

Descriptive Indexicals and Indexical Descriptions

Geoffrey Nunberg

1. Explaining the Referential–Attributive Distinction

As Donnellan (1966) and many others have pointed out, a sentence like (1) has two readings:

- (1) The person who's parked in front of the restaurant is in a hurry.

On the attributive reading, the description *the person who's parked in front of the restaurant* is interpreted as a quantifier: it says that the unique person who is parked in front of the restaurant is in a hurry, with no implication that the speaker has a particular person in mind—maybe she simply inferred his haste from his choice of a parking spot. On the referential reading, the description picks out a particular person—say Jones, who is waiting for his car at the valet parking station with visible impatience.

The question is, is (1) semantically ambiguous? Not if you accept the Russellian view, which holds that the description in (1) is always interpreted as a quantifier, with the referential interpretation arising via a conversational implicature that arises from the hearer's assumption that the speaker is trying to say something relevant and has a particular person in mind. In that case the utterance does not literally say anything about Jones, though the speaker may succeed in identifying him by uttering it.¹

For referentialists, on the other hand, (1) is ambiguous. On the referential reading, *the person who's parked in front of the restaurant* functions basically the way an indexical does. This is the effect achieved by Kaplan's 'Dthat' (see Kaplan 1979), an operator that turns descriptive content into a function from contexts to individuals. In that case

For comments on this chapter and on earlier versions of this work, I am grateful to François Recanati, Cleo Condoravdi, and the two anonymous reviewers of this chapter.

¹ This is by no means the only possible view of the relation between the readings. Bezuidenhout (1997) and Recanati (1993) have suggested that both referential and attributive readings are generated pragmatically from an under-specified semantic representation. For my purposes, though, the only important distinction is between those who argue that referential uses contain something like a hidden indexical in their semantic representations and those who do not—on this cut of things, both Bezuidenhout and Recanati would figure in the latter group.

an utterance of (1) with the description used referentially is semantically identical to an utterance of a sentence like *He is in a hurry*—both utterances literally express a singular proposition, in which the content of the referring expression plays no role.²

2. The Argument from Misdirection

Is this an empirical dispute? It feels as if it should be, but there does not seem to be a lot of hard empirical evidence that bears on it, and the arguments that people have offered have not been conclusive, to put it mildly. In introducing the distinction, Donnellan leaned heavily on a line of argument that Neale (1990) calls the argument from misdescription and Recanati (1993) calls the argument from improper use: the observation that a speaker can sometimes succeed in identifying the individual she has in mind with a description like *the person who is parked in front of the restaurant* even if it should turn out that the description does not actually apply to him. So suppose that Jones is the visibly impatient person who is holding the parking stub for the car that is parked in front of the restaurant, but that Jones is in fact merely getting the car for his friend Smith, who is inside paying the bill in an unhurried way. Even so, Donnellan asserts, we might use the description to pick out Jones. In this sense the description is regarded as analogous to a demonstrative accompanied by an inaccurate gesture, which might nonetheless enable someone to pick out the object that the speaker has in mind. Whereas if the sentence merely says that there is a unique person who is parked in front of the restaurant and who is in a hurry, it is held to be mysterious as to how we might use it to get to someone who does not satisfy that description.

The argument is not very persuasive. For one thing, it rests on people's intuitions about what is actually said by an utterance of (1), which is never a very sound basis for analysis. (The more you listen to people's intuitive judgments about the 'literal meanings' of sentences, the more you come to think that the notion is more jurisprudential than linguistic.) And in any event, people do not always operate the way Donnellan assumes they do when it comes to interpreting these utterances. True, we often allow a speaker a certain descriptive latitude in order to get on with the conversation, but we are not obliged to do so, particularly if the content of the description is controversial. Suppose that a White House press officer says, 'The man who won the majority of votes in Florida will be visiting Tallahassee on Sunday', with the intention of informing

² For the present purposes, I shall assume that direct referentiality (or as I will be calling it, direct interpretation) is a semantic property of indexical expressions as such. In fact Roberts (2002) argues that the direct referentiality of demonstratives, at least, follows from their dependence on a completing demonstration (where 'demonstration' is construed broadly). She notes that a Kaplan-type (1989*a*) analysis of demonstratives cannot deal with various problems raised by Heim (1985), or with the use of demonstratives that refer to discourse antecedents, such as in *She ignored his warning, and that got her in trouble*. It is not clear, though, how or whether this analysis could be generalized to 'pure' indexicals like *I* or *today*.

the public about President Bush's schedule. My own guess would be that a good number of Americans would respond to that by saying, 'No, he'll be spending the weekend in New York with Tipper.' And from the point of view of a referentialist, it is hard to explain why anyone would have this reaction, since on his view the speaker has not really said that Bush won Florida.³

Then, too, as both Kripke (1977) and Neale (1990) observe, these cases of misdescription are not limited to the uses of definite descriptions. Kripke gives the example of someone who points at Jones under the misapprehension that he is Smith, and it is easy enough to construct analogous examples using quantifiers. Imagine an out-of-it dad who says at his 11-year-old daughter's birthday party 'Every girl who is wearing a button with a picture of Christina Aguilera is attending my daughter's party' when in fact the invitees are all wearing badges with pictures of Britney Spears. Even so, a hearer could doubtless figure out who's being talked about, though clearly there is no question of saying that the expression *Every girl who is wearing a button with a picture of Christina Aguilera* contains a concealed demonstrative.

