

# *The deferred interpretation of indexicals and proper names*<sup>\*</sup>

GEORGE POWELL

---

## Abstract

One of the central tenets of direct reference theory, that indexicals are semantically marked so as to give rise to singular propositions, has recently been challenged, by Geoffrey Nunberg among others, on the grounds that indexicals can, in fact, be shown to lead to the same range of interpretations as descriptions. In this paper I want to take a look at so-called *deferred* interpretations of indexicals and proper names from a cognitive perspective grounded in Relevance Theory, with the aim of exploring what rôles semantic decoding and pragmatic inference play in the retrieval of these interpretations.

## 1 Introduction

What do indexicals, expressions such as *I*, *you*, *here* and *now*, contribute to the propositions expressed by utterances in which they appear? What, for instance, does the expression *I* bring to a proposition? I imagine many of us, without having to give the question too much thought, would probably say something like: *I* just serves to pick out a particular person, it picks out the person who's speaking. So if I say *I hope I'm not boring you yet*, my use of *I* just picks out George Powell, in order that I can say something about him/me. And faced with most of the data this intuition seems pretty robust, so robust, in fact, that it provides the bedrock for accounts of the semantics of indexicals within the direct reference tradition. But is it really as robust as all that? Are there uses of indexicals which don't fit into this nice clear picture? Over the last few years there has been a growing number of voices claiming just this: that there are, in fact, a range of uses, involving so-called *deferred* reference, in which indexicals do not introduce particular individuals into the propositions expressed by the utterances in

---

<sup>\*</sup> I would like to thank Robyn Carston for her encouragement, her support and, above all, for her patience.

which they appear. Consider the following examples, taken from Nunberg (1993):

- (1) *Condemned prisoner*: I am traditionally allowed to order whatever I like for my last meal.
- (2) *Chess teacher giving an introductory lesson to a student who has just played 4. N X P ...*: According to all the textbooks, you often get in trouble with that move.
- (3) *Medical pathologist pointing to a spot on his own chest*: When a person is shot here we can usually conclude that it was not suicide.

In each of these examples it seems that, on at least one reading, an indexical is being used to introduce something other than an individual. Have a look at (1), for example. The claim is that it has two readings: a (nonsensical) reading on which *I* does indeed introduce an individual, and a preferred reading on which *I* introduces a definite description something like *the condemned prisoner*, used attributively in the sense of Donnellan (1966). In much the same way, *you* in (2) is taken to introduce a description something like *the player who plays 4. N X P...*, again attributively used. In (3) it's not entirely clear how we'd want to paraphrase the contribution *here* makes to the proposition expressed, but all we need be clear about is that it doesn't serve to introduce the particular spot on his own chest which the pathologist is pointing at, as our original intuition might suggest it should.

So where do these observations leave us? Should we abandon our original intuition entirely in the face of irreconcilable evidence? But if we do that how do we explain the undoubted strength of the intuition, or the fact that it can account for the majority of the data? It would surely be ill advised to ditch out of hand a theory which can do so much good work for us. Perhaps, then, we should just stick to our guns and assume that uses like those in (1), (2) and (3) are somehow deviant? It seems to me that this smacks of sticking our head in the sand; if our theory of indexicals is to be worth its salt it better have some kind of account for uses which, although perhaps rare, are certainly intelligible and attested. As I imagine the foregoing may have hinted, I hope to steer a middle course between these two approaches: I don't want to turn my back entirely on direct reference, but direct reference cannot hide from the data if it is to provide a satisfactory account of indexicals as they are used in natural language.

So far, then, I've introduced the notion of the deferred use of indexicals; how about the proper names I promised in the title? I don't want to say too much about the deferred use of proper names yet: I hope that, as we go on, it will become clear why the availability or unavailability of deferred readings for proper names is of central significance. As a rough first approximation, the issue looks like this: if indexicals are

alone among referring expressions in their capacity to give rise to deferred readings, then that seems like pretty good evidence that that capacity has something to do with the semantics of indexicals i.e. with their linguistically encoded meaning; if, on the other hand, proper names can give rise to deferred readings just as indexicals can, then we might well be looking towards pragmatics for an overall account of deferred use.

Before getting on to the meat of this paper, let me say a couple of words about how I propose to proceed. In the first half of the paper, I want to look at the competing camps in rather more detail, starting with direct reference-flavoured accounts and then turning to the challenges that those accounts have had to face. Then I want, essentially, to turn the question around: rather than approaching the issues raised by deferred reference from a philosophical perspective grounded in truth-conditional semantics, I want to see whether any interesting results fall out from adopting a cognition-based approach. My aim is to see whether looking at things from a new perspective might allow us to ask different (and perhaps more interesting) questions and so reach different (and perhaps more interesting) answers.

## 2 Indexicals, proper names and direct reference

At the heart of the theory of direct reference is the notion that there is a lexically encoded distinction between, on the one hand, definite descriptions and, on the other hand, directly referential terms such as indexicals and proper names. There are various ways of looking at this distinction but whichever way you look at it, it comes down to something like this: what descriptions contribute to a proposition is some version of a Fregean sense; what directly referential expressions contribute to a proposition just is their referent. Kaplan (1977) is careful to distinguish this from the claim that directly referential expressions lack descriptive content; it's just that what descriptive content they do have plays a fundamentally different role in the interpretation process. Instead of forming part of the proposition expressed, it simply serves to indicate a referent and then falls away. To put this another way, utterances containing descriptions give rise to general truth conditions, utterances containing directly referential expressions give rise to singular truth conditions.

To capture this insight, Kaplan (1977) introduces a distinction between two different types of meaning: character and content. I don't want to go into this no doubt familiar distinction in any great depth, but as a rough rule of thumb an expression's character can be equated with its linguistic meaning and its content with what it contributes to the proposition expressed. In Kaplan's terms, then, the key thing about directly referential expressions is that their content just is their referent. So what *I* or *John*, for instance,

contribute to a proposition just is whoever they're being used to refer to. The content of a description, by contrast, is some kind of conceptual complex, some kind of Fregean sense. And the character of directly referential expressions? For Kaplan, the feature that distinguishes indexicals from other expressions is that their character and content differ. Let's take *I* again: we've just seen that the content of *I* will be an individual; what is its character? Well it's going to be something like *the speaker of this utterance*. So the meaning of *I* on this analysis just serves to pick out the speaker of the particular utterance and then falls out of the picture.

There's some debate about the character of proper names: for Kaplan proper names have identical character and content. For others, such as Recanati (1993), proper names have differing character and content; they are, in other words, a form of indexical. We need not, however, be too concerned here with the character of directly referential expressions; it is, rather, their content, what they contribute to propositions, that lies at the heart of the debate over deferred reference.

### 3 Geoffrey Nunberg and the challenge of deferred reference

In a series of papers (e.g. Nunberg (1990) and Nunberg (1993)), Geoffrey Nunberg has raised objections to the approach outlined above which one of its leading proponents has called 'the most serious which the theory of direct reference has ever faced' (Recanati (1993, p301)). The main thrust of Nunberg's argument is simple: at the heart of direct reference is the claim that there is a semantically encoded distinction between directly referential expressions and descriptions which leads to a difference in truth conditions; if, therefore, it can be shown that indexicals are capable of giving rise to the same range of interpretations as descriptions, then the theory of direct reference will be seriously weakened, at least so far as indexicals are concerned. And, for Nunberg, examples such as (1), (2) and (3) show just this.

But what lessons does he draw from these data? Let's just remind ourselves of the two main planks of the direct reference account of indexicals: the meaning (character) of indexicals does not enter into the proposition expressed, in other words the meaning of indexicals is indicative; and what indexicals contribute to propositions (their content) just is an individual. Where do these claims stand in the face of Nunberg's examples? Well the first claim remains unaffected: Nunberg sees no reason to dispute the notion that the linguistic meaning of indexicals is indicative. There is, after all, abundant evidence to illustrate this point. To take just one example, Kaplan (1977) points out that, if *the speaker of this utterance* were the content as well as the character of *I*, then (4) would express a necessary truth:

(4) If no one were to utter this token, I would not exist.

Kaplan (1977, p520)

‘Beliefs such as [(4)]’, says Kaplan, ‘could make one a compulsive talker’.

