SENTENCE MEANING

A thesis submitted by Brian Loar
to the Board of the Faculty of
Literae Humaniores for the
degree of Doctor of Philosophy



ACKNOWLEDGEMENTS

I am extremely indebted to Stephen Schiffer, for years of very enjoyable and fruitful conversations on these matters. I have profited a great deal from conversation about meaning and language with H.P. Grice, Marsha Hanen, Steven Stich, P.F. Strawson, and Robert Ware.

CONTENTS

ABSTRACT iv

INTRODUCTION ix

I THE MEANING OF SIMPLE SIGNALS 1
1. Some preliminaries about Signals and about Meaning 1
2. Schiffer on the Meaning of Simple Signals 5
3. Convention 10
4. The Convention which Governs Meaning 11

II REFERENCE AND PROPOSITIONAL ATTITUDES 20
Introduction 20
1. A Theory of Quine's 21
2. Referential but Non-extensional Occurrences of Singular Terms . 27
3. Reference and Satisfaction Conditions 30
4. An Account of the Dual Role of Some Singular Terms 33
5. Uniqueness Properties and Self-ascriptive Attitudes 38
6. The Extension to Propositional Attitudes in General 47

III REFERENCE AND THE MEANING OF SIMPLE SIGNALS 49
1. Non-propositional Signals 49
2. Qualified Indexicality 53
3. What Reference Is 55
4. The Meaning of Signals which Mean The E is F 60
5. Reflexive Signals 64
6. Referential Qualifiers 71
7. How to Refer on the Conventional Utterance of a Signal . . . 75
8. What is Meant on the Conventional Utterance of a Signal . . . 82
9. Signals with Imperative Meaning and Other Illocutionary Forces. 91

IV SENTENCE MEANING 99
1. The Convention which Governs the Meaning of Sentences . . . 99
2. A Problem Concerning the Nature of Meanings 110
3. Propositions and Intensions 117
4. Intensional Structures 121
5. Frege's Theory Incorporated 130
6. Word Meaning and Logical Syntax 140

Works Referred To 150

ABSTRACT

Introduction. Grice's first definition of speaker's meaning is described and stated. An important class of counterexamples lead to adopting a definition by Schiffer, which employs the notion of mutual knowledge.

I.1 Signals have no semantical structure, but have meaning. The intention is announced of defining 'z means p in G' in terms of speaker's meaning.

I.2 A signal's meaning consists in its being governed by a certain kind of convention. Schiffer's' analysis of 'x means p in G' is given.

I.3 An account of convention by Schiffer is stated. His two definitions together do not provide a convention which governs meaning.

I.4 A form of convention is proposed which is adequate for defining 'x means p in G'. It is argued that there are specifiable circumstances in which, if one utters a signal, then one's utterance accords with the convention only if one means such and such.

II Introduction. Not all signals are propositional; hence, to give their meaning, one must quantify into 'to mean that', about which there are difficulties.

II.1 A theory of Quine's is described in which he accounts for a limited kind of quantification into belief contexts. It is argued that Quine's

fails to make sense of a certain valid inference.

II.2 It is claimed that, contrary to entrenched assumptions, there are occurrences of singular terms which are both referential and non-extensional, that is, which do not satisfy what some have called 'Leibniz' Law', which is here called the "substitution principle".

II.3 The substitution principle is not a priori true. A general explanation of exceptions to it is offered.

II.4 The logical form of sentences containing non-extensional but referential occurrences of singular terms in belief contexts is stated. The theory requires that, on such occurrences, singular terms have a dual role.

II.5 Several points here. First, belief of things is under some uniqueness conception of them. Secondly, this cannot in general be construed in terms of uniqueness properties. Thirdly, what is required is sometimes that the believer's conception of something is as being uniquely related to him. Fourthly, a problem about self-ascriptive beliefs, and a way of representing their logical form.

II.6 An apparent difficulty in extending the account to 'means' is dealt with.

III.1 An account is given of certain signals which have pure indexical meaning — that is, signals which mean it is F - extending the device of Chapter I.

III.2 Qualified indexicality, as in signals which mean the E is F,

requires some new apparatus, in particular, the notion of referring.

III.3 A definition of referring qua, produced jointly with Schiffer, is given.

III.4 'z means <u>the E is F</u>' is defined, extending the kind of definition used for 'z means <u>it is F</u>', and using the notion of referring qua.

III.5 An account is given of the meaning of signals which are (as it were) "token-reflexive".

III.6 How to represent complex referential embedding in the specifying of the meaning of signals. "Referential qualifier" is defined.

III.7 A general account is given of what conventional restrictions there are on the references which are made on the utterance of a signal. The notion "super-reference" is defined; this captures how one's primary references are related to one's n-ary references.

III.8 Incorporating super-reference, and earlier points aboutthe relation between meaning and reference, we arrive at a general account of what is meant on the utterance of any indicative signal, regardless of referential complexities. The definition of meaning for all indicative signals.

III.9 An account of the logical structure of speaker's imperative meaning. How to ascribe meaning to imperative signals, and to signals which have specific illocutionary forces. The final definition of meaning for all simple signals.

IV.1 It is argued that the assumption that there are separate conventions for word meanings is superfluous. The convention which governs sentence meaning incorporates a complex abstract relation which associates each sentence with each of its meanings. A general definition is given of what it is for a sentence to mean such and such.

IV.2 A problem is raised about the nature of the entities which are to serve as meanings (propositions and propositional functions.) It is claimed that the unanalyzed notion of proposition will not enable us to state recursively the relations which relate sentences to their meanings.

IV.3 An identification of propositions with certain intensions — namely, functions from possible worlds to truth values — is rejected from counterexamples. Carnap's notion of intensional isomorphism serves as a constraint of any adequate account of propositions.

IV.4 Propositions are identified with intensional structures, certain hierarchically ordered entities which contain both intensions and structural functions, which are functions from intensions to intensions. Propositional functions and referential qualifiers are similarly construed.

IV.5 Frege's theory that the denotatum of a that-clause is the sense of its embedded sentence is shown to be compatible with the requirement that the meaning function for a language be recursively specifiable.

IV.6 It is claimed that, with an important qualification which is consistent with the general program, the notion of word meaning is derivative from that of sentence meaning. An example of how this may be shown for a particular language is given. A similar claim is made about the

derivativeness of logico-syntactic vocabulary, and an example of how to do it for a simple language is given. The concluding words express a certain unsupported bias concerning the nature of grammar.

INTRODUCTION

In 1957, H.P Grice published the article "Meaning", which has since had a most important influence on philosophical writing about language and speech acts. To describe it at its most general level, Grice's account makes utterer's meaning, or, as I shall henceforth call it, speaker's meaning, the conceptually basic notion in the theory of meaning. Persons produce sounds, gestures and marks in the presence of others, and mean something by those utterances. Grice proposed a definition of the notion whereby it is not a necessary condition of a person's meaning something that what he produces has any conventional meaning whatever. At the same time, the notion defined is intended to be fully general, and to cover all communication, from a caveman's tentative grunts to the orations of Cicero.

More specifically, Grice proposed that, in the fundamental sense of speaker's meaning, a person A meant something by his utterance x provided that "A intended the utterance of x to produce some effect in the audience by means of the recognition of this intention". It is not, as earlier causal theories suggested, sufficient for meaning something (nor is it necessary) that one produces an utterance of a type which has a tendency to produce certain effects in others. One must intend to produce the effect. Nor is that sufficient, since, as Grice pointed out, secretly planting evidence in order to affect the beliefs of another hardly counts as meaning something. So one must also intend that one's intention be

recognized. But that is not enough, for it would allow as an instance of meaning, for example, the informative display of a photograph, where one's informational intent is intended to be recognized. There is a difference between "deliberately and openly letting someone know" and "telling". Grice locates the difference in there being in the latter case an utterer's intention that the hearer's recognition of the primary intention should be operative in producing the primarily intended effect. That is what is special about telling.

Now as Grice showed in detail in a subsequent paper[1], there are crucial counterexamples to this definition. One learns a great deal about the nature of communication from these cases, but I shall not describe any here, except for one most important class. In that paper, Grice also gives a new kind of analysandum, as well as a new analysis. The early account was of meaning something or other, but that notion rests on the notions of meaning that p, and meaning that one's hearer is to do a certain thing. In the former case, the intended response is the belief that p, and in the latter, a certain action, or intention.

Throughout Grice's new analysis of these concepts, the original definition remains the heart of the matter. The central element in the concept of speaker's meaning is the intention to produce a certain kind of response in one's audience, a belief, intention, or action. That this be brought about by means of recognition of that intention is a require-

---

[1] "Utterer's Meaning and Intentions". For details on works mentioned, see Works Referred To.

ment for only certain kinds of case in Grice's revised account.

There is a certain class of counterexamples to the first account, presented and discussed by Grice in the later paper, which raises a fundamental problem in explicating what meaning is. P.F. Strawson pointed out that the following case is a counterexample to the definition's providing sufficient conditions for a speaker's meaning something.[2] S is arranging "evidence" that p, in a place where A will see it. Moreover, S knows that A is watching him do this, and that A does not know that S knows that A is watching him. S intends A to think that S does not wish A to have false beliefs concerning p (A has to be helped, a bit unscrupulously, to see the truth), and so intends A to believe that p out of trust for S in this respect, despite his crafty devices. Hence, where the intended response is the belief that p, Grice's conditions are satisfied. But, as Strawson writes, "this is not a case of attempted communication in the sense which...Grice is seeking to elucidate."[3] There is in the case a lack of explicitness of intention which seems crucial to the notion of speaker's meaning. Strawson suggests the addition of a further condition to Grice's definition, a condition represented in the following by 4):

1) S intends (i_1) to produce by uttering x a certain response r in A; and

2) S intends (i_2) that A should recognize S's intention (i_1); and

3) S intends (i_3) that A's recognition shall function as (part of)

---

[2] "Intention and Convention in Speech Acts"

[3] Ibid., p.447

A's reason for producing r; and

4) S intends (i_4) that A should recognize his intention (i_2).

As it happens, there are (as Strawson suggested there might be) yet further cases which show that these conditions are not sufficient, and in these cases there is the absence of a further intention (i_5) that intention (i_4) be recognized.[4] The device of adding further such conditions up to some (i_n) is unattractive, since there is a principle in the kind of counterexample involved which suggests that a counterexample could be produced at each level.[5] But, equally importantly, there is an element in normal communication which no such set of conditions could capture, which may be stated thus: all of the relevant communicative intentions are out in the open. Stephen Schiffer, whose theory of speaker's meaning I shall presuppose in the following chapters, has pointed out that there is a psychological element present in many interpersonal situations including, but by no means confined to, communication situations, which provides us with what we need to express the requirement of openness of communicative intent. He calls the element mutual knowledge.

Two persons, S and A, mutually know that p provided that

1) S knows that p;

2) A knows that p;

3) S knows that A knows that p;

4) A knows that S knows that p;

and so, ad infinitum.

---

[4]For a counterexample of this kind see Grice, op. cit., p. 155

[5]See Schiffer, Meaning, II.1

If two normal persons are present together at an explosion, and are aware of each other in the normal way, they mutually know that there was a bang. Most pairs of us mutually know that the Pope is Catholic.

Roughly speaking, Schiffer's use of the concept of mutual knowledge is this: in meaning something, S intends it to be mutual knowledge between himself and A that he intends to produce in A a certain response. S could not intend his intention to be more out in the open than that. The other novel features of the following definition of Schiffer's can be fully appreciated only by seeing how they deal with the counterexamples which prompt them. I refer the reader to chapters II and III of Meaning. Here is Schiffer's account of indicative meaning in the presence of an audience.[6] He also provides an account of imperative meaning, and a fully general account which covers cases of meaning where there is no specific intended audience, as in the posting of notices.[7]

> S meant that p by (or in) uttering x iff S uttered x intending thereby to realize a certain state of affairs E which is (intended by S to be) such that the obtainment of E is sufficient for S and a certain audience A mutually knowing (or believing) that E obtains and that E is conclusive (very good or good) evidence that S uttered x with the primary intention
> (1) that there be some r such that S's utterance of x causes in A the activated belief that p/r(t);
> and intending
> (2) satisfaction of (1) to be achieved, at least in part, by virtue of A's belief that x is related in a certain way R to the belief that p;
> (3) to realize E.

One feature of the definition needs special explanation. The variable 'r' ranges over reasons, and the notation 'the activated belief that p/r(t)'

---

[6] Ibid., p. 63

[7] Ibid., pp. 75-76

indicates that the audience is intended to have the belief that p for some reason of a truth-supporting, as opposed to a practical or moral, kind. In saying "you should, on moral grounds, believe that your wife is faithful" one would not thereby mean that his hearer's wife is faithful, since it is not necessarily intended that the belief be grounded in truth-supporting reasons.

There are two further elements in Schiffer's overall theory of meaning which should be mentioned. First, he claims that to mean something is, for some p or some X, either to mean that p, or to mean that one's hearer is to do X.[8] In support of this, he points out that in those cases of meaning in which there is some intended effect apart from a belief or an action — for example, anger or distress — there is also an intended belief or action whose specification is, in effect, the specification of what is meant. He also, in support of the claim, argues that for any English sentence, any true report of what is meant by a literal utterance of that sentence entails, for some p or X, either that the utterer means thereby that p, or means that someone is to do X.

Secondly, Schiffer offers an analysis of the nature of illocutionary acts[9] in which, roughly (a) performing an illocutionary act is identified with meaning something, and (b) types of illocutionary acts are classified and defined solely in virtue of various aspects of i) what is meant by the utterer, and ii) what type of reasons his hearer is intended to have for producing the intended response. If the account is correct, then the

---

[8] Ibid., ch III, sec. 5
[9] Ibid., ch IV, sec 2

concepts of the central kinds of illocutionary act — those fundamental to communication — presuppose in their analysis nothing of convention or of linguistic meaning, but presuppose only those concepts of intention and belief which are already present in the concept of speaker's meaning.

CHAPTER I

THE MEANING OF SIMPLE SIGNALS

1. Some Preliminaries about Signals and about Meaning

It is not only sentences and bits of language which have meaning. There are also simple signals, like certain gestures, the white flag of truce or surrender, trail-markers, air-raid sirens, and much of the paraphernalia of traffic control. Such utterance types are semantically unstructured. Their meaning is not a function of their parts or subsidiary features; they are, rather, characteristically learned and retained one by one, and not as the result of learning the meaning of more basic signalling-elements. In no really interesting sense do simple signals belong to systems, although they may belong to families of signals whose physical forms and whose meanings belong to narrow ranges, as with certain sets of traffic signals.

We should expect that, if such problems are solvable at all, an account of the meaning of simple signals will be more easily found than the corresponding account of sentence meaning. The relative tractability of the problem is not its only merit as a starting point, for, as we shall see, the essential feature of its solution will also determine the subsequent account of sentence meaning. With simple signals we may concentrate on the question of meaningfulness in abstraction from the compositional nature of sentence meaning, and, by freeing the

the notion of meaningfulness from the notion of language, we may elucidate the latter in terms of the former. Sentence meaning and the meaning of simple signals are much more closely related than they may seem—indeed, there exists a unified account of the two.

To give a quick preview of this similarity—a simple signal's having a certain meaning among members of a group will be determined by the group's having a certain kind of convention. Similarly, the meaning which a sentence has in a group is determined by the existence of a parallel convention—not, in this case, one which governs the sentence directly, but, rather, one which governs an infinite set of sentences to which the sentence belongs. In both cases the effect of the convention is the same; sentences and simple signals are alike in being conventional devices for communication—in a sense, this is all their having meaning consists in.

The fundamental concept of the theory of meaning is that of speaker's meaning. The account which I shall offer of sentence meaning will depend upon this concept, and on another key concept—namely, convention. Speaker's meaning, as we have seen, has an explication in wholly non-semantical terms, and, as we shall see, so does convention. It is psychological notions—principally those of belief and intention—which underlie our two key concepts. So, in effect, what I hope to have shown in what follows is that all semantical facts, including facts about word meaning, are in a way psychological facts, facts about the communication practices of a group. For many, the general idea will be already obvious, but it is by no means a trivial question what kind of

psychological fact a semantical fact is.

So far I have posed no precise question. We need an explicandum. In order to get a firm question with which to begin, I shall immediately introduce a restriction. Simple signals, like sentences, have "moods"; they may conventionally have the force of indicatives, or of imperatives. For the moment, I shall consider only the question what it is for a simple signal to have a certain indicative meaning in a certain group, and this we may represent by the form: x means p in G[1], where 'x' and 'G' range over simple signal types and groups of people, and 'p' indicates a propositional context.

Perhaps no distinction is more important for avoiding confusion concerning our explicandum than the following. We are concerned with utterance-types and their "timeless" meaning, that is, what one attributes to an utterance type x in answer to the question "What does x mean among members of G", abstracting from any particular occasion of the utterance of x. This is to be contrasted with the occasion meaning of an instance of x, and with the applied timeless meaning of an instance of x[2]. The former of these is, simply, what a speaker means by a particular utterance of x, in the fundamental sense of speaker's meaning; this may or may not be identical with one of the timeless

---

[1] "means" is, for now, to be read, "means qua simple signal". Subsequently this restriction will be unnecessary, for the final account will do equally for simple signals and for sentences.

[2] These notions I borrow wholly from Grice; see "Utterer's-Meaning, Sentence-Meaning, and Word-Meaning". pp. 225-227.

meanings of the utterance-type, according as the utterer is or is not, as it were, speaking literally. The applied timeless meaning of an instance of x is that one of its timeless meanings which the speaker intends it to have on the occasion; again this may or may not be identical with the occasion meaning of that instance of x. When someone utters something which has more than one timeless meaning, but utters it metaphorically, he will normally intend his real meaning to be apprehended <u>via</u> one of the timeless meanings of what he utters; in that case the utterance's occasion meaning and its applied timeless meaning are not identical.