In fact, as Wettstein (1981), Recanati (1993), and others have pointed out, the argument from misdescription is not really germane to the larger issue of whether sentences like (1) are semantically ambiguous. Wettstein maintains that it is not essential to the referentialist view to accept that referentially used descriptions can function semantically to pick out an individual who does not satisfy them. If (1) is used in a context in which Jones is holding Smith's parking stub, you could say that the description literally picks out Smith, while noting that the hearer can nonetheless recover the speaker's intended reference, and use that as a basis for further action if it is appropriate. In short, the argument from misdescription hardly gives us any grounds for resolving the controversy one way or the other.

3. The Argument from Incompleteness

This takes me to the second argument that has been appealed to here, which has its origin in Strawson's observation that we can use a noun-phrase like *the table* to identify a particular table even though it is not the only table in the world. On a Russellian analysis, the argument goes, this is puzzling unless we can provide a completing sense to distinguish the table from all others, and that creates a number of problems—how do we know which completing sense is pertinent, for example? Whereas if we treat the definite description *the table* on the model of a demonstrative description like *that table*, we can understand how it can be used to identify a particular salient object.

³ Or sometimes we can nod to the intended reference without endorsing the content of the description. In the movie *Quo Vadis*, there's a scene where a Roman general (Robert Walker) takes a Christian hostage (Deborah Kerr) to a banquet at Nero's palace. 'Have you ever seen your god up close?' he asks her? 'No,' she answers, 'I have never seen Nero up close', deftly demurring from the content of the description.

I confess that I have never been able to make much sense of this argument. It seems plain that when someone in an office says *The table is covered with books*, we interpret the description *the table* as applying to the unique table in the domain of things that the speaker is talking about. I do not mean to suggest that we have a wholly adequate account of domains of discourse, but this ground has been covered in an extensive literature, particularly as regards the interpretations of definite noun-phrases.⁴ And whatever form that account takes, it will have to explain the use of ‘incomplete’ quantifiers, as well, as Neale (1990), Salmon (1991), and others have pointed out. For example, we often say things like *Everyone left*, where there is no question of proposing a demonstrative reading.⁵

What is more, it is hard to see how the hidden indexical analysis of referential descriptions is going to let us off the hook in cases like this. It may be that the people who have taken this view have been misled by the familiar examples of referential descriptions—phrases like Donnellan’s *the man with the martini* or Strawson’s *the table*. These typically involve situations where the intended referent is salient in the immediate context, so that you could imagine referring to it with a demonstrative phrase like *that man* or *that table*. But ‘incomplete’ referential descriptions can be used to pick out remote things as well, where no demonstration would be possible, and in those cases it is hard to see how they could succeed in referring unless some uniqueness condition was satisfied. Say one of the Duke of Plaza-Toro’s footmen says to another *The Duke is visiting Venice*, using the definite description in a referential way. If it was not clear that there was a unique duke in the domain of discourse, how would it help to assume that the description contained a hidden indexical—how would you determine which duke the speaker had in mind without requiring a completing description or a suitable restriction on the domain of discourse?⁶ In short, the argument from incomplete descriptions simply has no bearing on the status of these expressions.

⁴ In this connection, see among many others Heim (1982), Kempson (1986), Löbner (1987), Kadmon (1987), Neale (1990), Birner and Ward (1994), and Roberts (1995, 2002). Among other things, these sources all discuss the possibility that a proper analysis of definites requires criteria other than or in addition to uniqueness, such as familiarity. For our purposes, though, it does not matter which analysis of definites we adopt, so long as definite descriptions come out as quantifiers.

⁵ Ostertag (1998: 26) argues that this approach does not explain why descriptions like *the record* do not seem to be equivalent to descriptions like *every record* when used in incomplete contexts: [A]n utterance of “the dog is barking” will generally be interpreted relative to a very small domain, while an utterance of “every dog” is barking will be assigned to a larger domain. . . . If the equivalence cannot be assumed, then the theory that treats “the *F*” as a quantifier. . . . is not Russell’s theory of descriptions’. I have no brief to make for the classical Russellian theory of descriptions, but I do not see the justification for Ostertag’s conclusion, since the presuppositional difference between ‘The dog is barking’ and ‘Every dog is barking’ seems to have a straightforward Gricean explanation. In general, we are obliged to couch our utterances in such a way as to accommodate as best we can the actual presuppositions of the context. For example, if it is common knowledge that Cleo is not an American citizen, I could felicitously say (i) but not (ii): (i) If Cleo were an American, she wouldn’t need a Green Card; (ii) If Cleo is an American, she doesn’t need a Green Card. It is not that (ii) would not be truthful, but in the circumstances it would be misleading, since it ignores what is commonly known. The difference between ‘The dog is barking’ and ‘Every dog is barking’ is perfectly analogous to this: while the version with *every* is consistent with the existence of a single dog, it is not a felicitous thing to say when it is common knowledge that there’s only one dog around.

⁶ See Roberts (2002) for reasons for assuming that a uniqueness presupposition is semantically associated with the uses of demonstratives.

4. The Argument from Methodological Preliminaries

Given the inconclusiveness of these arguments, it might seem that the status of referentially used descriptions is not really empirical at all.⁷ That seems to be the position of Kripke (1977), who argues that the issue should be thought of in methodological terms. He asks us to imagine a stipulated variety of English in which definite descriptions could have only quantifier interpretations. Even in that case, he argues, we can see how referential readings could arise on Gricean principles. And for that reason, according to Grice's 'modified Occam's razor' principle ('senses are not to be multiplied beyond necessity'), there is no ground for introducing a semantic ambiguity.