The second claim of direct reference, on the other hand (the claim, that is, that indexicals are semantically marked so as to give rise to singular propositions), seems seriously threatened by Nunberg’s examples, if we accept that the propositions expressed by those examples are general<sup>1</sup>. The insight that underlies this is that there is no necessary link between the two claims of direct reference: just because the meaning of an expression is indicative, that doesn’t necessarily mean that what that expression contributes to a proposition just is an individual. It might contribute a property other than that encoded by its linguistic meaning. And it’s this insight that leads Nunberg to propose an alternative account of the semantics of indexicals, an account grounded in the notion of deferred reference.

The question Nunberg sets out to answer is, essentially, this: if indexicals can contribute properties other than those encoded by their linguistic meaning, through what mechanism do these properties get into propositions? And, in rough terms, the answer he gives is that they get in because indexicals have the semantics they do. Nunberg starts by looking at the meaning of *we*: he observes that, although it may be so interpreted<sup>2</sup>, *we* is not standardly used as the plural form of *I*. It is, rather, used to mean something like *the group of people instantiated by the speaker or speakers of the utterance* (Nunberg (1993, p7)). Now as Nunberg points out, what this really comes down to is the idea that the interpretation of *we* involves a two-stage (i.e. deferred) process: first the hearer has to find out who the speaker of the utterance is and secondly he must decide which group instantiated by the speaker is being picked out by this use of *we*. With this in mind, Nunberg divides the meaning of *we* into three separate components: a deictic component, which picks out a particular element of the context of utterance (an *index*, in Nunberg’s terms): in the case of *we*, the deictic component will pick out the speaker; a classificatory component, which serves to lay constraints on the eventual interpretation: for *we*, the classificatory component will contain the features plural and animate, on Nunberg’s analysis; and a relational component which determines what relation must hold between the index and the interpretation: for *we* this

---

<sup>1</sup> As we shall see there are good reasons to believe that, contrary to appearances, at least some of Nunberg’s examples do, in fact, express singular propositions.

<sup>2</sup> As, for instance, when a group of authors jointly drafting a document start with the phrase *we the undersigned*.

will be something like *the index must instantiate the interpretation*. The key point is that, for Nunberg, the semantics of *we* is such that all three components must be resolved in the interpretation process: a hearer must go from the expression to its speaker and then from the speaker to a plural, animate interpretation instantiated by the speaker. Now that's not to say that there is no rôle for pragmatics in this process: clearly working out which plural, animate group is being referred to on a given use of *we* involves pragmatic inferencing. That pragmatic inferencing, in fact, is mandatory given what, for Nunberg, *we* semantically encodes.

The interpretation of *we*, then, standardly involves deferred reference. But it's not at all clear that direct reference theorists would have much trouble with the outline of this account of *we*: so long as the ultimate interpretation is singular rather than general, then all's well. They would, and in the case of Recanati do, however, have a problem with Nunberg's next move. Nunberg's claim is that the interpretation of indexicals standardly involves a process of deferred reference; in other words, indexicals in general work just like *we*. Let's have a look at *I* to see how this might work. Nunberg's proposal looks something like this: *I* has a deictic component identical to that of *we*, i.e. the deictic component of *I* picks out the speaker as index; it has a classificatory component which requires that the interpretation should be 'an individual or an individual property, which moreover can only apply to an animate' (Nunberg (1993, p20)); and it has a relational component which, again, determines that the index should instantiate the interpretation. With this in mind let's have another look at (1), repeated here as (5):

- (5) *Condemned prisoner*: I am traditionally allowed to order whatever I like for my last meal.

On Nunberg's analysis, the deictic component of *I* will, in this case as in every case, pick out the speaker as index. The hearer will then look for an animate individual, or 'individual property', instantiated by the speaker. Clearly, so long as the interpretation of *I* is an individual, the predictions of Nunberg's account will coincide with those of the direct reference account; after all, there's only one individual which is identical to any speaker, and that is the speaker herself. But in this case it is an individual property, that of being a condemned prisoner, which enters into the interpretation. As we've already observed, whether it will be an individual or an individual property which enters into the interpretation, and, if the latter, which individual property, will be down to pragmatic inference<sup>3</sup>, but that pragmatic inference is mandatory given Nunberg's

---

<sup>3</sup> In an earlier analysis (Nunberg, 1978), Nunberg saw this process in terms of the selection of a

semantics for indexicals.

To summarise, for Nunberg the direct reference account of indexicals must be wrong, given the evidence that indexicals can give rise to the same range of interpretations as descriptions. He thus proposes a semantics for indexicals which accounts for the transfer from individual to property evident in some of his examples as the upshot of a mandatory step of the interpretation process. For Nunberg, an indexical leads to its interpretation by pointing out an object, the index, which relates to that interpretation. Looking at it from the other end, deferred reference depends on the availability of an index: ‘in the absence of deixis, deferred interpretation is not possible’ (Nunberg (1993, p35)).

But why, given that there’s so much pragmatics going on in the step between index and interpretation, does Nunberg want to develop a semantic machinery to account for this step, rather than just leaving the whole thing to pragmatics? Well he has two answers to this question. Firstly, if these interpretations were just the product of some pragmatic tinkering with a literal meaning, you’d expect the same effect to crop up in the same contexts with other referring expressions, such as proper names. But there are at least some contexts in which a deferred reading is available for indexicals but unavailable for proper names. Compare (6) and (7) with (8) and (9):

- (6) *Spoken by Justice O’Connor, a member of the US Supreme Court:* We might have been liberals.
- (7) *Spoken by Bill Clinton:* The Founders invested me with sole responsibility for appointing Supreme Court justices.
- (8) O’Connor, Rehnquist, Thomas etc. might have been liberals.
- (9) The Founders invested Bill Clinton with sole responsibility for appointing Supreme Court justices.

For Nunberg (6) and (7) have general readings on which the indexicals introduce attributive descriptions; we in (6) would be equivalent to something like *the Supreme*

---

referring function. In broad terms the idea is that in any act of indexical reference the index and interpretation are linked by a referring function; for communication to succeed the intended referring function must be that with highest ‘cue validity’ in the context. For those examples in which indexicals do, in fact, turn out to refer to what Nunberg would call their index (such as most uses of *I*), the referring function with highest cue validity simply is the identity function (see also Berckmans (1990) and Fauconnier (1994)).

*Court justices, and me in (7) to the President of the US. (8) and (9), by contrast, appear to have only de re readings.*

The second reason Nunberg gives for rejecting a pragmatic account is that, if these general readings are pragmatically derived, then there must be some coherent literal interpretation for them to be derived from; but, in at least some of the examples, there is no coherent literal interpretation. Consider another of Nunberg's examples:

(10) Tomorrow is always the biggest party night of the year.

*Always* must quantify over instances, therefore, Nunberg argues, this utterance has no coherent interpretation on which *tomorrow* is taken to refer to a particular day, as direct reference would have it do. Rather, the only coherent interpretation has *tomorrow* standing in for some description such as *the Saturday before classes begin*. Given the lack of a coherent literal interpretation this general reading cannot, for Nunberg, be pragmatically derived. If, then, we want to argue for a pragmatic account of any of Nunberg's general interpretations, we're going to have to overcome these two obstacles.

#### 4 Some problems with Nunberg's account

Before I come on to that, however, I just want to highlight a couple of problems which might make us think twice before accepting Nunberg's account wholesale. Let's just briefly remind ourselves of the core of Nunberg's claim. Essentially what he is arguing amounts to this: an indexical receives its interpretation through a process of deferred reference. It can do this due to a tripartite semantic structure consisting of a deictic component, a classificatory component and a relational component. The deictic component points to an object, the index, and the combination of index, classificatory component and relational component lead to an interpretation. On this analysis, as we have seen, deixis is crucial to deferred reference: no deixis, no deferred reference.

The first problem with Nunberg's account is that it involves a great deal of redundancy. As we have seen, Nunberg analyses the meaning of *we* as: deictic component picking out the speaker; classificatory component constraining the interpretation to plural and animate; and relational component requiring that the index instantiate the interpretation. But how would we want to flesh out the semantics of other indexicals, *they* for instance? Let's look at one of Nunberg's examples of deferred reference involving *they*:

- (11) *Spoken by a teacher pointing at a little girl:* They haven't signed the permission form yet.

