Reference and Context

"I proceed. 'Edwin and Morcar, the earls of Mercia and Northumbria, declared for him; and even Stigand, the patriotic archbishop of Canterbury, found it advisable—'"

"Found *what*?" said the Duck.

"Found *it*," the Mouse replied rather crossly: of course you know what 'it' means."

"I know what 'it' means well enough, when *I* find a thing," said the Duck: "it's generally a frog or a worm. The question is, what did the archbishop find?"

— Lewis Carroll, *Alice's Adventures in Wonderland*, chap. 3

593. A main cause of philosophical disease — a one-sided diet: one nourishes one's thinking with only one kind of example.

— Ludwig Wittgenstein, *Philosophical Investigations*

1

Singular reference, in the narrow sense, is a certain kind of connection between a token of an expression of a language — a "singular term" — produced by a speaker on a given occasion, and a particular thing — the "referent" of the term. A speaker "makes a reference" or "refers" when he establishes such a connection by the act of producing the singular term (uttering it, inscribing it, etc.). A singular term itself "refers" if the speaker has made a reference by producing it.

In the broad sense, singular reference is any such connection

AUTHOR'S NOTE: Earlier versions of this approach to the theory of reference were presented in papers read to philosophy colloquia at the University of Minnesota and Rockefeller University in 1968 and at Duke University, the University of Toronto, and Macalester College in 1970, and to a conference on the philosophy of language held at the Minnesota Center for the Philosophy of Science in August 1968. I owe a number of useful suggestions to participants in those events, especially Saul Kripke

holding between a particular thing and something which is either a singular term in a language or else is an element in some system which is language-like in the relevant ways. Such "quasi-linguistic" entities include: symbols on maps, figures in paintings, statues, scale models, memory images, the components of one's visual field, thoughts, and so on. Each of these is an element in what I will call a *context*, and it is by virtue of referential connections between their elements and particular things in the world that contexts are anchored down to the things and situations which they are about. I use 'context' as a technical term to cover a class of things for which there is no convenient label,[1] some examples of which I will now discuss briefly, postponing further clarification of the notion until later.

A linguistic discourse is a context whose elements are words, phrases, sentences, paragraphs, etc., produced by one or more speakers. A map is a context whose elements are cartographic symbols signifying buildings, bridges, forests, rivers, and so on. A picture is a context consisting of such elements as blotches of paint on canvas, ink marks on paper, silver deposits on photographic printing paper, patterns of light projected on a screen, etc. My visual field is a context consisting of elements commonly called "visual sensations." In general anything which has *content* is a context, as I use the term. Anything that has *meaning* or *sense* is a context. Anything which *expresses* something or *represents* something is a context.

Some elements of contexts are or can be referentially connected with particular things outside the contexts: the word 'Fred' which I utter in a conversation refers to a friend of mine, this blue line on the map

and Hilary Putnam. Since I first started thinking about these problems I have benefited greatly from discussions with Keith Gunderson, and also from discussions at various times with Paul Benacerraf, Keith Donnellan, David Kaplan, and David Lewis. I also want to thank John Olney of System Development Corporation, Santa Monica, California, under whose auspices — in 1966, as a consultant on a project on discourse analysis and anaphora — I began to consider the relation between discourse structure and singular reference.

[1] Nelson Goodman's notion of a "symbol" comes close, but includes some things I exclude (e.g., most "serious" music) and excludes some things I include (e.g., photographs). (See his *Languages of Art* (Indianapolis and New York: Bobbs-Merrill, 1968), *passim*.) For further explanation, see section 11. The theory I am going to present concerns the reference of concrete linguistic and quasi-linguistic tokens, not the corresponding types; singular terms, linguistic contexts (= discourses), other singular elements, and the contexts of which they are part are all concrete individuals, actual or possible. For brevity, I will often not make this explicit in what I say.

refers to the Mississippi River, that variegated blotch of paint refers to the Duchess of Alba, this sensation of a white disk on a dark background refers to the moon, and so on. These are, of course, uncommon uses of 'refer' (except for the first one); I will apply this expression not merely to singular terms in a language but to anything which is like a singular term in the relevant ways. I will call such entities "singular elements" of contexts.

When a singular element is not referentially connected with anything outside the context there is a *failure of reference*: singular terms which name nothing, maps with imaginary rivers and mountain ranges, paintings of fictitious scenes, hallucinations, dreams, and so on. It is the business of a theory of reference to explain both failure of reference and successful reference, to give a systematic general account of how singular elements in contexts hook up with items in the world. Such a theory must deal first with linguistic reference, since we understand the linguistic cases best; many important facts about them are open to public observation, and we already have a good deal of knowledge about the syntax, and some knowledge about the semantics, of natural languages. Previous theories have generally tried to explain the connection between a singular term and its referent as a function of the meaning of the term and the properties of the referent, paying little or no attention to the circumstances in which the term is uttered. The most influential theory of this sort has, of course, been Russell's Theory of Descriptions. Strawson's revision of Russell[2] paid conspicuous lip service to the importance of "context" but didn't include a systematic theory of what a "context" is or how contexts affect reference.[3] The first significant advance beyond Russell was made by Keith Donnellan, whose paper "Reference and Definite Descriptions"[4] contained several striking examples of singular reference which could not in principle be explained in Russellian or Strawsonian terms. The theory I will present here is the result of an attempt to do justice to Donnellan's cases while retaining what still seems to be true in the Russell-Strawson tradition. I think that Don-

[2] First presented in his "On Referring" (*Mind*, 1950; reprinted in Antony Flew, ed., *Essays in Conceptual Analysis* (London: Macmillan, 1956), pp. 21–52; my references will be to the latter.)

[3] Except for such very general remarks as the following: ". . . by 'context' I mean, at least, the time, the place, the situation, the identity of the speaker, the subjects which form the immediate focus of interest, and the personal histories of both the speaker and those he is addressing." ("On Referring," p. 42.)

[4] *Philosophical Review*, 75 (1966):281–304.

nellan's work shows that a "pure" or "autonomous" theory of linguistic reference is impossible, and that the central concepts ('singular term', 'refer') will have to be generalized far beyond their original application to spoken and written language. I believe that a theory of linguistic reference will have to be combined with a systematic account of certain internal states of the speaker — his thoughts, beliefs, perceptions, memories, and so on — which are, so to speak, the intermediate links connecting the singular terms he utters with their referents out in the world. These intermediaries can themselves be understood only if we treat them as being quasi-linguistic in structure and content — as contexts, in my sense of the term — and as containing elements analogous to singular terms which can be referentially connected with things in the world outside the speaker's skin.

2

Since linguistic reference is a kind of connection between singular terms and things in the world, the first thing a theory of reference has to do is define the class of singular terms. Quine offers this definition: "A term is singular if it purports to name an object (one and only one), and otherwise general." General terms don't purport to name anything at all: "The general term may indeed 'be true of' each of *many* things, viz., each red thing, or each man [he is discussing 'red' and 'man'], but this kind of reference is not called naming: 'naming', at least as I shall use the word, is limited to the case where the named object purports to be unique."[5] What is "purporting"? It is not a matter of whether the term is in fact true of just one object:

> For 'Pegasus' counts as a singular term though true of nothing, and 'natural satellite of the earth' counts as a general term though true of just one object. As one vaguely says, 'Pegasus' is singular in that it *purports* to refer to just one object, and 'natural satellite of the earth' is general in that its singularity of reference is not something *purported* in the term. Such talk of purport is only a picturesque way of alluding to distinctive grammatical roles that singular and general terms play in sentences. It is by grammatical role that general and singular terms are properly to be distinguished.
>
> The basic combination in which general and singular terms find their contrasting roles is that of *predication* . . . Predication joins a general term and a singular term to form a sentence that is true or false accord-

[5] W. V. Quine, *Methods of Logic* (New York: Holt, 1959), p. 205.

ing as the general term is true or false of the object, if any, to which the singular term refers.[6]

To paraphrase this in my terminology: singular terms are the ones whose role is to be referentially connected with objects, and general terms are the ones whose role is to describe or characterize the objects with which the singular terms are connected. Singular terms can purport to refer without actually referring, for purporting is a purely intra-linguistic affair — this is what Quine is getting at when he says that singular terms are identified in terms of their "grammatical role."

I think that Quine is essentially right[7] — his remarks are, indeed, no more than a sophisticated gloss on sound linguistic common sense — and I will shortly take up the question of exactly what the grammatical role of singular terms is and how it enables them to purport singularity of reference. But first I want to consider an objection put by Strawson, who thinks Quine's approach is inadequate and advocates a different one, based on the notion of "identification." He points out that we can distinguish singular terms from general terms by the consequences of their failing to apply to anything. ('Apply to' means 'either refer to or be true of'.) If I say 'The captain is angry' and 'the captain' applies to somebody, then the failure of 'angry' to apply to him results in the sentence being false. If 'the captain' does not apply to anybody the result is not a false sentence but one that has no truth-value at all — there is a so-called "truth-value gap." In general: "Whether the sentence is true or false depends on the success or failure of the general term; but the failure of the singular term appears to deprive the general term of the chance of either success or failure."[8] So Strawson's version of Quine's distinction is this:

Singular terms are what yield truth-value gaps when they fail in their role. General terms are what yield truth or falsity, when singular terms succeed in their role, by themselves applying, or failing to apply, to what the singular terms apply to. This is more or less what we have. It scarcely seems enough. We want to ask 'Why?'[9]

[6] W. V. Quine, *Word and Object* (Cambridge, Mass., and New York: Technology Press of M.I.T. and Wiley, 1960), pp. 95–96.

[7] Though his notion of a singular term is broader than mine, since it includes the variables of quantification.

[8] P. F. Strawson, "Singular Terms and Predication," in P. F. Strawson, ed., *Philosophical Logic* (London: Oxford University Press, 1967), p. 72. See also Quine's reply to this in *Synthese*, 19 (1968):292–297.

[9] Strawson, "Singular Terms and Predication," p. 73.

He tries to answer this question by starting with what he regards as the central cases of singular reference, namely those predications "in which singular and general term alike may fairly be said to be applied to a single concrete and spatio-temporally continuous object." In these cases, at least, the "characteristic difference" between the two kinds of terms

is that the singular term is used for the purpose of *identifying* the object, of bringing it about that the hearer (or, generally, the audience) knows *which* or *what* object is in question; while the general term is not. It is enough if the general term in fact applies to the object; it does not also have to identify it.

But what exactly is this task of identifying an object for a hearer? Well, let us consider that in any communication situation a hearer (an audience) is antecedently equipped with a certain amount of knowledge, with certain presumptions, with a certain range of possible current perception. There are within the scope of his knowledge or present perception objects which he is able *in one way or another* to distinguish for himself. The identificatory task of one of the terms . . . is to bring it about that the hearer knows *which* object it is, of all the objects within the hearer's scope of knowledge or presumption, that the *other* term is being applied to. This identificatory task is characteristically the task of the definite singular term.[10]

In general, the identificatory task of the singular term

is successfully performed if and only if the singular term used establishes for the hearer an identity, and the right identity, between the thought of *what-is-being-spoken-of-by-the-speaker* and the thought of some object *already within the reach of the hearer's own knowledge, experience, or perception*, some object, that is, which the hearer could, in one way or another, pick out or identify for himself, from his own resources. To succeed in its task, the singular term, together with the circumstances of its utterance, must draw on the appropriate stretch of those resources.[11]

If the identificatory task is botched and nothing correctly identified, the singular term has failed in its mission and there is a truth-value gap, since nothing has been correctly or incorrectly described or characterized by the speaker for the hearer. Thus we can understand what singular terms really are: they are the terms which are used to identify things for hearers. And we can understand what singular reference is: it is the way singular terms apply to objects.

But there are complications. Not every failure of a singular term to

[10] *Ibid.*, pp. 74–75.

[11] *Ibid.*, p. 78; emphasis in original.

perform the identificatory task results in a truth-value gap. Only the "radical" failures do so, and these occur when there is nothing there to be identified, when "there just is no such particular item at all as the speaker takes himself to be referring to."[12] And this is obviously correct. For example, if I want to tell you that my cat died, and I say 'The cat died', the singular term 'the cat' might fail to identify my cat because you didn't hear what I said, or because you didn't understand it, or because you thought I was talking about the neighbors' cat, and so on. Still, I was talking about her and 'the cat' *did* refer to her — the failure of the singular term was not a *radical* failure, and my assertion has a truth-value. If this is so — if reference can occur without identification — why does Strawson insist that identification is *the* characteristic task of singular terms, that identification is what they are really *for*, and that understanding this enables us to distinguish singular terms from general terms? Why not just say that singular terms are the ones whose characteristic task is singular reference and be done with it? Strawson rejects this because he thinks that 'refer to' is just as ill-understood as 'singular term' and just as much in need of explanation.[13] The same goes for such cognate notions as 'specify'; he says, for example, that a remark by Quine that a singular term "is used purely to specify its object for the rest of the sentence to say something about" is "unsatisfactory, since 'specify' by itself remains vague. To remove the vagueness we need the concept of 'identifying for an audience' which I have just introduced."[14]

I don't believe that the introduction of this notion removes the vagueness. According to Strawson, I am in a position to identify an object for someone only if that object is in fact being "spoken of" by me; that is, only if it is the one "in question"; that is, only if the general term I use is "being applied to" that object rather than some other. My hearer must be able to "distinguish for himself" or "pick out" the object on the basis of some "thought of" it.[15] What is it to "speak of" an object? For it to be the one "in question"? To "think of" it? "Pick it out"?

[12] P. F. Strawson, "Identifying Reference and Truth-Values," *Theoria*, 1964, p. 103. (This paper is an amplification of the account of identification given in "Singular Terms and Predication.")

[13] See "Singular Terms and Predication," p. 74.

[14] *Ibid.*, p. 76.

[15] All the quoted words and phrases come from the passages from Strawson quoted above.

These are at least as vague as the notion of "specifying" which Strawson criticizes Quine for relying on, if not more so. They are themselves instances of what I earlier called singular reference in the broad sense. The connection that holds between a person and an object when he is thinking of the object, or when he distinguishes it or picks it out, or when he intends to speak of it or apply a general term to it, or when it is the one in question, is similar in the relevant respects to the connection that holds when the person has specified or referred to the object by using language. In each case we must take one thing — a spoken remark or an unspoken state of mind — and relate it uniquely to something else: the referent, the thing specified, the object of thought, the thing distinguished or picked out, the thing involved in the speaker's intentions, the object in question. I hope to show that a theory which accounts for such referential connections will in the last analysis have to use the non-linguistic ones to explain the linguistic ones (to this extent I agree with Strawson), but as things presently stand the non-linguistic cases are considerably more obscure and of little use in clarifying linguistic reference. They are part of the problem, not part of the solution.

3

I said earlier that whether and what a singular term "purports" is a purely intra-linguistic affair, and that this was what Quine was getting at when he said that singular terms are distinguished by their grammatical role. In English, for example, which expressions are the singular terms and what is grammatically distinctive about them? Suppose we say tentatively that at least the following kinds of expressions are singular terms:

(a) proper names ('George', 'Leon Trotsky', 'Brazil', etc.);
(b) so-called "definite descriptions" ('the man', 'the man wearing a plumed hat', 'the cat', 'my cat', etc.);
(c) personal pronouns ('he', 'it', 'they', etc.);
(d) so-called "demonstratives" ('this', 'this man', 'that', etc.).

This is more or less the standard roster of expressions used in making singular references. We are already perfectly familiar with these expressions, not just in the sense that we've seen them all before, but, more important, in the sense that we would be able to recognize wholly novel

expressions falling under the same categories. New proper names appear continually, and previously unuttered definite descriptions get uttered all the time; it would even be easy, if we wished, to introduce new pronouns (say, 'ger' for old people) or new demonstratives (say, 'thot' for things heard but not seen). Our ability to produce and comprehend new singular terms could no doubt be explicitly represented by a grammar of English, which would enumerate all the definite descriptions, tell us what sentence-frames accept proper names, and so on, but present purposes will be served well enough by exercising the inexplicit ability which such a grammar would represent.

Consider one of the expressions on the list, say 'the man'; since it is a singular term, it purports to refer to exactly one thing; thus it purports to refer to exactly one thing in the following context:[16]

D1: #A man was sitting underneath a tree eating peanuts. A squirrel came by, and the man fed it some peanuts.#

What point are we trying to make when we say that 'the man' in D1 "purports to refer" to exactly one thing? We are, I think, trying to say something about how that remark is to be taken — specifically, we are trying to say that it is appropriate to look for something (a man, presumably) which 'the man' refers to (= which the speaker refers to by uttering 'the man') and which is truly or falsely described by the rest of the sentence. It has often been pointed out that some expressions are grammatically "singular" (at least on a superficial analysis) but couldn't conceivably refer to anything about which something is being truly or falsely said. For example, 'Nothing is under the bed' — 'nothing' doesn't count as a singular term, despite the syntactic similarity of 'Nothing is under the bed' and 'The cat is under the bed.' Singular terms are those which at least *could* be referentially connected with something, and 'nothing' couldn't.[17]

Is it mere membership on the list given above that makes it appropriate to look for something which an expression refers to? No, for 'the man' is on the list and yet it could have no referent in this context, in which D1 is embedded:

D2: #Say, let me tell you a funny story I just made up. A man was

[16] '#' marks the boundary of a discourse.

[17] Unless it is used as an ad hoc substitute for a singular term, say as part of a code. That sort of thing is always possible, and I won't bother to point it out henceforth.

sitting underneath a tree eating peanuts. A squirrel came by, and the man fed it some peanuts. Then the squirrel said . . .#

The possibility of a referential connection between 'the man' and anything outside the context is explicitly canceled by the speaker's announcement that what he is about to say is only a story he made up. That suspends the normal purport which 'the man' has in isolation, which it has in the sentence 'The man fed it some peanuts', and which it retains in some supra-sentential contexts, such as D1. It is a matter of context whether 'the man' purports to refer to anything, and consequently it is a matter of context whether 'the man' is a singular term. The same goes for all the other things on the list, which I will call "singular expressions." An expression is a singular term *in* or *relative to* a given context. Now the question is: When is a singular expression a singular term? In other words: In which contexts does a singular expression purport to refer to exactly one thing? Or: When is it possible for there to be a referential connection between a singular expression in a context and some object outside the context?