5. The Descriptive Uses of Indexicals

I have no quarrel with the form of Kripke's argument, but it is always more satisfying to have data to fall back on. In this section I shall discuss some linguistic observations which seem to favor the Russellian analysis of referential descriptions, but which up to now have not played any role in these discussions. The phenomenon I have in mind involves what I shall call the descriptive uses of indexicals (the term was suggested by Recanati (1993) to describe some examples I gave in Nunberg (1993) in the service of another point):

- (2) *Today* is always the biggest party day of the year.
- (3) (Pointing at a TV that is showing a lopsided Stanford–Berkeley basketball game) *This* is usually an exciting game, but not tonight.
- (4) Condemned prisoner: I am traditionally allowed to order whatever I want for my last meal.

Unlike indexicals of the more frequently picked-over variety, the ones in (2)–(4) do not refer to individuals, but rather seem to contribute properties to the utterance interpretations. In (2), we take the interpretation of *today* as some property that is instantiated by the day of utterance—of being November 1, the day that fall midterms are over, the day following the homecoming game, or whatever. (Note though that it is not necessary that the hearer be able to identify the particular property the speaker has in mind—for his purposes, it might be sufficient to know that it is not a good day to try to go for a quiet drink at a bar near campus.)

⁷ I am ignoring here another set of arguments that people have offered in favor of the thesis that sentences like (1) are ambiguous. These have to do with the interpretations of pronouns in examples like (i): The horse which I bet on won. Hans had foreseen it. As the argument goes, only if we treat the subject of the first sentence in (i) as having a reading as a singular term can we explain the reading of the utterance in which Hans had foreseen that Secretariat would win without knowing that the speaker had bet on that horse. For an account of the deficiencies of this argument, see Heim (1991).

Example (3) works in the same way. The demonstrative *this* cannot refer to the very contest the speaker is pointing at, which is plainly boring; instead it contributes the property of being a Stanford–Berkeley game. And in (4), the pronoun *I* doesn't refer to the speaker, since obviously there couldn't be any traditions that deal specifically with *his* last meal; instead it refers to the role he exemplifies.

In calling these uses of indexicals 'descriptive', I do not want to give the impression that their interpretations are equivalent to nonindexical descriptions like *the day of our conversation* or *the game on TV* on their normal Russellian readings. To see the difference, we have only to recast the sentences so that they contain a modal or counterfactual context:

- (5) If we were talking to each other on November 2 instead of now, the day of our conversation would be the biggest party day of the year.
- (6) If we were talking to each other on November 2 instead of now, today would be the biggest party day of the year.
- (7) If you had turned to Channel Four as I asked, the game on the TV would be exciting.
- (8) If you had turned to Channel Four as I asked, this game would be exciting.

The descriptions in (5) and (7) can be re-evaluated relative to the contexts established by the conditional, but the indexicals in (6) and (8) cannot: *today* is still anchored in the actual day of speaking, and *this game* is the one whose current episode is on the channel we are watching now, not the one playing on Channel Four. In this regard, the indexicals here are no different from indexicals that pick out specific days or individuals, which have to be evaluated relative to the context of utterance. So whether or not these indexicals here are 'directly referential' in the strict sense, they are nonetheless what we can think of as directly interpreted, and for this reason they cannot be scoped by other operators.⁸

For the present purposes, we could describe the interpretations of the indexicals in (2)–(4) in any of several ways—as properties, as higher-order entities of some kind, or as under-specified discourse referents or constituents of abstract situations. I am willing to be agnostic about that—my chief interest is in the way these utterances contrast with utterances that contain descriptions. For example, suppose that tomorrow is also my daughter's eleventh birthday. Then I could use the description *Sophie's eleventh birthday* to refer to that day, in what would clearly be a referential use of the expression:

- (9) Sophie's eleventh birthday falls on a Tuesday.

⁸ The fact that indexicals are directly interpreted even when they contribute properties seems to pose a problem for the analysis of Bezuidenhout (1997), who offers a symmetrical treatment of descriptions and indexicals: both, she says, begin with under-determined representations from which either 'attributive' or 'referential' readings can be derived. In that case, though, I am not sure how the differences illustrated in (5)–(8) could be explained.

But I cannot substitute that description for the indexical in (2):

- (10) ?Sophie's eleventh birthday is always the biggest party day of the year.

That is, the descriptor *Sophie's eleventh birthday* can only apply to a unique date, which is inconsistent with the use of *always*. And by the same token, suppose that the prisoner who utters (4) has already had his head shaved in preparation for his execution. Then his guard could say (11) but not (12):

- (11) The prisoner with the shaved head gets the *spaghetti alla Carbonara*.
 (12) ?The prisoner with the shaved head is traditionally allowed to order whatever he wants for his last meal. (\neq 11)

That is, the description cannot contribute a property of the person it applies to—other than the very property that is explicit in its descriptive content, of course.

Suppose that we accepted the referentialist line and said that the definite descriptions in these utterances contain Kaplan's *Dthat* operator or something along those lines. Then why do they not have the same range of readings that true indexicals like *I* and *that* do? In each case you would have an expression containing a (real or crypto-) indexical that picks a certain individual out of the context and creates a singular proposition. But when you use a description like *the prisoner with the shaved head*, you are required to stop with the individual that the description applies to—you cannot go on to take one of his salient properties and make it the interpretation of the expression.

On the other hand, if we assume a Russellian analysis of the definite description in a sentence like (11), then it is not odd that the mechanism that produces a descriptive reading does not come into play. The individual that the description applies to is not a constituent of the proposition, so you do not have access to any of his properties, apart from those that are made explicit by the description itself. On the face of things, this seems a strong empirical argument for saying that referentially used descriptions have a Russellian interpretation, and that the sentences that contain them have the same truth conditions as sentences containing attributively used descriptions.