The idea is that *they* in this utterance is interpreted as something like *her parents*. But how does this interpretation proceed? Firstly, what is the deictic component of *they* which leads us to the index, in this case to the little girl? The answer Nunberg gives is, essentially, that there is no encoded deictic component, or, at least, no deictic component which alone identifies an index: *they* relies for deixis on an associated demonstration. And what of the relational component? As Nunberg himself points out, there seems to be no particular relation that need hold between the index and the interpretation in cases such as this.

Now Nunberg is entirely aware of these differences; they in fact form the basis of his distinction between two classes of indexicals: *participant terms* and *non-participant terms*. Participant terms are those nominal indexicals which, in Nunberg's terms, have as their index a participant in the exchange; that is *I* and *we*, which both have the speaker as index, and *you*, which has the addressee as index. Non-participant terms are, essentially, all other nominal indexicals: *he*, *she*, *they*, *this*, *that* etc. For Nunberg the interpretation of participant terms involves the cashing out of all three semantic components, whereas non-participant terms 'have no explicit relational component' (Nunberg (1993, p9)). In addition, 'with demonstratively used pronouns, the deictic component is provided entirely by demonstration' (Nunberg (1993, p23)), so for the majority of non-participant terms there is neither a semantically encoded deictic component nor a semantically encoded relational component. It seems, then, that much of the semantic machinery Nunberg has developed is actually redundant in many of the examples he seeks to account for.

The second problem with Nunberg's account is pointed out by Recanati (1993). Recanati argues that, in his distinction between index and interpretation, Nunberg has conflated two separate distinctions: that between index and referent, a distinction which Recanati grants is crucial to the interpretation of *we*, and the distinction between trigger and target in the terminology of Fauconnier (e.g. Fauconnier (1994)). Have a look at the following example:

- (12) We are in red brick.

Now the interpretation of (12) will, Recanati claims, look like:

- (13) The house we live in is in red brick.

But, Recanati asks, what is the process which underlies this interpretation? The answer seems to be something like: from *we* to the speaker, from the speaker to the group of people living in the house and from the group of people living in the house to the house itself. In Fauconnier's terms, the trigger is the group of people living in the house and the target is the house itself. As Recanati points out, the distinction between index and referent must be separate from that between trigger and target 'for the obvious reason that the index is neither the trigger nor the target' (Recanati (1993, p312)). It seems to me that this objection is crucial. However Recanati has allowed Nunberg an easy escape: example (12) is of a type which Nunberg expressly excludes from his analysis of deferred reference<sup>4</sup>. We can, however, see the same phenomenon at work in some of the examples cited by Nunberg himself:

There was an even clearer example of this type [the type of example (6)] in a recent cartoon by William Hamilton, which showed a group of conservative middle aged businessmen sitting around a boardroom table as one of them says: "in a couple of years we'll probably all be women."

Nunberg (1993, p14 fn20)

As Grimberg (1996) points out, the index in this example cannot be intending to include himself in the interpretation, since the interpretation is an all-female group. Grimberg takes this as evidence that there is no relational component encoded by *we*; as I shall argue later I believe this conclusion to be a bit hasty. It does, however, illustrate just the situation Recanati claims for (12): the index, here, is, as always, the speaker, the trigger is the group of businessmen around the boardroom table and the target will be something like *the board members* attributively used. What this shows, I shall be arguing, is that there are two entirely separate processes of deferral at work in Nunberg's examples, one semantically motivated and one entirely pragmatic.

Finally, before launching into the second half of this paper, I'd like to mention another issue raised by Nunberg's account. As I've stressed, one of Nunberg's key claims is that 'in the absence of deixis, deferred interpretation is not possible': the process of deferral relies, for Nunberg, on the identification of an index, and that can only be achieved through deixis, either in the form of an encoded deictic component or a demonstration. In the light of this, consider these two examples, both from Nunberg (1993):

---

<sup>4</sup> Nunberg actually accounts for uses such as this in terms of property transfer e.g. for (12) the property of being in red brick would transfer from house to inhabitants. For more on predicate transfer, see Nunberg (1995).

(14) *Spoken while pointing at a painting*: Now *he* knew how to paint goats.

(15) *We are walking through the Taj Mahal*: Gee, he certainly spared no expense.

(14) has, *ex hypothesi*, an associated demonstration; (15), on the other hand, has neither an associated demonstration nor an encoded deictic component (since *he* is a non-participant term and so lacks deixis). As we already know, no deixis, no deferred reference; so, if Nunberg is right, there must be a different kind of interpretation process at work in (14) and (15). Nunberg concludes that (14) and (15) contain examples of different uses of *he*: in (14) *he* is *deictic*, in (15) it is *contextual*. In a good discussion of Nunberg's distinction between deictics and contextuels, Grimberg (1996) points out that this result is highly counterintuitive: there seems little intuitive basis for assuming that the Taj Mahal is playing a fundamentally different role in the interpretation of (15) from that played by the painting in the interpretation of (14). Grimberg also observes that the use of *he* in (15) carries the hallmark of deferred reference: it can give rise to an interpretation on which the indexical appears to stand in for an attributively used definite description, just as it can in (14). In (14) *he* can either be interpreted as Chagall, say, or as *the man who painted that painting*; in (15) it can be interpreted either as Shah Jahan or as *the man who built this building*. There seems little more evidence to support a distinction between these two uses than that Nunberg's account begins to look a bit shaky without it. Nunberg himself, in fact, appears to realise that this is, to say the least, a murky area: he concedes that 'it is not always easy to tell when the context provides an implicit demonstration (Nunberg (1993, p 34)).

## 5 Deferred reference - a cognitive approach

5.0. At the start of this paper I suggested that by approaching questions of deferred reference from a cognitive perspective, interesting results might fall out. The cognitive perspective I intend to adopt is grounded in Relevance Theory, providing, as it does, useful insights into the semantics-pragmatics distinction, a distinction which will prove central to the following discussion. It seems to me that, looking at deferred reference from a relevance theoretic perspective, we're likely to find the question we're asking more or less entirely reversed: rather than asking, as Nunberg does, why deferred readings *are* available for indexicals, the key question from within a relevance theoretic framework, as I hope to show, is likely to be why deferred readings *aren't* available for other referring expressions, such as proper names, where they are available for indexicals. But before looking at that question we need to do a bit of preliminary scene-

setting.

## 5.1 Individual and general concepts

What kind of thing are we talking about when we talk about the interpretations of indexicals or, for that matter, of any referring expressions? I hope I'm saying nothing too controversial when I claim that, from a cognitive perspective, there can be only one general answer to this question: we're talking about concepts. Concepts, after all, must, in one form or another, be the building blocks of propositions on any representational account. It might seem, however, that, in accepting this notion, I've prejudged the entire issue against the direct reference camp; after all, it is, as we've seen, one of the central claims of direct reference that what directly referential terms contribute to propositions is nothing but their referent. I don't believe that I have: direct reference, as originally formulated by Kaplan, is a semantic theory, not a cognitive one; Kaplan is not attempting to discover what goes on in the mind, but rather what goes into the truth conditions.

The direct reference account, however, is not incompatible with a cognitive approach, as has been demonstrated, for instance, by Recanati (1993). Recanati draws a distinction between *descriptive* concepts and *de re* concepts. Descriptive concepts are the sort of thing that go to make up Fregean senses; *de re* concepts, by contrast, are, for Recanati, the cognitive corollary of directly referential terms. Their defining characteristic is that their content is truth conditionally irrelevant, leaving only their referent to appear in the truth conditions. Recanati envisages *de re* concepts as being of two types: *egocentric* and *encyclopaedic*. Egocentric concepts are the sort of thing we standardly entertain when we understand indexicals; they are temporary dossiers dominated by non-descriptive (perceptual) information. Encyclopaedic concepts, by contrast, are stable, long-term dossiers of predominantly descriptive information, the sort of thing we entertain when we understand proper names<sup>5</sup>.