The smallest context which a singular expression can have is just the expression itself, in isolation. A singular expression standing alone purports to refer to exactly one thing, by virtue of its "grammatical role," that is, because of the way it enters into predications — so far this just recapitulates Quine's view. Proper names, definite descriptions, pronouns, and demonstratives are therefore all singular terms when uttered in isolation. When such an expression is embedded in a wider context, it may still purport to refer or it may cease to do so. In the latter case, I will say that its purported reference has been *canceled* — in other words, it cannot be referentially connected with anything beyond the boundaries of the context. Purported reference can be canceled by a variety of devices, one of which we have already seen (I will discuss the others later, in section 6). In D2 the announcement at the beginning of the discourse that what followed was fiction canceled the purported reference of the singular expression 'the man', even though it was not canceled by anything in D1. 'The man' is a singular term in D1 but not in D2, and perforce not in any wider context in which D2 might be embedded. (Of course, a fictional context doesn't have to be explicitly marked as such, as long as it is understood in the same way as overtly fictional contexts like D2; and the device which marks a context as fictional doesn't have to be part of the text, properly speaking,

but may instead be attached to it in some conventionally understood way, which is what happens when the disclaimer is built into the title (e.g., *'Bertrand and Ludwig: A Novel'*), or when the text is preceded by some such warning as: "All the characters in this book are fictitious and any resemblance to actual persons, living or dead, is purely coincidental.")

However, it is not the mere presence of a singular expression within the scope of a remark like 'Say, let me tell you a funny story I just made up' that cancels its purported reference. For example:

> D3: #Say, let me tell you a funny story I just made up. Richard Nixon was playing ping-pong with Mao Tse-tung and . . .#

(As with D2, the reader may complete the story however he wishes.) In this case, a piece of fiction is wrapped around two singular references to actual people, and 'Richard Nixon' and 'Mao Tse-tung' still purport to refer as usual. It is only when a singular expression is *introduced* within the scope of a fiction-indicating device that its purported singular reference is canceled. 'The man' in D2 is understood on the basis of the earlier expression 'a man', and the latter is, in a familiar sense, its *antecedent*. If we ask *who* fed the squirrel, *which* man did it, one perfectly correct answer would be that it was the man previously mentioned; it is on the basis of that previous predication involving 'a man' that we identify the purported referent of 'the man' — he is not just any man, but, specifically, one sitting underneath a tree eating peanuts. But *that* man is merely a figment of our tale: he, the tree, the squirrel, and the peanuts were all *first* mentioned in such a way as to cancel the purported reference of all subsequent mentions of them. We can make what Strawson calls "story-relative identifications"[18] within the context, but we cannot go outside the context and identify *that* man, *that* tree, *that* squirrel, *those* peanuts as the ones the story was about.

Before going any further I want to introduce some new terminology. Grammarians sometimes talk about "anaphora," by which they mean the kind of relationship that holds between, for example, a pronoun and its antecedent.[19] Let an *anaphoric chain* be a sequence of singular

[18] See his *Individuals: An Essay in Descriptive Metaphysics* (London: Methuen, 1959), p. 18.

[19] The notion seems to be rather loosely employed, but the basic idea is that expressions in different sentences (or, sometimes, in different parts of the same sentence) are related anaphorically if one of them somehow helps us understand or interpret the other.

expressions occurring in a context, such that if one of them refers to something then all of the others also refer to it.[20] 'A man — the man' in D1 and D2 is an anaphoric chain. So are 'A squirrel — it' in D1 and D2, and 'A squirrel — it — the squirrel' in D2. To consider a real-life example:

> D4: #At eleven o'clock that morning, an ARVN officer stood a young prisoner, bound and blindfolded, up against a wall. He asked the prisoner_1 several questions, and, when the prisoner_2 failed to answer, beat him_1 repeatedly. An American observer who saw the beating_1 reported that the officer "really worked him_2 over." After the beating_2, the prisoner_3 was forced to remain standing against the wall for several hours.#[21]

This passage, which occurs in the middle of a book, contains the following anaphoric chains:

I. 'that morning';
II. 'an ARVN officer — he — the officer';
III. 'a young prisoner — the prisoner_1 — the prisoner_2 — him_1 — him_2 — the prisoner_3';
IV. 'a wall — the wall';
V. 'an American observer who saw the beating_1';
VI. 'the beating_1 — the beating_2'.

Unlike the others, (I) is part of a longer chain which begins outside the quoted text in an earlier part of the book. ('That morning' refers, in fact, to the morning of January 8, 1967.) (II), (III), and (IV) all begin with so-called indefinite descriptions and continue with definite descriptions and pronouns. (V) begins the same way and clearly could continue much as they do ('an American observer who saw the beating_1 — the American observer — he — the American', for example). (VI) begins with a definite description, 'the beating', but it is clear that the beating in question is the one first mentioned in the second sentence of the text; we can therefore think of the first link of (VI) as being an indefinite description — perhaps 'a beating of the young prisoner by

[20] Zeno Vendler has an interesting discussion of anaphoric relations between singular terms (though he doesn't call them "anaphoric") in his "Singular Terms," which is chap. 2 of his *Linguistics in Philosophy* (Ithaca, N.Y.: Cornell University Press, 1967). His "chains of identification" (see p. 63) resemble but are not the same as my "anaphoric chains."

[21] From Jonathan Schell's *The Village of Ben Suc* (New York: Random, Vintage Books, 1968), p. 54. I have added subscripts to keep track of expressions that occur more than once; they should not be considered as part of the expressions.

the ARVN officer' — which would appear after a deeper analysis of the second sentence.

The difference between an indefinite description and a definite description is, in a sense, merely stylistic. An expression like 'a young prisoner' can occur only at the beginning of an anaphoric chain, not at any later point in it. For example:

> D5: #At eleven o'clock that morning, an ARVN officer stood a young prisoner$_1$, bound and blindfolded, up against a wall. He asked a young prisoner$_2$ several questions, and, when a young prisoner$_3$ failed to answer, beat him repeatedly.#

(This is just the first two sentences of D4 with 'the prisoner' replaced by 'a young prisoner'.) If we want to understand this as normal English we will have to conclude that 'a young prisoner$_1$', 'a young prisoner$_2$', and 'a young prisoner$_3$' do not purport to refer to the same thing and do not together form an anaphoric chain. The reason for this is not that indefinite descriptions cannot refer at all and for that reason cannot be links in anaphoric chains, but rather that whenever an indefinite description enters an anaphoric chain it can enter only as the first link. The difference between indefinite descriptions and definite descriptions is that the former can only be used to initiate anaphoric chains and the latter only to continue them. If we find a definite description, say 'the prisoner', occurring in a text and we want to know what is being talked about, we search the preceding text for an expression which will serve as the *anaphoric antecedent* of 'the prisoner' — 'a prisoner', 'a young prisoner', or what have you. If we don't find an antecedent we say that one is "presupposed" or "understood" or "implicit" in the use of the definite description.[22]

Indefinite descriptions should, therefore, be added to the list of singular expressions, since in many contexts they purport to refer and thus count as singular terms. There is, however, a popular view, first stated by Russell,[23] which denies that indefinite descriptions are on a par with definite descriptions and holds that they cannot refer or purport to refer. David Kaplan argues for this view in the following way:

For example:

> A senator from New York is supporting Rockefeller. (5)

Now (5) certainly has subject-predicate grammatical form in English, but if you feel that its logical form is the same as

> Jacob Javits is supporting Rockefeller. (6)

you can quickly disabuse yourself by comparing:

> A senator from New York is supporting Rockefeller, and a senator from New York is not supporting Rockefeller. (7)

with

> Jacob Javits is supporting Rockefeller, and Jacob Javits is not supporting Rockefeller. (8)

Sentence (8) is a contradiction, but (7) is true. In fact, isn't it obvious that indefinite descriptions do not even purport to denote a unique object as names do?[24]

No, it isn't. All that (7) shows is that each occurrence of an indefinite description starts a new and different anaphoric chain, like the different occurrences of 'a young prisoner' in D5. It does not show that an indefinite description occurring in its customary place at the beginning of an anaphoric chain does not refer or purport to refer on a par with the definite descriptions occurring later on in the same chain. Since the two occurrences of 'a senator from New York' manifestly belong to different anaphoric chains, they do not purport to refer to the same thing, and it is for that reason that (7) is not a contradiction. (8) is a contradiction precisely because the two occurrences of the proper name 'Jacob Javits' do count as links in the same anaphoric chain, according to our rules for the use of proper names, and thus they purport to refer to the same thing.

Quine argues along similar lines:

> The difference between . . . indefinite singular terms [by 'term' he means roughly what I mean by 'expression'] and the ordinary or definite ones is accentuated when repetitions occur. In 'I saw the lion and you saw the lion', we are said to have seen the same lion; indeed 'it' or 'him' could just as well have been used in place of the second occurrence of 'the lion'. But in 'I saw a lion and you saw a lion' there is no such suggestion of identity. . . . There is no one thing named by the indefinite singular term 'a lion'; no one thing even temporarily for the space of the single sentence.[25]

[22] If an anaphoric chain begins with a definite description which is a singular term in that context, then the context is in a sense incomplete or elliptical. Most linguistic contexts that actually get uttered are of this sort.

[23] Perhaps the clearest exposition of it is in his *Introduction to Mathematical Philosophy* (London: Allen and Unwin, 1919), chap. 16. Part of what I say about indefinite descriptions is adapted from Strawson's *Introduction to Logical Theory* (London: Methuen, 1952), pp. 186–187.

[24] David Kaplan, "What is Russell's Theory of Descriptions?", in D. F. Pears, ed., *Bertrand Russell: A Collection of Critical Essays* (Garden City, N.Y.: Doubleday, Anchor Books, 1972), pp. 230–231.

[25] *Word and Object*, p. 113.

Again, this only shows that an indefinite description can't be repeated without starting a new anaphoric chain; it doesn't show that an indefinite description at the beginning of an anaphoric chain cannot refer or purport to refer. The tacit assumption underlying Quine's and Kaplan's arguments is that all singular terms in English behave like proper names or like individual constants in formalized languages, which can be repeated, at least within a single sentence, without changing their purported reference. This assumption is false. The difference between 'a lion' and 'the lion' is like the difference between 'The lion' and 'the lion' in written English; the latter never occurs at the beginning of a sentence and the former never occurs anywhere else, but that proves nothing about their status as singular terms.

Quine has another argument for the same conclusion:

> In 'I saw the lion', the singular term 'the lion' is presumed to refer to some one lion, distinguished from its fellows for speaker and hearer by previous sentences or attendant circumstances. In 'I saw a lion', the singular term 'a lion' carries no such presumption; . . . 'I saw a lion' counts as true if at least one lion, no matter which, was seen by me on the occasion in question.[26]

He apparently assumes that there are only two alternatives: either 'a lion' refers to some one lion, distinguished from its fellows for speaker and hearer by previous sentences or attendant circumstances, or it doesn't refer at all. But there is a third possibility: that 'a lion' refers to some one lion which is not distinguished from its fellows *for the hearer* by previous sentences or attendant circumstances, but which is distinguished from them *for the speaker*. In D4, for example, 'a young prisoner' does not refer to a young prisoner distinguished from his fellows for the hearer (or reader) by attendant circumstances or by previous sentences in that discourse; but it does refer to a young prisoner whom the author of the book was able to distinguish for himself and whom he introduces into the book at that point by using the expression 'a young prisoner'.

Russell presented this argument:

> . . . no one could suppose that "a man" was a definite object, which could be defined by itself. . . . when we have enumerated all the men in the world, there is nothing left of which we could say, "This is a man, and not only so, but it is *the* 'a man,' the quintessential entity that is just an indefinite man without being anybody in particular." It

[26] *Ibid.*, p. 112.

> is of course quite clear that whatever there is in the world is definite: if it is a man it is one definite man and not any other. Thus there cannot be such an entity as "a man" to be found in the world, as opposed to specific men. And accordingly it is natural that we do not define "a man" itself, but only the propositions in which it occurs.[27]

We would define "a man" itself only if it were a name (by which Russell meant something that "directly designates" an individual, i.e., refers to it). If it were a name it would have to name an indefinite man; since there are no indefinite men, it is not a name.

But what is it for a singular term to name or refer to (or purport to name or refer to) some definite thing? A "definite" thing is not, after all, a *kind* of thing, like a red thing or a spherical thing. 'Definite' and 'indefinite' have to do with the *way* an expression purports to refer, and that is a matter of context. What is it for an expression to purport to refer "definitely" in a given context? Consider the anaphoric chain (III) in D4. Its second link is 'the prisoner', which in that context is a singular term that purports to refer to some prisoner. Which prisoner? If we ask this question *the context will supply us with a definite answer* — for example, "The young prisoner just mentioned, the one that the ARVN officer stood up against a wall." In general, the sort of reference which is purported by a singular term occurring in an anaphoric chain can be made more definite (more specific, more exact) by an appeal to earlier links in the chain and what is predicated of them; but we cannot in that way make the purported reference of the *first* link more definite. To do that we have to go outside the context altogether. Since indefinite descriptions like 'a young prisoner' and 'a lion' and 'a man' can occur only as the initial links in anaphoric chains, we can see why their purported reference cannot be made more definite within the context and why they have the quality of "indefiniteness" which has led some people to deny that they purport to refer.

I am not claiming that indefinite descriptions are *always* singular terms, purporting to refer; like the other singular expressions, including definite descriptions, they qualify as singular terms in some contexts but not in others. When an indefinite description does not purport to refer, then the sentence containing it can be taken as equivalent to an existentially quantified sentence.[28] 'A senator from New York is sup-

[27] *Introduction to Mathematical Philosophy*, pp. 172–173. See also Strawson's criticisms of Russell's view about indefinite descriptions in "On Referring," p. 49.

[28] Sometimes but not always. See below, section 6.

porting Rockefeller' sometimes *is* to be read as equivalent to '$(\exists x)$(x is a senator from New York & x is supporting Rockefeller)', but it doesn't always have to be paraphrased this way. This is perhaps easier to see in the case of 'A senator from New York is supporting Rockefeller' than it is when we consider some of its stylistic variants, such as 'There is a senator from New York who is supporting Rockefeller', since the latter is closer syntactically to the standard English way of reading the corresponding existential quantification ('There is an x such that x is a senator from New York and x is supporting Rockefeller'). Concerning such sentences, Davidson argues:

> We recognize that there is no singular term referring to a mosquito in 'There is a mosquito in here' when we realize that the truth of this sentence is not impugned if there are two mosquitos in the room. . . . We learned some time ago, and it is a very important lesson, that phrases like 'a mosquito' are not singular terms, and hence do not refer as names or descriptions do.[29]

If what Davidson says holds for 'There is a mosquito in here' uttered in isolation, it should also hold for the same sentence when embedded in a larger discourse, say this one:

> D6: #There is a mosquito in here. You can hear it buzzing. See, it just landed on my left arm. Now it's biting me. [*The speaker swats the mosquito.*] Not much left of it now, is there!#

How do we understand the occurrences of 'it' in such a context? It is clear from the structure of the discourse that all the occurrences of 'it' refer to the same thing if they refer to anything at all.[30] It is also clear that sometimes they do refer, as in this situation: the speaker sees a mosquito and hears it buzzing (there is another mosquito present but he doesn't notice it), and he says: 'There is a mosquito in here. You can hear it buzzing.' Then he sees it land on his left arm and says: 'See, it just landed on my left arm.' Then he feels the bite and says: 'Now it's biting me,' then he swats it, looks at the remains, and says: 'Not much left of it now, is here!' It would be absurd to maintain that on the four occasions he uttered the word 'it' the speaker was not referring to that mosquito, talking about it, making remarks about it, commenting on its activities, describing what it was doing, and so on. There

[29] Donald Davidson, "The Individuation of Events," in Nicholas Rescher, ed., *Essays in Honor of Carl G. Hempel* (Dordrecht: Reidel, 1969), p. 220.

[30] Unless the setting is such as to force an ad hoc interpretation on what is said.

is (in my jargon) a referential connection between that mosquito and each of those tokens of 'it'. Does such a connection also hold between 'a mosquito' and that mosquito? Davidson denies that it does, on the grounds that the first sentence of D6 is true even though there is a plurality of mosquitoes present, from which it follows that 'a mosquito' is not a singular term. Why does the extra mosquito matter? There are three readings of the sentence to consider. First, we can take it as an existential quantification, to be paraphrased as '$(\exists x)$(x is a mosquito & x is in here)'. In this case the sentence is true and 'a mosquito' is not a singular term.[31] Second, we can read it as initiating an anaphoric chain ('a mosquito — it — it — it — it'), in which case it is a singular term referring to the mosquito seen, heard, and swatted by the speaker, and the sentence is true if and only if *that* mosquito was in the room. The presence of an additional mosquito not noticed or referred to by the speaker does not falsify the sentence, nor does it prevent him from referring to that one. If it had turned out that the mosquito heard and seen by the speaker was not inside the room but, say, just outside it, the presence of the unnoticed mosquito inside the room would not have saved the sentence from falsity.[32] Third, we can understand the sentence as saying that *exactly* one mosquito is in the room. Now this *is* falsified by the presence of the second mosquito, and it seems to be this case that Davidson had in mind when he argued that 'a mosquito' was not a singular term. But the claim that it is a singular term does not depend on taking the sentence that way.[33]

[31] And neither are the four occurrences of 'it', which are syntactically linked with 'a mosquito' and stand or fall with it. If 'a mosquito' is no more closely connected with the perceived mosquito than with the other one, the occurrences of 'it' are in the same position and do not refer uniquely to the one the speaker perceived. The only way to avoid this obviously false conclusion is to hold that the syntactic connection with 'a mosquito' is of no significance. But that won't do either. The occurrences of 'it' purported to refer to a *mosquito* only because of the anaphoric chain linking them with 'a mosquito'; if we deny this linkage, we can't account for their purported reference, and the first sentence might as well not be in the discourse.

[32] It is easy to overlook this because of the fact that if 'There is a mosquito in here' is true on the reading that takes 'a mosquito' as a singular term, it will also be true on the reading that takes the sentence as an existential quantification, provided that the referent of 'a mosquito' is a mosquito, as it will be in normal cases. (The abnormal cases are the ones pointed out by Donnellan; see below, section 9.)

[33] You are likely to take it that way only if you make both of the following assumptions: (i) if 'a mosquito' is a singular term, it must be understood the same way as 'the mosquito'; (ii) 'the mosquito' must be understood according to Russell's

4

Sentences containing indefinite descriptions are ambiguous. Sometimes 'A mosquito is in here' and its stylistic variant 'There is a mosquito in here' must be taken as asserting merely that the place is not wholly mosquito-less, but sometimes they involve an intended reference to one particular mosquito. Their disambiguation depends on how the speaker intends the contexts containing them to be related to other contexts.