Both utterance-types and utterance-instances may be called "ambiguous". In the former case, what is meant is that the utterance-type has more than one timeless meaning; in the latter case, what is meant is, perhaps, either that one cannot infer which is the applied timeless meaning of the instance, or which, among various candidates, is its occasion meaning. "Is that sentence ambiguous?" is not a determinate question, since "that sentence" may refer either to a type or to an instance.

There is a tendency among some writers on language to say such things as "The meaning of a natural language's sentences depends heavily on context." Such statements are best seen, at this point, as confusions. A sentence's meaning out of context—its properties as a member of the language, so to speak—is a matter quite distinct from both the occasion meaning and the applied timeless meaning of its utterances. And this tripartite distinction, which I here make in

terms of sentences, applies equally to simple signals.[3]

2. Schiffer on the Meaning of Simple Signals

The conventionality of meaning is the chief clue to its explication. We may take the truism that meaning is a matter of convention in a quite strong sense. It is not merely that meanings may be changed by the adoption of new conventions, but rather that, for a signal to have the meaning it has, there must exist certain conventions about it. To be more precise, I shall maintain that there is a certain kind of convention, indicated here schematically, which is such that:

x means p in G if, and only if, there exists in G the convention that...x...p.

Now if this is to be at all useful in achieving our reductionist goal, these two things must be the case:

a) the concept of convention must itself require for its explication no semantical notions (apart, perhaps, from that of speaker's meaning), and b) the specification of the content of the convention must itself require no semantical notions, apart from, of course, speaker's meaning. Moreover if the latter notion is to have the central position I have claimed for it, one would expect it to figure prominently in the specification of the appropriate kind of convention.

---

[3] I have spoken of timeless meaning as being of utterance types in order to emphasize the above sharp contrast. But one might, of course, speak of the timeless meaning of actual utterances and make the contrast nevertheless. Then, however, one must be careful actually to make the right distinctions, for example, among a) x means (timelessly) p, b) x means (here) p, and c) he meant by x that p. Ontological qualms about types may sometimes be stilled by distinguishing among the properties of their instances.

Schiffer, in Meaning[4], has offered us an account of the meaning of simple signals which comes quite close to explicating the sense in which meaning is conventional. However, his account does not entail that, if x means p in G, there exists a convention in G with regard to x. This is particularly unsatisfying since he also has given us a persuasive account of what conventions are. I shall now outline how Schiffer arrives at his account of the meaning of simple signals, and indicate its striking parallels with his subsequent account of convention.

The point of having signals with a conventional meaning is to facilitate communication. So it might seem that the timeless meaning of a signal consists in its being an effective communication device—if a member of G can effectively communicate p to most members of G by uttering x then, perhaps, x means p in G. Another, related, possibility may seem to be this: if there exists in G a practice, or regularity, of uttering x to mean that p, then x means p in G. Both these suggestions can be seen to be defective from the following considerations.

It is not a necessary condition of someone's meaning something by y that y itself have semantical properties. For y to be an effective communication device, all that is necessary is that y have some feature(s) from which the utterer's audience may infer that he uttered y with certain intentions, these being constitutive of his meaning such-and-such. As Schiffer puts it, y may have some natural feature—that is, one which is independent of anyone's ever having previously meant, or agreed to mean, anything by it, but which, nevertheless, makes it a suitable means

---

[4] pp. 174-202.

for communicating that p. A purring sound might, in certain circumstances, be an effective device for requesting milk—yet its effectiveness will not depend upon its having meaning. So, clearly, y's being, in G, an effective means for meaning that p is not sufficient for y's meaning p in G.

Similar considerations defeat the suggestion that y's meaning p in G consists in the practice, or regularity, of members of G uttering x to mean that p. There being such a practice is consistent with utterers of x intending their meaning to be inferred, in each case, from some appropriate natural feature, and not, for example, from the existence of the practice. We need a much stronger definition.

One would expect that, if x means p in G, the feature which makes x effective for communicating p "must not be one that is true of x independently of people's meaning something by x." (Schiffer).

If x means p, then, standardly, if someone utters x and means thereby that p, he will intend his audience to infer that <u>he</u> means that p at least partially from the fact that <u>x</u> means p.

An independent and equally important constraint on any adequate account is that, whatever constitutes x's meaning p in G, mutual knowledge of the relevant conditions among members of G will be necessary. Speaking roughly, if x means p in G, this will be mutual knowledge among the members of G. Nothing short of the requirement of mutual knowledge captures fully the element of <u>publicity</u> which is attendant on a signal's having meaning. Any lesser conditions will permit the construction of counterexamples not unlike those which make mutual

knowledge an essential feature of the account of speaker's meaning.

Amalgamating these requirements, we arrive at this intermediate definition:

> x means p in G if, and only if, it is mutual knowledge among the members of G that
>
> a) there is a precedent (i.e., a past practice) in G of uttering x to mean that p, and
>
> b) members of G are able to mean that p in G by virtue of its being mutually known in G that a) obtains.

But these conditions, in their turn, are neither necessary nor sufficient. That they are not sufficient may be seen from this case. There is a certain subgroup of G (a club, say, or a religious sect) in which x means p, while it fails to be the case that x means p in the principal group G. Clearly the above conditions may be satisfied consistently with its being false that x means p in G. But to stipulate that, for example, <u>most</u> members of G have uttered x to mean that p would be too strong, for that would rule out the perfectly possible case in which not most people have yet been in a position to mean p, but in which, nevertheless, we should wish to say that x means p. Schiffer suggests that what the definition lacks is the requirement that there exists, for all or most members of G, a certain conditional probability—namely that, if any member of G utters something and means thereby that p, then it is not unlikely that what he utters is x. Revising the definition to incorporate this condition, we get:

> x means p in G if, and only if, it is mutual knowledge among the

members of G that

a) there is a precedent (past practice) of members of G uttering x to mean that p, and

b) because of a) any member of G who means that p <u>might</u> do so by uttering x.

One final revision is needed, since the conditions just given are too strong. It may be that no one in G has ever uttered x to mean that p, and that, nevertheless, x means p by virtue of an agreement, or stipulation, in G to that effect (an agreement, perhaps, which there has been no occasion to implement.) So Schiffer's final definition is this:

x means p in G if, and only if, it is mutual knowledge among the members of G that

a) there is a precedent of members of G uttering x to mean that p, or an agreement among members of G that x may be uttered to mean that p, and

b) if any member of G utters x to mean that p, he will intend his hearer to infer his meaning, in part, from the fact that a) holds, and

c) should any member of G utter something and mean thereby that p, then what he utters might be x.

This account almost captures the <u>conventionality</u> of meaning, as may be seen from the following account of convention. Nevertheless, the definition fails to provide, or to entail, an actual convention which relates x and p, and which constitutes the former as meaning the latter.

3. Convention

A convention requires that a certain action be performed whenever, or only when, such and such circumstances obtain. So let us standardize reports of conventions in this fashion:

There is in G the convention of doing A whenever (only when) C.

Conventions are no more constituted by regularities than is the meaning of simple signals. The bare regularity lacks the essential element of deliberate concordance among the actions of the members of G. In short, it must be mutual knowledge among the members of G that the regularity obtains, and members of G must perform the action because they are expected to. As in the case of meaning, the general expectation of conformity which convention entails may be based either on a precedent (past practice), or on an agreement. So Schiffer offers the following definition of convention:[5]

There is in G the convention of doing A whenever (only when) C if, and only if, it is mutual knowledge among the members of G that

(1) there is a precedent in G of doing A (an agreement among members of G to do A) whenever (only when) C;

(2) on the basis (in part) of (1) almost everyone in G expects almost everyone in G to do A whenever (only when) C, and

(3) because of (2) almost everyone in G does A whenever (only when) C.

A signal's having meaning will, by Schiffer's definitions, have something striking in common with the existence of a convention. In

---

[5] Meaning, pp. 228-229.

both cases it will either be mutual knowledge among the members of the group that a certain practice obtains in the group (or would obtain if the circumstances were actualized) and that, when members perform the action they do so because it has been the practice; or it will be mutual knowledge that a certain agreement has been reached by members of the group, and that were members in situations appropriate for implementing the agreement, they would do so because of the agreement.

Regardless of their closeness, the definitions of convention and of x's meaning that p do not imply the existence of a convention which directly relates the utterance of x with meaning that p. The two most obvious candidates would be these: a) to utter something and to mean thereby that p only when what one utters is x; and b) to mean that p whenever one utters x. That the obtaining of neither kind of convention is necessary for x's meaning p is shown by (among other clear and conclusive considerations), in the case of a), the fact that two signals may have the same meaning, and, in the case of b), the fact that a signal may have more than one meaning.

It would be highly satisfying to have an explication of the <u>conventionality</u> of meaning, in terms of the existence of <u>conventions</u> which relate a signal and its meaning. Not only would this make clear the sense in which meaning is a matter of convention, but it would remove an impediment to a general identification of conventionality with the existence of conventions. There is also the straightforward consideration that, could we specify an adequate convention, we would thus understand better the relation between timeless meaning, and

occasion meaning. In what follows I shall propose a not too complex kind of convention which relates x and p, and whose obtaining may plausibly be taken as equivalent to x's meaning p.

4. The Convention Which Governs Meaning

Why do we have utterance types with conventionally determined meanings? In any act of communication, uptake or understanding is secured just in case the hearer has correctly inferred the utterer's intentions from his utterance. So achieving uptake is made surer to the extent that the inference from utterance to meaning is made easier and more certain. Consequently, if there are conventions which constitute utterance-types as being meaningful, one should expect that the existence of these conventions facilitates the inference.

Obviously, the kind of convention which would be maximally effective towards achieving this would be one which required members of the group to refrain from producing a certain signal unless they meant a certain specific proposition. That is, if it were the case that:

(C_1) There is in G the convention that whenever a member of G utters x, he means that p

effective communication of p would be secured, provided that x were easily produceable. If the convention obtained, x would mean p in G.

But, as we have seen, this convention's obtaining is not necessary for x's meaning p in G. It is too cautious and rigorous to be necessary for making x a maximally effective device for communicating p. The entrenchment of x as meaning p is not threatened by the fact that

speakers may often produce x without meaning p in circumstances in which it is clear to the audience that the speaker does not mean that p. There are, it would seem, easily recognized circumstances in which one is not to be taken as meaning what one's utterance timelessly means.

A closely related point is this. Signals may have more than one meaning—which our rigorous convention would make impossible. Although ambiguity is a defect in a technical language, whose sentences may often be uttered against a background which fails to disambiguate them, it is no great defect in a signal which would normally be uttered in circumstances which disambiguate it.

The two points are connected—the fact that a signal which means that p may be uttered without meaning that p (in violation of no convention which is required for keeping x as timelessly meaning p) is precisely what explains the possibility of x's having more than one meaning. Disambiguation will occur when circumstances make it clear that the one thing is not meant, and that the other is meant.

On the other hand, x's timelessly meaning p facilitates the communication of p only if in certain circumstances the hearer can depend on the inference: x was uttered, so p was meant. If such circumstances could be generally described (let's call them "K"), we could then suppose that x means p in G just in case

(C_2) There is in G the convention that, whenever any member of G utters x in circumstances K, he means that p.

When x means p in G, a member of G may utter x in the presence of another member of G without fear of being taken to mean that p, when (a)

he knows that the hearer will prefer some other explanation of his utterance—for example, that he is joking, speaking hyperbolically, implying something else, uttering x with one of its other meanings, singing a song, acting, exercising his vocal chords, or testing a microphone, or (b) he knows that the hearer would not, in the circumstances, take any utterance of his to mean that p—because, for example, p is extremely improbable, known to be false, or not within the ken of the speaker.

Clearly, such lists of circumstances will be of no use. But there is a pattern to them, which, if we can capture it, will give us our convention. For then the convention will require one to refrain from uttering x in circumstances which are not of the above kind unless one means that p.

Let us say that:

> the hypothesis H overrides the hypothesis H' as an explanation of an action A for a set of persons P
>
> if, and only if,
>
> members of P believe that
>
> i) it is not unlikely that H is true;
>
> ii) H is a substantially better explanation of A than is H'; and
>
> iii) either H and H' are incompatible in the circumstances, or the truth of H would make the supposition of H' otiose.

What the various circumstances listed above in a) have in common, then, is this: they are all cases in which there is an hypothesis

which overrides (for the audience), as an explanation of the utterance of x, the hypothesis that its utterer meant thereby that p.

Notice that we could now get a much more realistic and less rigorous convention, which would still provide a sufficient condition for x's meaning p—namely,

(C_3) There is in G the convention that whenever a member of G utters x in the presence of an audience for whom no explanation of his utterance overrides the hypothesis that he means thereby that p, then he means that p.

I must emphasize that this convention is sharply to be distinguished from: "whenever a member of G utters x in the presence of an audience for whom the best explanation of his utterance is that he means that p, then he means that p." That there is no hypothesis which overrides H as an explanation of A clearly does not imply that H is itself a good explanation of A.

That C_3 is not trivial may be seen from this consideration. Characteristically a group will have in its repertory of utterance-types only, say, sounds or gestures or inscription-types which have no other function in their activities. If a particular such utterance type had lacked the meaning it in fact has, any utterance of it would have been one which had no obvious explanation (in terms of its utterer's reasons); hence there would have been no hypothesis which overrode the hypothesis that, e.g., p was meant. The introduction of a convention like C_3 would then have substantial consequences; for one's audience could now infer from one's <u>otherwise unexplained</u> utterance of x that

one means that p.

Interestingly enough, some kinds of explanation which may override the hypothesis that p is meant are possible only because the convention exists. I have in mind here joking, hyperbole, acting, and other such types of utterance as are parasitical upon the existence of literal meanings.

But, most importantly, such a convention accommodates a signal's having more than one meaning. Suppose that x means p, and also means q. If one then utters x and means that q, one will expect the circumstances to be such that the hearer will take the hypothesis that one meant that q to override the hypothesis that one meant that p. So there is no conflict in there being both a convention of this kind which relates x and p and another which relates x and q. Notice that, if one utters x ambiguously (e.g., in circumstances in which the hearer cannot tell which of x's timeless meanings is its occasion meaning), then one has violated the conventions. For, an utterance is ambiguous between p and q if a) no hypothesis overrides the hypotheses that p is meant or that q is meant, and if b) neither overrides the other, and if c) it cannot be that both are meant. It is not an unsatisfying result that ambiguous utterances violate the conventions which govern meaning.

C_3 is still not a necessary condition of x's meaning p in G, for it remains too restrictive. It is not only the availability of an alternative explanation which permits one to utter something which means p without fear of being taken as meaning that p. Certain

utterances will not, in certain circumstances, be taken seriously because of the improbability that the speaker should wish to communicate that proposition by any means in the circumstances. Clearly this may be so even in the absence of an alternative explanation, and one who produces such a non-serious utterance violates no convention.

It may appear that the last condition and the condition of C_3 do not exhaust the cases in which one may utter a signal which means p, without meaning that p. Consider this case: I may in normal current circumstances utter the sentence "Mao has married Svetlana" and, regardless of whether my hearer has available some alternative explanation, he will not take me to mean that Mao has married Svetlana. But it is not true that I would not be taken to mean that proposition by any utterance in the circumstances. For, clearly, I might produce a small discourse, with a suitable build-up and an urgent or ponderous tone, and then, uttering that sentence, be taken literally.

What the example shows though is that 'in the circumstances' must be carefully construed. For it seems perfectly clear that in the two cases the circumstances are not relevantly the same as regards the question of whether the speaker means that Mao has married Svetlana. For, the build-up provides a different setting for the sentence in the second case.

But there are possible cases which differ only in this: one utterance is produced in an urgent tone, while the other is not, where the tone makes the difference between the one utterance's being literal and the other's not. How might this be represented as being a

difference in the circumstances? It seems to me that a good account of the function of tone in such cases is this: by one's tone one sometimes intends to convey a kind of higher order proposition—something like: '(contrary to expectations) my utterance is to be taken literally.' So we can represent the tonal difference as a difference in circumstances by supposing that a minimal condition on the circumstances in two cases being relevantly the same as regards the question of whether p is meant, is that nothing logically independent of p be meant in one case, and not in the other. Clearly, further conditions will be in order.

It will useful to have an abbreviated phrase which captures this condition together with the condition of C_3. So let us say that:

x's utterance of z is free, in circumstances C, for meaning
that p, audience y,
if, and only if,
i) x and y mutually know that x produced z;
ii) there exists for y no hypothesis which overrides as an explanation of x's utterance the hypothesis that x means by z that p, audience y;
and iii) it is not, for y, so improbable that (w) if x uttered w in C, x would mean thereby that p, audience y, that y could not take any utterance of x's in C as meaning that p, audience y.

On this basis, we may now define what it is for a signal to mean timelessly a certain proposition in G.

x means p in G

if, and only if,

(C_4) there is in G the convention that (x)(y) if x and y are members of G, and x utters z in circumstances in which x's utterance is free for meaning that p, audience y, then x, by his utterance of z, means that p, audience y.

CHAPTER II

REFERENCE AND PROPOSITIONAL ATTITUDES

Introduction

The account I have just given of the meaning of simple signal fails, in a crucial respect, to be fully general. The definiendum was 'x means p', which asserts a relation between a signal and a proposition, but there are simple signals, just as there are sentences, whose meaning cannot be identified with any proposition. What proposition, for instance, does the sentence 'He arrived on the last train', express? Its meaning is better identified with a propositional function than with a proposition. Similarly, the meaning of the white flag of surrender is something like "We surrender", and is, therefore, not to be identified with a proposition.