Having explained (albeit very sketchily) Recanati's distinction between *de re* and descriptive concepts, I now propose to replace it or, rather, to modify it. Recanati's distinction is firmly grounded within semantic notions of propositional singularity, direct referentiality and so on; it seems to me, however, to cross-cut the key cognitive

---

<sup>5</sup> That is not to say that egocentric concepts cannot contain descriptive information or that encyclopaedic concepts cannot contain perceptual information; it is rather a claim about which type of information predominates in each type of concept. For a more detailed account, see Recanati (1993, Chapter 7).

distinction. The question we need to ask ourselves is this: when can we think of an individual under a de re concept? In what relation do we have to stand to something in order to open a de re concept corresponding to it? Clearly to open an egocentric concept we need direct epistemic contact of some variety. But is the same true of encyclopaedic concepts? Recanati's answer is no:

For the referent to be thought of non-descriptively, it seems to be sufficient that the thinker acquire what Grice calls a 'dossier' for the definite description used to fix the reference.

In Grice's framework, 'X has a dossier for a definite description  $\delta$  if there is a set of definite descriptions which includes  $\delta$ , all the members of which X supposes...to be satisfied by one and the same item' (Grice, 1969, p.139)

Recanati (1993, p 109)

Now what does this criterion actually mean? Well it means something like: to have a de re thought all you need is enough information all of which you believe applies to a unique individual. But how much information is enough? Where between one and one million pieces of information can we safely draw the line between de re and descriptive concepts? From the perspective of direct reference, concerned as it is with truth conditions, there's a very good reason for drawing the line, wherever it may be that you decide to draw it; but from a cognitive perspective, concerned only with what's going on in the minds of speaker and hearer, there seems to me to be no good reason to draw what must in the end be an arbitrary line somewhere on this continuum<sup>6</sup>.

If that's right, then, where does it leave us? Let me propose the outline of an answer to that question: in cognitive terms the crucial difference between concepts is whether or not we believe they are representations of an individual, not how much uniquely identifying information they contain. What do I mean by that? Essentially I mean that there are two distinct types of concept: first there are those which, whether they contain one piece of information or a million, we believe correspond to an individual in the world (from now on I propose to call these *individual* concepts); and there are those which, again regardless of how much information they may contain, we believe do not uniquely represent any individual (henceforth *general* concepts). I would, for instance, have an individual concept of Robyn Carston, which will include a bunch of information gained through direct epistemic contact as well as otherwise; another of the

---

<sup>6</sup> Recanati does, to an extent, accommodate this insight through his notion of provisional de re concepts. A provisional de re concept is, essentially, a de re concept in search of a referent.

current president of Indonesia, containing information gained, I guess, through indirect epistemic contact as well as otherwise; and yet another of the oldest person in the world, which will, as far as I can see, contain very little information other than *is the oldest person in the world*. The crucial thing is that it is my belief that each of these concepts relates to a single individual. On the other hand each of these three individual concepts will contain general concepts: my concept of the current president of Indonesia, for instance, will contain the general concepts *president* and *Indonesian*; it will, in fact, contain the general concept *president of Indonesia*. Equally my individual concept of the oldest person in the world will contain general concepts such as *old*, *person*, *world* etc.

This distinction between individual and general concepts is little more than a suggestion; I haven't, as you will have noticed, offered much in the way of supporting argument. On the other hand, when we come to look at Nunberg's examples of deferred reference from a cognitive perspective later in this paper, I hope to persuade you that the distinction is a useful one<sup>7</sup>.

While reformulating Recanati's distinction between descriptive and de re concepts, I want to maintain, within the class of individual concepts, his distinction between egocentric and encyclopaedic concepts, or something very like it. Recanati (1993) provides extensive and, as far as I can see, very convincing arguments for such a distinction and, when we come to look at contexts in which deferred readings are apparently available for indexicals but not for names, the egocentric-encyclopaedic distinction is, I hope, going to do good work for us.

## 5.2 Indexicals and the conceptual-procedural distinction

Since Blakemore (1987) proposed that the meaning encoded by non-truth conditional expressions such as discourse connectives might be of a fundamentally different kind from that encoded by other expressions, there has been a great deal of research conducted within the relevance theoretic framework into the distinction between *conceptual* meaning and *procedural* meaning. Blakemore's original insight was, roughly, this: that, since utterance interpretation involves two radically different kinds of process, first the retrieval of a logical form and then the manipulation of that logical form through pragmatic inference, we might well expect to find two distinct kinds of

---

<sup>7</sup> Of course, if this is right, it may have interesting consequences for Donnellan's referential-attributive distinction. For lack of space, I shall have to leave an examination of those consequences for another time.

meaning corresponding to these processes. And this, for Blakemore, is just what we do find. The kind of meaning encoded by descriptions, for instance, is representational: it encodes concepts. The kind of meaning encoded by discourse connectives is, by contrast, computational: it encodes procedures which lay constraints on the retrieval of implicatures.

For Blakemore, then, procedural meaning constrains the retrieval of implicitly communicated meaning. Wilson and Sperber (1993) and others have extended Blakemore's original notion to the retrieval of explicitly communicated information; specifically, Wilson and Sperber suggest that pronouns, while contributing to truth conditions, might encode procedures. On the face of it, as Wilson and Sperber admit, this is little more than a reformulation of Kaplan's distinction between character and content: for Kaplan the character of *I* is something like *the speaker of this utterance*; for Wilson and Sperber *I* encodes 'an instruction to identify its referent by first identifying the speaker' (Wilson and Sperber (1993, p20)). The conceptual-procedural distinction does, however, have a cognitive grounding which the character-content distinction lacks. For Kaplan, character and content are, at least in one sense, the same *kind* of meaning: they are both descriptive. The failure of the character of indexicals to show up in truth conditions, then, must be accounted for by an extra (and ad hoc) piece of encoding, something like the REF feature of Recanati (1993). By contrast there is a fundamental difference between conceptual meaning and procedural meaning corresponding to the cognitive difference between representation and computation. The upshot of this is that procedural meaning is not excluded from truth conditions by an ad hoc encoding, it is excluded because it's simply not the sort of thing that appears in truth conditions i.e. it's not representational<sup>8</sup>.

It seems to me, then, that there may well be good reasons to suppose that indexicals encode something like procedures, although I'm not sure that a great deal of what follows will hang on this. But, if indexicals do encode procedures, which procedures do they encode? Let's take Sperber and Wilson's suggestion for the procedure encoded by *I* as a starting point: *identify its referent by first identifying the speaker*. In other words, the semantics of *I* takes the hearer as far as the speaker (or, in the light of the discussion above, as far as an individual concept of the speaker) and then leaves the rest to pragmatics. I imagine we'd want to say much the same about the encoded meaning of *you*: something like *identify its referent by first identifying the hearer*. In other words, *you* takes the hearer to his own concept of himself and then leaves him to work the rest

---

<sup>8</sup> This is clearly no criticism of Kaplan's approach, since providing a cognitive account of the interpretation of indexicals is no part of his aim.

out pragmatically.

So much for Nunberg's two singular participant terms; how about the plural participant terms? As we've seen, Nunberg (1993) launches his proposal off the back of a discussion of the semantics of *we*. The reason he gives for doing so is that the interpretation of *we* uncontroversially involves deferred reference. And intuitively this seems right; the interpretation of *we* does seem to involve finding the speaker and then finding a group of which the speaker is a member. In procedural terms, it looks like *we* encodes a two-part procedure, the first part leading to the speaker and the second leading to a group instantiated by the speaker<sup>9</sup>. But maybe this analysis of *we* is not as uncontroversial as it seems. Grimberg (1996) has argued that, for all indexicals, Nunberg's relational component should be regarded not as semantically encoded but, rather, as a metaphor for the operation of pragmatic processes. In other words for Grimberg the semantics of indexicals consists of a deictic component leading to an index, a classificatory component constraining the interpretation and no more. So the meaning of *we* would simply lead a hearer to the speaker and then determine that the interpretation should be plural and animate. Grimberg takes as evidence for this the fact that *we* can give rise to interpretations which are not instantiated by the speaker. Have a look again at (12), repeated here as (16):

(16) We are in red brick.

Now, as we've already seen, for Recanati, (16) shows that there must be a difference between Nunberg's distinction between index and interpretation and Fauconnier's between trigger and target: the group of people living in the house, while neither index nor interpretation, is the trigger in Fauconnier's terms. Recanati's view, then, is that *we* may well encode something like Nunberg's relational component, but that this semantics cannot account for deferred uses of *we*, which are the result of sub-propositional pragmatic processes. Grimberg, however, says:

It seems to me that the inferencing process from index to interpretation is essentially the same as that from trigger to target, with the 'trigger' in Recanati's example representing no more than an intermediate stage in that process.