For example, suppose that I am reading the morning newspaper and I come across the following story:

> D7: #Houston, Texas, March 10 (UPI) — Dr. Michael DeBakey stated at a press conference today that an artificial heart could be developed within five years. The famed Baylor University heart surgeon said that such a development would make heart transplants unnecessary.#

I then report this fact to you by saying:

> D8: #A doctor in Texas claims that artificial hearts will be developed within five years.#

Is 'a doctor' in that token of D8 a singular term? Is it possible to trace a referential connection between that expression and a particular person, such that what I said is true if and only if *that* person claimed that artificial hearts will be developed within five years? Or am I merely asserting that the class of Texas doctors claiming that artificial hearts will be developed within five years is non-empty, as the existential quantification reading of D8 would have it? In that case what I said would be true even if the news report about DeBakey were wholly erroneous and DeBakey had never made any such claim but some other doctor in Texas had, say in a private conversation, unknown to the reporter who wrote the story. Which reading is the correct one in this case? Imagine how the conversation might continue: you ask 'Who said that?' and I answer 'Dr. Michael DeBakey.' Or perhaps: you say 'I can't quite believe that' and I say 'Well, it was DeBakey who said it and he ought to know. He's a famous surgeon.' Or perhaps: you say 'What's his name?' and I say 'Michael DeBakey.' Such continuations would be unintelligible on the existential quantification reading, for they pre-

Theory of Descriptions (or something very much like it). (ii) is false, as I will argue below, in section 9.

suppose that one and only one person is being said to have claimed that artificial hearts will be developed within five years; they presuppose that there is a unique referent of 'a doctor' whose name can be requested by asking 'Who?' or 'What's his name?' and who can be identified by saying 'Michael DeBakey'.

Compare that with:

> D9: #A doctor normally makes at least $75,000 a year.#

Here the referential purport of 'a doctor' is canceled, and it would be a misunderstanding of the context to ask 'Which one?' or 'What's his name?' There is and could be no particular one. (But notice that 'There is a doctor who normally makes at least $75,000 a year' could be taken either as an assertion about DeBakey, say, who normally makes at least $75,000 a year, or as an assertion that there is at least one such doctor, based, perhaps, on a statistical analysis of medical incomes. D9 would not be naturally read in either of these ways, but rather as an assertion about the normal doctor's income.) The fact that 'a doctor' in D8 was intended to be taken as purporting a singular reference could have been signaled explicitly, for I could have said any of the following instead:

> D10: #A doctor in Texas — DeBakey, the heart surgeon — claims that artificial hearts will be developed within five years.#
>
> D11: #A doctor in Texas — namely, Michael DeBakey — claims that . . . #
>
> D12: #A doctor in Texas, Michael DeBakey in fact, claims that . . . #
>
> D13: #A certain doctor in Texas claims that . . . #

And so on. In each of these contexts it is made clear by the speaker that he intends 'a doctor' to be taken as a singular term. In D10, D11, and D12 the appearance of the proper name shows that the discourse is to be understood as being linked with another discourse not uttered by the speaker, which contains the proper name (in this case the newspaper story D7). D13 also implies such a linkage but does not present a name taken from the other context.

When a singular term in one context is connected with another the way 'Dr. Michael DeBakey' in D7 is connected with 'a doctor' in D8,[34] I will say that the two expressions are *referentially linked* in a referen-

[34] And with 'a doctor' and 'DeBakey' in D10, with 'a doctor' and 'Michael DeBakey' in D11 and D12, and with 'a certain doctor' in D13.

tial chain. If a singular term in a given context is not referentially linked to a singular term in another context, I will call it *referentially isolated.*[35] If the structure and content of the context in which a singular expression appears cancels its referential purport, thereby insuring that the expression will not be referentially linked with another expression in another context, the expression is not only referentially isolated but *referentially segregated* — its isolation is not merely an accidental result of the circumstances in which that token is produced, but is, so to speak, intrinsic to it. 'A doctor' in D9 is an example of this. The fact that the expression is segregated need not be indicated on the surface.

So far, then, we have anaphoric chains within contexts, such that if one expression in the chain refers to a given thing then so do all the others, and we also have referential chains between contexts, for which the same condition holds: if one expression in the chain refers to a given thing then so do all the others. Since the links in referential chains connecting different contexts are also links in anaphoric chains within those contexts,[36] we can see how very lengthy referential chains can be constructed: a singular term T_1 in a context C_1 is anaphorically linked with another singular term T_2 in C_1, which is in turn linked with T_3 in C_1, which is referentially linked with T_4 in another context C_2, which is anaphorically linked with T_5 in C_2, which is referentially linked with T_6 in C_3, and so on, until we come at last to some singular term T_n which refers to some object, which thus counts also as the referent of all the singular terms along the whole chain from T_1 to T_n.[37]

Thus the theory of singular reference is concerned with three broad topics:

(i) *Contexts.* What kinds of contexts are there, linguistic and non-linguistic? When does a singular expression in a linguistic context, or a singular element in a non-linguistic context, purport to refer? How does the content of a context affect the purported reference of a singular term or other singular element? When do singular elements form anaphoric chains? When is referential purport canceled, segregating the context?

(ii) *Connections Between Contexts.* Under what conditions are referential chains formed? When is a context referentially isolated, and when is it linked to another context?

(iii) *Connections Between Contexts and Referents.* When do the links in a referential chain actually refer to something? Which terms are connected with objects in the world "directly" rather than by way of other links in the chain? Under what conditions does this take place? What sorts of objects must the referents be if such connections are to be established?

Before sketching a theory of how such questions should be answered for contexts in general, I will deal first with the problem of linguistic reference, examining some popular theories and attempting to demonstrate their inadequacy. For the time being, then, 'context' will mean 'linguistic context' unless otherwise indicated.

[35] If D8 had been uttered "out of thin air," idly, with no such background as D7 provides, the occurrence of 'a doctor' which it contains would have been referentially isolated. This is not to say that it could have had no referent, but merely that it would not have been connected with its referent via any referential chains. See below, section 8.

[36] Since any singular expression is a link in at least one anaphoric chain, namely the one consisting of that expression itself.

[37] Thus the whole referential chain includes all the anaphoric chains as segments; every anaphoric chain is a referential chain, as I am using the phrase, but not vice versa.

5

So far my discussion has been mostly concerned with anaphoric chains beginning with indefinite descriptions, and referential chains in which a singular term T_1 in a context C_1 is referentially linked with an indefinite description T_2 in a context C_2. Now I want to consider other varieties of anaphoric chains and other modes of referential linkage.

First, proper names. An obvious characteristic of proper names is that generally all occurrences of them in a given discourse produced by a single speaker belong to the same anaphoric chain, unless there is something in the context which explicitly indicates the contrary. In a novel, say, or a biography, there will be a number of occurrences of names of characters — 'Humbert Humbert', 'Lolita', 'Winston Churchill' — and unless the reader is told otherwise he is entitled to assume that, e.g., in *Lolita* all the occurrences of 'Humbert Humbert' belong to one anaphoric chain (including other expressions as well, such as, in this case, a number of occurrences of 'I', since the novel is narrated in the first person) and all the occurrences of 'Lolita' belong to another (which would also include a number of occurrences of 'Dolores Haze', 'Dolly',

'Lo', some but not all occurrences of 'she', and so on). The ability to comprehend a novel, a biography, or even a single paragraph presupposes the ability to keep track of who's who by assigning a given singular expression to the right anaphoric chain, where the latter will normally be a mixture of proper names and other singular expressions. This means, for example, knowing when an occurrence of 'she' in *Lolita* belongs with the occurrences of 'Lolita' and when it belongs with the occurrences of the names of the other female characters in the book. The ability to identify anaphoric chains of this sort is obviously very complex: we employ our knowledge of the syntax and semantics of the language, plus our knowledge of how discourses are constructed, plus our knowledge of whatever special literary or scholarly or other conventions pertain to the genre in question, plus our knowledge of what the writer is likely to have meant, and so on.[38] I think it quite unlikely that anaphoric chains in any moderately lengthy discourse are formed according to rules in the way that grammatical sentences are formed according to rules — that is, in such a way that, for a given discourse, the rules generate all and only the anaphoric chains in it — but whether or not this is so is irrelevant to the present discussion. It is more to the point to consider how proper names get introduced into a discourse, that is, how the first occurrence of a proper name comes about. Various devices are in current use in spoken and written English, including the following (this list doesn't pretend to be exhaustive):

D14: #A man named Fred Schultz runs a liquor store around the corner from here. He was held up yesterday, and . . .#

D15: #A man known to the underworld as Joe the Snake was found dead in a vacant lot this morning. Police spokesmen stated that . . .#

D16: #George Frisbee, a friend of mine from college, just inherited a million dollars. He always used to say that . . .#

D17: #A certain bookie I know whom I'll just call "Harry" told me that . . .#

D18: #Stella Houston (a fictitious name) underwent a sex-change operation six months ago. "She" told this reporter that . . .#

[38] The difficulties involved in the disambiguation of even a single sentence taken by itself are great enough; see the discussion of this point by Katz and Fodor in their "The Structure of a Semantic Theory," in J. A. Fodor and J. J. Katz, eds., *The Structure of Language: Readings in the Philosophy of Language* (Englewood Cliffs, N.J.: Prentice-Hall, 1964), especially pp. 486–491.

D19: #It turns out that a man I saw at the beach is a spy. He's been working for the Russians for years. Who is it? Bernard J. Ortcutt.#

D20: #Stately, plump Buck Mulligan came from the stairhead, bearing a bowl of lather on which a mirror and a razor lay crossed. . . .#

The simplest way to introduce a proper name into a discourse is to just start using it, as in D18 and D20. Another way is to make a statement of identity, in one guise or another, utilizing a singular term that has already appeared in the context (D19 is of this sort), or one that is itself used there for the first time (D16 and D17).[39] Sometimes it appears that a proper name is introduced *de novo* into a context, in such a way that it is not referentially linked to another occurrence of the same name in an antecedent context. This is indicated explicitly in D17 by the phrase 'whom I'll just call "Harry" '; the pseudonym is introduced simply in order to have a name *in that context* for the person whom the speaker is referring to by another singular term ('a certain bookie I know'), and previous occurrences of the name in other contexts are irrelevant. It doesn't matter whether anyone has ever referred to that person by the name 'Harry' before, and even if someone did it would have no bearing on the present case. (This is the natural way to take D18 also.) Sometimes the situation is just the opposite: the first occurrence of the name in the context harks back to previous occurrences in other contexts, as with D14, D15, and D19. In D19 the point of what the speaker is saying is lost unless the token of 'Bernard J. Ortcutt' which he utters is referentially linked with another token of 'Bernard J. Ortcutt' in some previous context known to him. Otherwise the speaker would not be *identifying* the man seen at the beach as Ortcutt, for identification requires that the terms in the identity statement be independently connected with the referent.[40] The same goes

[39] In effect, D16 asserts 'George Frisbee = a friend of mine from college' and D17 asserts 'Harry = a certain bookie I know'. D14 and D15 are more complicated; D14 asserts that the man who runs the liquor store is actually *named* 'Fred Schultz' and presumably known by this name to persons other than the speaker; and D15 asserts that the dead man was known to certain people by the nickname 'Joe the Snake'. It is not always clear how the falsity of such assertions would affect the reference of the name. See below, section 14.

[40] Some identity statements don't identify anything; they serve instead to introduce new singular terms, such as 'Harry' in D17. Notice also that in D19 it is plausible, contrary to Strawson's view, to say that the speaker has identified the

for D11 and D12: 'namely, Michael DeBakey' and 'Michael DeBakey, in fact' indicate linkages with antecedent contexts containing 'Michael DeBakey'. Sometimes it isn't clear whether such a referential linkage is present. In D20 the speaker might be initiating an entirely new referential chain; he might be continuing one whose last link was 'Buck Mulligan'; and he might be continuing one whose last link was some other singular term (possibly even another proper name).

Proper names pose no special problems when they hark back to a previous occurrence of the same name in an antecedent linguistic context. Here the relationship is like that between proper names forming an anaphoric chain within a single context. Which other singular terms in antecedent contexts can be linked to proper names? As far as I can see, any category of singular term will do. I said in section 3 that indefinite descriptions appear only as the initial links of anaphoric chains. Chains so initiated can continue in various ways, as D4 showed: 'an ARVN officer — he', 'a young prisoner — the prisoner', and so on. Like definite descriptions, demonstrative phrases can have indefinite descriptions as antecedents. For example:

> D21: #I know a guy who likes to eat grasshoppers. Really! Sometimes this guy will eat five or six at a time. He puts vinegar on them.#

'This guy' has whatever reference 'a guy' has, since it belongs to the anaphoric chain 'a guy — this guy — he'. 'This' — which is sometimes erroneously believed to be a "pure" demonstrative which indicates its referents "directly" or "ostensively" — can function the same way:

> D22: #I've just spotted a large boulder. It seems to contain mica. This is the most interesting rock I've seen so far.#

D22 might be a running report by an explorer communicating by radio with his base camp, for example; in such a situation the token of 'this' which he utters cannot pick out anything "ostensively" for his absent audience, but they would know nevertheless that it had whatever referent 'it' and 'a large boulder' had, by virtue of their common membership in the anaphoric chain 'a large boulder — it — this'.[41] Pronouns with

spy as Ortcutt even though the hearer will not be helped by this if 'Bernard J. Ortcutt' doesn't appear in any antecedent context familiar to him.

[41] Sometimes, of course, 'this' would not continue the anaphoric chain; for example, the speaker could suddenly change the topic: '. . . seems to contain mica. [*He suddenly shifts his attention to another rock.*] *This* is the most interesting rock . . .' But when this doesn't happen — that is, when 'a large boulder' is an anaphoric ante-

the appropriate syntactic characteristics may be added to any anaphoric chain: 'a man — he', 'Bella — she', 'this guy — he', 'this — it', 'the table — it', and so on. Each of the singular expressions considered so far can initiate an anaphoric chain in a given context, and each can be referentially linked with other contexts.

Obviously there is a great deal more to be said about how anaphoric chains are formed in English discourses, but I don't have a theory that would do justice to this vast subject, and even if I did it would not advance my purpose to present it here, since I am concerned with singular reference in general, not just singular reference in English. My examples are from English because that is the only natural language I am competent to theorize about, but I am concerned not with the details of English discourse structure but with the broad outlines which it shares with other languages: anaphoric chains linking singular terms in a context, however they may be constructed by the speaker or comprehended by the hearer.

6

The referential purport of a singular expression can be canceled in a variety of ways, corresponding to the variety of purposes for which discourses can be produced. Referentially segregated contexts may be involved in all of the following.

(i) *Fiction.* If a discourse is produced as fiction *with respect to a singular expression E* contained in it, then: (a) the speaker produces it with the intention that no referential linkage shall hold between *E* and any singular term in an antecedent context,[42] and (b) he intends that

cedent of 'this' in D22 — *and* the context is a running report in which 'a large boulder' is uttered earlier than 'this', so that the former has a reference before the latter is even uttered, it seems clear that 'this' gets whatever reference it has from the earlier phrase, not "ostensively." (Later on I will present an analysis of ostension, in terms of perceptual contexts; see section 13.)

[42] This condition insures that *E* will not be connected with an extra-linguistic referent by way of an expression in another context; you might object that it does not rule out a "direct" connection with a referent. Later on I will argue that in most cases of linguistic reference it must always be possible to establish a referential linkage with a singular element in an antecedent linguistic or *non-linguistic* context. See section 9.

You might also think that this formulation ignores cases in which the speaker intends a linkage with another *fictional* context, as when I now say 'Superman is faster than a speeding bullet'. But the referential purport of 'Superman' would be canceled in such antecedent contexts — since they are fiction — and it would not be a singular *term*. If I had been unaware that the Superman stories are fiction,

this should be evident to his audience. If a discourse is produced as *pure* fiction, it is produced as fiction with respect to each of the singular expressions in it.[43] These are necessary, not sufficient, conditions. They don't distinguish between fictional uses of contexts and the other non-factual uses (suppositional, modal, etc.) discussed below; all such contexts meet conditions (a) and (b), and further differences among them must be explained in terms of the purposes for which they are produced rather than the status of the singular expressions they contain. Any such distinctions — e.g., that fiction is used to entertain or edify, that discussions of hypothetical cases are used to illustrate general principles, and so on — are likely to be rather vague, as these are, and we can get along well enough without them. As I pointed out in section 3, the fact that a discourse is fiction can be signaled by putting 'a novel' in the title and by stating that all of its characters are fictitious. The same effect can be achieved by 'Once upon a time . . .', 'Let me tell you a story . . .', and 'Have you heard the one about the traveling salesman who . . .', and the like, depending on the prevailing conventions. Often the non-linguistic circumstances will make it evident to the hearer that the speaker intends what he says as fiction even though that is not indicated openly.

(ii) *Supposition.* The following contexts are referentially segregated:

D23: #Assume that a man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually. After three years, how much money will the man have in his account?#

D23′: #Assuming that a man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually, how much money will he have in his account after three years?#

D24: #If a man were to deposit $5,000 in a savings account at 5¼% interest, compounded semiannually, how much money would he have in his account after three years?#

D25: #A man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually. After three years, how much money does he have in his account?#

D25′: #If a man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually, how much money does he have in his account after three years?#

then (a) would not hold, but then also I would not have produced my remarks as *fiction*.

[43] Most fiction is not pure. Historical novels, jokes about actual people, political satires, etc., are very impure; even novels whose characters are all fictitious may contain references to actual places, say, and be somewhat impure. Even if the characters of a discourse are all actual people, as in D3, some of the *events* may be fictitious (e.g., 'the ping-pong game' might appear later in D3).

These could be used to state a problem in a textbook, or as part of an exposition of the laws of compound interest (in which case the questions would be "rhetorical"). They are available for such purposes because in each case it is explicitly indicated that the referential purport of 'a man — the man' or 'a man — he' is canceled, so that these expressions cannot be referentially linked to singular terms in other contexts. In each case it would be incorrect for a hearer to ask 'Who?', 'Which man?', 'Is that him over there?', etc., and it would be incorrect for the speaker to continue the discourse by identifying the man as this one or that one. The segregation of the context is accomplished by different syntactic forms, but the effect is the same in each case: a supposition is made that such-and-such a person exists and does so-and-so, and a question is asked about the consequences.

There are very similar contexts in which the question of referential linkage can only be settled by the speaker's intentions:

D26: #A man deposited $5,000 in a savings account at 5¼% interest, compounded semiannually. Now, after three years, how much money does he have in his account?#

D27: #A man is depositing $5,000 in a savings account at 5¼% interest, compounded semiannually. How much money will he have in his account after three years?#

These could serve as textbook examples or exam questions, but they also could under certain circumstances be used to talk about particular people. (Imagine D26 continuing: # . . . He asked me to work it out for him. He doesn't trust the people down at the bank.#)

In a similar fashion, referentially segregating the context enables one to make suppositions and then state rather than ask about their consequences:

D28: #Assuming that a man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually, after three years he will have $5,432.10 in his account.#

D29: #If a man were to deposit $5,000 in a savings account at 5¼% interest, compounded semiannually, after three years he would have $5,432.10 in his account.#

D30: #If a man deposits $5,000 in a savings account at 5¼% interest, compounded semiannually, after three years he has $5,432.10

in his account.# (D30 is discussed further in (iv) below.)