6. A Pragmatic Explanation?

Let me deal now with some potential objections to this argument. In the first place, someone might say that the descriptive reading available for sentences like *Tomorrow is always the biggest party day of the year* is somehow a marginal or secondary phenomenon. This is the line taken by Recanati (1993), who suggests that the readings of indexicals in sentences like *Today is always the biggest party day of the year* are derived via a conversational implicature from literal readings in which they have their 'normal' references, referring to days, persons, or whatever.

That approach runs into a number of difficulties, though. For one thing, it is hard to see how a sentence like *Today is always the biggest party day of the year* could have any kind of literal meaning in which *today* picks out a unique day, say November 1, 2000, and says of it that it is always the biggest party day of the year—there simply is not any coherent proposition that we could associate with such an interpretation.⁹ And for that reason, it is not possible to suspend the descriptive reading of indexicals in sentences like these, which is usually taken to be a requirement for postulating a conversational implicature in the first place:

- (13) Some of the boys left—in fact, all of them did.
- (14) My love is a red, red rose, and I mean that literally.
- (15) Jack took off his trousers and got into bed, but not in that order.
- (16) ?Today is always the biggest party day of the year, but only if it's November 1, 2000.

In any event, all of this is beside the point. Even if we did say that the reading in a sentence like (2) arose via a conversational implicature, we would be left with having to explain why a similar implicature does not arise when we replace the indexicals like *I* and *today* with referentially used descriptions like *the prisoner with the shaved head* and *Sophie's eleventh birthday*. On any version of the referentialist hypothesis, after all, the indexicals and descriptions would have exactly the same literal or primary interpretations: each would contribute the individual it applies to.

7. The Maxim of Manner to the Rescue?

At this point, a referentialist might suggest that the difference in the implicatures available for indexicals and referentially used descriptions follows somehow from the intervention of Grice's maxim of manner. That is, while the two types of expressions have identical semantic import, there is something about the fact of choosing one or the other that either makes the descriptive reading available only for indexicals, or blocks it for referential descriptions.

It is true that conversational implicatures can attach to the choice of a description rather than an indexical as a means of identifying something, as in (17)–(20):

- (17) (Spoken on April 9) The doctor can see you tomorrow.
- (18) (Spoken on April 9) ?The doctor can see you on April 10.

⁹ It may be that we can interpret Recanati's analysis as primarily a thesis on the way these utterances are processed. That might be true, but there is no reason to assume that the semantics of the sentence should recapitulate the psycholinguistic mechanisms that speakers bring into play to figure out what it means.

- (19) (Spoken by an American) When we Americans travel, we like to have all the comforts of home.
- (20) (Spoken by an American) ?When Americans travel, they like to have all the comforts of home.

Both pairs demonstrate a general conversational principle to the effect that when an indexical *can* be used to identify something, we expect that it *will* be used. It is generally odd for someone speaking on April 9 to refer to the following day as ‘April 10’, when the use of *tomorrow* would pick out the same day, though certainly we would not claim that the sentence was false if she did. And similarly, it is a reasonable inference that the speaker of (20) is not an American, since if she were she would have been expected to use a noun-phrase containing *we*. But even so, the truth-value of the sentence does not stand or fall on the nationality of its speaker. And when the speaker flouts this rule, that choice usually occasions a conversational implicature. When somebody refers to himself by his name or with a description like ‘your little brother’, we usually assume that he has some ulterior motive for doing so.¹⁰

We will see in the final section that this principle does in fact have some bearing on explaining the descriptive interpretations of indexicals. By itself, though, it will not explain the difference we are interested in—not if indexicals and referential descriptions have the same type of interpretations, as referentialists claim. To invoke the maxim of manner, after all, we have to be able to assume that the speaker has deliberately chosen one form of expression over another that would have the same truth-conditional meaning. But most referential uses of descriptions cannot be paraphrased by indexicals, particularly if the intended referent is not actually present or salient in the context. The footman’s announcement ‘The Duke of Plaza-Toro will arrive at noon’ has no equivalent using an indexical, so there could be no implicature associated with his choosing a description to refer to his master. (There might be an implicature associated with using one particular description rather than another, of course, but that would not be relevant to explaining the behavior of descriptions as a class.)

What is more, when we look at cases where a referential description and an indexical phrase can both be used to pick out the same entity, it is hard to see how we could appeal to any difference in mode of expression to explain the differences in available interpretations. Suppose you have been visiting me at my house and need to pick up some cash on the way home. You could ask me either (21) or (22), to pretty much the same effect:

- (21) Is there an ATM around here?
- (22) Is there an ATM in the neighborhood?

But now suppose we are trying to get some cash in an unfamiliar neighborhood, and we notice an off-track betting parlor. I could then say (23):

¹⁰ I say ‘we usually assume’ because there are a few situations where the rule that requires the use of indexicals is suspended—e.g. in news stories in which reporters are expected to refer to themselves as ‘this correspondent’.

(23) There is usually a cash machine around here.

In this context, (23) means roughly the same as ‘there is usually a cash machine in the vicinity of an off-track betting parlor’. But note that this interpretation would not be available if I said (24) instead:

(24) There is usually a cash machine in the neighborhood.

That is, the description *the neighborhood* cannot contribute a property of the place we are in, but only the place itself. But the only difference between the two utterances that could explain this fact is that in one case the place is identified descriptively, while in the other it is identified indexically (see Appendix). And if there is no semantic difference between these two modes of referring, what could give rise to a conversational implicature via the maxim of manner?