---

<sup>9</sup> On this analysis, the use of *we* to refer to a group of speakers/writers arises where the optimally relevant group associated with the speakers simply is the speakers themselves. In other words, once the hearer has found the speaker(s), he has already found a plural group instantiated by the speaker(s), i.e. the speakers themselves. He should accept this interpretation in the absence of good pragmatic reasons not to do so.

Grimberg (1996, p 80)

As I understand it, then, Grimberg sees the process of interpretation that leads from this utterance of *we* to the interpretation *the house we live in* as essentially seamless: the same processes take us all the way, guided by the deictic and classificatory components and by general pragmatic principles. But surely there is a problem here? What's happened to the classificatory component? The ultimate interpretation in this example is a house, i.e. it is neither plural nor animate. As we've already seen, however, Nunberg might well not account for this example in terms of deferred reference. Perhaps, then, we might want to say something like: the standard interpretation of *we*, rather than those cases involving phenomena such as predicate transfer, involves deferred reference guided by a deictic component and a classificatory component but no relational component. Now there seem to me to be good reasons to suppose that this is wrong. Grimberg conducts her discussion off the back of:

(17) In a couple of years we'll probably all be women.

in which the interpretation of *we* does indeed comply with her analysis; interpretation, that is, goes via the speaker to a plural animate group but, apparently, to one not instantiated by the speaker. But consider a parallel example. Imagine a discussion between a group of print workers in Fleet Street some time around 1982:

(18) In ten years time we'll probably be one man and a computer working out of an office in Wapping.

Following Grimberg's argument, the speaker of (18) is, presumably, not intending to include himself in the interpretation of *we*, therefore the interpretation does not appear to be constrained by Nunberg's relational component. But, then, the interpretation doesn't seem to be constrained by Nunberg's (or Grimberg's) classificatory component either: you might, I suppose, want to say that *one man and his computer* denotes a plural entity, but it surely doesn't denote a plural animate entity. The only way that I can see out of this, arguing that (17) and (18) involve different types of interpretation process, seems to me to be a non-starter.

The picture that seems to be emerging here is, as Recanati has argued, of two distinct processes at work: on the one hand the interpretation of *we* standardly involves deferred reference, and it is at the level of the output of this process of deferral that the interpretation conforms with both the classificatory and relational components posited by Nunberg. This type of deferral is *semantic*, by which I mean that it is the result of a semantically mandatory phase of the interpretation process. On the other hand, once

this mandatory process has taken place, there is then a separate *pragmatic* deferral, constrained by general pragmatic principles, which leads from the output of semantic deferral as trigger to the eventual interpretation as target.

In passing, I just want to mention one more piece of evidence which seems to me to support a pragmatic deferral account of examples such as (17) and (18). Such examples typically have, if not a humorous feel ((17) is, after all, taken from the caption of a cartoon), at least a slightly anomalous feel. And that, it seems to me, is, in relevance theoretic terms, pretty well the hallmark of extra effort yielding extra effects. If, on the other hand, this interpretation were simply the upshot of standard, semantically guided processes, we'd be hard pushed to explain the anomaly.

On this analysis, then, we encode a two-step procedure something like *first find the speaker(s) and then find a group instantiated by the speaker(s)*. And it seems inevitable that much the same analysis must apply to plural *you*, i.e. plural *you* must encode something like *first find the addressee(s) and then find a group instantiated by the addressee(s)*<sup>10</sup>. I don't propose to provide any arguments for this analysis over and above those I've already provided for *we*.

That deals, then, albeit fleetingly, with Nunberg's participant terms; how about non-participant terms? Let's start with *he* and *she*. It seems to me that there is no need to propose anything like a deictic component for *he* and *she*; after all, even for Nunberg these are not semantically encoded but supplied by an associated demonstration. Nunberg also agrees that *he* and *she* lack relational components, so what does that leave us with? In Nunberg's terms it leaves us with a classificatory component consisting of the features male and singular in the case of *he*, and female and singular in the case of *she*; in procedural terms, it leaves *he* encoding something like *find an individual concept with a male referent*, and *she* encoding *find an individual concept with a female referent*. The rest can be left to pragmatics.

I'll come on to *they* in a moment, but first I want to have a look at *this*, *these*, *that* and *those*. It seems to me that *this* and *that* must, at least in some sense, encode a two-step procedure. Take *this*; in intuitive terms *this* is used to refer to an object relatively near to the speaker, so the hearer must find the speaker before he can find the object referred to. In procedural terms, *this* must encode something like *find the speaker and then find an object near the speaker*. The same goes, *mutatis mutandis*, for *that*. How about *these* and *those*? There might, at first, be a temptation to suggest that *these* just encodes something like *find a group relatively near the speaker* and *those*, *find a group*

---

<sup>10</sup> This is, of course, only inevitable if we assume a homonymy account for *you*; i.e. if we assume that singular *you* and plural *you* are separate lexical items. I don't have the space here to go into an examination of this issue.

*relatively distant from the speaker*. But Nunberg produces good reasons why this can't be right or, rather, why this can't be the whole picture. Consider the following example from Nunberg (1993):

- (19) *I point in sequence at two sample plates in my china shop, the first sitting in front of me, the second on a table across the room: These are over at the warehouse, but those I have in stock here.*

As Nunberg points out, in (19) the referent of *these* is further from the speaker than the referent of *those*, whereas the index of *these* is nearer the speaker than the index of *those*. What this comes down to is that the interpretation of *these* and *those* standardly involves deferral just as for *we*. In procedural terms, *these*, for instance, must encode something like *find an object near the speaker and then find a plural referent relating to it*<sup>11</sup>.

How about *they*? Well, *they* is a little problematic. The problem comes down to this: does *they* encode a plural, genderless equivalent of *he* or *she* (i.e. *find something plural*), or does it encode something more like *we* (i.e. *find a he or she and then find a group relating to it*)? There is no way that I can see to decide this: since *they* doesn't encode anything like the proximality/distality of *these* and *those*, both analyses lead to the same results. Given this, I don't propose to spend too much time worrying about the problem.

Finally I want to look at indexicals of time and place i.e. things like *here*, *now*, *tomorrow* etc. I don't want to suggest anything particularly revolutionary; in fact, standard accounts of these expressions seem to translate well into procedural terms. On a straightforward analysis, *here* would encode *find the place of utterance* (or, more accurately, *find an egocentric concept of the place of utterance*); *now*, *find an egocentric concept of the time of utterance* and *tomorrow*, *find a concept of the day after the day of utterance*. This is all very rough and ready but, as I stressed earlier, it does not seem to me that a great deal hangs on the minutiae of this procedural account.

### 5.3 Deferred reference and relevance

I have suggested, then, that there may be two separate processes of deferral at work in Nunberg's examples: what I have been calling semantic deferral and pragmatic deferral.

---

<sup>11</sup> Again, in the cases where *these* is used to refer to an ostended group, that group will satisfy the full encoding and will thus be accepted as the intended interpretation unless pragmatically blocked.

Now I want to persuade you that the availability of pragmatically deferred readings is to be expected within a relevance theoretic approach. Why should this be so? Relevance Theory embodies a particular conception of the nature of communication: the task facing a hearer is not that of decoding an encoded message, but rather that of inferring speaker intentions. Under this conception of communication the role of linguistic meaning is restricted to providing inputs to the inferential process which will ultimately deliver hypotheses about speaker intentions; or, in other words, linguistic meaning provides evidence for speaker intentions and no more. One upshot of this insight is that so long as the linguistic meaning of an utterance provides sufficient evidence for a hearer to infer the speaker's intention, it should have served its purpose; the linguistic meaning of an utterance can be as minimal as you like so long as it places adequate constraints on the interpretation process. Now what does this mean in practice? At the heart of Relevance Theory is the notion that interpretation is constrained by a presumption of optimal relevance; the hearer, that is, can assume that:

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

and can conduct the process of interpretation on the strength of these assumptions. If a speaker wants communication to be successful, she had better make sure that the content of her informative intention coincides with the optimally relevant interpretation of her utterance. If it does, communication should succeed, if it doesn't communication should fail.