Proper names can figure in suppositional contexts:

D31: #If John Smith were to deposit $5,000 in a savings account at 5¼% interest, compounded semiannually, after three years he would have $5,432.10 in his account.#

Depending on the speaker's intentions, 'John Smith' may be referentially linked with another context (and perhaps ultimately with an actual person), or its referential purport may be canceled. In the former case it would still involve a supposition, of course; not a supposition to the effect that there exists such-and-such a person, but a supposition that an actual person does so-and-so.

(iii) *Modality*. Referentially segregated contexts can be used to make points about what might, or could, or ought to be the case. For example:

D32: #A visitor might trip over that cord. Then he could sue us.#

D33: #A rattlesnake can kill you.#

D34: #You ought to buy a new car. The bank will loan you the money for it.#

The segregation of 'a new car', 'a rattlesnake', and 'a visitor' in D32–D34 is explicit in their structure; normally understood, they cannot be referentially linked to other contexts via those expressions. But this is not due to the syntactic position occupied by 'a new car' and the rest, as we can see by making some substitutions:

D35: #Uncle Bert might trip over that cord. Then he could sue us.#

D36: #That rattlesnake over there can kill you.#

D37: #You ought to buy this new Mazda. The bank will loan you the money for it.#

'Uncle Bert', 'that rattlesnake over there', and 'this new Mazda' could each be referentially linked with a term outside the context and ultimately with an extra-linguistic referent. To use the indefinite description instead is to indicate an intention to block the linkage. Yet the possibility of such a hook-up can be reinstated if the right sort of indefinite description is used, as is obvious from the following:

D38: #A nearsighted man who comes in here occasionally might trip over that cord. Then he could sue us. He's very litigious.#

D39: #A rattlesnake I just saw can kill you. He wasn't defanged like the others here in the zoo.#

D40. #You ought to buy a new car I saw today. It's got a little scratch on the top and the dealer is willing to come down on the price.#

Thus it is not the mere presence of an indefinite description within the range of a term like 'might' or 'can' or 'ought' that cancels its purported singular reference. It is much more complicated than that, and I believe that the preceding examples show that any simple account is likely to fail.

(iv) *Universal Quantification*. D30 can be paraphrased as a universal quantification:

(1) (x) ((x is a man & x deposits $5,000 in a savings account at 5¼% interest compounded semiannually) ⊃ x has $5,432.10 in his account after three years)

Apparently D30 must always be paraphrased as a universally quantified sentence, in which the singular expressions disappear ('a man' in (1) should be considered part of the predicate). If D30 had contained genuine singular terms, they could not have been brushed away by paraphrase; a singular term possesses at least the *possibility* of referential linkage with other terms in other contexts, but no expression in (1) does so. Thus the fact that (1) is evidently a correct paraphrase for any token of D30 shows that D30 itself does not contain 'a man' or 'he' as singular terms. If the latter had counted as singular terms, then the paraphrase would have had to look like this:

(2) a deposits $5,000 in a savings account at 5¼% interest compounded semiannually ⊃ a has $5,432.10 in his account after three years

with 'a' as an individual constant. (1) could also have been reached from

D41: #The man who deposits $5,000 in a savings account at 5¼% interest, compounded semiannually, has $5,432.10 in his account after three years.#

D42: #He who deposits . . . #

D43: #If John Smith deposits . . . #

D43, at least, is ambiguous between the two readings, and would sometimes have to be paraphrased as (2).

Sorting out when a discourse has the referential purport of some of

its singular expressions canceled by universal quantification[44] is not a simple matter, if you want general principles which, given the syntactic and semantic structure of the discourse, will tell you whether a given singular expression is referentially segregated. As in the case of fiction, supposition, and modality, I offer no such principles, nor do I think they will be easy to find.

(v) *Existential Quantification.* Sometimes the existential quantification

(3) $(\exists x)$ (x has been eating my porridge)

is a correct paraphrase of

D44: #Someone has been eating my porridge.#

and sometimes it is not, as we can see by considering

D45: #Someone has been eating my porridge. She says her name is "Goldilocks." Here she is. What are we going to do with her?#

The same goes for

(4) The window has been forced & $(\exists x)$ (x is a burglar & x has been in here)

D46: #The window has been forced. A burglar has been in here.#

D47: #The window has been forced. A burglar has been in here. You can see his footprints quite clearly. It's Joe the Snake again — I can spot those shoes of his every time.#

and also for

(5) $(\exists x)$ $(\exists y)$ $(\exists z)$ (x is an enchanted evening & y is a stranger & z is a crowded room & you will see y on x across z)

D48: #Some enchanted evening, you will see a stranger, across a crowded room.#

D49: #Some enchanted evening, you will see a stranger, across a crowded room. She will be Agent 99, and she will give you the recognition signal. You will leave the room with her immediately and . . . #

The fact that D44, D46, and D48 can be expanded into D45, D47 and D49, which cannot take the quantificational paraphrase, shows that their

[44] A useful but inaccurate phrase to apply to sentences or discourses which would be paraphrased by universal quantifications. No sentence of English could actually be universally quantified, since English does not possess quantifiers or bindable variables.

referential segregation[45] is not something built into the structure of the discourse.

The following, however, are segregated no matter what the circumstances:

D50: #Has a burglar been in the house?#

D51: #Buy me an ice-cream cone!#

D52: #I promise to give you a bicycle for Christmas.#[46]

D53: #I dreamed that a new continent arose out of the Pacific Ocean.#

D54: #Imagine a mountain. Think of it as a place where the cops have wooden legs, and all the dogs have rubber teeth, and . . . #

Each should be thought of as involving an existential quantification along roughly the following lines:

(6) Has the following been the case? — $(\exists x)$ (x is a burglar & x is in the house)

(7) I dreamed that the following was the case: $(\exists x)$ (x is a new continent and x arises out of the Pacific Ocean)

and so on. Putting other singular expressions in place of the indefinite descriptions in D50–D53 may restore the possibility of referential linkage: e.g., 'the burglar' or 'Joe the Snake' or 'he' for 'a burglar', 'that ice-cream cone' for 'an ice-cream cone', 'that bicycle' for 'a bicycle', 'Australia' for 'a new continent'.[47]

(vi) *"Generic" Reference.* Consider:

D55: #A humpback whale is a whalebone whale with a rounded back and long, knobby flippers. It is on the verge of extinction.#

D56: #The dodo lived on Mauritius. It is extinct.#

[45] I.e., their referential segregation with respect to the singular expressions in question; for brevity I will often not state this qualification.

[46] This is borderlinish. I think that #I promise to give you a bicycle for Christmas. It's that one over there.# (indicating a particular one, not a brand or model) is incorrect, but others may think otherwise. It isn't essential that everyone agree with all my examples as long as enough of them are accepted to establish the general claims I am making, for these claims concern the range of possibilities inherent in the structure of the discourses we utter, not the details of how these possibilities are realized.

[47] But the place occupied by 'a mountain' seems absolutely closed; if we try to substitute a proper name or a demonstrative phrase, the result seems somehow malformed: #Imagine Mao Tse-tung. . . . # Imagine him doing or being what? I can't imagine just *him*, for he is already a real person.

'A humpback whale — it' and 'the dodo — it' cannot be referentially connected, directly or by way of referential linkages with singular terms in other contexts, with any particular whale or dodo, for no particular whale or dodo could be extinct — 'extinct' applies to the species as a whole or not at all. For the same reason they cannot be read quantificationally, if the variables range over concrete objects only. They must be taken as referring to the species *Megaptera novaeangliae* and *Raphus cucullatus*, respectively, if they refer to anything. The normal referential purport of the singular expressions has been canceled — they can no longer take concrete things as referents — and another involving abstract objects put in its place.

As with fiction, supposition, etc., some contexts are ambiguous between the two kinds of readings; for example:

D57: #The whale is a mammal.#

Strawson would disagree:

It is obvious that anyone who uttered the sentence, 'The whale is a mammal', would be using the expression 'the whale' in a way quite different from the way it would be used by anyone who had occasion seriously to utter the sentence, 'The whale struck the ship'. In the first sentence one is obviously *not* mentioning, and in the second sentence one obviously *is* mentioning, a particular whale.[48]

But consider:

D58: #No, ladies and gentlemen, not every creature you see here today is a fish. The lobster — he's perched on that rock in the middle of the tank — is a crustacean, and the whale is a mammal. But all the others are fish.#

Unsegregated occurrences of 'The whale is a mammal' may occur less frequently than segregated ones, but that fact is of no theoretical interest.

7

When a singular expression is referentially segregated, we can say that it is to be read "internally" or given an "internal" reading, and that the discourse itself is to be read internally with respect to that expression. When an expression is not segregated, we can say that it is to be read "externally" and that the discourse is to be read externally with respect to it. Thus, for example, 'a mountain' in D54 must be read internally,

[48] "On Referring," p. 21.

and D54 must be read internally with respect to 'a mountain'. I use 'internal' and 'external' instead of the current '*de dicto*' and '*de re*' because expressions read externally need not refer to a *thing* — be "*de re*" for some *re* — at all, or be intended as such; they may only hark back to a singular expression in another context which itself fails to refer and which may itself be referentially segregated. And I use 'internal' and 'external' instead of Quine's 'referentially opaque' and 'referentially transparent' because his distinction rests on the notion of an expression's having "purely referential position," defined as follows:

When a singular term is used in a sentence purely to specify its object, and the sentence is true of the object, then certainly the sentence will stay true when any other singular term is substituted that designates the same object. Here we have a criterion for what may be called *purely referential position*: the position must be subject to the *substitutivity of identity*.[49]

But substitutivity of identity doesn't even hold in general for the expressions in a single anaphoric chain, as I have already argued, since e.g., 'a man' and 'the man' can be linked in an anaphoric chain and will refer to the same thing if to anything but will *never* be interchangeable (see section 3). Therefore they are not used "purely to specify" their referents and are thus referentially opaque. But among them some will be read externally and some internally. Thus this distinction cuts across his.

The cases of referential segregation discussed in this section suggest that the unit of analysis — *if we are considering a natural language* — is at the very least a whole discourse, not an isolated single sentence. Those examples showed the great variety of ways in which and purposes for which referential purport can be canceled. By expanding the discourse in one way or another, one or another option is opened or closed, and it seems obvious that we can place no limit on the length of expansions which may turn out to be relevant. It is better to consider the single sentence as a nest of possibilities for anaphoric connection within a larger discourse and for referential linkage to other contexts, with the realization of these possibilities depending on subtle and complex structural properties of the discourse plus the speaker's intentions concerning inter-contextual linkage. (The relation of the sentence to the discourse is analogous to the relation of the word to the sentence.) This makes our problem harder if we want to find the "logical form" of single

[49] *Word and Object*, p. 142.

sentences by trying to show how they should be paraphrased in a canonical language, since we will typically end up with several disparate logical forms for each sentence. Quine says in *Word and Object* that

> it would be folly to burden a logical theory with quirks of usage that we can straighten. It is the part of strategy to keep theory simple where we can, and then, when we want to apply the theory to particular sentences of ordinary language, to transform those sentences into a "canonical form" adapted to the theory. If we were to devise a logic of ordinary language for direct use on sentences as they come, we would have to complicate our rules of inference in sundry unilluminating ways.[50]

It does not conflict with Quine's principles, though it does conflict with his practice, to attempt the paraphrases only at the level of whole discourses, including some very long ones. Simple sentences in isolation such as 'I saw a lion' will usually be ambiguously paraphrased. To find out whether this token uttered by this speaker on this occasion goes over to '$(\exists x)$ (I saw x & x is a lion)' or to 'I saw *a*' we must find out what discourse it is part of and how it is linked to other contexts. Once we have that information we can paraphrase the whole discourse, if that is what our purposes require.

I suggest that we take as basic the notion of an extended discourse containing one or more anaphoric chains and having a certain "content." Discourses constructed out of different sentences arranged in different ways will normally contain different anaphoric chains, but may nevertheless have the same content, like D23–D27 above. A discourse with a given content may be used for various purposes: to assert straightforwardly what the speaker takes to be ordinary fact about individual things, or to propound fiction, or to make general statements, or modally, or suppositionally, or what have you. I am not concerned with these other uses here, since they all involve cutting off the possibility of singular reference.

8

If a discourse is not referentially segregated, when and how do its singular terms refer? If T_1 in C_1 refers because it is linked to T_2 *in* C_2, that only postpones the question. How is it that T_2 refers? How does a referential connection with an extra-linguistic object make its first entry into a linguistic context?

[50] *Ibid.*, p. 158.

If T_2 in C_2 uttered by Y is referentially linked with T_1 in C_1 uttered by X, and X and Y are different people, and T_1 refers to O, and Y's utterance of T_2 is preceded by X's utterance of T_1 and is thus dependent on it for its referential connection with O, then Y in uttering T_2 has made a *secondary reference* to O. For example, taking the reporter who wrote D7 as X and the speaker who uttered D8 as Y, and assuming that X referred to DeBakey in writing 'Dr. Michael DeBakey', we can see that in uttering 'a doctor' Y made a secondary reference to DeBakey. If someone refers but does not make a secondary reference, then he has made a *primary reference*.[51]

What I just said about secondary reference doesn't cover cases in which what X says now is linked "externally" with what he said earlier, as would happen if he developed amnesia and then read things he had written before, or if his memory was bad but not wholly destroyed and he relied heavily on what he himself had previously written or otherwise recorded, and so on. Such cases shade over into ordinary ones familiar to everyone, since nobody's memory is perfect. A complete definition of 'secondary reference' would have to clear up such cases,[52] but that can be ignored here. Secondary reference is explained by primary reference. What accounts for primary reference? I will begin by discussing theories of primary reference for definite descriptions.

If a definite description in a given discourse is a genuine singular term, possessing the possibility of referential linkage outside the discourse, then it cannot be paraphrased away a la Russell, since the paraphrase will not contain any expression capable of being linked with another context; it will only contain quantifiers, bound variables, predicates, and so on, which cannot sustain the linkage. Strawson's account of definite descriptions does not brush them away; they remain singular terms (at least in most cases: it isn't clear to me whether he believes that Russell's analysis *never* holds). Given that T is a definite description which is not referentially segregated and which does not get its reference secondarily, when does it refer and what does it refer to?

[51] I found it necessary to make this distinction because of an important class of cases pointed out to me by Saul Kripke, commenting on an earlier statement of this theory that covered primary reference only. See below, section 14.

[52] By expanding 'X and Y are different people' into 'either X and Y are different people or else they are the same person and X at t_1 stands at Y at t_2, with respect to the utterance of T_1 and T_2, *as if* they are different people', or something of the sort. The 'if . . . then' could then be replaced by 'if and only if'.

I now want to introduce the notion of the "descriptive content" of a singular term in (or relative to) a given discourse.[53] Again:

D1: #A man was sitting underneath a tree eating peanuts. A squirrel came by, and the man fed it some peanuts.#

The descriptive content of 'the man' in D1 is: man who ate peanuts while sitting underneath some tree. The descriptive content of 'it' is: squirrel which came by some tree underneath which some man was sitting and eating peanuts. I will say that a singular term *denotes* something if that thing satisfies its descriptive content. Thus 'the man' in D1 denotes any man who ate peanuts while sitting underneath some tree, and 'it' denotes any squirrel which came by some tree underneath which some man was sitting and eating peanuts. A singular term *uniquely* denotes something if it denotes it and nothing else.

Thus the descriptive content of a singular term in a discourse is, roughly speaking, what sort of thing the discourse says the referent is supposed to be, what properties it is supposed to have, what sorts of other things it is supposed to be related to and in what ways, and so on. The term not only purports to refer but purports to refer to a thing of a specified kind. Do other links in the same anaphoric chain have the same descriptive content? Not if 'a man' and 'a squirrel' are assigned descriptive content in the way that 'the man' and 'it' were, for the assignment of content to 'the man' and 'it' did not include what was predicated of them at the point at which they appeared in the context. We did not include 'and who fed peanuts to some squirrel which came by' in the content of 'the man' or 'and which was fed peanuts by the man' in the content of 'it', for these were being *asserted* to hold for the referents of 'the man' and 'it', and such assertions are intended to say something new, not merely unfold what is already contained in the terms. Parallel to this, the descriptive content of 'a man' would just be 'man' and that of 'a squirrel' merely 'squirrel', and these terms would denote anything which was a man or a squirrel, respectively.

There is an issue which I have put off considering up to now. Any discourse can appear part by part, one sentence being uttered and then another and so on, possibly separated by periods of silence of various lengths. D6 (the mosquito story) and D22 (the explorer's report) would normally occur that way. On the other hand, a discourse can be produced and presented "all at once," as a single unit, like a book or a newspaper story. The audience must take it in part by part, but from the point of view of the speaker it is all of a piece: the order of the sentences and the arrangement of the anaphoric chains partly depend on contingencies of exposition which select one from among a number of different discourses having the same content, such as:

D59: #A squirrel came by a tree. A man fed it some peanuts. He had been eating them. The man was sitting underneath the tree.#

D60: #A man was eating peanuts. He was sitting underneath a tree. He fed some peanuts to a squirrel which came by.#

The same information about this situation is introduced in a different order in D1, D59, and D60, but what we make of it is the same in the end: man-sitting-underneath-tree-eating-peanuts-and-squirrel-coming-by-and-man-feeding-peanuts-to-squirrel. As far as content goes, D1, D59, and D60 are equivalent. They do not contain the same anaphoric chains — D1 has 'a man — the man', D59 'a man — he — the man', and D60 'a man — the man — he' — but they do contain the same "characters": a man, a tree, a squirrel.[54]

The singular terms in a temporally extended discourse like D6 or D22 must, I think, be assigned descriptive content in the way I just described, since "what sort of thing the referent is supposed to be" depends on what the speaker meant by what he said, how he was putting sentences together in order to construct a story with a certain content. If in the mosquito story the descriptive content of the first occurrence of 'it' is taken to include not only 'mosquito in here' but what was predicated at later links in the chain ('lands on left arm', 'bites', etc.), then the descriptive content has nothing to do with what the speaker meant and what sort of thing he was purporting to refer to, since he did not foresee the events which led him to make those remarks later in the discourse.

[53] This is an adaptation of Paul Ziff's notion of "information-content"; see his *Semantic Analysis* (Ithaca, N.Y.: Cornell University Press, 1960), pp. 97–101.