This complicates matters considerably. In my definition of the meaning of a signal, I used the propositional attitude 'to mean that'. Now, as everyone knows, there are serious difficulties with quantifying into propositional attitude contexts. But that is just what we shall have to do, if we are to generalize the account of the meaning of simple signals to cover those signals which lack propositional meaning. Consider, for example, a signal which means "I surrender". What is meant by the production of the signal is that the producer of the signal surrenders; clearly, in expressing this, we have quantified into the 'that'-clause.

So in this chapter I propose an account of the logical form of

propositional attitude contexts which represents the referentiality of singular terms within them, and makes sense of quantification into them. I take as a paradigm of the class of propositional attitude verbs, the verb "believes".

1. A Theory of Quine's

Frege held that names and other singular terms do not have their normal reference in "oblique" contexts. Since Frege, weighty arguments have been advanced and accepted, and the view is well entrenched. But it is a most implausible claim, not easy to reconcile with certain obvious and ordinary facts. Quine[1] has come as close as anyone to recognizing this in print, but even he appears to accept a variant of Frege's theory. There is, though, in Quine's qualified theory, a foothold from which one might arrive at a more acceptable view. So I shall begin by describing his theory.

Suppose that

1) Ralph believes that a certain person is a spy.

This clearly supposes rather more than that Ralph believes that spies exist. It appears to be a quantification into an oblique context, and hence by the usual conventions of mixing quantificational notation and English it should be equivalent to

2) (Ex) Ralph believes that x is a spy.

---

[1] "Quantifiers and Propositional Attitudes".

Quine points out that, despite its initial appeal, the significance of this sentence is doubtful. Here is a story he tells.

> There is a certain man in a brown hat whom Ralph has glimpsed several times under questionable circumstances on which we need not enter here; suffice it to say that Ralph suspects he is a spy. Also there is a gray-haired man, vaguely known to Ralph as rather a pillar of the community, whom Ralph is not aware of having seen except once at the beach. Now Ralph does not know it, but the men are one and the same.[2]

Now let us suppose that Ralph, if required to give an opinion, would sincerely deny that the man seen at the beach is a spy. We then have a case in which the man in the brown hat is the man seen at the beach, and in which it is true that

3) Ralph believes that the man in the brown hat is a spy

but in which it is false that

4) Ralph believes that the man seen at the beach is a spy.

Quine draws from this failure of implication the conclusion that 3) does not assert a relation between Ralph and some man. His assumption appears to be that if a relation were thus truly asserted the same relation would be truly assertable between Ralph and the man under any description. As a result, 2) is held by Quine to make no sense, for it clearly requires that Ralph stand to someone in the putative relation. Since the point about 3) is quite general, no relation is expressed by the context

5) x believes that...y...

---

[2] Ibid., p. 185

If Quine is right, the man whose spyhood is in question 'does not receive reference' in 3) 'because of referential opacity'. Thus Frege's theory is upheld. To accommodate this alleged fact, Quine proposes that we take 3) as asserting a relation between Ralph and a proposition. So we have:

6) B* (Ralph, [the man in the brown hat is a spy]).[3]

The irreferentiality of singular terms in belief contexts is thus to be explained by their occurring in designation of their sense.

But how are we then to construe the acceptable sentence 1)? Quine's suggestion is that, just as Ralph may stand in the required relation to a proposition, he may similarly be related to an intension (propositional function) and some item of which, as it were, he believes the intension to be true. Thus it may be true that

7) B (Ralph, a[a is a spy], the man in the brown hat).

The problematic term has here a normal referential occurrence, and together with the identity 7) will imply

8) B (Ralph, a a[is a spy], the man seen at the beach).

Both of these are to be taken as true in Quine's story, though of course Ralph will not know that his beliefs thus relate him to the man seen at the beach. (One may read 8) as "Ralph believes of the man seen at the beach that he is a spy.") Using this device

---

[3] Quine first fancifully presents his theory in terms of intensions, then reveals his preference for sentences and open sentences.

The notation is not Quine's. 'B' may be read "believes-of", and B* believes. The notation for intensions (propositions and propositional functions) is this. A sentence in square brackets denotes the proposition it expresses. An n-adic open sentence in square brackets with quasi-operator-variables prefixed denotes the relevant n-ary propositional function.

we may now represent the structure of 1), for we may generalize 7) to get:

9) (Ey) B (Ralph, a[a is a spy], y).

This, then, is Quine's version of 1).

An attractive feature of the analysis is that it enables us to represent concisely a certain not uncommon way of being mistaken. In the case described, it may happen that

10) B (Ralph, a[a is not a spy], the man seen at the beach).

The joint truth of 8) and 10) convicts Ralph, not of irrationality, but merely of fallibility. For it does not follow that

11) B (Ralph, a[a is a spy and a is not a spy], the man seen at the beach).

To make it plausible that this is a non-sequitur and thus to defend both Ralph and Quine, I suggest the following. If 8) does what we intuitively require, it will ascribe to Ralph that state of mind he would be in if he were sincerely to assert, referring to the man in the brown hat (who is the man seen at the beach), "He is a spy." 10), in parallel fashion, will ascribe to Ralph that state of mind required for sincerely asserting "He is no spy", referring to the same man. Even though Ralph has revised no judgments, he is not prepared sincerely to assert of that man, "He is a spy and yet no spy." This is not to suggest that belief has a satisfying analysis in terms of sincere assertion, but merely that the connection between the two makes it possible to point up features of this crucial notion "belief-of".

Before we proceed, there is a complication which needs to be dealt

with. As we have treated it so far, 4) is false in Quine's story; but another way of reading it makes it true, namely as meaning something like "Ralph believes of the man seen at the beach that he is a spy." (We are at the beach and I say 'Ralph thinks that that fellow is a spy', proceeding with the story.) This is, in itself, no problem for Quine, since 8) captures it. But 3) will then also have two readings, a strong interpretation (the one we have given it so far) and a weak interpretation, on which it is rendered by 7). This characterization ('strong' and 'weak') is given on the assumption that the former implies the latter, but not conversely; that is, if Ralph believes that the man in the brown hat is a spy, then Ralph believes of him that he is a spy. The result of this assumption is that 3), on its strong interpretation, together with the identity, implies 4) on its weak interpretation. The upshot is that there is a reading of 4) on which, given the identity, it follows from both interpretations of 3). Although this requires a qualification of the claim that substitution of the co-referential term in 3) fails to preserve truth, the point remains that if the result of the substitution (viz., 4) is taken on its strong interpretation truth is not preserved; indeed this failure to be implied is the mark of the strong interpretation.

My assumption that the strong interpretation implies the weak has, in effect, been maintained by Quine, for he holds with regard to a pair of sentences parallel to 3) and 7) that "the kind of exportation which leads from ⌜3⌝ to ⌜7⌝ should doubtless be viewed in general as implicative."[4] The device of considering would-be sincere assertions may make

---

[4]Ibid., p. 188

this clear. If 3) is true, Ralph is in that state he would be in if, referring to the man in the brown hat, he were sincerely to assert 'The man in the brown hat is a spy.' In that case, he would be in a state appropriate to the sincere assertion of 'He is a spy', referring to that man. The latter state, though, is what is reported by 7).

The obtaining of this implication, however, leads to serious difficulties with Quine's analysis, as follows. Recall that Quine holds that 3) (on our strong interpretation) does not assert a relation between Ralph and some man; in fact the adequacy of 6) as a final analysis depends on it. But this is hardly compatible with 3)'s implying 7), for the latter asserts some such relation. So, if the implication holds, it would be natural to suppose that 'the man in the brown hat' occurs with its normal reference in 3), on its <u>strong</u> interpreration, since a statement of its truth conditions would introduce, to speak roughly, the man in the brown hat himself, and not merely the sense of 'the man in the brown hat.'

But how is this possible? If the relevant singular term occurs with its normal reference on the strong interpretation of 3), the same should hold of 4). But then if 4), on its strong interpretation, asserted a relation between Ralph and some man (in consequence of the man's being referred to), it would seem that 4), on that interpretation, should be true if that relation were truly asserted between Ralph and that man under any other description of him. But that would be so only if 4), on its strong interpretation, followed from 3) with the identity, and that, as we have seen, fails to be so. So it would

appear that certain natural assumptions lead us to a contradiction.

2. Referential but Non-Extensional Occurrences of Singular Terms

I shall try to strengthen one of our incompatible assumptions, as a preliminary to an attack on the other. Let us say that, on the weak interpretation of sentences like 3) and 4), the relevant singular term occurs extensionally, and that, on the strong interpretation, it occurs non-extensionally. Now there are independent reasons for supposing that singular terms may occur in belief contexts non-extensionally, and yet with their normal reference.

There is a tendency in some quarters to ignore pronouns and demonstratives when considering the behavior of singular terms. This neglect, I think, is partially responsible for Frege's and Quine's view of singular terms in belief contexts.

Suppose that we are at a costume ball, and I say, pointing to a certain man

12) Michael believes that that masked man is a diplomat.

Now in some circumstances it would be perfectly natural to take this on the strong interpretation. That is, 12) may be asserted not merely on such grounds as that the masked man is Ambassador Brown and that Michael believes that Ambassador Brown is a diplomat. But even on this non-extensional occurrence of the term, it seems plausible that I (and my words) referred to the masked man, and that, if you like, I asserted a certain relation to hold between Michael and him. I offer the following considerations in support of this.

i) There seems, in this case, to be that need for specification from the context which is characteristic of normally referential occurrences of demonstrative phrases. Thus the speaker may be asked 'To which masked man are you referring?' with the same appropriateness as if he had said 'That masked man is a diplomat.'

ii) It seems not far-fetched that, in uttering this sentence, I asserted a certain relation to hold between Michael and the masked man. For, in showing that what I said is false, one might quite aptly point out that certain relations fail to hold between the two, or that certain other relations do hold. Thus:

a) Michael wasn't in his vicinity all evening. Perhaps you mean someone else.

b) Michael was heard to say, pointing at the fellow, 'He's obviously a gate-crasher.'

c) The two were seen deep in conversation and the masked man as we know, speaks ungrammatically.

d) The masked man is Congressman Smith, whom Michael would recognize anywhere.

iii) Consider the inference:

Whoever Michael believes is a diplomat is a diplomat.

Michael believes that that masked man is a diplomat.

So, that masked man is a diplomat.

My intuitions have it that this is a perfectly acceptable inference. The most natural way of construing its validity is to take it as having the form '$(x)(Fx \rightarrow Gx)$; Fa; so Ga'.

Of course there may be more complex ways of representing its validity while denying the relevant singular term in the minor premiss a normal reference. But on the obvious way of doing it the term has to be treated as having a normal reference.

iv) The analysis of certain complex sentences would be simplified if we could suppose such terms to have a normal reference. For example,

13) Michael thinks that that masked man is a diplomat, but he obviously is not.

A very plausible thing to say about the role of the pronoun 'he' is that it anaphorically picks up the reference of 'that masked man'.

These considerations do not prove the point, but they suggest, I think, that our theories would be simpler and more realistic if we could allow a normal reference to some singular terms which occur non-extensionally. But to do so requires the denial of a principle which is so deeply entrenched that it would seem eccentric to deny it. It is called by some "Leibniz's Law", but I shall refer to it as "the substitution principle." It may be formulated like this:

> If S' arises from S by substitution of a singular term N for one or more occurrences of a singular term N in S, and if N and N' have the same reference in S and S', then S' is true iff S is true.

The application of the principle to Quine's case is this. If, on the strong interpretation of 4), the term 'the man seen at the beach' occurred with its normal reference it would have the same reference as the corresponding term in 3). Hence 4) should be true if 3) is. But that is not so, hence the term does not occur with its normal reference in 4).

I shall, in the next section, show how the substitution principle might be false, and then, in the following section, how to construe non-extensional but referential occurrences of singular terms in belief contexts.

3. Reference and Satisfaction Conditions

Let us consider some general facts about the truth conditions of sentences. Utterances of sentences containing genuine singular terms are true or false depending on whether the items designated by the singular terms (on the occasion of utterance) satisfy or fail to satisfy certain conditions associated with, or expressed by, the sentence. Each such sentence, in other words, has satisfaction conditions such that its utterance is true if, and only if, the appropriately ordered n-tuple of its references satisfies those conditions. So, for example, an utterance of "She shot the armadillo" is true if, and only if, the indicated female shot the armadillo.

By a sentence's frame I shall mean the result of deleting its singular terms, with small letters inserted to mark the positions of deletion. So, the frame of 'She shot the armadillo' is 'x shot y'. Now on the basis of the above fact about satisfaction conditions, it is usually assumed that if two sentences have the same frame they have the same satisfaction conditions. For normal well-behaved sentences this does seem to be true: for example, for all sentences whose frame is 'x shot y'. An utterance of any such sentence will be true just in case the item designated by the term in the 'x'-position shot the

item designated by the term in the 'y'-position (shelving irrelevant complexities introduced by the possible ambiguity of 'shot').

Now it is not difficult to see how this general assumption might be false, and, most importantly, false in cases which make the substitution principle false. That is, it is a perfectly coherent possibility that there are pairs of sentences with the same frame (where the respective singular terms are co-referential) which do not have the same satisfaction conditions. For suppose that singular terms in certain sentences had a dual role of the following vaguely described kind. Not only do they designate, or refer to, items which must satisfy the sentence's satisfaction conditions if the sentence is to be true, but they also contribute to the very conditions which must be satisfied. In other words, still keeping in mind the distinction between the references of an utterance and the satisfaction conditions of the sentence uttered, it may be the case that there are sentences whose frames do not alone determine their satisfaction conditions, but do so only in conjunction with what singular terms occur. It would, in that case, be possible for two sentences with the same frame, and the same set of references, to differ in satisfaction conditions and, hence, to differ in truth value.

Unfortunately there are few clear examples in English apart from those psychological contexts I wish ultimately to account for. But Quine, in an earlier paper, called attention to a type of sentence which will serve nicely as an example of my generally described possibility. It is true that

14) Giorgione was so called because of his size.

Despite the fact that Giorgione is clearly here referred to, truth is not preserved by the substitution of another name of him, e.g., 'Barbarelli'; for it is false that

15) Barbarelli was so called because of his size.

The falsity of 15) is due to the fact that Barbarelli was not called 'Barbarelli' because of his size, though he was, of course, called 'Giorgione' because of his size. This, at least, is independent of how we refer to him. So the satisfaction conditions of 14) and 15) do not depend uniquely on their frame ('x was so called because of his size'), but also depend on what singular term occurs. The man Giorgione-Barbarelli satisfies the conditions of 14), and fails to satisfy the conditions of 15). Hence the term 'Giorgione' in 14) is to be taken as making a two-fold contribution to the sentence; on the one hand, it has its expected role in locating the reference of the utterance, and, on the other hand, it contributes to the satisfaction conditions which the reference of the utterance must satisfy if 14) is to be true. The substitution of co-referential terms, then, may sometimes fail to preserve truth, for such substitution may give the resulting sentence satisfaction conditions which differ from those of the original sentence. Hence utterances of the two sentences may have different truth conditions; in our examples, we may say that different properties are in the two cases predicated of the same person, despite the identity of those sentences' frames.

There is a Leibnizian principle salvageable from this—a principle

which is fully general, appropriately trivial, and which does not pretend to be an a priori constraint on the behavior of singular terms. It is this.

If the sequence $\langle x_1, \ldots, x_i, \ldots, x_n \rangle$ satisfies a sentence S, and if $x_i = y$, then $\langle x_1, \ldots, y, \ldots, x_n \rangle$ satisfies S.

4. An Account of the Dual Role of Some Singular Terms

Recognizing the vague possibility that singular terms may make what I called a "dual contribution" to the truth-conditions of sentences does not of itself solve any problems raised by singular terms in belief contexts. What is needed is some specification of that additional contribution singular terms make to belief sentences which accounts for the non-extensionality of their occurrence.

But first it may be useful to introduce a refinement into the straightforward analysis of extensional occurrences of singular terms in belief contexts. The relation B holds among a believer, an intension and an item of which the believer believes the intension to be true. There is, in this relation, a crucial lack of generality. How, for example, are we to account for the extensional reading of

16) William believes that Cynthia will marry Herbert?

If we render it as

17) B(William, ab[a will marry b], Cynthia, Herbert) we have a new 'B', one which should express a tetradic relation. But we can not introduce a new, ad hoc, relation of different degree at each level,

since there would then be no effective procedure for determining the semantical contribution of the verb 'to believe'.

To meet this difficulty, I propose the simple maneuver of taking the 'belief-of' relation to hold among persons, n-adic intensions (properties and relations, or, if you like, propositional functions), and n-ary sequences of items. Thus 7), for example, becomes

7') B(Ralph, a[a is a spy], ⟨the man in the brown hat⟩),[2]

and 17) becomes

17') B(William, ab[a will marry b], ⟨Cynthia, Herbert⟩).

Let us consider 12) again. On the interpretation which gives the phrase an extensional occurrence, Michael and the masked man are asserted to stand in a relation which we may represent as

18) B(y, a[a is a diplomat], z).

I maintained that on the interpretation of 12) which gives the term a non-extensional occurrence, Michael and the masked man are asserted to stand in a certain relation. But this relation cannot be represented adequately by 16), since something stronger is being asserted, something whose truth depends upon Michael's identifying the man as such and such and believing of him under that description (or, as I shall say henceforth, under that property) that he is a diplomat.

Now, to identify something as such and such in the sense required is merely correctly to believe of it that it is such and such. So,

---

[2] Henceforth I shall omit brackets for one-membered sequences.

for example, the identification involved in the non-extensional interpretation of 17) is:

19) B(Michael, a[a is that masked man], that masked man).