Now how do these observations tie in with deferred reference? Well, so long as the deferred interpretation of a referring expression is the optimally relevant one, then it should be the accepted interpretation. It seems to me that this condition involves two separate requirements: firstly, that the intermediate stage in the interpretation, the trigger, should not, itself, be optimally relevant; and secondly that there should not be any more relevant stimulus which the speaker could have produced (compatible with her abilities and preferences). So long as these conditions are met, however, pragmatic deferral is to be expected under a relevance theoretic conception of communication.

## **5.4 Pragmatic deferral and indexicals**

Let's briefly summarise what we've seen in the last three sections. Firstly we've

proposed a rough and ready semantics for indexicals on which indexicals encode procedures which lead to individual concepts; and secondly we've argued that some kind of pragmatic deferral is no more nor less than we'd expect within a relevance theoretic framework. But what do we get when we put these two ideas together? Well, if the procedures proposed above are anything like right, many of the examples Nunberg seeks to account for with his tripartite semantics for indexicals must, on our account, be the product of pragmatic deferral. And on the face of it, this has a pretty plausible feel to it. To convince ourselves of this, let's return to a couple of Nunberg's examples; let's look at examples (7) and (10), repeated here as (20) and (21):

(20) *Spoken by Bill Clinton*: The Founders invested me with sole responsibility for appointing Supreme Court justices.

(21) Tomorrow is always the biggest party night of the year.

On the semantics for indexicals proposed above, *me* in (20) will initially lead the hearer to an individual (egocentric) concept of the speaker, i.e. to an individual concept of Bill Clinton. Why, then, is this individual concept not accepted as the eventual interpretation? We'd like to be able to say that it's not accepted because it's not optimally relevant. And why might that be? One possibility is this: put in pretheoretical terms, there's something bizarre about the idea that the Founders of the United States, in the 18th century, could have made a decision concerning an individual born in the 20th century. Finding a context in which this interpretation will have sufficient contextual effects will put the hearer to a great deal of processing effort; the interpretation cannot, therefore, be optimally relevant. With this interpretation pragmatically blocked, the hearer should, as always, try out the next most accessible interpretation. He will, in this case, look within his individual concept of the speaker where he will find the general concept *the President of the United States*. Finding that an interpretation on which *me* is interpreted as equivalent to *the President of the US* provides enough effects for not too much effort, he will accept this interpretation as optimally relevant. The hearer can thus arrive at the deferred interpretation through standard pragmatic inferencing.

Much the same line of argument can be used to account for the interpretation of (21): *tomorrow* leads to the hearer's concept of a particular day  $\delta$ ; that concept is pragmatically blocked as the eventual interpretation since *always* must quantify over instances; the hearer thus hunts through information stored with his concept of  $\delta$  for another interpretation and comes up with something like *the Saturday before classes begin*.

It therefore seems that, on the face of it, we might be able to account for some of Nunberg's examples in terms of pragmatic deferral, i.e. in terms of the free operation of pragmatic processes. But, as we already know, things are not that straightforward: if we want to argue for a pragmatic account we're going to have to tackle Nunberg's two objections. I imagine that it may already be clear how I intend to deal with Nunberg's second objection, that deferred readings cannot be pragmatically derived because there is often no coherent literal interpretation for them to be derived from. Recanati (1993) has mounted what seems to me an entirely convincing defence against this objection. Nunberg appears to be working on the assumption that any pragmatic account must be based on the retrieval of implicatures; that is, the pragmatic inferencing involved must operate once a global utterance interpretation has been achieved<sup>12</sup>. But why would we necessarily want to assume this? There is, after all, abundant evidence to suggest that pragmatic inferencing is not restricted to the retrieval of implicatures but plays a crucial role in retrieving the proposition expressed<sup>13</sup>. As soon as we assume that the inferencing involved in pragmatic deferral takes place locally rather than globally, as Recanati explicitly does, there's no need for any global interpretation to be retrieved before deferral takes place: it is, in fact, precisely because no 'literal' global interpretation is available, for Recanati, that deferral does take place.

What of Nunberg's first objection, then? If I'm right, and examples such as (20) and (21) really do involve pragmatic deferral, why can we not get the same readings in the same contexts with proper names, if that is the case? Now it seems to me that there may well be no one standard answer to this question, but a whole range of factors which together may prevent deferred readings of proper names from achieving optimal relevance in particular contexts. There are, however, a couple of general patterns which I'd like to draw out, in order to show that we don't necessarily need to resort to a semantic account of the type adopted by Nunberg.

First, however, we need to establish in which contexts deferred readings really are available for indexicals but blocked for proper names. I want to start by weeding out one group of Nunberg's examples which seem to me to involve no deferral in any useful sense. Consider (22), (14), repeated here as (23), and (24):

(22) *Spoken while pointing at a girl child to identify her father:* He is in real estate.

---

<sup>12</sup> Nunberg does grant that, for Recanati, the pragmatic processes involved in deferral contribute to the proposition expressed, but he does not seem to have appreciated the significance of this to Recanati's account.

<sup>13</sup> See, for instance, Carston (1988, 1996).

(23) *Spoken while pointing at a painting*: Now he knew how to paint goats.

(24) *Spoken while pointing at a book to identify its author*: She was my chemistry teacher.

On Nunberg's analysis, the interpretation of each of these examples involves finding an index, as indicated by the demonstration, and then finding an interpretation suitably associated with that index. As we've already seen, this forces Nunberg to the conclusion that (22)-(24) involve a different type of interpretation process from (15), repeated here as (25):

(25) *We are walking through the Taj Mahal*: Gee, he certainly spared no expense.

Now on the face of it that seems a pretty counterintuitive result; as I suggested earlier, the Taj Mahal in (25) seems to be playing much the same role in interpretation as, for instance, the painting in (23). Moreover, we could, without any problem, rework (23) so that the painting in question was the only one in the room and (23) was uttered without any associated demonstration<sup>14</sup>. Evidence such as this has led Kaplan (1989) and Bach (1992) to conclude that the role of demonstrations is no more than context shuffling in order to reveal speaker intentions: a speaker, intending to refer to something which is not already maximally contextually salient, performs a demonstration; since demonstrations are intentional acts, the hearer will ask himself what the demonstration is for, and will come up with the answer that it is to make something in the context more salient. Interpretation will then proceed according to this new hierarchy of contextual salience. Now if Kaplan and Bach are right, there is no need for the speaker to point at the actual referent; after all, in (22)-(24) this wouldn't be possible. All the speaker need do is point at something which makes the intended interpretation optimally relevant in the revised context. On this analysis, the interpretation processes involved in (22)-(24) are no different from that involved in (25): each involves assigning reference to an indexical according to context.

Ah, says Grimberg (1996), but there must be some deferral, because there is an obvious way in which these examples can give rise to general interpretations, the very hallmark of deferral. Consider, for instance, if the speaker of (23) has no idea who the goat-painter is, or the speaker of (25), who built the Taj Mahal. The interpretation in

---

<sup>14</sup> Nunberg might argue that there was an implicit demonstration in such a case but, given that, as we've seen, Nunberg himself admits that it's sometimes hard to tell when there is an implicit demonstration and when not, I can't see that this would be a very fruitful approach.

each case must, then, be equivalent to a definite description attributively used. Now it seems to me that the distinction drawn earlier between individual and general concepts can help us to overcome this objection: in each case the hearer simply retrieves an individual concept, as required by the semantics of the indexical. So the hearer of (23) will simply retrieve, or open, an individual concept containing the information *is the painter of that painting*, and insert that concept into his interpretation; in much the same way the hearer of (25) will retrieve or open a concept containing the information *is the person who built the Taj Mahal*. The key point is that, in each case, the hearer is accessing a concept which he believes to correspond to a unique individual in the world, albeit an individual with whom he may have no direct epistemic contact.

Finally, before addressing Nunberg's second objection to pragmatic accounts, there are a couple of pitfalls we would do well to be wary of. Recanati (1993) argues that many of the general interpretations that Nunberg cites for indexicals are simply illusory. Let's just have a look at one of the examples he tackles, taken from Nunberg (1990):

- (26) In the movie *The Year of Living Dangerously*, Mel Gibson plays a reporter in Sukarno's Indonesia who is looking for a shipment of arms destined for the local communists; who will kill him if they find out he is on to them. He is interviewing a warehouse manager, who tells him, 'I have seen no such shipment. And you should be careful; I might have been a communist'.