[54] But we might also have had this, which clearly has the same content as the others: #A feeding of peanuts took place. It was done by a man who had been sitting underneath a tree eating them, to a squirrel which came by.# 'A feeding of peanuts — it' appears here but not in D1, D59, or D60, which do not on the surface contain singular terms purporting to refer to events. Are we to say that such terms are somehow there anyway, that the event is also part of the story, or what? I don't believe that my account of reference depends on resolving this issue one way or the other, so for simplicity I will ignore this kind of problem and assume that all the relevant singular terms are there on the surface and that each discourse has an anaphoric chain for each character in the story. For a discussion of this problem, see Donald Davidson's "The Individuation of Events."

It will, I think, introduce avoidable complications to make allowances for discourses in which the descriptive content of the anaphorically connected singular terms accumulates over time, so henceforth I will assume that every discourse is of the simultaneous variety, in which each term in an anaphoric chain has the same descriptive content as well as the same reference as all the others. I am also going to ignore discourses uttered by more than one speaker, discourses uttered by vacillating speakers who introduce discrepancies into the content of what they are saying by retracting their previous remarks or contradicting themselves, discourses involving questions or commands or other non-assertive speech acts, and so on[55] — everything but the most humdrum factual narration.

How is descriptive content comprehended? How does the hearer manage to grasp the story and identify the characters, assigning each one to the singular terms purporting to refer to him? Like the ability to recognize anaphoric chains (which is, in fact, one of its components), this ability is subtle and complex and difficult to theorize about. Nevertheless, discourses *do* have content, the stories they tell *are* identifiable, the descriptive content of their singular terms *can* be discerned, and we can safely proceed by dealing with clear examples even though we lack an analysis of what makes them examples.

9

The orthodox view of primary reference is a theory which I will call *denotationism*. First I will discuss the denotationist analysis of definite descriptions and later on consider denotationist accounts of other singular terms. Take first the case in which a speaker utters a short discourse of the following sort:

#The king of France is wise.#

The denotationist view — as expounded by Strawson, who discusses this example in "On Referring" — is that in uttering 'the king of France' the speaker is referring to whatever thing happens to be uniquely denoted by 'the king of France'. Thus if he is referring to anything at all he is referring to something which uniquely satisfies the descriptive content of the term, i.e., some person who is sole king of whatever country 'France' refers to. If nothing is denoted by the term, or if more than

[55] It is the existence of such discourses that makes it necessary to regard multi-sentence discourses in general as sequences of sentences rather than conjunctions.

one thing is, he has failed to make a reference. Unique denotation is held to be both a necessary and a sufficient condition for primary reference in all cases involving unsegregated definite descriptions.

Denotationism is false, as Keith Donnellan has shown. He introduces a distinction between the "attributive" and the "referential" uses of definite descriptions, as follows:

> A speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so. A speaker who uses a definite description referentially in an assertion, on the other hand, uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing. In the first case the definite description might be said to occur essentially, for the speaker wishes to assert something about whatever or whoever fits that description; but in the referential use the definite description is merely one tool for doing a certain job — calling attention to a person or thing — and in general any other device for doing the same job, another description or a name, would do as well. In the attributive use, the attribute of being the so-and-so is all important, while it is not in the referential use.

To illustrate this distinction, in the case of a single sentence, consider the sentence, "Smith's murderer is insane." Suppose first that we come upon poor Smith foully murdered. From the brutal manner of the killing and the fact that Smith was the most lovable person in the world, we might exclaim, "Smith's murderer is insane." I will assume, to make it a simpler case, that in a quite ordinary sense we do not know who murdered Smith (though this is not in the end essential to the case). This, I shall say, is an attributive use of the definite description.

The contrast with such a use of the sentence is one of those situations in which we expect and intend our audience to realize whom we have in mind when we speak of Smith's murderer and, more importantly, to know that it is this person about whom we are going to say something.

For example, suppose that Jones has been charged with Smith's murder and has been placed on trial. Imagine that there is a discussion of Jones's odd behavior at his trial. We might sum up our impression of his behavior by saying, "Smith's murderer is insane." If someone asks to whom we are referring, by using this description, the answer here is "Jones." This, I shall say, is a referential use of the definite description.

That these two uses of the definite description in the same sentence are really quite different can perhaps best be brought out by considering the consequences of the assumption that Smith had no murderer (for example, he in fact committed suicide). In both situations, in using the definite description "Smith's murderer," the speaker in some sense presupposes or implies that there is a murderer. But when we hypothe-

size that the presupposition or implication is false, there are different results for the two uses. In both cases we have used the predicate "is insane," but in the first case, if there is no murderer, there is no person of whom it could be correctly said that we attributed insanity to him. Such a person could be identified (correctly) only in case someone fitted the description used. But in the second case, where the definite description is simply a means of identifying the person we want to talk about, it is quite possible for the correct identification to be made even though no one fits the description we used. We were speaking about Jones even though he is not in fact Smith's murderer and, in the circumstances imagined, it was his behavior we were commenting upon. Jones might, for example, accuse us of saying false things of him in calling him insane and it would be no defense, I should think, that our description, "the murderer of Smith," failed to fit him.[56]

Thus 'Smith's murderer' used referentially in a situation where Jones did not murder Smith does not denote Jones but refers to him, and does not refer to the person who murdered Smith (if there is such a person) even though it denotes him. Unique denotation is neither necessary nor sufficient for reference.

Donnellan's notion of referential use involves a fairly complex situation in which a speaker is attempting to get a message across to a hearer by any means necessary, including non-denoting descriptions. One might think that his point holds only for such situations, but that would be a mistake, as we can see by considering what might happen in a similar case in which the speaker asserts 'Smith's murderer is insane' purely for his own benefit. He might be writing or dictating notes on the trial, or merely expressing aloud his opinion about Jones's mental condition.[57] In uttering 'Smith's murderer' he refers to Jones even though the definite description does not denote Jones. How is this possible? If the referential connection between 'Smith's murderer' and Jones is not established as a result of Jones being denoted by 'Smith's murderer', how is it established?

So far we have been considering only those linguistic contexts which are written or uttered aloud or otherwise publicly produced. What about those which are not? Many if not most people think many if not most of their thoughts in the words of the language they speak, and these words form sentences which form discourses which are in the relevant

[56] "Reference and Definite Descriptions," pp. 285–286.

[57] It is not essential that the speaker is making an assertion. He might instead be asking himself a question: 'Is Smith's murderer insane?'

respects the same as the ones we have been discussing. Call these *covert* discourses. Most of the discourses I have discussed so far could appear covertly as well as overtly, could *be* thoughts as well as express them (that is, some tokens of them could).[58] Like overt discourses, covert discourses consist of expressions some of which will be singular expressions connected together in anaphoric chains. Some of these will count as singular terms which could be referentially linked with singular terms in other contexts. In particular, they could be linked with singular terms in overt discourses produced by the same person. Thus we could have the following:

> D62 (covert): #Jones is on trial here for the murder of Smith. He sure looks like a criminal. He's got beady eyes. He must be guilty. He murdered Smith. He's behaving very strangely. He's staring like a madman. Smith's murderer is insane. That's why Smith's body was so horribly mutilated.#
>
> D61 (overt): #Smith's murderer is insane.#

D61 is what the speaker says, D62 is what he thinks (and, in a sense, what he "means" by D61).

In a case in which D61 is not uttered idly but expresses the speaker's thoughts, the relation between D62 and D61 will be like the relation between D7 and D8. D62 contains the anaphoric chain 'Jones — he — he — he — he — he — he — Smith's murderer', and the token of 'Smith's murderer' in D62 is referentially linked with the token of 'Smith's murderer' in D61. The following referential chain extends over the two contexts:

$$\underbrace{\text{Jones} - \text{he} - \ldots - \text{he} - \text{Smith's murderer}}_{\text{D62}} - \underbrace{\text{Smith's murderer}}_{\text{D61}}$$

'Smith's murderer' in D61 refers to whatever 'Smith's murderer' in D62 refers to, and the latter refers to whatever 'Jones' refers to. In this case 'Jones' refers to Jones the defendant, and so does 'Smith's murderer',

[58] This is hardly a novel suggestion. As Plato put it:

> *Socrates.* . . . And do you accept my description of the process of thinking?
> *Theaetetus.* How do you describe it?
> *Socrates.* As a discourse that the mind carries on with itself about any subject it is considering. You must take this explanation as coming from an ignoramus; but I have a notion that, when the mind is thinking, it is simply talking to itself, asking questions and answering them, and saying Yes or No. . . . So I should describe thinking as discourse, and judgment as a statement pronounced, not aloud to someone else, but silently to oneself. (*Theaetetus*, 189E–190)

even though Jones did not murder Smith. The actual murderer is *denoted* by 'Smith's murderer' — whose descriptive content in D61 is 'murderer of Smith' — but is not its referent, since he is not the referent of 'Jones'.

This kind of example makes it possible to understand how a definite description can have a referent even when it obviously would never denote anything uniquely.

D63 (overt): #The cat is on the mat.#

'The cat' and 'the mat' denote all the cats and all the mats and hence (in this world) denote no cats or mats uniquely. Instead, they get their reference via linkages with singular terms in covert discourses.

I don't claim that a definite description *never* refers by virtue of uniquely denoting something, for we have such examples as these:

D64: #Shoot the first man who comes through that door!#

D65: #Our one-millionth customer will receive a month's free groceries and a check for $100.#

D66: #The last word uttered by the last human being will be "rosebud."#

D64 might be uttered by someone leading a defense against a lynch mob about to invade a building. Although the speaker doesn't know who the first one through the door will be — that is, at the time of utterance he can refer to that person *only* by uttering 'the first man who comes through that door' — it is *that* man, the one uniquely denoted by the definite description, who is to be shot when the time comes. He, whoever he is, is the subject of the command, and whether or not the command is obeyed is a matter of whether or not he is shot. Therefore he is the referent. D65 is similar. Whoever turns out to be uniquely denoted by 'our one-millionth customer' will get the groceries and the money if the pledge is kept, and whether or not the pledge has been kept is a matter of whether that person gets them. D66 is idle speculation, but it is true or false nevertheless, depending on whether 'the last word uttered by the last human being' uniquely denotes something which is as described. That thing will be the referent because it is the thing whose properties determine whether the statement is true or false.

I believe that Donnellan would disagree with this, for he says:

denoting and referring should not be confused. If one tried to maintain that they are the same notion, one result would be that a speaker might

be referring to something without knowing it. If someone said, for example, in 1960 before he had any idea that Mr. Goldwater would be the Republican nominee in 1964, "The Republican candidate for president in 1964 will be a conservative," (perhaps on the basis of an analysis of the views of party leaders) the definite description here would *denote* Mr. Goldwater. But would we wish to say that the speaker had referred to, mentioned, or talked about Mr. Goldwater? I feel these terms would be out of place. Yet if we identify referring and denoting, it ought to be possible for it to turn out (after the Republican convention) that the speaker had, unknown to himself, referred in 1960 to Mr. Goldwater. On my view, however, while the definite description used did *denote* Mr. Goldwater . . . the speaker used it *attributively* and did not *refer* to Mr. Goldwater.[59]

I agree, of course, that referring and denoting are distinct notions and I agree that they are not coextensive in any cases except the kind just discussed, but I think that it is possible to explain why one feels reluctant to say that the speaker is "referring" in those cases. For singular terms in general — with the sole exception of terms in referentially isolated linguistic contexts such as D64–D66 and the one Donnellan cites — primary reference can occur only if the speaker has *knowledge of* the referent (I will explain this notion later), which he clearly does not in those examples. I think there is a tendency to use 'refer' when applied to speakers in such a way as to rule out the cases where knowledge of the referent is lacking. For the sake of a uniform terminology I will use 'refer' there too, though it sounds somewhat unnatural, since I want to draw attention to what I think is the most important thing which they all have in common, namely the fact that with both kinds of singular terms there is one and only one thing the facts about which have a decisive bearing on the success or failure of the speech act, on whether the statement is true or false, the command obeyed or disobeyed, the promise kept or broken, and so on. If such a thing is related via a singular term to the speech act performed in producing the context, I will call it the *referent* of the term.

10

Unique denotation doesn't determine reference in the case of a definite description in an overt discourse which forms a referential chain with a singular term in an antecedent covert discourse. The reason for

[59] "Reference and Definite Descriptions," p. 293.

this is that the connection with the term in the other discourse provides a way for the description to be connected with a referent independent of whether its descriptive content happens to fit that referent. Denotation can fail and reference succeed because there is an alternative route to the thing referred to. If this alternative route is removed — that is, if the definite description happens not to be linked with a term in a previous context — then there is nothing left but denotation to fix the identity of the referent, and since the reference is supposed to be to one thing only the denotation must be unique. This explains why the denotationist analysis applies to some cases and also why reference is possible in the other cases. This approach can easily be extended to cover indefinite descriptions.

D67 (covert): #An insect is crawling up my arm. It looks as if it's going to bite me. It's a mosquito. There, it bit me. That really hurt.#

D68 (overt): #A mosquito just bit me.#

The referential chain is:

$$\underbrace{\text{an insect} - \text{it} - \text{it} - \text{it} - \text{it}}_{\text{D67}} - \underbrace{\text{a mosquito}}_{\text{D68}}$$

'A mosquito' refers to whatever 'an insect' refers to, and if the speaker is mistaken in identifying it as a *mosquito* then 'a mosquito' — and also any occurrences of 'the mosquito' anaphorically connected with it in expansions of D68 — would denote whatever mosquitoes there are in the world but would refer to whatever non-mosquito 'an insect' referred to: a horsefly glimpsed by a badly astigmatic speaker, a Disney-designed mosquito robot, or what have you. When a term is referentially linked with a previous context, denotation is neither necessary nor sufficient for reference.

This takes care of definite and indefinite descriptions in overt discourses. It obviously applies to pronouns as well:

D69 (overt): #It bit me.#

'It' could be linked to the occurrences of 'it' in D67; if so, it would refer to whatever they refer to. The same pattern of analysis applies to demonstratives and proper names which are referentially linked to antecedent contexts.

But this doesn't solve the basic problem: when and how do referentially unsegregated singular terms in linguistic contexts overt or covert make primary references? Singular terms like 'Jones' in D62 and 'an insect' in D67 could be referentially linked with other covert discourses, conscious or unconscious, which could themselves be linked with still others, and so on, but tracing the referential chains into the past in this way will, if we have a case of primary reference, lead us eventually to a covert discourse which contains a singular term which has a referent but is not referentially linked to any earlier singular term in any other discourse, overt or covert. When does such a term refer?

A covert discourse like D67 could occur as a kind of internal running commentary on what the speaker was (or thought he was) *perceiving*: something that looked like an insect, was seen and felt crawling up his arm, looked as if it were going to bite him, and bit him. In such a case there is no mystery about what object he was referring to: it is that object which he perceived as an insect, saw and felt crawling up his arm, and so on. Perhaps that object is really an insect, perhaps not; if it is an insect, perhaps it is a mosquito, perhaps not. Whatever it is, it is the thing he was perceiving, well or poorly, as it really was or as it really wasn't, and *it* is the referent of 'an insect' in the covert discourse and 'a mosquito' (or 'the mosquito') in the overt discourse. It is the referent because a perception of it was connected with the singular terms. If the speaker only *seemed* to perceive something but in fact perceived nothing at all, then there was no referent.

11

In section 1, I said that a context was something which had content, which could represent some thing or situation, and which could contain singular terms or analogous elements which could be referentially connected with objects in the world. Linguistic contexts are discourses, overt or covert. Other overt contexts belong to non-linguistic representational systems — maps, pictures, and so on — and these can appear covertly as well. Everything I regard as a context or an element of a context is a concrete individual object which is spatio-temporally or at least temporally located and which is capable of causal interaction with other objects.[60] As a materialist I hold that every context is *in fact* a physical

[60] Borderline case: an eternal God located outside space and time might think in words, i.e., possess covert linguistic contexts. (In which language? Hebrew?) Any-

thing of some sort,[61] but a non-materialist, a dualist for example, could still accept my analysis of reference, though we hold different opinions about what certain contexts are made of. He would regard covert contexts as immaterial, since he holds that all mental entities are immaterial, but I would regard them as states of or events in the speaker's brain. No matter how these issues are to be decided, their resolution cannot affect my theory of reference, since all I require is that contexts be concrete things of some sort.

Sometimes one context has the same content as another and can be used for the same purposes. Examples of this are the contexts D28–D30, all of which could be used to make the same point about compound interest (a false one, as it happens). They have the same content because, even though they are referentially segregated, *within* the context the singular expressions count as singular terms and these singular terms purport to refer to exactly the same kind of situation: some man depositing $5,000 in a savings account at 5¼% interest compounded semiannually and after three years having $5,432.10 in his account. With respect to content they are perfect paraphrases of one another. If someone uttered D28 and I then commanded 'Say it another way!' and he uttered D29, he would have said it another way, obeying my command. Paraphrases can be understood as translations within a given language instead of between languages, and like translations they may convey the content of the original well or badly, may stick as closely as possible to the original forms of expression or be highly non-literal, may carry over the content of the original wholly or only in part, and so on. Like translation, paraphrase can decrease the number of distinct anaphoric chains[62] but it ought not to introduce any new ones (at least, not any new ones that were not already somehow implicit in the original). For example, D1 could be paraphrased by

D70: #A man fed a squirrel some peanuts.#

which leaves the tree out of the story, but not by

D71: #A man fed a squirrel some peanuts while a frog croaked loudly.#

one who finds this intelligible will have to reject my claim that all contexts are concrete objects.

[61] Taking 'thing' in a very broad sense, so that it covers such things as events and discontinuous individuals consisting of scattered parts.

[62] That is, distinct maximal anaphoric chains — ones which could not have additional singular expressions already present in the context added to them.

because this introduces a new character into the story.

Paraphrase and translation are processes of transferring content from one linguistic context to another, and if accurate they do not alter the content substantially, except possibly to reduce it.[63] I want to generalize this notion to cover all cases in which, starting from an antecedent context with a given content, there is produced a second context which has all or part of the content of the previous context and which introduces no new content.[64] "Translation" in this extended sense is thus a process involving two temporally related concrete objects such that beginning with the first one a second one with the same or, within limits, partly the same, content is produced.

Examples:

(a) The overt discourse D68 is a (very) partial translation of the covert discourse D67, and similarly for D61 and D62.

(b) Starting from a map of the Northern Hemisphere with Mercator Projection, I draw a map of the Northern Hemisphere with Simple Conic Projection.[65]

(c) Starting from a map I describe its content in words:

D72: #In the center there is a river running roughly north-south.

[63] 'Substantially' is vague. Minute inaccuracies are frequent, e.g., in translations of books, but there are no general criteria setting off bad translations from bad attempts so inaccurate as to be nontranslations. We can live with this vagueness, and assume that in each case where accuracy is at issue some appropriate criteria are in force.