I defer the difficult question of what is packed into <u>being that masked man</u> until the next section; let us say for now that it involves just those properties necessary for Michael's making the implied identification.

But the conjunction of 19) and

20) B(Michael, a[a is a diplomat], that masked man).

is not sufficient to capture the non-extensional sense of 12), since 20) might be true by virtue of Michael's believing it of the fellow under another property. What is required is that Michael realize that one and the same person is that masked man and a diplomat, or, better, that he believe <u>of</u> that masked man that he is that masked man and a diplomat. I suggest, then, that to believe such and such of something under a certain description is to have a certain <u>conjunctive belief</u> with regard to it (rather than merely a conjunction of beliefs).

12), then, asserts the following relation between Michael and the masked man:

21) B(y, a[a is that masked man & a is a diplomat], z)

and its full rendition is

22) B(Michael, a[a is that masked man & a is a diplomat], that masked man).

So the phrase 'that masked man' makes two distinct contributions to the truth-conditions of 12) on its non-extensional interpretation. In one of these, the expression is referential in the normal way. In the other, the expression partially determines what conjunctive property is the middle

term of the triadic relation asserted by 12); or, looked at in a more illuminating way, it contributes to the satisfaction conditions of 12), with the belief-relation and the relevant intension generating a complex relation between Michael and the masked man.

We can now see why 12), in its strong sense, fails to follow from, say,

23) Michael believes that Ambassador Brown is a diplomat

together with the relevant identity. For this will have these two interpretations:

24) B(Michael, a[a is a diplomat], Ambassador Brown),

and

25) B(Michael, a[a is Ambassador Brown & a is a diplomat], Ambassador Brown).

The substitution of co-referential terms which leads from 23) to 12) clearly may change the satisfaction conditions of the sentence. For quite different relations may be asserted in 12) and 23), as may be seen by comparing 22) with 24) and with 25).

The path to this result passed through the observation that the inference from 3) to 7) appears to be valid. Can the validity of this inference now be represented? 3), on a non-extensional interpretation, becomes

26) B(Ralph, a[a is the man in the brown hat & a is a spy], the man in the brown hat).

Does it follow that

7') B(Ralph, a[a is a spy], the man in the brown hat)?

It is extremely implausible that there is a generally valid inference from 'x believes that p' and 'p entails q' to 'x believes that q'. Nevertheless, it does seem that in the specific case of

'conjunction elimination' (so to speak) we do accept such inferences from beliefs to beliefs. The only kind of alleged counterexample which I know of fails to be one. It is this. Someone may believe a conjunction but also believe that the truth of one of the conjuncts is a necessary condition of the truth of the other. But what this shows is not that there is no valid inference from 'x believes that p and q' to 'x believes that q', but rather that there is no valid inference from 'x believes that p and q' to 'x would have believed that q even had he not believed that p', which is obvious and irrelevant. So I suggest that we take the validity of the inference from 26) to 7) as imposing a condition on any adequate account of the relation expressed by "B".

If the inference holds, 26) implies

9) (Ey) B(Ralph, a[a is a spy], y).

That is, 3), even on its non-extensional interpretation, implies that Ralph believes that a certain person is a spy.

The trouble with the peculiar sentence

2) (Ex) (Ralph believes that x is a spy)

should now be clear. A term in the position of the variable may have more than a referential function. Consequently, there will be no single relation asserted by all sentences which have as their frame 'y believes that x is a spy'. One needs, or may need, more than the reference of the term in the 'x'-position to determine what is being asserted.

On the strong interpretation of 3), there are two things which may be called its existential generalization with regard to the term 'the man in the brown hat', namely 9), and the stronger:

27) (Ey) B(Ralph, a⌈a is the man in the brown hat & a is a spy⌉, y).

If what I have said about 'conjunction elimination' is right, then, agreeably enough, 9) follows from 27). That is, if there is something of which x believes such and such under a certain description, then there is something of which x believes such and such.

5. Uniqueness Properties and Self-Ascriptive Attitudes

There is a serious gap in this account. I have pretended that there is such a thing as the non-extensional interpretation of 12), as if, in the non-referential contribution it makes the phrase 'that masked man' determines, simply by its meaning out of context, a certain property a⌈a is that masked man⌉. But the only property which is determined by the sense of the words 'that masked man' is the property a⌈a is a masked man⌉. That this cannot be substituted for 'a⌈a is that masked man...⌉' in 22) with a result which adequately renders 12) may be seen from the following case.

Suppose that the masked man is Armbruster, that Michael has currently seen Armbruster in a mask, and that Michael believes Armbruster to be a diplomat. It is then true that

28) B(Michael, a⌈a is a masked man & a is a diplomat⌉, Armbruster),

and therfore, that

29) B(Michael, a⌈a is a masked man & a is a diplomat⌉, that masked man).

Nevertheless, 12) may be uttered with the right references, and be false. It may, for example, be uttered with the implication that Michael believes it of the masked man as he is to be seen there dancing with a bear

('<u>that</u> masked man'). Michael may, in these circumstances, fail to recognize the normally dignified Armbruster, even though 29) is true.

On a normal non-extensional utterance of 12), what is meant is that Michael believes of the masked man that he is a diplomat under a certain property which individuates the man for Michael as <u>that</u> masked man. Let us say that F is a uniqueness property, just in case there is a G such that F is the property of being the only G, (i.e., $F = a[a = (\eta y)(Gy)]$.) For any non-extensional utterance of 12), then, there is a certain uniqueness property—namely, $a[a = (\eta y)$ (y is a G masked man)]—such that what is meant by that utterance of 12) is

30) B(Michael, $a[a = (\eta y)$(y is a G masked man) & a is a diplomat], that masked man).

Being <u>that</u> masked man, then, is being a masked man who has certain further properties which conjointly individuate him in the required way. 12) may be regarded as containing a kind of second-order indexical, where the relevant property is determined by the utterance. Thus the property $a[a$ is that masked man] of the previous section is whatever property of the form $a[a = (\eta y)$(y is a G masked man)] is meant by the utterance of 12).

This casual invoking of uniqueness properties will not satisfy all, nor should it.

It is a philosophical commonplace that our non-relational conceptions of objects do not always individuate them for us, and, moreover, that at least sometimes one's individuating an object requires relating it to

some item which has a special relation to oneself, for example, this time and place, this sensation, me.

Thus, imagine Jenkins' being in stereotyped surroundings, unaware of his location, and unaware of the time. There is before him a jug which, he has discovered, contains water. So it is true that

31) Jenkins believes that the jug contains water.

Clearly, 31) may be true in a non-extensional sense, but under what kind of uniqueness property might Jenkins believe what he does of the jug? There is nothing about it which he regards as unique, except its being currently related in such and such a way (spatially, perceptually) to him.

Let us term "self-ascriptive" any belief a person has about himself—whether it be the belief that he possesses a certain property, or the belief of some items that they stand in a certain relation to him. How shall we give the form of self-ascriptive beliefs, using the apparatus of intensions and relational belief?

What is it for a man to believe himself to be F? It may seem that there is a simple answer to this—namely: x believes himself to F if and only if

32) B(x,F,x).

That 32) fails to provide a sufficient condition for a belief's being self-ascriptive may be seen from this case. Jones looks through a mirror-device (a kind of periscope), and sees a man looking through a mirror-device. Though it is himself he is seeing (there is no need to go into details here), he believes it to be someone else, and, furthermore,

believes that the man is a spy. So we have:

33) B(Jones, a⌜a is a spy⌝, the man seen through the periscope).

and, therefore,

34) B(Jones, a⌜a is a spy⌝, Jones).

It is, nevertheless, false that Jones believes himself to be a spy.

What is immediately suggested is this. A person x's self-ascriptive belief is not only x's belief of x, but x's belief of x under a property which, as it were, identifies x to himself as himself. More precisely (the suggestion is), for each person x there is a property G such that x believes himself to be F if and only if

35) B(x, a⌜a = (ιz)(Gz) & Fa⌝,x).

But this fails to provide a necessary condition of having a self-ascriptive belief; it also fails to provide a non-circular sufficient condition. Let us take the latter point first.

We want a general account of x's believing himself to be F. Clearly not <u>any</u> proposition of the form 35) will be sufficient for that; the same cases which show 32) to be insufficient show that arbitrary instances of 35) are insufficient. Rather, it is suggested that for each person there is <u>some</u> particular G which makes the truth of 35) sufficient for x's believing himself to be F. But this provides no <u>analysis</u> of self-ascription, unless we can specify a condition which makes a particular property x's own self-ascriptive property (so to speak). It is extremely doubtful that we can specify a completely general condition other than: x believes G to be true of himself. For whatever G is, we might imagine bizarre cases in which x fails to realize

that it applies to himself, though he believes it to be true of some person who is in fact identical with him.

But, most importantly, it is not a necessary condition of x's believing himself to be F that something of the form of 35) be true. For, a person may lose enough of his memories, and be in a sufficiently undistinguished environment for there to be no uniqueness property which he believes to be true of himself, and yet still have beliefs about himself. He may, of course, believe himself to be the person who is R to that (some object of his perceptions). However, this can provide him with a unique self-characterization only on the assumption that he has available some uniqueness property which applies to it without relating it to himself, which is contra hypothesem.

It may be thought that, in the case of one's sensations, one may be aware of them "directly" and not under some uniqueness property—that if x believes F of a sensation and it is one of his sensations, then (necessarily) x believes that his own (such and such) sensation is F. But this is false. For example, it is not impossible that sensations are neural events, and hence, in a way, public, in which case a man may have a belief of a sensation of his without believing it to be his own.

The problem of self-ascription, then, is this. It is not sufficient for x's believing himself (his own sensation) to be F, merely that he believes of x (x's sensation) that he (it) is F. For x may have such a belief of x (of x's sensation) under a property which fails to identify it to him as himself (as his own). But we cannot get an account of self-ascription by requiring that x have the belief of himself under

such and such property, for there will be cases in which there is no uniqueness property which identifies x for himself, though clearly in such cases x may have beliefs about himself.

Let us return to the case which 31) describes. Let us suppose that the only uniqueness property available to Jones (a propos the jug) is: being the only jug which is R to him. Just as we failed to capture simple self-ascription with 32), so we cannot capture Jenkins' belief in:

36) B(Jenkins, ab[b = (ιz)(Rza & z is a jug) & b contains water], ⟨Jenkins, the jug⟩).

Cases analogous to the case which shows the insufficiency of 32), will show that 36) fails to capture the self-ascriptive character of the property under which Jenkins believes what he does of the jug.

If we are to have a fully general account of the logical form of belief sentences, we must be able to represent self-ascriptive beliefs, both for their own sake, and because an important class of beliefs which do not seem at first glance to be self-ascriptive are, in fact, indirectly so. I do not know how to represent such beliefs using the triadic 'B' and the apparatus of intensions. Nevertheless, we can capture their logical form at a less than completely analytical level, but still in a manner consistent with my current purpose—namely, getting an account of reference and quantification in belief contexts.

If a person believes himself to be F, he stands in a certain dyadic relation to F. If a person believes himself to be related by the n-adic

relation R to the sequence of items $\langle x_1,...x_{n-1}\rangle$, he stands in a certain triadic relation to that n-adic relation and that (n-1)-ary sequence. Thus, if Jones believes of x and y that he, Jones, is better at badminton than x is, but worse than y is, then he stands in (what we might call) the self-ascription relation B^+ to the triadic relation (intension), abc[a is better at badminton than b & a is worse at badminton than c], and the binary sequence $\langle x,y\rangle$. The notational device makes the first slot (the 'a' slot) in the relational expression the one which is, as it were, applied to the self-ascriber.

In general,

37) $B^+(y,a_1...a_n\lceil Ra_1...a_n\rceil, \langle z_1...z_{n-1}\rangle)$

gives the logical form of

38) y believes of $z_1,...z_{n-1}$ that he is R to them.

One further modification is needed. The uniqueness property under which one has a belief of something may be one which relates the thing to oneself spatio-temporally. Now whether the relevant temporal relation is one of simultaneousness with the belief, or priority to it, we must suppose that there is some reference to (as it were) the time at which the belief obtains. (Jones believes, or believed, that the F thing which is spatially related to him in such and such a way at this or that time is such and such). Accounting for beliefs about one's <u>now</u> is as difficult as accounting for beliefs about oneself. (The '<u>here</u>' is not a separate issue; here is where I am now). There are a few self-reflexive devices which might solve this problem along B^+ lines. One possibility is to take the subject of B^+ as a temporal segment of a

person, a person-at-time. A more complicated, and perhaps more attractive, possibility is this. There are good reasons for supposing that the analysis of verb constructions will require quantifying over events and states. If so, Jones' having believed such and such is to be construed as there having existed a certain belief-state which was Jones'. So, if Jones believed that he was F at <u>that</u> time, there was a belief-state x such that there obtained among x, Jones and F the triadic relation: <u>x</u> is a believing by Jones that he is F at the very time of <u>x</u>. More precisely,

39) $(Ex)\ B^{o}(Jones,x,ab\lceil a \text{ is F at } b\rceil)$.

However we come to represent the logical form of such self-ascriptive belief attributions, we can now give a good approximation to an adequate representation of believing something under a self-ascriptive uniqueness property. Suppose that Vincent believes of that vase that it is Ming, under the (quasi-) property: being the only vase which is R to me. This is true if and only if

40) $B^{+}(Vincent,ab\lceil b = (\eta z)(z \text{ is a vase } \& Rza) \ \&\ b \text{ is Ming}\rceil,$ that vase).

We have on our hands an embarrassment of belief-relations. There is the original B, which holds among believers, n-adic intensions, and n-ary sequences. There is B^{+}, which holds among believers, n-adic intensions, and (n-1)-ary sequences. (Or perhaps, in place of B^{+}, something like B^{o}.). And what about cases in which reference is made to no particular object of belief, as in

41) William believes that what goes up must come down.?

This appears to assert a dyadic relation between William and a proposition—

42) B*(William, ⌈What goes up must come down⌉).

And what about the self-ascription of a property which involves no relation to a referred-to item, as in

43) Harvey believes that he is a paragon ?

This, by analogy with former cases, asserts a dyadic relation between Harvey and a property:

44) B'(Harvey, a⌈a is a paragon⌉).

We need not be distressed by this inflated notation. The use of the four relational expressions 'B', 'B^{+}', 'B*', and 'B' ', does not commit one to their irreducibility, or to supposing that the English verb 'to believe' is four-ways ambiguous. It is my concern now to represent at a certain level the logical form of belief-sentences, and not to give an analysis of 'believes'. It seems to me undeniable that, for example, some belief-contexts have to be construed as containing referential positions and others as not containing them. This has an effect on logical form.

There are reductivist possibilities of an unattractive kind. One might, for example, take propositional belief (B*) to be just the triadic relation B, holding among a person, a 0-ary intension (proposition) and, perhaps, the null sequence. But this is a spurious economy, which throws no light. The parallel method of reducing 'B' ' to 'B^{+}' is scarcely more attractive.

A question that should be asked is this: would the need for assigning several relations of different degree to the verb 'to believe', in order to account for the logical form of belief-sentences, have the consequence that 'to believe' is ambiguous, in the intuitive sense in which it seems not to be?

Let me emphasize that my current concern is with logical form; but it would be very agreeable if this stock of belief-relations could be reduced to one.

6. The Extension to Propositional Attitudes in General

What I have suggested concerning the logical form of belief-sentences I intend to be applicable to propositional attitude verbs in general. There is, though, a rather striking apparent difficulty in extending the account to a certain class of verbs which includes the verb 'to mean', my ultimate concern in these pages.

Consider:

45) (In saying what he said), Robert meant that the man in the brown hat is a spy.

To achieve unity of exposition, let us take 'B' as representative of 'belief-of', and pretend that non-self-ascriptive uniqueness properties are always available. If we then adopt 'M' to express 'meaning-of', the analysis of 45) which parallels our rendition of the corresponding belief sentence, with the singular term occurring non-extensionally, is:

46) M(Robert, a⌜a = (ιy)(Fy & y is a man in a brown hat) & a is a spy⌝, the man in the brown hat).

But this appears to imply that

47) M(Robert, a[a = (ιy)(Fy & y is a man in a brown hat)], the man in the brown hat).

And it may seem clear that, in the kind of case which 45) would typically report, it would be false that Michael meant by his utterance (compare 'asserted') that a certain person was the man in the brown hat. What he meant was that that man is a spy. Another example: if I say "He is a spy", it will not be the case that I mean that the person referred to is male. (Again, compare "assert").

Happily, it may be granted that in uttering 45) one will not normally be implying something of the form of 47). This is quite compatible with 46)'s being the correct analysis of 45), because, as may be seen from a moment's reflection, 46) does not imply 47), contrary to first appearances. If I intend to cause you to believe of a certain person that the conjunctive property F & G applies to him, it may still be false that I intend to get you to believe of him that the property F applies to him—for, I may know that you are already in this state. The point is that my intention may be to get you to believe of the thing that F and G are conjointly instanced in it.

'Intends' and 'hopes' share with 'intends that x believes' the feature that, where the second term of the relation is a conjunctive intension, the elimination of one of the conjuncts does not always preserve truth. In this respect such appetitive verbs differ from 'believes'.

CHAPTER III

REFERENCE AND THE MEANING OF SIMPLE SIGNALS

1. Non-Propositional Signals

The locution which concerned us in the first chapter was 'x means p'. But, as I have already indicated, its analysis does not provide a fully general account of the meaning of simple signals. For, like sentences, simple signals may have indexical aspects, and hence the conventions which govern some of them do not fully determine what is meant on their conventional, or (so to speak) literal, employment. When a Roman spectator at the circus turned his thumb up, he meant that a certain one of the performers was to be spared. But it was not part of the conventional meaning of that gesture that Marius was to be spared.