In the light of examples like these, it is hard to imagine how pragmatics could come to the rescue of a hidden indexical analysis. At this point, then, the burden of proof is clearly on the referentialist: as Grice points out, the definitive test for a conversational implicature is a demonstration that it be capable of being worked out.

If the difference between the available readings of indexicals and referentially used descriptions cannot be explained pragmatically, it must rest on a semantic difference between the two types of expressions. But where could this come from? I suppose a referentialist could still try to claim that the hidden indexical that figures in a referential use of the description *the man with the martini* was linguistically marked with a feature that arbitrarily blocks the availability of a descriptive reading but otherwise left the interpretation untouched. (We don’t know a lot about hidden indexicals, after all—nobody has ever seen one.) But this is a desperate expedient. It amounts to saying that there is no good explanation for the difference between the available readings of indexicals and referentially used descriptions—that is just how nature made them.

8. Where Do Descriptive Readings Come From?

The observations I have made up to here are sufficient to show that definite descriptions simply do not behave like indexicals and demonstratives, and that the hidden-indexical analysis cannot be sustained. But this leaves us with an unanswered question: why do only demonstratives and indexicals permit descriptive readings, and permit them, it seems, in every language we look at? Or conversely, why can descriptions not have readings like these?

To answer these questions I should say something about how demonstratives and indexicals acquire these readings. In Nunberg (1979, 1993), I described examples like these as instances of ‘deferred indexical reference’, a notion based on Quine’s ‘deferred ostension’. The idea is that reference in these cases proceeds via a two-stage process. We

first identify the *index* of a term—that is, either the demonstratum of a demonstrative or the contextual element that an indexical picks out in virtue of its linguistic meaning—and then proceed to identify the referent of the term, which is something that stands in a salient functional relation to the index. So if you point at a copy of a newspaper and say ‘That was bought by Rupert Murdoch for five hundred million dollars’, you can succeed in referring to a publisher in virtue of the functional relation between newspaper companies and their publications.¹¹

I still think this is the right structural approach to explaining how references are identified in these cases of demonstrative metonymy—that is, cases where you point at one thing to identify something that stands in a relation of contiguity to it. But I have come to think that it is misleading to use the term ‘deferred reference’ to describe these examples. It seems just as straightforward to assume that the remote referents picked out by the demonstratives here are themselves made present in the context by a demonstration that accompanies the use of a demonstrative. (I am taking ‘demonstration’ in a broad way, as something implicit or explicit in the context or in the mention of the demonstrative term itself.) When you call to someone’s attention a copy of a newspaper, that is, you are also calling to his or her attention a newspaper publisher. And when you point at a picture of Carnap (to take an example used by Kaplan), you are also pointing at Carnap himself.¹² In this sense there is nothing ‘deferred’ about these uses; they are more on the order of other cases of ambiguous or under-determined demonstrations.

One important reason for preferring this way of accounting for demonstrative metonymy is that the phenomenon does not extend to the uses of demonstratives whose antecedents are elements introduced into the discourse by earlier referring expressions, rather than things present in the context of utterance itself. For example, I can point at Tiger Woods and say (25):

(25) That’s what I want to take lessons in.

But this use of the demonstrative does not have a parallel in (26):

(26) ?Whenever Mary sees Tiger Woods on TV, she wants to take lessons in that.

When we refer to Tiger Woods by name, that is, there is no demonstration, explicit or implicit, that can pick out the virtual golf bag that he is carrying. And for the same reason, these metonymic references are not available for pure indexicals, either. Tiger Woods cannot use the pronoun *I*, for example, in a way that parallels the use of *that* in (26):

(27) ?Mary wants to take lessons in me.

¹¹ There are further conditions on this kind of reference that I will not go into here; see Nunberg (1993) for a discussion of these.

¹² It is notable how often philosophers use examples of people pointing at photographs to identify the people they picture or of maps to identify places, almost always without seeming to notice the apparent indirection of reference in these cases.

Or suppose two teachers at a nursery school are talking about how accomplished the mothers of their charges are. One of them can get away with pointing at a little boy and saying, ‘She’s a famous athlete.’ But when the mother visits the school and wants to brag about her son, she cannot say, ‘I am one of your brightest students.’¹³ Here again, then, we see that metonymic references require a demonstration, whereas *I* can only pick out something in virtue of the role it plays in the utterance.

So it would seem as if the notion of deferred reference is a mistake, if by that we mean that indexicals or demonstratives can sometimes refer to something other than what we can get to directly via their linguistic meanings or the demonstrations associated with them. But in that case, where do the descriptive readings of indexicals like *I* and *today* come from? The only remaining possibility is that the interpretations of these uses of indexicals are the very things that their linguistic meanings pick out of the context. When someone says ‘Today is always the biggest party day of the year’, that is, we have to assume that the interpretation of *today*—a day-type or day-property, say—is directly picked out by the linguistic meaning of the expression.

9. Granularization of the Context

In order to make sense of this conclusion, I should say something about the way we talk about contexts and the meanings of indexical expressions. There is a more-or-less standard view of this which cuts across theoretical differences, and which goes something like this. Every utterance is associated with a certain set of elements or coordinates that bear a certain relation to it—its speaker, addressee, time, place, world, salient objects, and so on. And the linguistic meanings of indexical expressions like *I*, *here*, and *today* provide descriptions that pick out those elements in virtue of their relation to the utterance. The content of those meanings may not figure as part of the content of the utterances that contain the expressions, but we evaluate it relative to a particular context in the same way we might evaluate the use of any other description.

