Reprinted in M. Loux (ed.), 1979, *The Possible and the Actual*, 225-234. Ithaca: Cornell University Press.
- Stanley, J. 2000. Context and logical form. *Linguistics and Philosophy* 23, 391-434.
- Taylor, K. 2000. Emptiness without compromise: a referentialist semantics for empty names. In A. Everett and T. Hofweber (eds.), *Empty Names, Fiction, and the Puzzles of Non-Existence*, 17-36. Stanford: CSLI.
- Wettstein, H. 1981. Demonstrative reference and definite descriptions. *Philosophical Studies* 40, 241-257. Reprinted in Wettstein 1991, 35-49.
- Wettstein, H. 1984. How to bridge the gap between meaning and reference. *Synthese* 58, 63-84. Reprinted in Wettstein 1991, 67-85.
- Wettstein, H. 1986. Has semantics rested on a mistake? *The Journal of Philosophy* 83, 185-209. Reprinted in Wettstein 1991, 109-131.
- Wettstein, H. 1991. *Has Semantics Rested on a Mistake?* Stanford: Stanford University Press.
- White, S. 1982. Partial character and the language of thought. *Pacific Philosophical Quarterly* 63, 347-365.
- Wiggins, D. 1975. Identity, designation, essentialism and physicalism. *Philosophia* 5, 1-30.

Wilson, D. and D. Sperber. 1981. On Grice's theory of conversation. In P. Werth (ed.) *Conversation and Discourse*. London: Croom Helm, 155-178.

Wilson, D and D. Sperber. 1992. On verbal irony. *Lingua* 87, 53-76.

Wilson, D. and D. Sperber. 1993. Linguistic form and relevance. *Lingua* 90.1/2, 1-25.

Wilson, D and D. Sperber. 2002. Relevance Theory. *UCL Working Papers in Linguistics* 14, 249-287.