Recanati (1993, p 301)

Nunberg points out that there are two possible readings of *I might have been a communist* in (30): on one reading the warehouse manager is saying, of himself, that he might have been a communist (given a different upbringing, say); on the other, the preferred reading in the context, he is saying something like *The person you are talking to might have been a communist*. For Nunberg the first reading is singular, since *I* contributes an individual, and the second is general, since *I* contributes a concept (or *property* in Nunberg's terms). Recanati starts by asking the following question: what thought is it that a hearer must think in order to have understood an utterance of (26)? And the answer he comes up with is that the hearer must think something like *That man might have been a communist*, with *that man* corresponding to an egocentric concept of the warehouse manager. But surely that leads us straight back to our first interpretation, i.e. to the interpretation on which the speaker is saying of himself that he might have been a communist? No, says Recanati, and this becomes clear as soon as we understand the distinction between metaphysically possible worlds and epistemically possible worlds. For Recanati, the speaker of (26) is not saying that there is a metaphysically possible world at which he is a communist; rather, he is saying that there is an

epistemically possible world at which this is the case. To put this another way, in order to understand (26), a hearer must entertain a thought which looks something like:

(27) For all I know, that man might have been a communist.

The proposition the hearer entertains in understanding (26), then, does, just as Nunberg claims, contain a concept, but, crucially, it is a *de re* concept and thus contributes nothing to truth conditions beyond its referent. Translating this argument into the terms we've been developing in this paper, the interpretation of (26) just involves retrieving an individual concept of the speaker, as determined by the semantics of *I* and goes no further. For Recanati, many of Nunberg's examples are just such sheep in wolves' clothing.

The second potential pitfall which we should be aware of is illustrated by (1). In his analysis of (1), Nunberg assumes an interpretation of *traditionally* which requires that it should quantify over instances; and, given this assumption, it follows naturally that *I* cannot have a *de re* interpretation. But are we sure that Nunberg is right about *traditionally*? Consider the following paraphrase of (1):

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

Now I hope you agree that, in (28), there is no reason to assume that *by tradition* quantifies over instances. Given this, how can we be sure that *traditionally* in (1) is not interpreted as equivalent to *by tradition* in (28)? And, if that's right, there's no reason to suppose that *I* receives anything other than a perfectly straightforward *de re* interpretation in (1). Moreover, Nunberg himself seems to have doubts about the interpretation of *traditionally* in this example. In a discussion of (1) 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 claimed deferred reading of (1) becomes markedly less accessible when we replace *traditionally* with *usually*:

(29) *Condemned prisoner*: I am usually allowed to order whatever I like for my last meal.

## 5.5 Pragmatic deferral and proper names

With those caveats under our belt, the time has come to look at Nunberg's second

objection to pragmatic accounts: that deferred readings are unavailable for proper names in the same contexts in which they are available for indexicals. The first question we're going to need to ask is whether Nunberg is right; whether it is always the case that deferred interpretations are unavailable for proper names. It would, after all, be quite surprising if deferred readings of proper names were never available, given that, as we've seen, pragmatically deferred readings are to be expected under a relevance theoretic conception of communication. And there do, indeed, seem to be cases in which proper names can give rise to deferred interpretations just as indexicals can. Imagine yourself in Allied Command Headquarters, some time around May 1944; you overhear someone saying:

- (30) I've looked through the meteorological charts and D-Day has been the hottest day of the year every year for the last decade

How are you going to interpret the expression *D-Day* in this utterance? It seems to me you'll probably interpret it as *June 6th*. The interpretation process, that is, will go via an individual concept of June 6th 1944 to a general, or generic, concept of June 6th.

There do seem, however, to be certain contexts in which, as Nunberg argues, deferred interpretations are available for indexicals but blocked for proper names, and these examples provide Nunberg with the core of his argument against pragmatic accounts of deferred reference. Consider the following examples, some of which we've already seen:

- (31) *Spoken by Justice O'Connor, a member of the US Supreme Court:* We might have been liberals.
- (32) O'Connor, Rehnquist, Thomas etc. might have been liberals.
- (33) The Founders invested me with sole responsibility for appointing Supreme Court justices.
- (34) The Founders invested Bill Clinton with sole responsibility for appointing Supreme Court justices.

Now Nunberg's claim is that (31) and (33) have two readings each whereas (32) and (34) have just one each. Whether or not he is right that deferred readings of (32) and (34) are entirely unavailable, there does seem to be a marked disparity in the availability of such readings between the examples using indexicals and those using proper names.

And, on the face of it, this looks like strong evidence in favour of Nunberg's account, or something like it; after all, the examples are set up so that the different referring expressions appear in precisely the same contexts.

Perhaps, then, Nunberg is right and the procedures proposed above are far too simple; perhaps we do need to build into the semantics of many more indexicals elements of meaning leading to an index and other elements of meaning leading from index to interpretation. But before we give so much away, could there be any other reasons for the unavailability of deferred readings of proper names in these contexts? It seems to me that there could.

Let's first consider (31) and (32). Now, to be honest, I'm not entirely sure whether or not (32) does have only a *de re* reading as Nunberg claims; it seems to me that in certain lights I can get an attributive reading. But if Nunberg is right and there is no attributive reading, there may well be a good reason for it without having to resort to a semantic deferral account. As we've already seen above, the interpretation of *we* does standardly involve a two-stage process: from expression to speaker and from speaker to group instantiated by speaker. As part of the mandatory interpretation process, therefore, the hearer must retrieve some kind of property delimiting the group which, in Fauconnier's terms, provides the trigger for the pragmatic deferral in these cases. This property (in the case of (31) the property of being Supreme Court justices) is, therefore, made highly contextually salient and it is a small (and, in processing terms, uncostly) leap from this property *indicatively* used to the same property *attributively* used. In (32), by contrast, interpretation does not involve the retrieval of any joint property: the proper names in (32) present their referents individually, not jointly. It must therefore be a more costly step, in processing terms, to retrieve the intended property in (32) than in (31); (32), therefore, cannot be the most relevant stimulus the speaker could have produced, given that (31) (or the equivalent using *they*, *you* etc) offers the same effects for less processing cost.

How about (33) and (34) then? I want to suggest a possible account for these two examples which I hope might well shed light on a wide range of the cases in which there truly does appear to be a disparity between indexicals and proper names. Indexicals and proper names differ in two significant (although non-semantic) ways. Firstly the linguistic meaning of indexicals, the procedures they encode, are, in obvious ways, bound to the context of utterance. This has the by-product that the link between an indexical and its referent is typically short-term: I can only refer to here as *here* while I am actually here, I can only refer to you as *you* while I am addressing you. This is simply a result of the fact that the kind of properties encoded by the meaning of indexicals, properties such as being the speaker or the addressee, are temporary. The link between a name and its referent, by contrast, remains stable across contexts: I can

refer to Tony Blair by uttering *Tony Blair* regardless of where or when I am speaking (so long as my addressee knows who *Tony Blair* names or, rather, so long as he has an individual concept relating to Tony Blair which contains the information *is called Tony Blair*). Moreover, any number of other people can also refer to Tony Blair by uttering *Tony Blair*. The other difference between indexicals and proper names we've already seen: understanding an indexical standardly involves retrieving an egocentric concept, whereas understanding a proper name standardly involves retrieving an encyclopaedic concept. And, as we know, egocentric concepts tend to be repositories of short-term, perceptual information, whereas encyclopaedic concepts generally store long-term, propositional information<sup>15</sup>.

Now how might these differences help us explain the data we're trying to explain? Consider (33) and (34) again: why can *Bill Clinton* in (34) not lead to the same general reading as *me* in (33)? Saul (1997) points out that failure of substitutivity of co-referential names (or something that looks very much like it) is not limited to those contexts usually thought of as referentially opaque. Consider the following examples:

(35) Clark Kent went into the phone booth, and Superman came out.

(36) Clark Kent went into the phone booth and Clark Kent came out.

Although *Clark Kent* and *Superman* are co-referential, there nevertheless appears to be a difference between what is expressed by (35) and (36). Much the same effect can be shown to arise in the linguistic context of examples (33) and (34). Consider:

(37) The Mayor of Metropolis invested Superman with the power to arrest criminals.