In this sense, at least, the context of utterance is a domain like any other, and the question of determining what entity or element is picked out by the linguistic meaning of an indexical should be subject to the same considerations that we apply in determining the references of other descriptions. That is, there is no appreciable difference between evaluating the meaning of *today* in an utterance of (28) and evaluating the meaning of ‘the day of the last Giants game’ in an utterance of (29):

¹³ Just to fill in the pattern, note that pure indexicals can also have descriptive readings when they have discourse antecedents. It is not easy to construct examples that make this point unequivocally, but it can be seen in (i): Whenever Prisoner 28528 and I have tried to order the last meal that tradition specifically allows *us*, we have been rebuffed, and so have all the other prisoners on death row. The token of *us* in (i) does not pick up its reference directly from the context (among other things, that hypothesis would leave us unable to how to explain the sloppy reading of the VP ellipsis, which would require that *us* be treated as a variable). But the descriptive reading is available nonetheless.

- (28) Today is cloudy.
 (29) The day of the last Giants game was cloudy.

In particular, the evaluation in both cases requires that we restrict the domain appropriately. One important consideration, obviously, is that the domain is usually restricted to a set of relevant entities—with (29), for example, we take the domain to be the set of all Giants games played in the contextually relevant year, not all the Giants games that will ever be played.

But we may also want to restrict the domain in other ways, to include only the relevant properties of the individuals in question. Reimer (1998*a*) makes this point by way of explaining the interpretation of a sentence like (30), as uttered by a newscaster describing a local Daughter's Day event:

- (30) Every daughter present was accompanied by a parent.

Clearly, Reimer says, an assertion of (30) can be true even if some of the parents are females, and hence daughters themselves. As she puts it (p. 103):

After all, if the context can restrict the domain so as to include only the 'relevant' individuals, then why can't it similarly restrict the properties of those individuals, so as to include only their 'relevant' properties? . . . [T]o say that certain individuals in the domain are assigned some of their properties, but not others, is not to say they don't have those other properties. Rather, it is to say that those other properties are contextually irrelevant, and therefore are not assigned to the individuals in question.

Once we acknowledge that only certain conversationally relevant properties of individuals figure in the domain, it follows that the conversational purposes can determine what counts as an individual, as well. That is, there can be only as many individuals in the domain as are individuated by the conversationally relevant properties—the context is subject to what we can think of as a contextual granularization.

This sort of restriction of domains is very familiar, even if it is described in various ways. Take a sentence like

- (31) The burglars came in through the window.

At first blush, (31) might seem analogous to Strawson's *The table is covered with books*, where the 'incomplete' description applies to the unique table in the domain. But a speaker can utter (31) without being in a position to provide an identifying description of the particular window that the burglars entered through, nor does anyone care which one it was. In fact, (31) would make a true assertion even if it should turn out that there were two burglars who entered through different windows. Some people have taken sentences like (31) as exceptions to the uniqueness condition for definites, but I think this misses the mark.¹⁴ Rather, we would say that for conversational

¹⁴ Kadmon (1990) argues that the referent of a definite description must be 'not relevantly differentiable' in context, and Birner and Ward (1994) offer a similar analysis of the italicized NPs in (i) and (ii): (i) As soon as my cousin arrived in Santiago, she broke her foot and had to spend a week in *the hospital*; (ii) Your ten o'clock

purposes, the properties that distinguish one window from another in the house in question are simply irrelevant—the only important thing is that the burglars chose this mode of entry rather than the door or the chimney.

Examples like (31) have been widely discussed, though often with a sense that the phenomenon is restricted to a relatively small class of words and that it may involve some idiomaticity (see e.g. Abbott 1999). But in fact the phenomenon is a great deal more general than that. Take the uses of the descriptions in (32) and (33), which are suggested by some of the examples of descriptive indexicals that I gave earlier:

(32) The Stanford–Berkeley game is usually exciting.

(33) The biggest party day usually comes after midterms are over.

It is odd to suppose that (32) and (33) involve exceptions to the uniqueness criterion for definites, as if the words *game* and *day* are invariably restricted to referring to particular contests and particular days. Instead, we would say that the properties that distinguish one particular game or day from another simply are not represented in the discourse model—that is, that there is no way in terms of the contextually relevant properties to distinguish the Stanford–Berkeley game that occurred on January 10, 2001 from Stanford–Berkeley games that occurred on other dates. In the discourse model determined by the context, there is only a single thing to which the description *the Stanford–Berkeley game* can apply. As I suggested earlier, it does not matter for the present purposes whether we describe this interpretation as a set of properties or generalized quantifier, or as the constituent of an abstract situation. The important thing is that these elements are simply available in the discourse models for certain contexts in the same way that particular games and days are available in the models relevant to other contexts.

In this regard, the context is no different from other domains: we recognize only as many distinct contextual elements as are individuated by the conversationally relevant properties. This point may be clearest when we consider the way we interpret indexicals like *here* and *now*. Suppose you are at a reception after a conference talk and someone asks you:

(34) Is John here? I could use a ride back to my hotel.

On the natural interpretation of (34), the reference of *here* is the room or immediate area where the speaker and hearer are located—a place small enough so that it would not be too inconvenient to get hold of John before he drives away. That is, the conversational purposes impose a certain spatial granularity on the domain of discourse, such

appointment said he'd be late because he had to stop first at *the bank*. Both of these analyses assume that these examples do involve some weakening of the condition of uniqueness. Other analyses follow lines more like the one I am taking here; see e.g. Epstein (2000), who analyzes sentences like these by reference to 'roles', and in particular Löbner (1987), who takes these NPs as referring to unique constituents of abstract situations, an analysis that is equivalent to the one I am assuming here. In any event, none of these points is crucial to the argument I am making in this chapter, since the crucial point for me is only that all definites be analyzed as quantifiers of one sort or another.

that two places are distinct only when the difference between being in one or the other of them has significant consequences for getting a ride home from somebody.