[64] I call this process *translation* in order to underline the essential similarity to ordinary translation, which is itself a consequence of the essential similarity between linguistic and non-linguistic contexts. This is the sort of thing Wittgenstein described in the *Tractatus* (4.0141): "There is a general rule by means of which the musician can obtain the symphony from the score, and which makes it possible to derive the symphony from the groove on the gramophone record, and, using the first rule, to derive the score again. That is what constitutes the inner similarity between these things which seem to be constructed in such entirely different ways. And that rule is the law of projection which projects the symphony into the language of musical notation. It is the rule for translating this language into the language of gramophone records."

[65] The same job could be done by a machine. Human abilities like conscious deliberation are not essential to inter-contextual translation; all that is required is some law-governed process which, operating on things with content, consistently yields other things with content, within prescribed limits of accuracy and reliability. Furthermore, contexts can originate mechanically. Content-bearing objects containing elements capable of singular reference (in the broad sense), and capable of referential linkage with elements in other contexts, are, for example, produced by cameras. Once produced, they can be translated mechanically as well; e.g., a camera in a satellite photographs part of the earth's surface, the negative is processed automatically, the picture is transmitted to earth by radio, and ultimately an image appears on your television screen: hurricane in the Caribbean.

A large swamp lies to the west. It's shaped like an inverted iota. A heavily forested area lies east of the river.#

Most of the information in the map is lost in D72, so unlike (b) it is only a partial translation.

(d) Looking at a painting I describe the scene it depicts by saying:

D73: (overt): #A firing squad is about to execute several prisoners. The soldiers are pointing rifles with fixed bayonets at them, and they have terrified expressions on their faces.#

(e) In the same situation, D73 occurs covertly.

(f) I picture again to myself something I once saw and translate into words:

D74: #The apparition of these faces in the crowd; petals on a wet, black bough.#

This could be either overt or covert.

In all these cases we have antecedent contexts containing singular elements which are referentially linked to singular elements in contexts derived from them by translation. In (c), for example, the map contains, say, a blue line, and this is referentially linked with 'a river' and 'the river' in D72; a portion of the map shaded green is linked with 'a heavily forested area'; and so on. The map could (and typically would) contain singular elements not linked to anything in D72, such as symbols for buildings, roads, etc. With a different selection from the content of the map, different singular terms would have been used in the second context ('a building', 'a road'). In (d) the painting contains a blotch of paint which is linked to 'a firing squad' and 'the soldiers',[66] and in (f) there is some component of the memory image which is linked to 'these faces in the crowd'.

A singular element in the second context refers to whatever the corresponding element in the first context refers to. It is obvious that the denotationist analysis doesn't apply to such non-linguistic elements as maps and pictures.[67] Denotationism holds that they refer to — i.e., are

[66] "Plural terms" like 'the soldiers' can be referentially linked to a plurality of elements in an antecedent context. If C_1 has, e.g., # . . . a soldier . . . another soldier . . . another soldier . . . #, C_2 might contain 'several soldiers' or 'the soldiers' or 'the group of soldiers' or 'the group' or 'the squad' and so on. But a detailed analysis of plural terms would involve complications not worth pursuing here. (It should be noted that they count as singular terms in my sense, for they can be referentially linked with other terms in other contexts and ultimately with referents.)

[67] Of course, they can sometimes be referentially isolated and used in much the same way as the definite descriptions in D64–D66.

maps of or pictures of — whatever they uniquely denote, and not to anything else. They uniquely denote something if it and it alone is accurately represented by them (satisfies their descriptive content). Thus a picture of a face[68] uniquely denotes someone if and only if it shows what he looks like ("resembles" him, we usually say), and a map uniquely denotes a tract of land if and only if the latter is as represented by the map: river here, forest there, and so on. Yet clearly we could have a picture of someone painted by an incompetent (or playful, or nearsighted, or avant-garde) portrait artist which failed to denote its subject, much less denote him uniquely; despite the poor likeness, it would still be a portrait of him — would "refer" to him, in my sense. Even if someone else — say an inhabitant of another galaxy who won't be born for a million years — happens to be denoted by the painting, he will not be its referent and it will not be a portrait of *him*.[69] The same goes for maps: an incompetent cartographer might produce a highly inaccurate map of *this* territory, and it would only be a coincidence that it fits *that* one perfectly.

Case (f) involves an antecedent context of the kind I now want to discuss, namely non-linguistic covert contexts involved in memory, imagination, and perception. So far I have stressed the similarity between overt non-linguistic contexts, like pictures and maps, and linguistic contexts which have the same or partly the same content; we have already accepted covert linguistic contexts, the existence of which few would seriously dispute,[70] but covert pictures and maps — "mental images," "mental pictures" — are generally thought to be unfit for regular philosophical service. They are rejected for various reasons: because they are supposed to be essentially "private" and unknowable by anyone other than the person whose mind they are in; because there are no clear criteria for telling whether somebody else has one; because a brain surgeon rummaging through your skull will see nerve tissue but no pictures or

[68] Taking 'picture of a face' as a characterization of its content, with no implication that it is referentially connected with this face or that. 'Of' is ambiguous in this respect, and where necessary I will indicate which reading I intend. Goodman would signal the content-characterization reading by the expression 'face-picture'; for a discussion see *Languages of Art*, pp. 21–31.

[69] The distinction is especially apparent for photographs; see David Kaplan, "Quantifying In," *Synthese*, 19(1968): p. 198.

[70] Even behaviorists, abandoning their principles, have been known to accept them. Gilbert Ryle, for example, speaks of "your silent colloquies with yourself" which "I cannot overhear" (*The Concept of Mind* (London: Hutchinson, 1949), p. 184).

maps or images. I will take up these points one by one. First, the privacy objection. As a materialist I hold that covert contexts are physical objects — brain states or events — which are no more private or unknowable than any other physical objects, even though they cannot be casually detected by other people. If a dualist accepts everything else I have said but claims that covert contexts are immaterial entities made of "spook-stuff,"[71] that they are private and unknowable by others, that they are somehow indelibly subjective, then *he* will have problems, but they won't be problems for me.[72]

Second, the "criteria" objection. Suppose you are looking at a painting and I ask you to tell me what you see. It is a picture of a man wearing a plumed hat, and you say 'It's a man wearing a plumed hat.' If there is an epistemological gap between what you saw and what you said — if it was pure coincidence that you uttered words whose content matched that of the painting — then certainly you didn't *report* to me what its content was, and what you said was no indication that you knew what its content was. Conversely, if we concede that you were telling me what you *knew* about its content, that you did make a genuine report, then it would be irrational not to conclude that, between the stimulation of your retina by light from the painting and the articulation of the sounds you uttered, there were inner states having the same content as what you saw and what you said. Such inner states I call "covert contexts." We can at present identify them only in terms of their content, not in terms of their physical structure, but they are not unique in that respect: if a television camera scanned the same painting and if on the screen of a television set hooked up to it by cable there appeared a picture of that painting, then — barring miracles — we would have to conclude that between the lens of the camera and the image on the screen there were inner states of the system which we who know no electronics can describe only as states having the same content as the painting and the picture of it on the screen.[73] If we were to look inside the camera, or dissect the cable, we wouldn't see any pictures, but we would see things — wires, transistors — which in fact had the same content (at that moment) as the pictures we saw, just as a surgeon who looks at your brain sees no pictures but does see things — masses of nerve cells — which at that moment have the same content as certain pictures.[74] Such objects are what I am calling "covert pictures," "covert maps," "memory images," and so on — all perfectly respectable physical things, no more mysterious than inscriptions on paper or sound waves in the air.

[71] As J. J. C. Smart calls it.

[72] What still is a problem is the alleged certainty or incorrigibility of first-person reports on the content of experiences, but the solution, whatever it turns out to be, has no bearing on my claims about reference.

[73] When I perceive my own mental states, I identify them only in terms of their content, with no knowledge of the physical structure which "embodies" or "encodes" it. It is as if I could grasp the content of a painting without seeing the paint.

Mental states such as memory, imagination, and perception have in each case a certain content.[75] If I remember, or seem to remember, the appearance of the house I lived in as a small child, then whether or not the house really looked like that, my present "memory image" (to use the everyday expression for it) gives the content of the memory: a one-story house, painted white, with a green roof, and so on. In general a necessary condition for remembering something is possessing a covert context with appropriate content, which can be translated, partially at least, into an overt context, either linguistic ('It was a one-story house . . .') or non-linguistic (e.g., by drawing a picture). I want to stress that I am concerned with memory contexts — whether they are linguistic or non-linguistic, visual, auditory, tactile, or whatever — only as objects having a certain content, and aside from that I make no claims about what their structure might or must be.[76] The same goes for imagination: any analysis will have to include the fact that a person who is imagining is imagining something of a particular sort, some possible thing or event or situation; that is, what he is doing involves a covert context with such-and-such content. If I imagine a golden mountain, then whatever else is going on in me, at least I possess a context whose content includes: golden mountain.

[74] An undeveloped negative has the same content as the print which will eventually be made from it, but you can't see what its content is; that is something you can find out only indirectly.

[75] That is, they have "intentionality." For a useful discussion, see the Chisholm-Sellars correspondence in Herbert Feigl, Michael Scriven, and Grover Maxwell, eds., *Concepts, Theories, and the Mind-Body Problem*, Minnesota Studies in the Philosophy of Science, vol. 2 (Minneapolis: University of Minnesota Press, 1958), pp. 507–539. For an interesting account of the views of Brentano and Meinong, see J. N. Findlay, *Meinong's Theory of Objects and Values* (Oxford: Oxford University Press, 1968), chap. 1, "The Doctrine of Content and Object," pp. 1–41.

[76] C. B. Martin and Max Deutscher call them "structural analogues"; their paper "Remembering" (*Philosophical Review*, 75(1966):161–196) presents a very interesting discussion of the problems involved in a causal analysis of memory.

Typically, when someone consciously perceives a physical thing or event, at least two things are happening: first, he is having an experience with a particular content; this determines what sort of thing he "seems to see" (hear, etc.). For example, if I perceive the full moon on a dark night in a cloudless sky, then normally I will possess a context whose content includes at least: white disk on a dark background. (It may, of course, include more, depending on how much detail I see, what other things are visible, what I make of the experience — e.g., not merely 'white disk' but 'moon' or 'goddess' — and so on.) Such contexts I will call *perceptual contexts* (subspecies: visual contexts, auditory contexts, etc.). 'Perceptual' is used noncommittally here: the person who possesses such a context may not be perceiving anything, but at least it appears to him that he is. Secondly, when someone perceives a physical thing or event there is a causal connection between his perceptual context and the thing he perceives: roughly, the occurrence of the former must be explained in terms of the latter. Grice has stated the case for a causal theory of perception,[77] and I won't repeat his arguments here, except to point out that a causal analysis is the only one likely to appeal to materialists. (Non-materialists may accept whatever theory appeals to them and still hold that perception involves having perceptual contexts, but they will have to find some other way of making intelligible the connection between the perceiver and what he perceives.) Again, I want to emphasize that the claims I make about perceptual contexts have to do only with their content and their relations with other contexts via translation and referential linkage, *not* with the ways they happen to be encoded in the brain (or, if you prefer, in some morsel of spook-stuff.) I don't have a materialist analysis of what it is for a perceptual context to have content, but perceptual contexts are no

[77] See H. P. Grice, "The Causal Theory of Perception," *Aristotelian Society Supplementary Volume* XXXV (1961); reprinted in Robert J. Swartz, ed., *Perceiving, Sensing, and Knowing* (Garden City, N.Y.: Doubleday, Anchor Books, 1965). For causal analyses of knowing in general, see Alvin I. Goldman, "A Causal Theory of Knowing," *Journal of Philosophy*, 64(1967):357–372; Peter Unger, "Experience and Factual Knowledge," *Journal of Philosophy*, 64(1967):152–173; and Peter Unger, "An Analysis of Factual Knowledge," *Journal of Philosophy*, 65(1968):157–170. For a non-technical survey of issues in the psychology of perception, see R. L. Gregory, *The Intelligent Eye* (London: Weidenfeld and Nicolson, 1970). Advanced discussions, from rather different points of view, are presented in James J. Gibson, *The Senses Considered As Perceptual Systems* (London: Allen and Unwin, 1968), and Ulric Neisser, *Cognitive Psychology* (New York: Appleton-Century-Crofts, 1967).

worse off in that respect than linguistic contexts; no one has a plausible materialist analysis of what it is for words and sentences to have content (sense, meaning, intension), but they still do.

Like other contexts, the non-linguistic covert contexts involved in memory, imagination, and perception[78] can contain singular elements which, like singular terms in discourses, can be referentially linked with other singular elements in other contexts. Suppose that I am now consciously remembering (or at least it seems to me that I am remembering) the house I lived in as a small child; there is a context which can be translated (fully or partially) into discourses like the following, which could appear either covertly or overtly:

D75: #It is a one-story house, painted white. It has a green roof. . . .#

D76: #The house has one story. It is painted white. It has a green roof. . . .#

D77: #There is a one-story house, painted white. It has a green roof. . . .#

'It' in D75, 'the house' in D76, and 'a one-story house' in D77 are all referentially linked to some singular element in the memory context which has the descriptive content: one-story house, painted white, with a green roof . . . This is parallel to cases (c), (d), and (e); a singular term in a linguistic context is linked with an element of an antecedent non-linguistic context and refers to whatever the latter refers to (taking 'refer' in the broad sense, of course). If in the memory context the element in question refers to object O, then O is the thing which is being remembered. It may or may not have had a green roof, may or may not have been painted white, may or may not have been a house: memories can be inaccurate, even highly inaccurate.[79] Denotation of O by a singular element in a memory context is not a necessary condition for that context's being (or containing, or being contained in)[80] a memory of O. And denotation, even unique denotation, is not a sufficient condi-

[78] The present discussion doesn't require any finer distinctions than these. A more detailed analysis would also have to take into account dreams, hallucinations, eidetic imagery, and so on.

[79] How inaccurate can they get without ceasing to be memories? A similar problem arises with perception; see below.

[80] Is the context identical with the memory or is it only one component of it? Can part of a memory context be a memory in its own right? It is not clear how these questions ought to be answered, so I will continue to use these terms rather loosely.

tion for memory: for the thing I remembered might in fact have had a brown not a green roof, or it might have been painted yellow not white, or it might have been a garage not a house, and there might have existed at the same time in another place one and only one house which I had never seen but which had just those properties, but nevertheless the latter would not have been the thing which I remembered. A memory must be causally connected with its object, which must, except in cases of time travel or precognition, exist prior to the memory.[81]

Although contexts involved in episodes of imagining, both linguistic and non-linguistic, can be referentially linked with later contexts — for I can remember and describe what I imagine — they cannot be linked with earlier ones (as long as it is an *object* which we are imagining, not just some fact about it).[82] Even apart from such linkages their singular elements cannot refer to any actual things.[83] Like fictional contexts, they are referentially segregated, and if a singular element happens to denote something uniquely, the resemblance is "purely coincidental."

This brings us again to perceptual contexts. Like all the others, they can be referentially linked to later contexts, overt and covert, linguistic and non-linguistic, including memory contexts. A person perceives something if and only if he possesses a perceptual context with a singular element which refers to that thing. You may take this as a definition: either of 'perceive', if you think you already understand 'refer', or of 'refer' (for perceptual contexts) if you think you already understand 'perceive'. In the example discussed earlier, I perceive the moon if and only if the white disk in my visual context — strictly speaking, the element in my visual context whose content is 'white disk'[84] — refers to the moon. Either side of this biconditional will be true if and only if there is a certain sort of causal connection between the thing perceived (= the referent) and the perceptual context. *Exactly* what sort of connection? As things stand at present with the theory of perception, one cannot say.[85]

Denotationism is, of course, false for perceptual contexts. What

[81] This doesn't cover every use of the words 'memory' and 'remember'; e.g., there is a sense in which I remember the meeting tomorrow if I remember that there will be a meeting tomorrow.

[82] See above, footnote 47.

[83] Leaving open the question of whether they refer to unactualized possibles.

[84] There would literally be a white disk in my visual context only if there were a white disk in my brain (or mind).

[85] See Grice, "The Causal Theory of Perception," V–VI.

counts is the causal pathway along which information passes from the object perceived to the perceptual context; it is this which determines the identity of the thing which is seen, heard, touched, smelled, etc. The information may be degraded or contaminated in transit or distorted by the perceiver, but still it is *that* object which is perceived and not some other one which, quite accidentally, happens uniquely to fit the content of the perceptual context. So, for example, if my visual context has the content white-disk-on-dark-background because I am looking through experimental lenses at a dark disk drawn on a piece of white paper, the latter is the thing I see, even though the way it seems to me is not the way it is; and even if the world happens to contain exactly one white disk on a dark background, which is not causally connected in the right way with my visual context, that thing, though it uniquely satisfies the content of my perceptual context, is not the thing I see. If in this situation I utter the singular term 'the white disk' referentially linked to the visual context, then what I am referring to is the dark disk. If a person utters a singular term which is referentially linked, directly or indirectly, to a singular element in an antecedent visual context, the questions 'What did he refer to?' and 'What did he see?' will have the same answer. In the case of hallucinations, dreams, and some kinds of optical illusion, the answer will be 'Nothing.' In other cases it won't be clear what the answer is. Here, for example, is a hard case described by G. E. M. Anscombe:

> An interesting case is that of *muscae volitantes*, as they are called. You go to the doctor and you say: "I wonder if there is something wrong with my eyes or my brain? I see" — or perhaps you say "I *seem* to see" — "floating specks before my eyes." The doctor says: "That's not very serious. They're there all right" (or: "You see them all right") — "they are just the floating debris in the fluids of the eye. You are a bit tired and so your brain doesn't knock them out, that's all." The things he says you see are not *out there* where you say you see them — *that* part of your intentional description is not true of anything relevant; but he does not say that what you are seeing is that debris *only* because the debris is the cause. There really are floating specks. If they caused you to see a little red devil or figure of eight, we should not say you saw them. It may be possible to think of cases where there is nothing in the intentional object that suggests a description of what is materially being seen. I doubt whether this could be so except in cases of very confused perception — how could a very definite intentional description be connected with a quite different material object of seeing? In such cases, if we are

in doubt, we resort to moving the supposed material object to see if the blurred, not colour-true, and misplaced image of it moves.[86]

An even thornier case is provided by the "fortification figures" experienced by people who have migraine headaches:

> The visual disturbance usually precedes the headache and can occur without any headache. It generally begins near the center of the visual field as a small, gray area with indefinite boundaries. If this area first appears during reading, as it often does, then the migraine is first noticed when words are lost in a region of "shaded darkness." During the next few minutes the gray area slowly expands into a horseshoe, with bright zigzag lines appearing at the expanding outer edge. These lines are small at first and grow as the blind area expands and moves outward toward the periphery of the visual field. The rate of expansion of both the arc formed by the zigzag lines and its associated band of blindness is quite slow: some 20 minutes can elapse between their initial appearance near the center of the visual field and their expansion beyond its limit. It is then that the headache usually begins, behind or above the eyes. It is the only unpleasant aftereffect of a spectacular visual display.[87]

The figures look about like this:

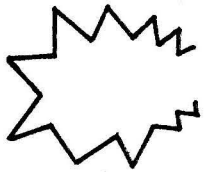

It is claimed that they are caused by a wave front of electrical activity moving across the visual cortex and interacting with neurons packed in a hexagonal lattice: "the advancing waves of disturbance draw continuous traces across the cortex and in less than half an hour reveal part of the secret of its neuronal organization."[88] I would count this as visual perception, since an element in a visual context is produced by an event in such a way as to convey information about it; the shape of the fortification figure and the speed at which it expands can be explained by the

[86] In her paper "The Intentionality of Sensation: A Grammatical Feature," in R. J. Butler, ed., *Analytical Philosophy, Second Series* (Oxford: Blackwell, 1965), pp. 178–179. An "intentional description" is one which would express the content of the perceptual context.