(38) The Mayor of Metropolis invested Clark Kent with the power to arrest criminals.

Again there seems to be a difference: the most accessible interpretation of (37) seems to involve the Mayor investing Superman, as Superman, with the power to arrest criminals, and the most accessible interpretation of (38) involves him investing Clark Kent, as Clark Kent, with this power. In other words, the mode of presentation seems to be getting into the proposition expressed somehow. With this in mind, let's return to (34): the most accessible interpretation of (34) is going to involve the Founders investing Bill Clinton, as Bill Clinton, with responsibility for appointing Supreme Court

---

<sup>15</sup> Again, for more extensive argument in favour of the egocentric-encyclopaedic distinction, see Recanati (1993, chapter 7).

justices. A hearer who wants to reach a deferred interpretation of (34), therefore, will have to go to the effort of eliminating the extra information provided by the mode of presentation. This extra effort, it seems to me, may bar (34) from being the optimally relevant stimulus for a speaker who wishes to communicate the deferred interpretation; he could, after all, simply say:

(39) The Founders invested the President with sole responsibility for appointing Supreme Court justices.

requiring no such extra processing effort from the hearer. Why, then, is there no such effect in the case of (33)? The answer to this seems to come in two parts: firstly, since the linguistic modes of presentation associated with indexicals are context-bound, there can be no implication that the Founders made their decision about Bill Clinton under the mode of presentation encoded by *I*; that is, the Founders cannot have decided to give the responsibility to Bill Clinton, under the mode of presentation *speaker of this utterance*. A hearer thus has to expend no extra processing effort, as he does for (34). But secondly, we must ask ourselves why (39), in which there is no deferral, is not a more relevant stimulus than (33), given that it would seem to require less processing effort. I guess the answer may look something like this: there is an expectation that a speaker, talking of a group or property which she instantiates, will use a first-person indexical; a failure to do so will raise expectations of extra effects for extra effort, for instance that the speaker is trying to distance herself from the group or property. Now I'm not suggesting that this is a hard and fast rule, but rather that it is one factor that may affect the relevance of particular stimuli in particular contexts.

The disparity between examples such as (33) and those such as (34) may, then, be the result of a parallel disparity between the stability across contexts of the modes of presentation linguistically encoded by indexicals and proper names. There seems to me to be another possibility, however: that this difference is not (or not only) down to the linguistic modes of presentation of indexicals and proper names but (also) to what Recanati (1993) calls their psychological modes of presentation, i.e. to the different types of concept they correspond to. Why might this be so? The linguistically encoded meaning of proper names is such that understanding a proper name consists in retrieving an encyclopaedic concept; that is, in retrieving a store of propositional information all relating to a single individual<sup>16</sup>. Now consider what that means for the interpretation of an utterance such as (34): if the speaker of (34) intends his addressee to

---

<sup>16</sup> For a more detailed account of the relation between names and encyclopaedic concepts, see Recanati (1993) and Powell (1997).

retrieve a deferred interpretation parallel to that available for (33), he is putting his addressee to gratuitous processing effort. The addressee must first access a body of propositional information all pointing at a single individual and then throw that body of information out again in order to retrieve the intended interpretation. A rational addressee will be left asking: why did the speaker use a proper name, and thus get me to retrieve all this information, if she didn't want me to use any of it? There are, after all, utterances such as (39) which would have required me to go to no such effort for the same effects. However, a speaker who utters (33), intending his addressee to retrieve a deferred reading, is putting his addressee to no such effort: as we've already seen egocentric concepts don't contain large amounts of propositional information, as encyclopaedic concepts do, but, rather, short-term perceptual information. The hearer doesn't, therefore, have to filter out lots of unnecessary information in order to retrieve the intended interpretation.

These two lines of argument are clearly closely allied but it seems to me that they are, nonetheless, distinct. As I hope I have made clear, one relies on the relative stability of the expression-referent link across contexts, while the other relies on the kind of information stored in psychological modes of presentation. It seems to me entirely possible that both of these factors may kick in to bar deferred readings of proper names in contexts such as (34). My only aim here is to demonstrate that there may well be good explanations for this data which do not require a semantics for indexicals such as that proposed by Nunberg; in other words, to show that deferred reference in these examples may well be the product of free pragmatic processes operating at a sub-propositional level.

## 6. Conclusion

I hope, then, that I have managed to convince you of the following: that deferred reference may well not be the single phenomenon Nunberg would have us believe; that there may well be two entirely distinct types of deferral, what I have been calling semantic deferral and pragmatic deferral, at work in Nunberg's examples; and that the arguments Nunberg uses in support of a semantic account for the entire range of his examples may not be quite so robust as they at first seem. There may, in fact, be perfectly good pragmatic accounts of deferral, grounded in a proper understanding of the cognitive background which underlies reference.

But we started this whole discussion with the idea that Nunberg's deferred uses pose a grave threat to the tenets of direct reference theory, at least so far as indexicals are concerned. Where, then, does direct reference stand at the end of our discussion?

Surprisingly, this seems to be the subject of considerable agreement between the opposing camps. Recanati (1993) argues that Nunberg's examples are only a threat to direct reference on a strong reading, a reading on which the distinction between directly referential expressions and descriptions is taken to be, essentially, that directly referential expressions lead to singular propositions while descriptions lead to general propositions. For Recanati, however, a weaker reading of direct reference remains unchallenged by Nunberg's data; on that reading, the difference resides not in directly referential expressions leading to singular propositions but rather in the fact that the process of interpretation of directly referential expressions must, mandatorily, pass via an individual, even if interpretation then continues, guided by general pragmatic constraints, past that individual.

And Nunberg finds much to agree with in this:

I think Recanati is right to insist that the central claim of direct-reference theories - "the thesis of the *asymmetry* between referential and nonreferential terms" - survives the challenges posed by the observations I have offered here.

Nunberg (1993, pp 32-33)

The effect of Nunberg's objections, then, has been this: to highlight that direct reference cannot be a theory of propositional content; but, after all, once we understand the integral role that pragmatic inference plays in retrieving the proposition expressed, we could never have expected direct reference to be any such thing. Direct reference should, rather, be taken as a theory about differences in encoded meaning and, as such, it remains, as far as I can see, unscathed by Nunberg's attack.

## References

- Bach, K. 1992. Intentions and demonstrations. *Analysis* 52, 140-146.  
 Berckmans, P. 1990. Demonstrative utterances. *Philosophical Studies* 60, 281-295.  
 Blakemore, D. 1987. *Semantic Constraints on Relevance*. Oxford: Blackwell.  
 Carston, R. 1988. Implicature, explicature, and truth-theoretic semantics. In R. Kempson (ed) *Mental representations: the interface between language and reality*, 155-181. Cambridge: Cambridge University Press.  
 Carston, R. 1996. Enrichment and loosening: complementary processes in deriving the proposition expressed. *UCL Working Papers in Linguistics* 8, 61-88.  
 Donnellan, K. 1966. Reference and definite description. *The Philosophical Review* 75, 281-304.

- Fauconnier, G. 1994. *Mental Spaces*. Cambridge: Cambridge University Press.
- Grice, P. 1969. Vacuous names. In D. Davidson and J. Hintikka (eds) *Words and Objections*, 118-45. Dordrecht: Reidel.
- Grimberg, M. L. 1996. *Against rigidity: an investigation of the semantics and pragmatics of indexicality*. Ph.D. thesis, UCL. Page references are to revised version.
- Kaplan, D. 1977. Demonstratives. In J. Almog, J. Perry and H. Wettstein (eds.) *Themes from Kaplan*, 481-563. New York: Oxford University Press.
- Kaplan, D. 1989. Afterthoughts. In J. Almog, J. Perry and H. Wettstein (eds.) *Themes from Kaplan*, 565-614. New York: Oxford University Press.
- Nunberg, G. 1978. *The Pragmatics of Reference*. Bloomington: Indiana University Linguistics Club.
- 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.
- Powell, G. 1997. Names. UCL MA Dissertation. Unpublished ms UCL.
- Recanati, F. 1993. *Direct reference: from language to thought*. Oxford: Blackwell.
- Saul, J. 1997. Substitution and simple sentences. *Analysis* 57.2, 102-108.
- Wilson, D. and D. Sperber. 1993. Linguistic form and Relevance. *Lingua* 90.1/2, 1-25.