But relative to other conversational purposes, the granularity would be different. Suppose that instead of (34), the speaker says:

(35) I need a ride home; the bus service here is terrible.

In this case, the properties relevant to individuating distinct places are those that might entail differences in bus services—being in different towns or neighborhoods, for example. And in other cases, *here* can refer to an area as large as a nation or even a planet, in the same way that *now* can refer either to an instant or an era.

This is not to say that the interpretation of *here* in (34) makes any mention of neighborhoods—the neighborhood just happens to be the smallest relevant chunk of the context that is available for reference. Of course (34) and (35) do not involve distinct uses or meanings of *here*. In both cases, *here* picks out its referent in virtue of the same linguistic meaning, ‘the place of utterance’. The only difference between the utterances is in that we appeal to different criteria to determine what counts as a distinct place.

Examples like these show that we have to evaluate the linguistic meanings of indexicals relative to a certain granularization of the context. In the case of (34) and (35), this is a question of individuating distinct locales in a continuous region.¹⁵ But we can invoke the same principle to explain the descriptive readings of indexicals. Take *Today is always the biggest party day of the year*. *Today* picks out the day of utterance, just as it always does. But in the relevant discourse model, there are only as many distinct days as are individuated by the properties relevant to the conversational purposes, just as there are in a sentence like *The biggest party day of the year is usually the day after midterms*. That is, there is no way to individuate the day of utterance from others that have the same relevant properties. In this sense, the reference here is not ‘deferred’: the linguistic meaning ‘the (calendar) day on which the utterance takes place’ is satisfied by the only day in the domain of discourse that corresponds to the time of utterance—but relative to the conversational purposes, that day simply does not have the properties that differentiate the actual day of the utterance from others that fall on the same date.

We will tell the same sort of story about an utterance like (4), *I am traditionally allowed to order whatever I want for my last meal*. Here, the domain of discourse from which the reference of *I* has to be drawn can contain only the sorts of entities that could be the subjects of traditional dicta—sets of personal properties, for example. And yet the analysis I am suggesting may seem more strained here than it does for the example involving *today* that I just mentioned. The reason is that the linguistic meaning of *today* includes the descriptor ‘day’, which applies as well to day properties or types of days as it does to particular days. That is, we could paraphrase the meaning of (2) as (36), provided we give wide scope to the definite description:

¹⁵ Note that the same inferences that determine the interpretation of *here* in examples like these are what determine the interpretation of a given use of a description like ‘the economy’ or ‘the weather’.

(36) The day of this utterance is always the biggest party day of the year.

Whereas however we describe the linguistic meaning of *I*—as ‘the speaker of the utterance’, ‘the agent of the context’, or whatever—it is not easy to think of this descriptor as applying naturally to the entity that is the subject of the traditional dictum about last meals. That is, the linguistic meaning of *I* picks out its contextual index *qua* its role as speaker, but that entity is being asked to figure in the content of the propositional content *qua* its status as a condemned prisoner. So (4) does not have a natural paraphrase as (37), even when we try to give wide scope to the description ‘the speaker of this utterance’:

- (4) I am traditionally allowed to order whatever I want for my last meal.
 (37) The speaker of this utterance is traditionally allowed to order whatever he wants for his last meal.

But this takes me back to the second question I began this section with: why can’t the referential uses of definite descriptions have these descriptive interpretations? If the description ‘the speaker of this utterance’ can take us to the speaker, why can’t we then throw out the content of that description and take that person himself as the interpretation? And if the properties relevant to individuating persons in the contextual domain do not permit us to differentiate the speaker from other persons who are also condemned prisoners, then why can’t the subject of (37) have a descriptive interpretation?

At this point I want to revisit a familiar assumption about the way the content of a description functions in these referential uses. As Wettstein (1981: 36) puts it:

Consider the referential use [of definite descriptions]; there are contexts in which a speaker wants to draw his audience’s attention to an entity, perhaps one visually present to both speaker and audience, in order to go on and, for example, predicate something of it. It is irrelevant to the purposes of the speaker, in many such cases, how the attention of the audience is directed to the referent. Pointing with one’s finger or uttering a proper name would do as well as some elaborate description.

In other words—and this view is generally shared by both Russellians and referentialists—the choice of descriptive content in a referentially used description has no conversational relevance over and above its utility in picking out the intended referent. And that is what permits the referentialist to further assume that we can discard the content when its identificational work is done.

But this is not quite true. The most obvious counter-examples are sentences like (38), say as the opening sentence of a newspaper story:

- (38) The contractor who contributed \$150,000 to the governor’s re-election campaign was awarded a \$10 million state contract yesterday.

The description in the subject NP of (38) is clearly used referentially—the speaker has a particular person in mind, and the story will most likely go on to identify him by name. But the content of the description is not functioning simply to identify the

individual, but to attribute some property to him that is relevant to the predication—the sentence implies that the contract was awarded because the contractor made a large contribution (or at least, it suggests that some people might infer that some funny business was afoot). And clearly the choice of another description that identified the same person would change what was said. Suppose, for example, that the contractor who contributed \$150,000 to the governor’s campaign was also his brother-in-law. Then on the referentialist view, there would be no semantic difference between saying (38) and (39):

- (39) The governor’s brother-in-law was awarded a \$10 million state contract yesterday.