Virtually all familiar simple signals have some degree of indexicality. For there are certain more or less determinate kinds of information which we wish to be able to convey about any members of indeterminate classes of objects, times, places and persons. Hence there is some utility in having signal-types which have conventionally determinate informational content, but which do not conventionally fix reference to determinate particulars. There are cases, of course, in which such openness of reference is no advantage. Conspirators in need of a signal for conveying information about definite particulars may agree that

their signal will mean p, and that may be sufficient for the signal's meaning p in that group, for the duration.

Though a signal (type) may lack propositional meaning, it will nevertheless have a determinate meaning. A ring painted around a tree may conventionally be a signal to the cutter that the tree is ready for cutting. Though there is no single proposition which every such painted ring conveys, there is a certain propositional function—namely, a[a is ready for cutting]—which in each case is meant to hold of the tree thus marked. So it seems that the timeless meaning of certain simple signals can be identified with a propositional function.

The sense in which a simple signal has a propositional function as its meaning is rather different, of course, from the sense in which a verb phrase does. The painted ring and the English verb phrase "is ready for cutting" have the same propositional function as their meaning, but they have significantly different semantical properties.

This difference will be adequately marked if we adopt as our current analysandum 'x means <u>it is F</u>', where 'F' ranges over propositional functions. I begin with this simple sort of simple signal in order to introduce a perspicuous basic schema for the subsequent general account. The immediately next sections are concerned with indicative meaning, but once an adequate overall account of that is reached we can deal with signals which have imperative force and with signals with more specific illocutionary forces.

The reference of any utterance of a signal which means it is F is left open by the convention which gives that signal its timeless meaning. Nevertheless it would appear to be an aspect of the timeless meaning of such a signal that, on its full conventional employment, its utterer refers to something or other—that is, means the relevant propositional function to hold of some contextually definite item. So we should expect that, whatever the convention is which governs a signal which means it is F, one utters the signal fully in accordance with the convention only if (Ez) one means by one's utterance that F of z. The 'free for meaning' proviso of the earlier definition is quite as relevant here as it was in chapter I. We need it, though, in a slightly revised form, so let us adopt the following:

x's utterance of z is free for meaning F of some n-ary sequence[1], audience y, if, and only if,

(i) it is mutually known by x and y that x uttered z;

(ii) there is, for y, no hypothesis which overrides as an explanation of x's utterance the hypothesis that x means, thereby, F of some n-ary sequence, audience y;

(iii) it is not, for y, so improbable that (w) if x uttered w in C, x would mean thereby F, of some n-ary sequence, that y could not take any utterance of x's in C as meaning F of some n-ary sequence.

---

[1] x means F of an n-ary sequence if F is an n-adic propositional function and $M(x,F,\langle z_1 \ldots z_n \rangle)$.

Putting this to work, we obtain the following account:

z means <u>it is F</u> in G

if, and only if,

there is in G the convention that

(x)(y) if x and y are members of G and x utters z in circumstances in which his utterance is free for meaning F of something, audience y, then x means thereby F of something, audience y.

The definition attributes to a signal which means <u>it is F</u> two important semantical properties. First, if someone utters the signal fully in accordance with the convention, then he is letting it be known that he is referring to something; and, secondly, he is letting it be known that, whatever that item of reference is, he means F of it. In other words, the convention determines not merely what the utterer means of whatever he is referring to, but it also requires that there be something or other to which he is referring. So if, in circumstances in which his utterance of a signal which means <u>it is F</u> is free for meaning F of something, the utterer fails to provide a reference for his utterance, he has violated the convention. This defect in an utterance is notably similar to ambiguousness. In both cases the circumstances are suitable for an informative utterance of a signal or sentence, but the speaker fails to provide a setting which, in the one case, disambiguates the utterance and, in the other, yields a reference for it.

Apart from this, there is no conventional restriction on what is referred to by an utterance of such a signal.

I shall return shortly to the question of what reference is.

2. Qualified Indexicality

A signal which means it is F has, as we might call it, unqualified indexicality. There is no conventional restriction on what sort of thing is referred to by its utterance, or, better, no conventional restriction on how its referent is referred to. But signals may have a kind of qualified indexicality—they may, for example, have the same meaning as an English sentence of the form 'The A is B'—that is, they may mean the E is F.[2] One utters such a signal fully in accordance with its conventionally determinate meaning only if one refers to something as being E.

There is a new kind of semantical property here—one which is not possessed by signals whose meaning can be expressed as it is G. The best way of showing this is to show that one cannot assimilate the case of signals which mean the E is F to the case of signals which mean it is G.

If someone utters a signal which means the E is F with its full conventional force, then, for a certain thing which is E, he means by his utterance that the E is F. By the account of propositional attitude constructions proposed in the last chapter, x means that the E is F if there is a D such that D is a uniqueness specification of E—that is, D is a uniqueness property and being D entails being E—and x means a[Fa & Da] of the E.[3]

---

[2] I shall use the letters 'D', 'E', 'F', 'G' as variables over properties, and, sometimes, relations.

[3] To keep matters simple, I ignore for the time-being self-ascriptive cases. The subsequent general account covers them.

From this, it might appear that a certain small modification of the definition of 'z means it is G' will be sufficient to accommodate signals which mean the E is F. For, consider what someone means on a normal utterance of a signal which (timelessly) means it is G. Clearly there will be some uniqueness property D' (to ignore for the moment questions about self-ascription) such that the speaker means a[Ga & D'a] of the item in question. For he will intend the hearer to believe of the item under a certain property it is G. So perhaps the only difference between a signal which means it is G, and a signal which means the E is F, is that there is no conventional restriction on what uniqueness property is meant on an utterance of the former (though there must be some such uniqueness property), while in the latter case there is the restriction that the uniqueness property which is meant on its utterance will be some (uniqueness) specification of E.

That this is not an adequate account of the semantical properties of a signal which means the E is F may be seen without having to consider an actual definition. The account treats the E aspect of 'x means the E is F' as making merely a contribution to what is meant on a conventional utterance of the signal. As I shall subsequently show, it is true that if someone utters such a signal with its full conventional force, then for some specification D of E, he does mean a[Fa & Da] of some item. But this does not exhaust the significant semantical features of such an utterance.

Compare the two English sentence forms 'The F is G' and 'It is the F and it is G'. On a conventional utterance of the former, one intends

the item's being the F (or being believed to be so) to let the hearer know that it is that item which one is asserting to be G. The speaker refers to the item as the F. On an utterance of the second kind of sentence, the speaker does not (conventionally) refer to something as the F, but rather asserts of something independently referred to that it is the F and is G.

An account of 'x means the E is F' which has the E aspect merely contributing to what is meant would not capture the referential contribution which the E aspect makes on a conventional utterance of the signal. But in order to capture this additional feature, we need an account of reference. More specifically, we need the notion 'x, by his utterance of z, refers to w qua D', 'D' being a variable over uniqueness properties.

3. What Reference Is

What is it to refer to something? The fact that we can, in the sense earlier described, quantify into the 'that'-clause in 'the speaker means that ...' may suggest that referring to x is to be identified with meaning something of x.

Meaning F of x, it is true, is a sufficient condition of referring to x. For even if the speaker does not produce something which has the separate function of referring to x, he still in some way makes known of which thing he means F to hold, and so has referred to that thing. But the distinctness of meaning something of x, and referring to x, is confirmed by the following trivial counterexample to the former's

being a necessary condition of the latter. Suppose a person begins a serious utterance by producing a name—say, "Horace Greeley"—and is then interrupted before he says anything more. He referred to Horace Greeley, even though there was nothing which he meant of Horace Greeley. He had intended to say something by which he meant that Horace Greeley is such and such, but in fact he did not succeed in doing so.

Indeed, the distinctness of referring and meaning is already implicit in the fact that (as I mentioned in the last section) the full semantical force of a signal which means the E is F is not identical merely with what is meant on its conventional utterance. Referring to something as the E is not the same as meaning of it that is the E.

A speaker refers to something qua its having a certain uniqueness property. Thus, on a normal utterance of "the cat is on the mat", one refers to the cat qua (the property of) its being the only cat which is such and such. In certain cases, the speaker will be taking advantage of the fact that the hearer already believes the referred-to item to be the only D, and intends this belief of the hearer's to be operative in the hearer's recognition that it is that item which the speaker intends to mean something of. But it often happens that, in referring to x qua D, we do not expect our hearer to have noticed x already, or to have been aware of x's being D. Rather, one will intend by one's utterance to produce this awareness in the hearer, and thereby to produce in the hearer the recognition that it is x about which one means something further.

The notion of "activating" a belief[4] is just right for expressing

---

[4]To activate a belief is either to produce a conscious belief or to make someone entertain an already held (dispositional) belief.

what these two cases have in common. Whether or not the hearer is already aware of x's being D, the speaker intends by his utterance to activate in the hearer the belief that x is D in order that the hearer will recognize that it is x of which the speaker means such and such by his (perhaps further) utterance.

It is not only qua some property that one may refer to something but also qua its being uniquely related in a certain way to certain items.

These several requirements may be combined to give us this preliminary definition of reference:

By uttering z, x refers to w qua D (or qua R to the sequence s),

audience y,

if, and only if,

x intends it to be mutual knowledge between x and y that x intends, by his utterance of z, to activate in y the belief that w is D (or R to s) in order that y will recognize that for some F (or for some T and s') x means that y is F (or T to s'), or x means that y is to make it the case that y is F (or T to s').

Several features of the definition require comment.

(a) Although I continue to defer the general question of imperative signals and signals with more specific illocutionary forces, it is convenient once and for all to have a definition of reference which covers all cases. Since in all illocutionary acts one either means that such and such, or that the hearer is to do such and such, this definition exhausts the cases.

(b) As the reader may have noticed, my assumption that all reference is qua a certain uniqueness property or uniqueness relation is not compatible with what I pointed out in chapter 2 regarding the fact that uniqueness properties (or relations) are not always available. I shall subsequently return to this question, when I discuss utterance-reflexive reference, but let it suffice now to notice that, just as there is the self-ascriptive belief relation B^+, so there is the self-ascriptive reference relation refers$^+$. Roughly speaking, one refers$^+$ to an item as being related in a certain unique way to oneself.

(c) The 'by uttering z' in the definiens refers to one's total utterance z. There is a more restricted sense in which one refers to something by an utterance—namely, the sense in which, in "George climbed the Matterhorn", one refers to George by "George", and to the Matterhorn by "The Matterhorn". It is an interesting question how one would capture this notion, but I do not need it for current purposes.

The above condition unfortunately fails to provide a necessary condition of referring to something qua D. For not only may one refer to something in order to mean something of it, but one may refer to something in the course of or in aid of referring to something else. So, if I say "the bride's sister fainted" I refer to the bride's sister and mean of her that she fainted, but I refer to the bride primarily in order to fix the reference to her sister, and it is incidental to this that I might also mean something of the bride.

There is no limit to the degree to which referential expressions

may be embedded within referential expressions. The form "the F which is R to the G which is S to the H which is ..." may be continued indefinitely. Consequently, if we are to give a general characterization of the contribution which a reference makes to an utterance we need a recursive definition.

Two preliminary points:

(1) Let us say that reference to x is primary reference when reference to x is made in order principally to assert something about, or to mean something of, x. If x is being referred to qua its being R to some sequence of items, we may say that reference to each of those items is, in its turn, secondary reference. And so on for n-ary reference The point is that there is no limit either to the number of things which may be referred to at any level, or to the number of levels at which things may be referred to.

(2) If an item of reference is not referred to qua its having a certain uniqueness property, then it is referred to qua its being uniquely related to certain (also referred to) items . This simple classification may seem not to accommodate the reference which is made in the use of expressions like "the cat which chased the mouse", for the referent is there referred to both as having a certain property, and as standing in a certain relation to a certain thing. The simple point, though, is that whatever is referred to by that phrase is referred to as having a unique relation to a certain thing (the mouse), where standing in that relation to something entails both being a cat and chasing that thing.

The previous definition of reference may be expanded to cover

recursively the case of embedded reference.

By uttering z, x refers to w qua D (or qua R to the sequence s), audience y,

if, and only if,

(i) x intends it to be mutual knowledge between x and y that x intends, by his utterance of z, to activate in y the belief that w is D (or R to s) in order that y will recognize that for some F (or some T and s') x means that w is F (or T to s') or x means that y is to make it the case that w is F (or T to s'); or

(ii) x intends it to be mutual knowledge between x and y that x intends, by his utterance of z, to activate in y the belief that w is D (or R to s) in order that y will recognize that for some u, T and s', w belongs to the sequence s' and x by his utterance of z is referring to u qua T to s'.[5]

Reference is made to some item either in order to mean something of it or in order to refer to something else

4. The Meaning of Signals Which Mean the E is F.

From the already mentioned complexity of referential embeddings, it will be apparent that the definiendum of current concern—namely, 'z means the E is F'—is far too narrow to cover the full range of simple signals. I shall, nevertheless, provide it with a definiens,

---

[5] This definition is a joint production with S. R. Schiffer, and emerged from combining our separate attempts.

for in doing so we may isolate and independently justify a certain important feature of the final account.

One thing should be clear: on a normal utterance of a signal which means the E is F, there will be a certain uniqueness specification D of E such that the utterer refers to something qua D. And the account may seem to be completed by the further condition that the utterer means of that referent that it is F.

But it is not a sufficient condition of uttering such a signal with its full conventional import that one refers by one's utterance to something qua (some uniqueness specification of) E and means of it that it is F. For there are describable cases in which one refers by an utterance to a certain thing qua (say) D, and means of it that it is F, but means this of it under some property other than D. Here is an example. Ralph has still failed to identify the man in the brown hat as the man seen at the beach, and he attempts to point up the radical difference between them. By his utterance, Ralph means that the man in the brown hat is a spy while the man seen at the beach is a patriot. So the fellow gets referred to by Ralph qua man in the brown hat, and Ralph means of him that he is a patriot (under another property). But it is false that Ralph, by his utterance, means what would be meant on a fully conventional utterance of a signal which means the man in the brown hat is a patriot.

What the missing element is should be clear. On a fully conventional utterance of a signal which means the E is F, there will be a certain uniqueness specification D of E, such that the utterer refers to something

qua D and means a[Fa & Da] of it. The repetition here is essential. To refer to something by the utterance of such a signal is to single it out as the subject of some current predication. But what one predicates is meant of the thing under the very uniqueness property qua which it was referred to. That is, in a successful communication of this type, two things are necessary: first, one gets one's hearer to have in mind that it is x to which one is referring qua D; secondly, one causes one's hearer to believe that x is F under the property D— in other words, to believe a[Fa & Da] of x.

The news in the utterance of such a signal (what is asserted rather than presupposed) is that part of what is meant which is independent of how the subject of predication is referred to. In this sense, what is asserted on the literal utterance of a sentence is not always identical with what is meant by the utterance of that sentence. I shall return to this question in the next chapter.

The earlier account of 'z means the E is F', which gave merely what is meant by an utterance of such a signal, was defective not because it failed correctly to specify what is thus meant, but because it left out the additional referential function of E. It is with respect to this latter feature that conventional utterances of "The F is G", and "It is the F and it is G" differ

So, we have the following definition:

z means the E is F in G

if, and only if,

there is in G the convention that (x)(y) if x and y are members of G

and x utters z in circumstances in which, for some D, D is a uniqueness specification of E and x's utterance of z is free for meaning a[Fa & Da] of something, audience y, then, for some D, D is a uniqueness specification of E, and x, by his utterance of z, refers to something qua D and means a[Fa & Da] of it, audience y.

A few comments:

(a) The case of z's meaning <u>it is F</u>, that is, of z's having unqualified indexicality, may be brought under the definition as the case in which E is a vacuous property—for example, the property a[a = a].

(b) I have not mentioned reference in the antecedent of the convention (perhaps the notion of an utterance's being 'free for referring to something' would be the appropriate one) since, as far as I can see, if an utterance is free for meaning a[Fa & Da] of something, then it will be 'free' for referring to it qua D, and vice-versa. Consequently reference comes into the convention only in what it requires when the antecedent condition obtains.

(c) The convention is not fully adequate, since it only covers cases in which uniqueness properties are available. But it is perfectly possible for there to be a case in which full communication exists between x and y on x's utterance of a signal which means <u>the E is F</u>, even though x's and y's conceptions of themselves, their environments, and the object of reference are, for whatever reasons, not sufficiently singular to provide a uniqueness property. This deficiency will be remedied in the final account of the meaning of simple signals.

I now take up a closely related matter—namely, the kind of

reference made on the utterance of signals which have reflexive aspects.

5. Reflexive Signals

Some signals are speaker-reflexive—that is, they mean I am F; some are hearer-reflexive—that is, they mean you are F; and some are utterance-reflexive. As regards this last class of signals, there is no English word or phrase whose function is specifically the focusing of reference on the utterance itself. But I have in mind here the tenses of verbs, and the occurrence of "here" and "now". I shall suppose that a signal may mean this^{+} is F, where the reference which is made on an utterance of such a signal is conventionally restricted to the utterance itself, just as in the former two cases reference is restricted to the speaker and to the hearer.

This class of signals is not a sub-class of the class of signals which mean the E is F. Here is a simple (no doubt, unnecessary) proof that there is no E such that a signal means I am F if and only if it means the E is F. Either E applies to at most one thing, or it does not. If the former obtains, at most one person could hope, without being deceived or intending deception, to refer to himself by a conventional utterance of the signal. Otherwise, E applies to more than one thing, in which case, as long as this is generally known, a speaker might utter the signal with its full conventional force and not refer to himself.