[87] Whitman Richards, "The Fortification Illusions of Migraines," *Scientific American*, May 1971, p. 89.

[88] *Ibid.*, p. 94.

properties of the waves and the way the neurons are packed together. Assuming that the migraine victim visually perceives events occurring in his visual cortex, does he literally see them? I am inclined to say that he doesn't. Seeing ought to involve light and eyes, or things functionally equivalent to them.

12

If the first (i.e., earliest) link in a referential chain is a singular element E in a perceptual context, and the last is a singular term T in an overt linguistic context, then the following are true under all the same circumstances:

(i) the speaker perceives, via E, an object O;
(ii) E refers to O;
(iii) T refers to O;
(iv) the speaker in uttering T refers to O.

Sometimes the perceptual context occurs immediately before the utterance of the linguistic context, and the latter serves as a kind of "running commentary" on it; this is how D6, D22, and D62 can be accounted for. Sometimes the perceptual context occurred at some time in the past and is linked by memory to the present utterance.

However, this does not account for all cases of primary reference,[89] for I can refer to objects which I have never perceived. In the remainder of this section I will present an analysis which I believe is adequate and show how it applies to the most important cases.

Excluding cases of secondary reference, the following condition holds: a singular element E in a context C possessed or produced by a person P refers to an object O if and only if either (i) E in C is referentially linked with an element E' in an antecedent context C' and E' in C' refers to O, or (ii) P has *knowledge of* O via E in C. If C is an overt linguistic context, then P in uttering E is referring to or making a reference to O. Clauses (i) and (ii) are not mutually exclusive; e.g., to remember something is to have knowledge of it, but the singular element in the memory context is linked to an element in an antecedent context (or else it wouldn't be memory). C can be either overt or covert.

What does (ii) mean? What is it to have "knowledge of" something?

[89] I.e., primary reference by singular terms that are not referentially isolated, excluding cases like D64–D66 in which the term refers to whatever it uniquely denotes. I will not state this qualification explicitly in the following discussion.

If someone perceives something, he has knowledge of it. If he detects or observes it, he has knowledge of it. If he introspects it, he has knowledge of it. If he becomes aware of it by clairvoyance or telepathy, he has knowledge of it. If he precognizes it, he has knowledge of it. If he is in mystical rapport with it, he has knowledge of it. And so on. The notion of someone possessing "knowledge of" something has been neglected in recent epistemological theorizing, which has tended to focus exclusively on "knowing how" and "knowing that"; but it has appeared before in the theory of reference. Russell distinguished between "knowledge by acquaintance" and "knowledge by description," which he characterized as follows:

> I say that I am *acquainted* with an object when I have a direct cognitive relation to that object, i.e. when I am directly aware of the object itself. . . . it is natural to say that I am acquainted with an object even at moments when it is not actually before my mind, provided it has been before my mind, and will be again whenever occasion arises. . . . When we ask what are the kinds of objects with which we are acquainted, the first and most obvious example is *sense-data*. When I see a colour or hear a noise, I have direct acquaintance with the colour or the noise. . . . we have also (though not quite in the same sense) what may be called awareness of *universals*. . . . It will be seen that among the objects with which we are acquainted are not included physical objects (as opposed to sense-data), nor other people's minds. These things are known to us by what I call "knowledge by description," which we must now consider. . . . What I wish to discuss is the nature of our knowledge concerning objects in cases where we know that there is an object answering to a definite description, though we are not *acquainted* with any such object. This is a matter which is concerned exclusively with *definite* descriptions. . . . I shall say that an object is "known by description" when we know that it is "*the* so-and-so," i.e. when we know that there is one object, and no more, having a certain property; and it will generally be implied that we do not have knowledge of the same object by acquaintance.[90]

This epistemological distinction runs parallel to an equally basic distinction in the theory of reference: *logically proper names* can be used simply and directly to refer to the objects of acquaintance, but everything else must be referred to by means of definite descriptions. A speaker can refer to something by uttering a logically proper name only if he has knowledge of it by acquaintance. Knowledge of the referent is a necessary condition for referring by means of logically proper names. What I am advocating is a kind of generalization of this principle which takes into account the facts about referential linkage and which applies to things which we cannot be acquainted with (in Russell's sense). What I call "singular terms" survive paraphrase as well as Russellian logically proper names do, and for the same basic reason: they refer, not because they uniquely denote their referent, but because they are connected with it in a more direct way, and this connection can hold even when denotation fails or when (as in the case of logically proper names) they have no descriptive content and hence no denotation. Any paraphrase which reveals their logical form will have to contain some expression which can sustain this kind of connection.

"Perception" is a loose notion. We can perceive physical objects and events external to us by seeing, hearing, smelling, tasting, and touching them; if I detect the presence of a burglar by seeing the marks he left on the window frame and smelling the smoke of his cigar, do I perceive *him*? Perhaps, perhaps not — but I do have knowledge of him, even if we don't know whether to call it perception. Seeing the marks is not seeing the burglar; smelling the smoke is not smelling the burglar; but if I see the marks and smell the smoke and say 'A burglar has just been here', then 'a burglar' and other terms anaphorically connected with it later on, such as 'the burglar' and 'he', refer to whoever left those marks and that smoke. If it turns out that they were not left by a burglar but by a policeman planting a hidden microphone in my house, then he is the referent; and he remains the referent even if there was also a burglar there who left no spoor. I *detected* the burglar — or whoever it was who was there — and even if this isn't perception in the ordinary sense it is close to it. (You might object that it differs from ordinary perception because there has been an inference from evidence; but ordinary seeing and hearing involve inferences, often highly complicated ones, though they are normally unconscious.[91])

Even if perception is taken broadly enough to cover such examples, there are others where it seems clear that nothing is literally perceived but there is — or at least there is alleged to be — knowledge of various objects. Russell claimed that we know universals by acquaintance ("though not quite in the same sense") and can on this basis refer to

[90] Bertrand Russell, *Mysticism and Logic* (Garden City, N.Y.: Doubleday, Anchor Books, 1957), pp. 202–207 *passim*.

[91] See the books by Gregory and Neisser in footnote 77. Analogous things can be said about 'observe', especially in scientific usage.

them by using logically proper names; but only metaphorically do we "perceive" them. I can't perceive something that doesn't exist yet, but sometimes I can have knowledge of an event which is about to happen. If I see a stick of dynamite attached to a burning fuse, I can by uttering 'the explosion which is about to occur here' refer to the explosion which occurs shortly thereafter, for it is within the scope of my knowledge. If the dynamite turns out to be a dud but a leaking gas main which I was completely unaware of explodes anyway, then even though that explosion is uniquely denoted by the definite description I was not referring to it: I had no knowledge of it.

The theory of reference merges on its periphery with the theory of knowledge. The denotationist analysis of reference is inadequate precisely to the extent that reference depends ultimately on knowledge of the referent by the speaker at the time of utterance, or at some earlier time, or (in cases of secondary reference) by someone else. One can have knowledge of something, by perception for example, even though one's conception of it — as given by the descriptive content of the singular element in the perceptual context — is highly inaccurate; a singular term which has that content will fail to denote the object in question, but it will still refer to it if it is referentially linked to the singular element by which the object is perceived.

There is some truth in denotationism — any popular philosophical theory has some truth in it — but for the most important uses of singular terms it fails to identify the principles by which the referent is selected. This inadequacy originated in Russell's epistemology: if the only concrete things which can be the referents of genuine singular terms are the objects of acquaintance, and acquaintance extends only to sense-data, then any non-sense-datum which is referred to at all will have to be referred to in some way that does not involve singular terms which survive paraphrase into a "logically perfect" language. The solution is the Theory of Descriptions: any sentence containing a definite description is understood as asserting that exactly one thing satisfies its descriptive content, that is, exactly one thing is uniquely denoted by it. We get at physical objects only by a semantic shot in the dark: we specify properties or relations and hope that they are uniquely exemplified. The limitations of this way of setting up referential connections were well understood by Russell:

When we, who did not know Bismarck, make a judgment about him,

the description in our minds will probably be some more or less vague mass of historical knowledge — far more, in most cases, than is required to identify him. But, for the sake of illustration, let us assume that we think of him as "the first Chancellor of the German Empire." Here all the words are abstract except "German." The word "German" will again have different meanings for different people. To some it will recall travels in Germany, to some the look of Germany on the map, and so on. But if we are to obtain a description which we know to be applicable, we shall be compelled, at some point, to bring in a reference to a particular with which we are acquainted. Such reference is involved in any mention of past, present, and future (as opposed to definite dates), or of here and there, or of what others have told us. Thus it would seem that, in some way or other, a description known to be applicable to a particular must involve some reference to a particular with which we are acquainted, if our knowledge about the thing described is not to be merely what follows logically from the description. For example, "the most long-lived of men" is a description which must apply to some man, but we can make no judgments concerning this man which involve knowledge about him beyond what the description gives. If, however, we say, "the first Chancellor of the German Empire was an astute diplomatist," we can only be assured of the truth of our judgment in virtue of something with which we are acquainted — usually a testimony heard or read. . . . All names of places — London, England, Europe, the earth, the Solar System — similarly involve, when used, descriptions which start from some one or more particulars with which we are acquainted.

. . . knowledge concerning what is known by description is ultimately reducible to knowledge concerning what is known by acquaintance.

The fundamental epistemological principle in the analysis of propositions containing descriptions is this: *Every proposition which we can understand must be composed wholly of constituents with which we are acquainted.*[92]

If we accept this principle we will have to find ways of translating each sentence used in talking about concrete objects into one containing no singular terms except logically proper names referring to sense-data. If this is impossible — and it is now generally agreed that it is — then the principle will have to be abandoned, and with it Russell's view that only sense-data can be referred to by paraphrase-immune ("logically proper") singular terms. If singular terms refer to things other than the objects of immediate acquaintance, we have to explain how this is possible. Strawson and many others believe that some version of denota-

[92] *Mysticism and Logic*, pp. 209–212.

tionism will do; I have argued that it will not, and the present approach is proposed as an alternative.

It does not contain, nor is it committed to, any particular theory of how we have knowledge of objects or any theory of what objects are there for us to know. Any theory which correctly accounts for ordinary perception and everything analogous to it, up to and including such things as clairvoyance and mystical insight, if there are such things, will be compatible with this theory of reference. Anyone who accepts a non-causal theory of perception (as I do not) can accept what I say about the connection between perception and reference, though he will reject some of my examples. Anyone who believes (as I do not) that we have knowledge of abstract entities can argue that that puts us in a position to make singular references to them, and what he says will be fully compatible with this analysis. That is why it would be incorrect to call this a "causal" theory of reference, as has been suggested to me. If it must be pigeonholed it would be more accurate to call it an "epistemic" theory; what gives it a "causal" flavor is the incidental fact that the processes of inter-contextual translation and the formation of referential chains are intelligible to most people only on a causal analysis, as are perception and memory; but that is not essential to the theory.

13

What about "demonstratives" like 'this' and 'that'? What about other indexical expressions like 'I' and 'here' and 'now'? We have already seen that 'this' sometimes has whatever reference it has because of an anaphoric connection with an antecedent singular term (see D21 and D22 in section 5). 'This' and 'that'[93] can also be referentially linked to previous contexts, linguistic or non-linguistic, overt or covert. The most important cases are those in which they are linked to singular elements in perceptual contexts. If I see a tree in front of me, my perceptual context containing some singular element whose descriptive content is, say, 'tree with white blossoms along its branches', then I can say or think 'This is the most beautiful thing I've seen today', where 'this' is referentially linked to that element. Since the latter refers to that tree (= I perceive the tree by way of it), 'this' refers to the tree also; it is the thing whose beauty or lack of it fixes the truth-value of my remark. The connection

[93] And their plurals 'these' and 'those'; see above, footnote 66.

between the word and the tree is mediated by that element in the perceptual context; if it were not, we would be unable to explain why the word hooked up with the tree rather than any of the other things present: birds, people, shrubbery, houses, or what have you.

'This' accompanied by pointing is more complex. I point only for your your benefit, in order to establish a referential link between the word 'this' which I utter and some element E in your visual context; if I am successful, my remark that "this is a so-and-so" will be understood as predicating so-and-so-ness of the thing you see via E. There will be a number of such elements at a given time, for at a given time you see, or seem to see, a number of things; to zero in on the right one I introduce into your visual context a new element, one by whose presence there you see my finger (or a stick, or whatever); this new element (think of it as the "finger-percept") is geometrically related to E in such a way as to get you to single out E as the recipient of the linkage with 'this'. There are more involved cases to which the same sort of analysis applies (e.g., I point not at the thing but at a picture of it and say 'this', referring to the thing *through* the picture of it, not to the picture itself), but I won't go into those, nor will I discuss the nonvisual analogs of pointing.

'This' without pointing can produce the right linkage if the circumstances are right. I say 'This is Fred Schultz'; no other person is in sight; you link your Schultz-percept with the token of 'this'. The less prominent the objects are, the chancier such linkages become. The theory of reference need not concern itself very much with these cases, nor with other problems of reference *to* or *for* another. I cannot refer to O for you unless I can refer to it for myself, and all the interesting problems are already posed by the latter.[94]

It is almost correct to say that 'I' refers to whoever utters or inscribes it. "Almost" because we have to rule out some obvious exceptions: numerous occurrences of 'I' in *Lolita* were inscribed by Vladimir Nabokov but none of them refer to him; a translation of a speech in a foreign language might contain occurrences of 'I', and they would refer not to the translator who uttered them but to the original speaker, who uttered not 'I' but 'ich' (or whatever); and so on. The first case can be disposed

[94] This is why I don't think that Strawson's notion of "identification for an audience" can play the leading role he has cast it in; to understand identification we must already understand reference. (See above, section 2.)

of by ruling out occurrences of 'I' which are referentially segregated, as fictional occurrences are; in the second case the word uttered by the translator is linked to the word uttered by the original speaker and for that reason refers to him, not the translator.

When 'I' is uttered unsegregated and without linkage to an antecedent context, then it does refer to the person who utters it. But who is that? Compare:

D78: #After I die, I will be buried.#

D79: #After I die, I will no longer exist.#

(These could be either overt or covert.) The statements made by someone who utters D78 and D79 can both be true only if the reference of 'I' is not the same in both discourses. Taking an individual to be an aggregate of temporal stages, the referent of 'I' in D78 must, and the referent of 'I' in D79 must not, include some corpse-stages. Descartes considered excluding everything which is not a thinking-being-stage: ". . . it might indeed be that if I entirely ceased to think, I should thereupon altogether cease to exist."[95] The scope of 'I' can vary from one context to another, depending on what descriptive content it has: 'I [body]', 'I [person]', 'I [thinking being]', and so on. Thus 'I' is similar to some definite descriptions used denotatively, namely the ones which get at their referents by building on another singular term. In D64, starting with the referent of 'that door', we get the referent of 'the first man who comes through that door' denotatively: it is whoever bears to that door the relation of coming through it first. Starting with the present temporal stage of the speaker, we get the whole referent of 'I' denotatively, according to the descriptive content of the term in that context, e.g., all body-stages causally related to the present one and satisfying certain criteria of continuity.

'Here' behaves in a similar way. If it is referentially segregated, as in a work of fiction, it refers to no place at all. If unsegregated, it can be referentially linked with a singular term in some antecedent context, perhaps a previous occurrence of 'here', perhaps some other term which can refer to a spatial location ('the Amazon basin', '50°N, 40°W', 'this room', 'the place where my body is presently located', etc.). The latter can, but need not, refer to the place where the speaker is presently located, or to something in his vicinity; e.g., I point to a map of Yosemite

[95] *Second Meditation.*

Valley and say 'Here is where I would like to be.' What if 'here' occurs neither segregated nor linked? A popular answer is to say that in such cases 'here' refers to "the place of utterance." But what place is that? Strictly speaking, the utterance is a process occurring in the speaker's mouth and throat, but that region obviously doesn't comprise the whole of the referent of 'here' in very many cases. What else is to be included? This will vary from context to context, depending on what further content attaches tacitly to 'here'. Suppose I am talking to you long distance and I say: 'It's very cold here. That's why I've spent the day indoors. But something's wrong with the thermostat — it's very hot here.' The content of the first 'here' includes something like 'this city' or 'this region' and that of the second includes something like 'this room' or 'this building'. The vagueness of the additional content insures that the regions referred to will have equally vague boundaries (just as the vagueness of 'body' and 'person' insure that the referent of 'I' — that is, the speaker himself — will begin and end fuzzily).

As with 'here', so with 'now'. Referential segregation is possible, in which case 'now' will name no time. Referential linkage is possible, in which case 'now' can name some time other than the time of utterance. If unsegregated and unlinked, 'now' names "the time of utterance," but the latter's identity depends on further conditions which vary from case to case and are likely to be vague.

When a token of 'I' or 'here' or 'now' is neither segregated nor referentially linked with an antecedent term, must the speaker have knowledge of the referent, as claimed in section 12? The notion of self-knowledge is certainly obscure, but it seems likely that any plausible analysis of it will hold that a competent speaker who consciously and deliberately utters 'I' at least knows of his own existence. Thus 'I' appears to be no exception to my claim. What about 'here' and 'now'? The difficulty is in deciding what counts as knowledge of places and times. Can I perceive or observe or detect or witness or experience or be aware of a place or a time other than by perceiving, observing, etc., some object or event located in it or at it? If knowledge of a place or time merely amounts to knowledge of something located there, then 'here' and 'now' are not exceptions to my claim; the speaker will always have knowledge of at least one thing which is located in the place and at the time in question, for he himself is located there. If knowledge of a place or time amounts to more than knowledge of something located there, then we will have

to wait and see what more is required before deciding whether 'here' and 'now' — in the relatively infrequent cases in which they are not referentially linked even to an antecedent perceptual context — can refer to things of which the speaker has no knowledge.