But obviously (39) has a very different conversational significance. I am not saying that sentences like these *entail* a causal relation between the content of the description and what is predicated of its referent, but they strongly invite that inference. And they do so on the assumption that the content of a referentially used description can have a conversational relevance over and above its usefulness in identifying the referent that the speaker has in mind.

In fact, these are only the extreme cases of a more general principle: following the maxim of relation, we tend to construe the content of *any* description as being relevant to the conversational purposes at hand, even when it seems to be functioning in a purely identificational way. Consider, for example, the difference between saying ‘please pass the potatoes’ and ‘please pass the red plate’. It may very well be that either description would be a perfectly reasonable way of identifying the thing that the speaker wants the hearer to pass him, but the second utterance is clearly an odd way to go about things, and it would probably trigger some sort of conversational implicature—maybe the speaker is interested in the plate *qua* plate (is he a china specialist?). And by the same token, we invite different inferences when we say ‘The man with the martini is Jones’, ‘The man in the ill-fitting suit is Jones’, and ‘The man who just yelled “blimey!” is Jones’, even if each of those descriptions is an equivalently effective way of drawing the hearer’s attention to the person that the speaker has in mind. Since the choice of descriptive content is always under-determined, that is, we assume that the speaker generally has some additional motive for choosing one description rather than another as a way of identifying someone—though it is true that sometimes not much attaches to the difference.

These observations sit very easily with the Russellian analysis of referential descriptions, which assumes that the content of the description is always in fact a part of what is said. By themselves, though, they may not rule out a referentialist account, since a referentialist could always argue that the inferences associated with utterances like these arise out of conversational implicatures via the maxim of manner. That is, the writer of (37) may not actually have *asserted* that the man who was awarded the state contract was a contributor to the governor’s campaign, but the fact that she chose that means of identification rather than some other might nonetheless be significant.

But at this point let us return to the observation that these referential uses of descriptions cannot have descriptive readings. Why should this be? After all, if the relevance of descriptive content is merely a conversational inference, rather than implicit in the semantics of the utterance itself, then we ought to be able to ignore that content when the conversational purposes require us to. Take an example like ‘Sophie’s eleventh birthday is always the biggest party day of the year.’ Suppose that the description ‘Sophie’s eleventh birthday’ takes us to November 1, 2001, and that the property of falling in one or another particular year is not conversationally relevant—in the domain, that is, there are only as many distinct days as are individuated by the date of the year they fall on. Then once we have identified the referent of the description, we should be able to toss its content aside and take the day-type as its interpretation, so that the utterance comes out as meaning something like ‘November 1 is always the biggest party day of the year.’

But that is not possible here. We are required to interpret the utterance as if the property of being Sophie’s eleventh birthday is relevant to what is being said, even if that leads us to an incoherent interpretation of the whole. As much as we’d like to, that is, we simply *cannot* throw out the content of the description—it turns up like a bad penny. And that is a fact that is simply incompatible with the referentialist account of these sentences.

This same observation explains why an utterance like (37) sounds odd to us: in choosing to refer to himself as ‘the speaker of this utterance’ rather than as *I*, the speaker conversationally implicates that he has some relevant reason for choosing that particular descriptive content—we would infer that he means us to believe that the tradition about last meals applies to people in virtue of their roles as speakers, even if that does not make a lot of sense. But there is no analogous implicature when a speaker refers to himself as *I*. Inasmuch as the linguistic meaning of the expression does not figure as part of the utterance content, we are not obliged to construe it as conversationally relevant. The fact that we have picked out an element of the context in virtue of its role as the speaker does not mean that that property is pertinent to its individuation.¹⁶

Appendix

I should say something about the difference between definite descriptions and demonstrative descriptions (or ‘complex demonstratives’) like ‘that guy’. It was suggested by

¹⁶ In fact there can almost never be a manner implicature associated with the use of an indexical, since as we saw that use is obligatory whenever it is felicitous. Since you are ordinarily obliged to refer to today as ‘today’, there is no significance attached to that choice. (There are some exceptions to this generalization: the choice of a demonstrative description can be significant when a definite description would do as well, as in ‘You’re not going to tell that story about your accident again, are you?’ But these are invariably tinged with an affective import, and do not have any bearing on the cases we are talking about here.)

one of the reviewers of this chapter that such expressions should be regarded as equivalent to definite descriptions. But there are good reasons for distinguishing between the two. For one thing, demonstrative descriptions cannot be re-evaluated relative to another hypothetical context:

- (i) If you had turned to Channel Four as I asked, the show would be more interesting;
- (ii) If you had turned to Channel Four as I asked, this show would be more interesting.

In (i), the description *the show* could refer to the show on Channel Four, but *this show* in (ii) can only be evaluated relative to the actual context—that is, it refers to the show that is currently visible, not the one on Channel Four. In this regard, *this show* behaves like a pure demonstrative, not like a description.

It should not be surprising, then, that demonstrative descriptions permit descriptive readings that are not easily available for definite descriptions. For example, you might ask either (iii) or (iv) to someone watching the Stanford–Cal game on TV:

- (iii) Are you enjoying the game?
- (iv) Are you enjoying this game?

And if the respondent is in fact enjoying the game that is on TV, he might answer with either (v) or (vi):

- (v) Yes, I'm enjoying the game a lot.
- (vi) Yes, I'm enjoying this game a lot.

But if the respondent wants to invoke a general disposition to enjoy Stanford–Cal games, he cannot felicitously answer with (vii), but must use (viii):

- (vii) ?Yes, I always enjoy the game.
- (viii) Yes, I always enjoy this game.

That is, only the version with the demonstrative description *this game* permits a reading where the description contributes a property of the game that is currently on view.