I shall begin with the traditionally most neglected of these three types—namely, simple signals which mean you are F. It should now be

evident that a fully conventional utterance of such a signal is not merely one on which the speaker means F of the person who happens to be the intended hearer. Rather, the speaker refers to the hearer *qua* his being the intended hearer of the utterance and means F of him, under the property of being the intended hearer. How might we capture this?

There is something to be learned from making explicit a certain feature of speaker's meaning which lies beneath the surface of Schiffer's definition. Very roughly speaking, someone means something by his utterance provided that he intends it to be mutual knowledge between him and a certain other person that he intends to produce in the latter a certain (appropriate) effect. Now the point is that the mutual knowledge which would thus obtain between speaker and hearer if the speaker's intention were fulfilled would be for the former *of* the latter under a certain property, and vice-versa, each being aware of the other's relevant conception of him, and aware that the other is thus aware, and so on, with no ignorance at any stage.

The situation is made more complex by the fact that in normal circumstances it is quite irrelevant to what is going on whether or not there actually are available to the speaker and the hearer uniqueness properties which might serve as their separate conceptions of each other. That is, in typical cases of communication, the speaker's conception of the hearer will be of him as uniquely related in a certain way to him (the speaker); the hearer will have a corresponding self-related conception of the speaker; and this will be mutually known by them. We can represent this more accurately as follows.

Let K^+ be the self-ascriptive knowledge relation, with the logical properties of B^+. If x and y mutually know that p, there will be uniqueness relations R and T such that:

(i) $K(x,p)$, and

(ii) $K(y,p)$, and

(iii) $K^+(x,ab[Rba \ \& \ K(b,p)],y)$, and

(iv) $K^+(y,ab[Tba \ \& \ K(b,p)],x)$, and

(v) $K^+(x,cd[Rdc \ \& \ K^+(d,ab[Tba \ \& \ K(a,p)], c)], y)$, and

(vi) $K^+(y,cd[Tdc \ \& \ K^+(d,ab[Rba \ \& \ K(a,p)], c)], x)$, and

⋮

There is an obvious rough formula for continuing this list. Even-numbered lines (except ii) report y's knowing of x, under the relation T to y-self, that x's knowledge of him is as reported three lines back; odd-numbered lines (except i) report x's knowing of y, under the relation R to x-self, that y's knowledge of him is as reported in the preceding line.[6]

If this infinite set of conditions hold, then x and y mutually know that p (as I shall say) <u>under</u> $\langle R,T \rangle$.

If someone utters a signal which means <u>you are F</u>, his intended conception of the hearer will typically be a relation (rather than a property) of just the kind which figures in these last conditions, the speaker's conception of the hearer being, typically, self-ascriptive. It will be useful to have a device for specifying the relation which

---

[6] There is an adequate recursive condition which correctly relates x, y, R, T and p by, as it were, summing up this infinite conjunction, but the list gives an accurate idea of the effect of such a condition.

serves as the speaker's conception of the hearer for any speaker, hearer, and utterance, and so we may define the following function, which takes relations as its values.

$$H(x,y,u) = R$$

if, and only if,

For some T, u is an utterance of x's and x intends it to be mutual knowledge between x and y <u>under</u> $\langle R,T \rangle$ that x intends by his utterance u to produce in y a certain belief or action.

The function H (whose definition here is only approximate) provides us with the means for stating what is meant on any conventional utterance of a signal which means <u>you are F</u>. If u is a fully conventional utterance by x of a signal which means <u>you are F</u>, then, for some person y and some relation R, $H(x,y,u) = R$, and x refers^{+} by his utterance of u to y qua R to himself, and x means^{+} thereby $ab[Fb \ \& \ Rba]$ of y.

This requires some elucidation.

(a) The notions <u>refers</u>$^{+}$ and <u>means</u>$^{+}$ inherit their self-ascriptiveness from the self-ascriptiveness of <u>intends</u>$^{+}$ as it figures in their analyses. In general, the intention that an item should be such and such will be of that item either under a certain uniqueness property, or under a certain unique relation to the intender himself.

(b) There is a consequence of the use of self-ascriptive reference and meaning relations which will seem strange at first. It may be brought out by contrast with cases of reference and meaning which are not self-ascriptive, (if there are any such). Suppose that x refers to w qua D, and means $a[Fa \ \& \ Da]$ of w — that is, he means that w is F, under the

property D. In general, one might assume, the content of what is meant is identical with the content of what is intended to be believed (or done). But this assumption is not compatible with the admission of a self-ascriptive meaning relation. For suppose that the utterer means^{+} ab[Fb & Rba] of w. What does he intend his hearer to believe? Clearly, he will not normally intend his hearer to believe^{+} ab[Fb & Rba] of w, for R is the unique relation in which the speaker regards w as standing to (the speaker) himself. Rather there will be a distinct uniqueness relation T such that the speaker intends the hearer to believe^{+} ab[Fb & Tba] of w.

This should be taken to show, I think, that one of our weak intuitions on these matters is mistaken. It is not true that the content of what is meant is always identical with the content of what is intended to be believed.

The relation ab[Rba] to himself under which the speaker means^{+} F of w is that relation to himself under which the speaker intends^{+} of w that, for a certain T, it is mutual knowledge between speaker and hearer that the former intends the latter to believe ab[Fb & Tba] of s. One should keep in mind that to mean^{+} something is to intend^{+} a certain thing.

Second-person utterances have a special feature. If x produces a conventional utterance u of a signal which means <u>you are F</u>, audience y, then, for some R, H(x,y,u) = R, and x means ab[Fb & Rba] of y. R is x's conception of y, under which x intends^{+} such and such. But there is no need for postulating a self-conception under which y is intended

to believe F of himself. What is intended to obtain is a case of simple self-ascriptive belief—namely, that B'(y,F).

For dealing with signals which mean <u>this^{+} is F</u>, we shall require a small further bit of apparatus. If x means something by an utterance u, then he will intend there to be mutual knowledge between him and his hearer, under a certain speaker-conception of the utterance u, and under a hearer-conception of u, that he intends of u that ... So we may define a function which yields a certain relation, for a given speaker, hearer, and utterance, which is such that the speaker's conception of the utterance is as having that relation to the speaker himself.

$$U(x,y,u) = R$$

if, and only if,

For some T, x utters u and intends it to be mutual knowledge between x and y—where x's conception of u is as R to himself, and y's conception of u is as T to himself—that x intends by uttering u to produce such and such (appropriate) effect in y.

On this basis we can say that a signal means <u>this^{+} is F</u>, provided that if u is one of its fully conventional utterances, for some y and R, U(x,y,u) = R and x, by his utterance of u, refers^{+} to u qua R to himself and means^{+} ab$\lceil$Fb & Rba$\rceil$ of u, audience y.

Signals which mean <u>I am F</u> are, in certain respects, the simplest to deal with, and yet they create a difficulty from the point of view of ultimately subsuming such signals under a general definition of the

meaning of simple signals.

What it is for a signal to mean I am F may be specified like this: if x utters such a signal with its full conventional meaning, then x refers^{+} to himself, and means^{+} F of himself. We may make more explicit what is concealed behind this, as follows.

With second-person and utterance-reflexive signals, the useful functions, H and U, both take as their values that relation which is the content of the speaker's conception of something—namely, of the hearer and of the utterance. But where first-person utterances are concerned, there is no speaker's conception of himself—self-identity singles out nothing. Nevertheless, it may be illuminating if we here isolate the hearer's intended conception of the speaker, by specifying the function S.

$$S(x,y,u) = T$$

if, and only if,

For some R, u is an utterance of x's and x intends it to be mutual knowledge between x and y under $\langle R,T \rangle$ that x intends by his utterance u to produce in y a certain belief or action.

S is H's mirror-image, and we may with its aid specify more fully what is meant by a first-person utterance—this time from the hearer's intended point of view. If u is x's fully conventional utterance of a signal which means I am F then, for some T, $S(x,y,u) = T$ and x intends it to be mutual knowledge between x and y that x intends y to believe^{+} $ab[Fb \ \& \ Tba]$ of x.

The difficulty which I referred to above is this. We should have

a quite general way of specifying the meaning of any signal. In the account developed subsequently, the canonical form for specifying the meaning of a second-person signal will not be 'z means <u>you are F</u>', but rather 'z means <u>F(H)</u>'. Utterance-reflexive signals mean <u>F(U)</u>. In both these cases, a signal's meaning what it does is its standing in a certain relation to a pair which consists of a property and a certain function,—where in both cases that function maps the utterance of the signal onto the speaker's conception of the referred to item. No such function is relevant to the case of first-person signals. So, in order to achieve a certain uniformity in the style of meaning attributions, I shall say that a first-person signal means <u>F(S)</u>. This introduces an asymmetry into the style of meaning attributions for reflexive signals, but the subsequent general definition of the meaning of simple signals does in fact assign to signals which are said to mean <u>F(S)</u> the conventional semantical properties of signals which mean <u>I am F</u>.

6. Referential Qualifiers

How much complexity of reference and of referential embedding may be packed into the meaning of a simple signal? The matter is obscure. On the one hand, a logical-atomistic insistence that the structure of a signal should be isomorphic to the structure of the signalled fact or would-be fact is at odds with real uses of signals. For example, the contrast between showing one lamp and showing two lamps does not mirror any structural difference between the British Army's approaching by land and their approaching by sea. But, on the other hand, there

seem to be limitations of some kind on what a simple signal might convey—could one communicate the whole message of the Book of Genesis with a wave, or with a yawn?

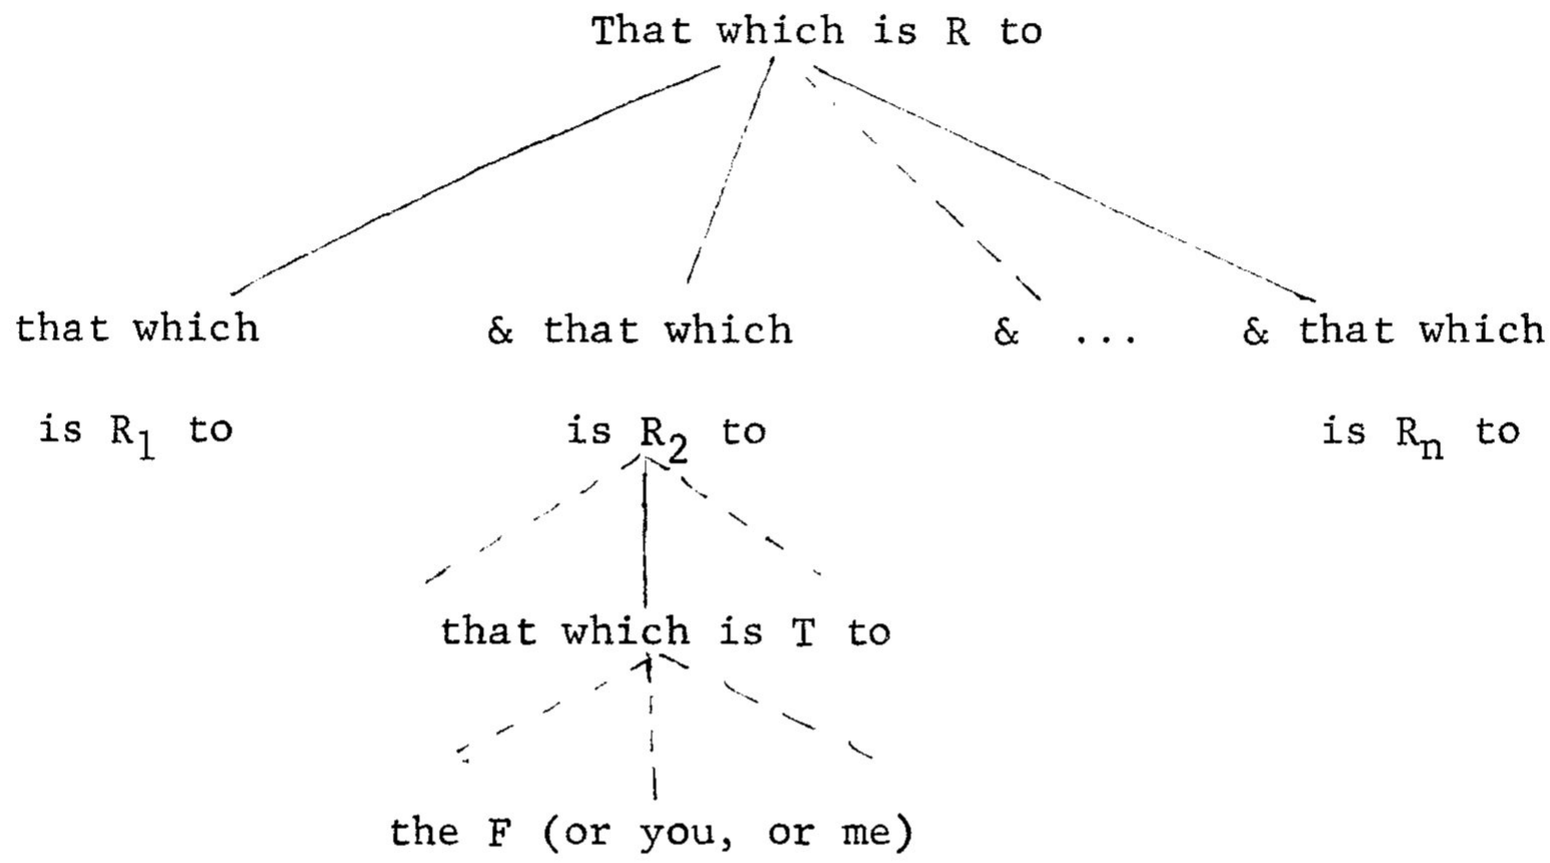
Probably the only real limitations are psychological ones, but I shall not be concerned to argue that point here. Rather, I shall make allowance for the possibility that anything that can be meant by a structured utterance type—by, say, a sentence—can be meant by a simple signal, and, of course, conversely. This procedure does not prejudice the above issue—for, if it is held that there is a limit to the complexity of meaning of a simple signal, then that is equivalent to holding that there is a limit to the existence of those conventions which by my definition constitute the meaningfulness of signals.

This liberal procedure has two distinct advantages. First, in getting a preliminary account of how generally to describe the meaning of signals, we may simply consider certain semantical properties of sentences. And, secondly, in subsequent dealings with sentence meaning, we may then use precisely the same manner of meaning attribution as we have used for simple signals.

The referential structure of English sentences is complex in two dimensions, neither of which is represented in the specification of a signal's meaning <u>the E is F</u>, or <u>you are F</u>. First, as I pointed out earlier, a referring expression may itself contain further referring expressions, and they in their turn may contain yet further such expressions, and so on indefinitely. Secondly, a sentence may contain any number of primary referring expressions, each with their separate

further embeddings. That is, a sentence may have (roughly speaking) the form "Relation R holds among the E (which is ...) and the F (which is ...) and the G and ..."

The structure of a referring expression may be presented in a tree, thus:



Each node which is not final is occupied by an (n+1)-adic relation, the node having n branchings. This relationship between the degree of the relation at a node, and the number of branchings there, merely represents the fact that in "that which is R to x_1 and ... and to x_n", one place in the relation R is applied to the item which is being principally referred to—that is, the referent of the whole expression. Final nodes are occupied either by properties ("that which is R to ... and to the cat"), or by one of the functions H, U, and S ("you", "me").

So far in this chapter, we have had reasonably perspicuous explicanda for our reductivist program for simple signals—namely, "z means <u>the E is F</u>", "z means <u>you are F</u>", "z means <u>I am F</u>" and "z means <u>this^{+} is F</u>". Since these styles of meaning-attribution do not cover all the cases, we have now to fashion a style that is fully general—that is, one which will cover signals with any referential structure, however complex. One of the preliminary problems in the theory of meaning is to invent a precise problem to solve.

In the case of signals which mean <u>the E is F</u>, or <u>F(H)</u> (i.e., <u>you are F</u>), let us call E and H <u>referential qualifiers</u>. Now, with regard to referring expressions in general—whose structures can be represented in a tree diagram of the kind shown—their essential semantical content consists in what relations, properties, and reflexive functions are embedded within the scope of what relations. Let us define an extended sense of 'referential qualifier' which captures this—one on which a referential qualifier mirrors the structure of the corresponding tree. The definition must, clearly, be recursive—thus:

e is a referential qualifier

if, and only if,

(i) e is a property, or

(ii) e is one of H, U and S, or

(iii) for some R, and $y_1 \ldots y_n$, R is an (n+1)-adic relation, and $y_2 \ldots y_n$ are referential qualifiers, and $e = \langle R, \langle y_1 \ldots y_n \rangle\rangle$.

Expressions like "the F", "you", and "me" have referential qualifiers of kinds (i) and (ii). The application of the recursive clause may be seen from an example. Suppose a referring expression has the form "that which is R to (a) that which is T to the D and the E, and to (b) that which is R' to the F and the G and you." The corresponding referential qualifier is: $\langle R, \langle\langle T, \langle D, E\rangle\rangle, \langle R', \langle F, G, H\rangle\rangle\rangle\rangle$. In a sense, every semantical respect in which two referring expressions differ may be represented by a difference in their referential qualifiers.

We may now adopt a quite general style of attributing indicative meaning to simple signals, and thus provide a definiendum for the over-all account of simple signals. For any indicative signal (and, for that matter, sentence) z, there is a property or relation F, and referential qualifiers $e_1 \ldots e_n$ such that z means $F(e_1 \ldots e_n)$. F is simply that n-adic relation which, on a fully conventional utterance of z, is asserted to hold among the n primary references of the utterance. And $e_1 \ldots e_n$ are the referential qualifiers which correspond to the referring expressions of some sentence which means the same as z.