14

We have already seen that a proper name can be introduced *de novo* into a linguistic context by means of an identity statement:

D17: #A certain bookie I know whom I'll just call "Harry" . . . #

'Harry' is anaphorically linked to 'a certain bookie I know' and has whatever reference it has. Denotationism cannot account for proper names which get their reference this way, since it cannot explain the reference of the definite descriptions, pronouns, etc., with which such proper names are linked. Proper names can, of course, be linked to definite descriptions used denotatively — e.g., in D64 we could have had something like: 'the first man through that door — call him "Primus" ' — just as they can be referentially segregated, as in fiction, but such uses don't pose any new problems; neither do examples of proper names with secondary reference, although the latter provide some especially striking counter-examples to denotationism. (See below.)

A proper name can also be linked to singular elements in earlier nonlinguistic contexts, for example perceptual contexts. I see the moon and call it 'Luna'; the name is linked to the singular element in my visual context by which I see the moon and which, say, has the descriptive content 'white goddess in the sky' (that is, I see the moon as a white goddess in the sky). 'Luna' refers to the object I see: the moon, not some goddess. The name 'Luna' denotes nothing, but refers to the moon, because it is referentially linked with an element in my perceptual context which is in turn causally related to the moon. (Referential linkages are themselves to be given a causal analysis, of course, but I have nothing useful to say on that topic.)

What about "pure" cases of proper names, ones which don't derive their reference from linkages with other singular terms? Several denotationist theories have been developed to account for them, and they all pursue the same general strategy: (i) find, or construct, a context containing all occurrences of the proper name which according to certain criteria are relevant to determining its reference; (ii) determine the content of the name as it appears in this context; (iii) find something which uniquely satisfies the content (or a certain part of the content) of the name — that unique *denotatum* will be the referent.

In an interesting paper entitled "Substances Without Substrata,"[96] Neil Wilson proposes the following procedure to find the referent of a proper name, as used by a given individual at a given time:

> Let us suppose that somebody (whom I am calling "Charles") makes just the following five assertions containing the name "Caesar." Let us suppose in addition that we know the significance which Charles attaches to expressions other than "Caesar" and that, in the beginning at least, we are ignorant of Roman history.
>
> (1) Caesar conquered Gaul.
> (2) Caesar crossed the Rubicon.
> (3) Caesar was murdered on the Ides of March.
> (4) Caesar was addicted to the use of the ablative absolute.
> (5) Caesar was married to Boadicea.
>
> . . . We have Charles' five assertions. We now conduct an empirical investigation, examining all the individuals in the universe. We might suppose that Charles intends the word "Caesar" to signify or designate Prasutagus (who, as every schoolboy knows, is the husband of Boadicea). On this supposition (5) could be called true and all the rest would have to be called false. Or we might suppose that "Caesar" signifies the historical Julius Caesar, in which case (1)–(4) could be called true and (5) would have to be called false. There do not seem to be any other candidates since any number of persons must have conquered Gaul and/or crossed the Rubicon and/or used the ablative absolute to excess. And so we act on what might be called the Principle of Charity. We select as designatum that individual which will make the largest possible number of Charles' statements true. In this case it is the individual, Julius Caesar. We might say the designatum is that individual which satisfies more of the asserted matrices containing the word "Caesar" than does any other individual.

In my terminology, this amounts to the following: (i) We construct a discourse consisting of all the assertions containing the name in question which the speaker would be willing to make (he may or may not have actually uttered this discourse), e.g.:

> D80: #Caesar conquered Gaul. He crossed the Rubicon. He was murdered on the Ides of March. He was addicted to the use of the ablative absolute. He was married to Boadicea.#

(ii) We decide what is the descriptive content of the name in that con-

[96] *Review of Metaphysics*, 12(1959):521–539; the quotes are from pp. 530–532.

text (the whole story about "Caesar"). (iii) We find out what individual in the universe satisfies more of the descriptive conditions than does any other individual. That individual need not be uniquely denoted by all of the conditions taken together, but he must be uniquely denoted by a greater portion of them than the nearest rival. If no individual satisfies any of the conditions, or if two or more individuals are tied for first place in satisfying them, then the name has no referent. (As Wilson points out, we might wish to weight the conditions for importance, or to count some of them as absolutely essential, but these complications don't affect the main point.)

This sort of approach makes possible a denotationist account of proper names in which they are not merely abbreviated definite descriptions but have an independent status. Proper names could thus be used to initiate referential chains as well as to continue those started by other terms. John Searle proposes a similar procedure:

> . . . though proper names do not normally assert or specify any characteristics, their referring uses nonetheless presuppose that the object to which they purport to refer has certain characteristics. But which ones? Suppose we ask the users of the name "Aristotle" to state what they regard as certain essential and established facts about him. Their answers would be a set of uniquely referring descriptive statements. Now what I am arguing is that the descriptive force of "This is Aristotle" is to assert that a sufficient but so far unspecified number of these statements are true of this object. . . . The question of what constitutes the criteria for "Aristotle" is generally left open, indeed it seldom in fact arises, and when it does arise it is we, the users of the name, who decide more or less arbitrarily what these criteria shall be. If, for example, of the characteristics agreed to be true of Aristotle, half should be discovered to be true of one man and half true of another, which would we say was Aristotle? Neither? The question is not decided in advance.[97]

No matter how ties are to be judged, the referent of 'Aristotle' must satisfy at least *some* of the agreed conditions:[98]

> . . . if a classical scholar claimed to discover that Aristotle wrote none of the works attributed to him, never had anything to do with Plato or Alexander, never went near Athens, and was not even a philosopher but was in fact an obscure Venetian fishmonger of the late Renaissance,

[97] "Proper Names," in Charles E. Caton, *Philosophy and Ordinary Language* (Urbana: University of Illinois Press, 1963), pp. 158–159.

[98] As Searle asserts in his *Encyclopedia of Philosophy* article on names (vol. 6, p. 490).

> then the "discovery" would become a bad joke. The original set of statements about Aristotle constitute the descriptive backing of the name in virtue of which and only in virtue of which we can teach and use the name. It makes sense to deny some of the members of the set of descriptions of the bearer of the name, but to deny them all is to strip away the preconditions for using the name at all.

Presumably, the referent of the name must also satisfy *more* of the conditions than any other individual (for then why wouldn't that other individual be the referent?) — as required by Wilson's Principle of Charity.

Similar views have been developed by P. F. Strawson[99] and Paul Ziff.[100] Can any such view be correct? Remember that such an analysis cannot account for proper names which are referentially linked to antecedent singular terms (or to singular elements in non-linguistic contexts), for those names take whatever referents their antecedents have, whether or not the referent satisfies their descriptive content, or satisfies more of it than any other object. If I christen the moon 'Luna' (e.g., link it to a moon-percept in a visual context) and then proceed to use the name with the content 'white goddess in the sky', it nevertheless names the moon, which is no goddess, and would not name any white goddess who happened to be lurking, unglimpsed by me, in the sky that night. Similar examples can be constructed in which one's use of the name is linked to a memory context which derives from some past encounter with the object but which has with time become very inaccurate. The present memory and the linguistic context based on it determine the content of the name, which because of the inaccuracy doesn't denote the object, but the connection with the past experience of the object makes it *that* object which is named by the name.[101] If we consider cases of secondary reference — which can be understood as a kind of social analog to individual memory — the same kind of examples crop up, such as this case described by Donnellan:

> . . . The sort of description generally mentioned as helping to pick out,

[99] See *Individuals*, pp. 190–192.

[100] See his *Philosophic Turnings* (Ithaca, N.Y.: Cornell University Press, 1966), pp. 94–95.

[101] For an example of how a proper name could be based on confused memory, see Keith Donnellan, "Proper Names and Identifying Descriptions," *Synthese*, 21 (1970):343. His discussion of denotationist theories of proper names based on what he calls the "principle of identifying descriptions" applies to proper names the insights about reference contained in "Reference and Definite Descriptions."

say, Thales, is such as 'the Greek philosopher who held that all is water'. Nothing is made of the fact that such descriptions are given by us derivatively. We might be pardoned if we supposed that the referent of 'Thales' is whatever ancient Greek happens to fit such descriptions uniquely, even if he should turn out to have been a hermit living so remotely that he and his doctrines have no historical connection with us at all.

But this seems clearly wrong. Suppose that Aristotle and Herodotus were either making up the story or were referring to someone who neither did the things they said he did nor held the doctrines they attributed to him. Suppose further, however, that fortuitously their descriptions fitted uniquely someone they had never heard about and who was not referred to by any authors known to us. Such a person, even if he was the only ancient to hold that all is water, to fall in a well while contemplating the stars, etc., is not 'our' Thales.

Or . . . suppose no one to have held the ridiculous doctrine that all is water, but that Aristotle and Herodotus were referring to a real person — a real person who was not a philosopher, but a well-digger with a reputation for saying wise things and who once exclaimed, "I wish everything were water so I wouldn't have to dig these damned wells." What is the situation then regarding our histories of philosophy? Have they mentioned a non-existent person or have they mentioned someone who existed but who did not have the properties they attribute to him? My inclination is to say the latter.[102]

And thus we can see how Aristotle *could* turn out to have been an obscure Venetian fishmonger of the late Renaissance, if enough documents we now think are genuine and accurate happen in fact to be spurious or mistaken. (Recall Russell's example of the world having

[102] "Proper Names and Identifying Descriptions," pp. 352–353. This kind of example, in which a remote historical figure turns out to be totally different from what we thought, was first pointed out to me by Saul Kripke, as a counter-example to denotationism in which the speaker does *not* have knowledge of the referent. His example concerned Jonah, who is supposed to have gone to sea, been thrown overboard, been swallowed by a great fish, etc., but who might in fact have been a landlubber about whom people told tall tales which were eventually recorded in the Book of Jonah and about whose life we actually know nothing at all. Kripke's approach to the theory of reference is presented in his "Naming and Necessity," in Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language* (New York: Humanities, 1972). (This paper was written before I had an opportunity to read "Naming and Necessity," and what I have said about Kripke's views derives from a conversation with him; I don't think, however, that there is any important difference between the former and the latter.) A similar "causal" theory of secondary reference is developed by David Kaplan in "Quantifying In."

Although examples of secondary reference would refute denotationism even if nothing else did, I don't believe that it poses any important problems not already present when the referential chains are confined to a single speaker.

been created five minutes ago, complete with prefabricated dinosaur bones, histories of philosophy, etc.)

As with perception and memory, names which hook up with their referents via secondary reference are counter-examples to any denotationist analysis. (Unlike the former, they do not satisfy the condition that the speaker have knowledge of the referent.) Is anything left for denotationism to be correct about?

First, we have to exclude names of historic figures like Thales and Aristotle, whom we can pick out only by making secondary references. (For the sake of brevity I am ignoring proper names of things other than people.) Second, we have to exclude many names of contemporary persons of whom the speaker has no knowledge but whose names he has heard or seen, often in a context which supplies them with minimal descriptive content (e.g., 'John Smith, San Francisco' at the bottom of a petition). Third, we have to exclude all names which are linked to present or past perceptual contexts containing singular elements referring to the persons named; this eliminates all names whose use is based on personal encounters with their referents. Fourth, we have to exclude all names which are linked to non-linguistic contexts like photographs, TV images, and the like. But we can stop here, for it is obvious that virtually all of the names uttered by speakers of natural languages have already been eliminated. Might there nevertheless be some examples, however remote, of proper names for which the denotationist analysis is correct? Perhaps something like the following: several outstandingly horrible murders are committed under similar conditions, with the same technique, against the same kind of victim. We begin to say that the city is menaced by "Jack the Ripper." We don't claim that Jack the Ripper committed *all* of the murders, for some of them might be due to coincidence, or to imitators of the earlier ones, or to persons with saner motives using the Ripper murders as a cover. We don't claim of any given one of the murders that *it* was committed by Jack the Ripper — not even the first one, for it is perfectly possible that somebody else did that one and thus unwittingly inspired Jack the Ripper to do all the others. But we do claim that Jack the Ripper committed a substantial number, probably most, of them, and that he certainly committed more of them, probably a lot more of them, than anyone else. If we now assert that Jack the Ripper is a raving lunatic, what we say is true if and only if there is someone sat-

isfying those conditions and *he* is a raving lunatic. If half the murders were committed by A and half by B, then we might refuse to concede Ripperhood to either of them. If the murders turn out to have been committed by a woman (Jill the Ripper?), we might say that there was no such person as Jack the Ripper, on the grounds that maleness is essential to him. The denotationist analysis of proper names fits this kind of example perfectly, but it is easy to see how rare this kind of example is, how rare it is for the referent of a proper name always to be lurking unobserved just beyond our view, never quite definitely pinned down to this or that alleged fact about him, never identified as someone already known to us. Is this an exception to my claim that primary references generally involve knowledge of the referent? What *do* we know about Jack the Ripper? Do we really know anything about him? I am inclined to say that if he exists at all then we have knowledge of him: we detected his presence among us even though we never actually observed him. (To simplify matters I assume that none of us ever observed at any other time the person who was Jack the Ripper.) It is like the example in section 12 of the burglar who leaves traces but whom I never actually perceive. The distinction between having knowledge of a thing and merely having knowledge of its effects is a vague one, as is the distinction between seeing a thing and merely seeing some of its effects (or some of its parts). The notions of perceiving, observing, detecting, "finding," or in some other way coming to have knowledge of an object may depend for their general utility on the fact that we generally can get a firmer grip on the objects we are interested in than was ever gotten on Jack the Ripper.

Finally, I must correct an oversimplification which I have indulged in for the sake of orderly exposition. All along I have spoken as if a singular term which is introduced by way of an identity statement linking it anaphorically with another term in the same context, or which is referentially linked to an antecedent singular term or other singular element in an earlier context, were linked with only one such antecedent, which would provide the only possible route connecting it with its referent. In other words I have assumed that all referential chains look like this

and that none look like this

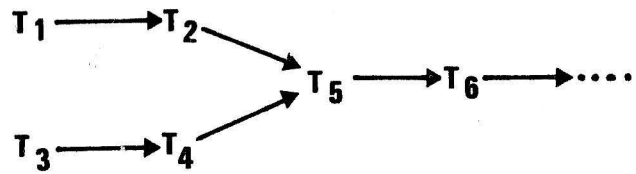

with two or more chains converging at a single point where a new term is introduced. Yet this assumption is obviously false. Suppose that my use of a proper name is linked to independent earlier uses of it by two different people, who I assumed were referring to the same person. If my assumption was correct, then all is well: both paths lead back to the same referent, which is thus the referent of the tokens of the name which I utter. But what if they lead back to different people? To whom am I referring? Both persons? Neither?

Consider an example not involving secondary reference. I see someone enter a building and a minute later I see someone who looks and is dressed exactly like him come out. Later on I utter (or think) the definite description 'the man I saw at City Hall', where this term is linked both to the visual context by which I saw someone entering the building and the visual context by which I saw someone leaving. (They are given equal weight.) If the two percepts hook up with the same person, then all is well: he is the referent. But what if the person I saw entering was the identical twin of the person I saw leaving? Am I referring to both? To neither?

If two or more referential chains converge at the point where a new term is introduced, I will call the latter *multiply linked*. Each of the terms to which the new term is linked could itself have been introduced by a multiple linkage, and multiple linkages obviously can involve many more than two antecedent terms, so that the possibility of there being some discrepancies in the pedigree of a singular term involving memory or secondary reference is very real indeed. Thus it is entirely possible that some name I now use with no qualms is connected with several different objects, as a result of my own or somebody else's past mistakes, but that I will never be in a position to discover this fact: skepticism about one's references is as easy to slide into as skepticism about one's claims to know facts about the things one is referring to. When such a mistake has been made, am I unwittingly

talking about two (or more) different people and perhaps uttering truths about one of them and falsehoods about the other? Or am I talking about neither and thus about no one at all?

If one found out about the mistake, then perhaps in most cases common sense would find a way of repairing the damage, e.g., by disconnecting my present token of 'Franklin Roosevelt' from a single maverick referential chain which leads back to Teddy Roosevelt. But what if the mistake is never discovered? It is irrelevant that *if* it were discovered we would know how to correct it and would subsequently be able to make a definite reference to a single object, for the problem is: what are we *now* referring to with the uncorrected version? Furthermore, some mistakes seem to be wholly incorrigible, no matter how much we might find out about the circumstances surrounding them (e.g., my example about 'the man I saw at City Hall'). Such attempted references would have to be abandoned altogether.

Do we live in a world in which such mistakes occur very often? Minor mistakes like the one about 'Franklin Roosevelt' can simply be ignored, even though it seems likely that they occur frequently, just as we ignore the fact that there are no perfectly straight lines or perfectly flat surfaces in the world and continue to employ geometrical concepts which presuppose that there are. Major mistakes, irreparable ones which spoil the reference beyond repair, may not occur very often in our world, at least among the more rational inhabitants of it, but if they do there are no fail-safe devices to insure that the references go through anyway. To suggest that a version of Wilson's Principle of Charity be employed here, that the referent of the term is the one which would make the greatest number of our assertions true, looks like an evasion of the problem, a desperate attempt to insure that despite our cognitive mishaps we are still referring to *something*. I see no reason why the theory of reference should be loaded in favor of our generally saying things which are true (or at least truth-valued), no more than the theory of knowledge should be loaded in favor of our generally remembering things as they were, seeing them as they are, or predicting them as they will be. Multiply linked singular terms are inherently risky, like memory, perception, and induction.

15

Let me conclude this paper by borrowing the words with which Carnap concludes *Meaning and Necessity*. It seems to me that his remarks, although written more than twenty-five years ago, still apply to the present situation (except, perhaps, that instead of 'the best method' I might prefer to say 'the true theory'):

Let me conclude our discussions by borrowing the words with which Russell concludes his paper ["On Denoting"]. It seems to me that his remarks, although written more than forty years ago, still apply to the present situation (except, perhaps, that instead of 'the true theory' I might prefer to say 'the best method'):

'Of the many other consequences of the view I have been advocating, I will say nothing. I will only beg the reader not to make up his mind against the view — as he might be tempted to do, on account of its apparently excessive complication — until he has attempted to construct a theory of his own on the subject of denotation. This attempt, I believe, will convince him that, whatever the true theory may be, it cannot have such a simplicity as one might have expected beforehand.'