7. How to Refer on the Conventional Utterance of a Signal

The timeless meaning of signals which mean the E is F was given by specifying a) how reference is made on their conventional utterance, and b) what is meant on their conventional utterance. This is just what must be done for the general case. In order to define what it is for the signal z to mean $F(e_1 \ldots e_n)$, we shall need a general characterization of both the kind of reference made on a conventional

utterance of z, and what is meant on such an utterance.

The properties and relations which are constituents of the referential qualifiers of sentences and signals are not characteristically uniqueness properties and relations. That this is so may be seen on a random inspection of the kind of definite descriptions and demonstrative phrases which occur in speech. Expressions which connote fully explicit uniqueness properties, like "the present Queen of England", are not extremely common among them.

Nevertheless, those properties and relations qua which actual reference is made are always uniqueness properties and relations. (Even in self-ascriptive reference, the referent is referred^{+} to qua being uniquely such and such to the speaker.) Uniqueness properties and relations, when they (as it were) cluster around a certain primary reference, may be grouped hierarchically in a way parallel to the arrangement of properties in a referential qualifier. Thus, if x refers to z and w qua F and G respectively, and refers to y qua R to z and w, we may take $\langle R,\langle F,G\rangle\rangle$ as, in its abstract way, capturing the structure of that piece of reference. The class of all such structures may be recursively defined as follows:

e' is a <u>specified qualifier</u>

if, and only if

(i) e' is a uniqueness property, or

(ii) e' is one of H, U and S, or

(iii) for some R and $e'_1 \ldots e'_n$, R is an (n+1)-adic uniqueness relation, and $e'_1 \ldots e'_n$ are specified qualifiers, and $e' = \langle R,\langle e'_1 \ldots e'_n\rangle\rangle$.

So, on the fully conventional utterance of a signal which means F(e), there are a) the referential qualifier e, which is part of the timeless meaning of the signal, and b) some specified qualifier whose constituents are the uniqueness relations and properties qua which the primary through n-ary referents of that utterance are referred to. Now, the question is: what constraints does e place on the relevant specified qualifier of a conventional utterance of a signal which means F(e)?

The problem may be clarified by showing that a certain straightforward answer is inadequate. It may seem that on a fully conventional utterance of a signal which means F(e), there will be a certain specified qualifier e' which is such that e' is isomorphic to e and each property or relation in e' is a uniqueness specification of the corresponding property or relation in e. But although this is sometimes true, it is not always true. Suppose e = $\langle R, \langle D, E \rangle\rangle$. It need not happen that on a fully conventional utterance of a signal which means F(e), reference is made to something qua some property which is a uniqueness specification of, for example, D. For the D aspect of the meaning of the signal may be here exploited by the utterer in these ways: a) There is a uniqueness relation T such that being T to something entails being D and the utterer refers^{+} to his (in this case, secondary) referent qua its being T to (the utterer) himself. For example, by uttering "the car" I may be referring^{+} to something qua its being my car. b) There is an (n+1)-adic uniqueness relation T' such that being T' to anything entails being D, and the utterer refers to his referent qua its being

T' to $y_1 \ldots y_n$ (or, perhaps, the utterer refers$^+$ to his referent qua its being T' to himself and $y_1 \ldots y_{n-1}$.) In this kind of case, there is the further complication that the additional items $y_1 \ldots y_n$ (here in tertiary reference) are being implicitly referred to qua their being such and such. I see no a priori reason to deny that such implicitly referred to items are referred to qua their being related to yet further items, which in their turn ..., and so on. Clearly, if there is no theoretical limit to the degree of embedding of such implicit references, then we shall have to allow for every possible degree in our account. Similar considerations are relevant, in our example, to the relation R. It may be that there is a uniqueness relation of greater degree than R being which to anything entails being R to something, and which is such that reference is made to something qua its standing in that relation to further items.

In these cases, the properties and relations in e contribute something to—that is, place some restrictions on—what the corresponding specified qualifier is. That relation between e and e' which holds when e' is the relevant specified qualifier of a conventional utterance of a signal which means F(e) is defined as follows.

e' is a specification of e

if, and only if,

(ia) (For some F and G) e = F, and F and G are properties, and G is a uniqueness specification of F, and e' = G; or

(ib) (For some F, R, and $e'_1 \ldots e'_n$) F is a property and e = F, and R is an (n+1)-adic uniqueness relation such that, for any x, $x_1, \ldots, x_n$, x's being R to $x_1 \ldots x_n$ entails x's being F, and $e'_1 \ldots e'_n$ are specified qualifiers and $e' = \langle R, \langle e'_1 \ldots e'_n \rangle\rangle$; or

(ii) e = H or U or S, and e' = e; or

(iii) (For some R, T, $e_1 \ldots e_n$, and $e'_1 \ldots e'_2$) $e = \langle R, \langle e_1 \ldots e_n \rangle\rangle$, and T is an (n+m+1)-adic uniqueness relation ($m \geq 0$) and, for any $x, x_1, \ldots, x_{n+m}$ x's being T to $x_1 \ldots x_{n+m}$ entails x's being R to $x_1 \ldots x_n$, and $e'_1 \ldots e'_{n+m}$ are specified qualifiers, and, ($1 \leq i \geq m$) e'_1 is a specification of e_i, and $e' = \langle T, \langle e'_1 \ldots e'_{n+m} \rangle\rangle$.

Some comments on the notion of a specification.

(a) Notice that any of H, U or S may occur in e' though it does not occur in e. S is particularly important: on an utterance of "the car" the relevant specification may be $\langle R, \langle S \rangle\rangle$, where R is the relation: "being the only car which is owned by".

(b) H, U and S are not, of course, relations qua which items are referred to; they will be cashed at the next stage.

This definition provides half of what we need for a fully adequate account of the referential aspects of a conventional utterance of a signal which means F(e). On any such utterance, there is a specification e' of e, which is such that reference is made to various items qua the properties and relations in e'.

But not only are these items referred to qua the properties and relations in e'; they are also (if e' is other than completely simple) referred to qua standing in those relations to each other. Suppose

$e' = \langle R, \langle F, G \rangle \rangle$. There are then items x, y, and z such that y and z are referred to qua F and G, and x is referred to qua R to y and z.

Notice that the items referred to in primary through n-ary reference may themselves be hierarchically grouped in a way which is (with the exception of one kind of case) isomorphic to the specified qualifier qua whose constituents they are referred to. So, for example, in the kind of case just mentioned, there corresponds to $\langle R, \langle F, G \rangle \rangle$, the (to adopt a convenient term) referential complex $\langle x, \langle y, z \rangle \rangle$.

The only cases in which a referential complex is not fully isomorphic to the relevant specified qualifier are all cases in which the latter includes S. The reason is simple. If I refer$^+$ to x qua R to myself, then the relevant specified qualifier is $\langle R, \langle S \rangle \rangle$, but the only relevant referent is x, and so the referential complex is merely x. For, 'y refers$^+$ to himself' is not construable as 'y refers to y', and, though the former may entail the latter, that is not explicit in the self-ascriptive 'refers$^+$'. In general, if $e' = \langle R, \langle e'_1 \ldots e'_n \rangle \rangle$ and, say, $e_i = S$, then the corresponding referential complex has the form $\langle x, \langle z_1 \ldots z_{n-1} \rangle \rangle$.

Our current concern is to specify, for any e, what references are made, and how they are made, on the conventional utterance of a signal which means F(e). So far we can say that, on such an utterance, there will be a specification e' of e and a referential complex g such that a) what the structure of g is, is determined by e', and b) how the constituents of g are referred to is determined by the constituents of e'. But how determined? It is this relation which we may now define;

I shall call it <u>super-reference</u>.

It will be useful to have a certain notion separately defined beforehand. Suppose $\langle R,\langle e_1 \ldots e_n\rangle\rangle$ and $\langle w,\langle g_1 \ldots g_k\rangle\rangle$ are a qualifier and referential complex which go together on a particular utterance. The leading item of the complex—namely, w—is referred to qua being R to certain items. But the items to which w is related will only sometimes be $g_1 \ldots g_k$ themselves, for each g_i may in turn be complex, and, in that case, w is being related, not to g_i itself, but to the leading member of g_i. The structure of the corresponding member of $\langle e_1 \ldots e_n\rangle$ determines whether it is g_i or the leading member of g_i to which w is referred to as being related. If the former is a property, H, or U, or $\langle R\langle S\rangle\rangle$, then g_i is itself the item in question; otherwise it is g_i's leading member. The function f captures this —

$$f(e,g) = z$$

if, and only if,

(i) if e is a property or H or U, or (for some R) $e = \langle R,\langle S\rangle\rangle$, then $z = g$;

(ii) if (for some $R,y,e_1 \ldots e_n$, and $g_1 \ldots g_k$) $e = \langle R,\langle e_1 \ldots e_n\rangle\rangle$ and $g = \langle y,\langle g_1 \ldots g_k\rangle\rangle$, then $z = y$.

The relation which obtains between a specified qualifier and the referential complex which corresponds to it on a particular utterance may now be defined.

x by his utterance u (audience y) <u>super-refers</u> to g qua e

if, and only if,

e is a specified qualifier, and

(i) e is a uniqueness property and x refers (by u etc) to g qua e; or

(iia) e = H, and (for some R) H(x,y,u) = R, and g = y,

and x refers^{+} to y qua R to x-self; or

(iib) e = U, and (for some R) U(x,y,u) = R, and g = u,

and x refers^{+} to u qua R to x-self; or

(iii) (for some R and $e_1 \ldots e_n$) R is (n+1)-adic, and e = $\langle R, \langle e_1 \ldots e_n \rangle \rangle$, and

a) if S is not among $e_1 \ldots e_n$, then (for some w and $g_1 \ldots g_n$)

g = $\langle w, \langle g_1 \ldots g_n \rangle \rangle$, and x refers to w qua R to $f(e_1,g_1) \ldots f(e_n,g_n)$,

and ($1 \leq i \geq n$) x super-refers to g_i qua e_i, and

b) if S is e_m (for some m from 1 to n, n>1) and if $\langle j \ldots k \rangle$ results

from deleting m from $\langle 1 \ldots n \rangle$, then (for some w and $g_j \ldots g_k$)

g = $\langle w, \langle g_j \ldots g_k \rangle \rangle$, and x refers^{+} to w qua R to x-self and to

$f(e_j,g_j) \ldots f(e_k,g_k)$, and (for each i in $\langle j \ldots k \rangle$), x super-refers

to g_i qua e_i, and

c) if S is the only member of $\langle e_1 \ldots e_n \rangle$, then x refers^{+} to g

qua R to himself.

The referential aspects of the conventional utterance of a signal which means $F(e_1 \ldots e_n)$ may be given thus: on any such utterance (for some $e'_1 \ldots e'_n$) $e'_1 \ldots e'_n$ are specifications of $e_1 \ldots e_n$ and (for some $g_1 \ldots g_n$) the utterer super-refers to $g_1 \ldots g_n$ qua $e'_1 \ldots e'_n$ respectively, except when e'_i = S, in which case the utterer super-refers to n-1 referential complexes qua the remaining members of $\langle e'_1 \ldots e'_n \rangle$.

8. What is Meant on the Conventional Utterance of a Signal

As we have seen, what is meant on the utterance of at least certain kinds of simple signals is partially determined by how the utterance's

referents are referred to. Suppose, for example, that D is a specification of E and that on a certain conventional utterance of a signal which means the E is F the utterer refers to something qua D. His meaning is, then, not merely that F of the referent, but more fully, that a[Fa & Da] of the referent.

What about those properties and relations qua which reference is made to non-primary referents? Are they also determinants of what is meant? They are—and an obvious corollary is that those items which are referred to in non-primary reference belong to the sequence of items of which what is meant by the utterance is meant. As we shall see, this still permits us to draw the distinctions (a) between the primary referents—or subjects of predication—of an utterance, and its non-primary referents, and (b) between what is asserted on an utterance, and how reference is made on that utterance.

I do not know how to prove that every property and relation qua which something is referred to in n-ary reference is part of what is meant. But I shall show that it is true for secondary reference, and I suspect that the reader will be as persuaded as I am that the considerations which are there relevant would apply at any level of reference, however embedded.

Suppose that z means the bride's sister is F, that D and R are uniqueness specifications of the property of being a bride and the relation of being a sister to, respectively, and that, on a certain conventional utterance of z, reference is made to y qua D and to x qua R to y. Now what is meant on that utterance is not merely that

F of x; rather, the utterer means of x that it is F and is R to y. For, exactly the same considerations as applied to the simple case of signals which mean <u>the E is F</u> apply here. What is meant, is meant of x under the relation to y qua which it is referred to. But that is to say that the utterer means ab[Fa & Rab] of ⟨x,y⟩. And from that it follows that something is meant of the secondary referent y.

Moreover, the property D, qua which the secondary referent y is referred to, cannot be left out of what is meant. For someone might mean ab[Fa & Rab] of ⟨x,y⟩, while what he means is not what can be meant on a conventional utterance of a signal which means <u>the bride's sister is F</u>. Here is an example. Suppose William thinks that the bride, the bride's sister, the first violinist, and the first violinist's sister are four people, whereas they are in fact two, the bride being the first violinist, and the bride's (relevant) sister and the first violinist's (relevant) sister being one and the same. Suppose that William, in his ignorance, says "the first violinist's sister is F, but the bride's sister is G", where F and G are independent. It clearly might be the case that the relation R which is, in the first reference, the relevant specification of <u>being a sister to</u> might also be the relevant specification in the second reference. But then William has meant ab[Fa & Rab] of ⟨x,y⟩, and he has referred on that utterance to y qua some specification of being a bride, and to x qua R to y, where R is a specification of being a sister to. Nevertheless, what William has meant here is not what could be meant on a conventional utterance of a signal which means <u>the bride's sister is F</u>.

The remedy should be clear; where the relevant specifications and referents are as above, the conventional utterer of such a signal means ab[Fa & Rab & Db] of ⟨x,y⟩. What is meant of the primary referent, is meant of it under that relation to the secondary referent qua which the former is referred to. And what is meant of the latter is meant of it under the property qua which it is referred to. I can think of no reason why it should not happen that, at each level of reference, similar considerations would show the necessity for (a) incorporating the item of reference into the sequence of items of which something is meant, and (b) incorporating the property or relation qua which that item is referred to into the whole propositional function which constitutes what is meant.

If we are to generalize this relation between how reference is made on an utterance and what is thereby meant, we shall need precise definitions of two relations—namely, first, the function which goes from the specified qualifiers $e'_1 \ldots e'_n$ of the conventional utterance of a signal which means $F(e_1 \ldots e_n)$ to the complex propositional function which is meant on that utterance and, secondly, the function which maps the referential complexes of an utterance onto the sequence of items of which what is meant thereby is meant.

These two functions must dovetail in such a way that the sequence of referred to items which is the value of the second function for an utterance parallels the propositional function which is the value of the first function for the predicated property and specified qualifiers of that utterance. Consider a conventional utterance of

"the bride's sister is R to the groom's brother", where F, T, G, and T' are appropriate uniqueness specifications, such that the first specified qualifier is ⟨T,⟨F⟩⟩, and the second is ⟨T',⟨G⟩⟩, and the two referential complexes are ⟨x,⟨y⟩⟩ and ⟨x',⟨y'⟩⟩ (that is, y is the bride, x is her sister, y' is the groom and x' is his brother). The utterer then means abcd[Rac & Tab & Fb & T'cd & Gd] of ⟨x,y,x',y'⟩. Clearly, it is by no means the same thing to mean the same thing of, say, ⟨x,x',y,y'⟩.

The subsequent overall procedure is simple, though the actual definitions are complex. First, P* is defined—it being a function from simple specified qualifiers to propositional functions which are conjunctions of the relations and properties which are embedded at various levels in the former. Secondly I define P, which uses P* and maps an utterance's asserted property or relation and its sequence of specified qualifiers onto the propositional function which is meant on that utterance. And then we have Q, which maps sequences of referential complexes (taking into account the corresponding qualifiers) onto the straight sequence of all the references of the utterance. On the basis of P and F we may then define 'z means $F(e_1 \ldots e_n)$ in G'.

There are two points, both concerning the reflexive functions (H, U and S), which may clarify certain complications in the definitions.

(1) P* is a function not merely from a specified qualifier to propositional functions, but from quadruples consisting of a specified qualifier, a speaker, a hearer and an utterance—hence, 'P*(e,x,y,u)'.

This is because specified qualifiers may contain H or U, and the contribution of these two qualifiers to what is meant on an utterance is determined by the speaker, hearer, and utterance, for which they give, respectively, the speaker's hearer-conception, and the speaker's utterance conception.

(2) Consider a conventional utterance of a signal which means F(H) i.e., you are F , where x, y and u are speaker, hearer and utterance. If H(x,y,u) = R, then, by u, x refers^{+} to y qua ab[Rba] to x-self, and means^{+} ab[Fb & Rba] of y. Notice that the leading operator variable in the latter continues to control the internal position it controls in the former. The reason is this: the self-ascriptive position in a propositional function which is the object of a self-ascriptive propositional attitude is—by the convention adopted in chapter 2—always the position which is controlled by the first operator variable. This introduces a (mainly superficial) complication into the definitions of P* and P. If H, U or S occurs at some level of embedding, then that position in its corresponding propositional function which is controlled by the first operator variable must continue to be controlled by the first operator variable in the larger propositional function into which the former is incorporated.

A) "P*(e,x,y,z)"

(i) If e is a property, then P*(e,x,y,z) = e;

(iia) (R) if e = H, and H(x,y,z) = ab[Rba], then P*(e,x,y,z) = ab[Rba] ;

(iib) (R) if e = U, and U(x,y,z) = ab[Rba], then P*(e,x,y,z) = ab[Rba] ;

(iii) $(R)(e_1)\ldots(e_n)$ <u>if</u> $e = \langle R, \langle e_1 \ldots e_n \rangle\rangle$, and $e_h \ldots e_k$ are those of $e_1 \ldots e_n$ which contain (at some level) H, U, or S but are not identical with S, and $e_r \ldots e_s$ do not contain H, U or S, and e_i = S, and

(q) if q is among h...k, then $P^*(e_q,x,y,z) =$

$ab_1 \ldots b_{m_q}[T_q(b_1 \ldots b_{j_q-1}, a, b_{j_q+1} \ldots b_{m_q})]$, and

(p) if p is among r...s, then $P^*(e_p,x,y,z) =$

$b_1 \ldots b_{m_p}[T_p(b_1 \ldots b_{m_p})]$,

<u>then</u>, $P^*(e,x,y,z) =$

$$acb_1^1 \ldots b_{m_1}^1 \ldots b_1^n \ldots b_{m_n}^n [R(c, b_1^1 \ldots b_1^{i-1}, a, b_1^{i+1} \ldots b_1^n) \;\&$$

$$T_h(b_1^h \ldots b_{j_h-1}^h, a, b_{j_h+1}^h \ldots b_{m_h}^h) \;\&\; \ldots \;\&\; T_k(b_1^k \ldots b_{j_k-1}^k, a, b_{j_k+1}^k \ldots b_{m_k}^k) \;\&$$

$$T_r(b_1^r \ldots b_{m_r}^r) \;\&\; \ldots \;\&\; T_s(b_1^s \ldots b_{m_s}^s)].$$

In definition A), the propositional function which is the value of P* in the recursive case (clause iii) has these features. The first conjunct ("$R(c,b_1^1,\ldots)$"), corresponds to the (n+1)-adic relation which is the lead in the specified qualifier. The 'c' position is the position which applies to the object which is the referent of the whole qualifier. In the same conjunct, the i-th variable position is controlled by 'a', the first operator variable of the whole, just in

case S is the i-th member of $\langle e_1 \ldots e_n \rangle$. In the conjuncts '$T_h(\ldots)$' to '$T_k(\ldots)$', '$_{j_h}$', '$_{j_k}$' number the positions which are controlled by the first operator variable 'a', this being determined by which position is controlled by the first operator variable in the propositional functions which are the values of $P^*(T_h,x,y,z),\ldots,$ and $P^*(T_k,x,y,z)$.

B) "$P(R,\langle e_1 \ldots e_n \rangle,x,y,z)$" may now be defined.

$(R)(e_1)\ldots(e_n)$ <u>if</u> $e_h \ldots e_k$ are those of $e_1 \ldots e_n$ which contain H or U or S but are not identical with S, and $e_r \ldots e_s$ do not contain H or U or S, and e_i = S, and

(q) if q is among h...k, then $P^*(e_q,x,y,z) =$

$ab_1 \ldots b_{m_q}\left[T_q(b_1 \ldots b_{j_q-1},a,b_{j_q+1} \ldots b_{m_q})\right]$, and

(p) if p is among r...s, then $P^*(e_p,x,y,z) =$

$b_1 \ldots b_{m_p}\left[T_p(b_1 \ldots b_{m_p})\right]$,

<u>then</u> $P(R,\langle e_1 \ldots e_n \rangle,x,y,z) =$

$$ab_1^1 \ldots b_{m_1}^1 \ldots b_1^n \ldots b_{m_n}^n \left[R(b_1^1 \ldots b_1^{i-1},a,b_1^{i+1} \ldots b_1^n) \;\&\right.$$

$$T_h(b_1^h \ldots b_{j_h-1}^h,a,b_{j_h+1}^h \ldots b_{m_h}^h) \;\&\; \ldots \;\&\; T_k(b_1^k \ldots b_{j_k-1}^k,a,b_{j_k+1}^k \ldots b_{m_k}^k) \;\&$$

$$\left.T_r(b_1^r \ldots b_{m_r}^r) \;\&\; \ldots \;\&\; T_s(b_1^s \ldots b_{m_s}^s)\right].$$

In definition B, the first conjunct, 'R(...)', is (speaking roughly) that part of the whole meaning which might be regarded as what is <u>asserted</u>. If S is the i-th among the primary specified qualifiers of the utterance, then the i-th position in 'R(...)' is controlled by the first variable of the whole.

The final preliminary to our definition of meaning for indicative signals is the definition of that function which, as it were, lines up all the referents of an utterance—that is, all constituents, at any level, of the referential complexes of the utterance—into the sequence of items of which what is meant by the utterance is meant. This is relativized to the corresponding specified qualifiers, for the simple reason that whether a referential complex is itself a referent or merely a collection of referents is determined by the qualifier. So we have the following definition of "$Q(\langle e_1 \ldots e_n \rangle, \langle g_1 \ldots g_k \rangle)$". The notation "$s_1+\ldots+s_n$" means "the sequence consisting of the members of s_1 in order, then the members of ..., and then the members of s_n in order."

C) $Q(\langle e_1 \ldots e_n \rangle, \langle g_1 \ldots g_k \rangle) = z$

if, and only if,

(i) $n = 1$, and e_1 is a property or H or U or (for some T) $\langle T, \langle S \rangle\rangle$

and $z = g$; or

(ii) $n = 1$, and $e_1 = \langle R, \langle e'_1 \ldots e'_q \rangle\rangle$ and $g_1 = \langle w, \langle g'_1 \ldots g'_k \rangle\rangle$ and

either (for some i) $e'_i = S$, and deleting i from 1...q leaves $m_1 \ldots m_k$, and $Q(e_{m_j}, g_j) = s_j$,

or no $e'_i = S$, and $q = k$, and $Q(e_j, g_j) = s_j$,

and $z = \langle w \rangle + s_1 + \ldots + s_k$; or

(iii) $n > 1$, and either (for some i) $e_i = S$, and deleting i from 1...n leaves $m_1 \ldots m_k$, and $Q(e_{m_j}, g_j) = s_j$,

or no $e_i = S$, and $n = k$, and $Q(e_j, g_j) = s_j$,

and $z = s_1 + \ldots + s_k$.

I offer the following definition of non-propositional indicative meaning for simple signals.

z means $F(e_1 \ldots e_n)$ in G

if, and only if,

there is in G the convention that (x) (y) (u) if u is x's utterance of z in circumstances in which there are specifications $e_1' \ldots e_n'$ of $e_1 \ldots e_n$, and a sequence s, such that u is free for meaning $P(F,\langle e_1' \ldots e_n'\rangle,x,y,u)$ of s (or if $e_1' = e_n' = S$, u is free for x's meaning$^+$ F of x-self); then there are specifications $e_1' \ldots e_n'$ of $e_1 \ldots e_n$ and (complexes) $g_1 \ldots g_k$ such that x super-refers to $g_1 \ldots g_k$ qua $e_1' \ldots e_n'$,[7] and means (or means$^+$) $P(F,\langle e_1' \ldots e_n'\rangle,x,y,u)$ of $Q(\langle e_1' \ldots e_n'\rangle,\langle g_1 \ldots g_k\rangle)$ (or, if $e_1' = e_n' = S$, x means$^+$ F of x-self).

9. Signals with Imperative Meaning and Other Illocutionary Forces

Some signals have the straightforward imperative meaning which characterizes most English sentences whose grammatical mood is the imperative. And there are signals whose conventional illocutionary force is more specific than that of either indicative or imperative signals; they mean the same as members of a certain important sub-class of performative sentences. Imperatives first.

---

[7] x super-refers to $g_1 \ldots g_k$ qua $e_1 \ldots e_n$ just in case a) if no $e_i = S$, then $k = n$ and x super-refers to $g_1 \ldots g_n$ qua $e_1 \ldots e_n$ respectively; or b) if $e_i = S$, then $k = n-1$, and x super-refers to $g_1 \ldots g_k$ qua $e_1 \ldots e_{i-1}, e_{i+1} \ldots e_n$ respectively.

On the conventional utterance of an imperative signal, the speaker means that his hearer is to do such and such. Whether the speaker advises, requests, entreats, orders, or exhorts the hearer is left undetermined by the conventional meaning of the imperative signal; but to do any of these things in uttering the signal would be to utter the signal with its full conventional force. Before we adopt a style for ascribing meaning to imperative signals, it will be useful to have a clear idea of the logical form of 'x means that his hearer is to do such and such.'

A speaker's meaning <u>that</u> F of the sequence s (speaker's indicative meaning) consists in there being a certain relation among a speaker, a hearer, an utterance, a propositional function, and a sequence (of referred-to items), this being representable as "M(x,y,u,F,s)". If it were necessary, the preceding definition of the indicative meaning of simple signals could be re-cast in terms of this prim notation.

Similarly, a speaker's imperative meaning is a relation among a speaker, hearer, utterance, propositional function and sequence. The relevant propositional function is of the kind expressed by action-verbs, and its first slot is, in an imperative utterance on which it is meant, applicable to the hearer—that is, to the intended agent. The latter plays a kind of dual role as hearer/intended agent, and in the following notation he gets a mention in both capacities—that is, as audience and as the first member of the intended sequence. The forms of imperative meaning attributions are, then, these:

$M!(x,y,u,a_1 \ldots a_n[Fa_1 \ldots a_n], \langle y,x_1 \ldots x_{n-1} \rangle)$, and

$$M!^{+}(x,y,u,\ ab_1\ldots b_n\lceil Fb_1\ldots b_{i-1},a,b_{i+1}\ldots b_n\rceil,\langle y,x_1\ldots x_{n-1}\rangle).$$

The second occurrence of 'y' has a mildly otiose air, since the nature of imperative meaning secures that it is the hearer to whom the agent-slot of the propositional function applies. Nevertheless, this unorthodox repetition insures that there will always be a final sequence for the relation, (as would not otherwise be so on utterances like "Jump!) and so allows a compensatory economy in the number of styles of imperative meaning attribution.

For an analogous reason, I shall carry over this repetitive feature into the following form for attributing timeless meaning to imperative signals, by keeping explicit the 'H', which might seem otherwise otiose. So we have: 'z means $(M!)F(H,e_1\ldots e_n)$'.

If this is true of some signal z, then that signal means the same as an English sentence of the form "Do F (with regard) to $e_1\ldots e_n$".

It is important to keep in mind that, as a propositional function, F may have any degree of complexity of internal structure; for example, it may be hypothetical, or it may be expressable as (roughly) abc⌈when b happens, then Rac⌋. I shall provide an account of the structure of propositional functions in the next chapter; the point now is that it is not necessary, in generalizations about the meaning of signals, which are semantically unstructured, to employ any but the current bland manner of referring to the propositional functions assigned to them one by one.

An agreeable feature of this manner of attributing meaning to

imperatives is that the sentence "Do F to $e_1 \ldots e_n$" means $(M!)F(H,e_1 \ldots e_n)$, and the sentence "you will do F to $e_1 \ldots e_n$" means $F(H,e_1 \ldots e_n)$. The propositional (as it were) part of their meaning is the same; they differ only in mood.

Now I shall not give a separate definition of the meaning of imperative signals, since, first, an adequate definition emerges from making the obvious changes in the preceding definition of meaning for indicative signals, and, secondly, imperative signals, like (I hope) all signals, are covered by the subsequent general definition.

Some signals are conventional devices for performing specific kinds of illocutionary act—for example, for promising that p, urging that q, and warning that r. Their conventional force is, roughly speaking, the same as that of corresponding performative sentences of the form "I hereby F that ____", where in the place of 'F' there occurs a verb which names the type of illocutionary act thereby performed. I shall consider the fact that a signal has such a conventional force an aspect of its timeless meaning.

Now it is not true of every signal which is a conventional means for performing a specific kind of illocutionary act that its timeless meaning is <u>fully</u> captured either in something of the form 'z means $F(e_1 \ldots e_n)$' or in something of the form 'z means $(M!)F(e_1 \ldots e_n)$'. To make this clear, let us first consider certain facts about illocutionary acts.

As Schiffer has argued,[8] every kind of illocutionary act is such

---

[8] <u>Meaning</u>, section IV. 2, pp. 92-104

that, in performing an act of that kind, one either means that such and such, or means that the hearer is to do such and such. Nevertheless, it is not true that each kind of illocutionary act A is such that one A's that, say, p, if and only if one means (or means!) such and such. To complicate matters, there <u>are</u> types of illocutionary act of which this does hold, and so it may be useful to consider an example of each—commanding, and objecting, respectively.

Consider x's commanding that the door be closed. Now there is no F such that x commands the door to be closed just in case x means! that his hearer is to do F; nor is there any bit of indicative meaning by virtue of which x's act is a command. As Strawson pointed out in "Intention and Convention in Speech Acts",[9] in commanding, ordering, requesting or urging that the door be closed, the same thing is meant! in each case. These acts differ not in what is meant but, as Strawson suggested,[10] and as Schiffer has subsequently argued,[11] in the <u>reasons</u> the hearer is intended to have for producing the directed action. It follows from these considerations that there is no F, nor any $e_1 \ldots e_n$, such that z is a conventional device for commanding that the door be closed if, and only if, z means $(M!)F(H,e_1 \ldots e_n)$. The reason is simple; if the opposite were true it would then be possible for a speaker to utter z — a conventional device for commanding — with its full conventional force, without commanding anyone to do anything. For,

---

[9] pp. 454-456

[10] Ibid.

[11] loc. cit.

to achieve the latter, not only must one mean! such and such, but one must also intend one's hearer to have certain kinds of reasons for conforming.

Not every type of illocutionary act is like commanding in this respect. Some—objecting, for example—are of the kind which, modifying a term of Schiffer's, we might call 'meaning-identifiable'. The matter is simple: by uttering z, I object that q, if, and only if, I mean, by uttering z, that q is a reason for not believing the relevant (entertained, suggested, asserted) proposition. That is, there is a certain propositional function F — ab[a is a strong reason for not believing that b] — such that one objects if and only if one means that F of certain things. Consequently, the full timeless meaning of a conventional device for objecting is captured by something of the form 'z means $F(e_1 \ldots e_n)$'.

Now it is essential not to confuse this last point with the quite distinct, and false, proposition that the meaning of a signal which has the same force as the performative "I (hereby) object that q" is to be located in a convention by virtue of which the conventional speaker <u>means</u> <u>that</u> he thereby objects that q. This is a familiar Austinian point, which is well supported by certain arguments of Schiffer's.[12]

Nevertheless, it does not follow that signals which are conventional devices for performing non-meaning-identifiable kinds of act are not correctly characterized by indicative or imperative meaning

---

[12] op. cit., pp. 107-110.

attributions. For example, it seems true that if z is a conventional device for commanding that the hearer is to do F with regard to (the appropriate referents of) $e_1 \ldots e_n$ then by virtue of that fact, z means $(M!)F(H,e_1 \ldots e_n)$. The latter does not give the whole of the timeless meaning of z in this case, but it applies to z by virtue of the kind of convention which governs a signal for commanding; such a convention requires, as it were, all that the simple imperative convention requires, and more.

This may seem to imply what is false—namely that "I hereby command you to do F to $e_1 \ldots e_n$" and "Do F to $e_1 \ldots e_n$" have the same meaning, since it is true of each that it means $(M!)F(H,e_1 \ldots e_n)$. The problem is only apparent, and may be resolved as follows. Though it is true that the former sentence means $(M!)F(H,e_1 \ldots e_n)$, it is not true that that is its <u>full</u> meaning, though it is the full meaning of the latter sentence. In the usual sense of sameness of meaning, two sentences have the same meaning just in case they have the same full meaning.

The <u>full</u> meaning of a signal makes it a conventional device either for performing some non-meaning-identifiable illocutionary act, or for meaning! that the hearer is to do such and such, or for meaning that such and such. These three exhaust the possibilities. Let 'I' range over the class which contains the relation M, the relation M!, and each relation I' such that $I'(x,y,u,F,\langle x_1 \ldots x_n \rangle)$ holds when x by his utterance u, audience y, A's that F of $x_1 \ldots x_n$, where A-ing is a non-meaning-identifiable kind of illocutionary act.

I shall then adopt the following completely general schemata for

ascribing meaning to simple signals: (1) z means (I)F($e_1 \ldots e_n$), and (2) z means <u>(I)p</u>. And if we let 'm' range both over propositions, and over pairs which consist of a propositional function and a sequence of qualifiers, we have this unified general style of meaning-attribution: z means <u>(I)m</u>.

The final, fully general, definition of both propositional and non-propositional meaning for simple signals is as follows.

z means <u>(I)m</u> in G

if, and only if,

either

(1) (for some F and $e_1 \ldots e_n$)m = $\langle F, \langle e_1 \ldots e_n \rangle\rangle$, and there is in G the convention that (x) (y) (u) if x and y are members of G and u is x's utterance of z in circumstances in which there are specifications $e'_1 \ldots e'_n$ of $e_1 \ldots e_n$, and a sequence s, such that u is free for I-ing P(F,$\langle e'_1 \ldots e'_n \rangle$,x,y,u) of s, audience y, (or if $e'_1 = e'_n$ = S, u is free for I^+-ing that F of x-self), then there are specifications $e'_1 \ldots e'_n$ of $e_1 \ldots e_n$ and complexes $g_1 \ldots g_k$ such that x super-refers to $g_1 \ldots g_k$ qua $e'_1 \ldots e'_n$ and x I's (or I^+'s) P(F,$\langle e'_1 \ldots e'_n \rangle$,x,y,u) of Q($\langle e'_1 \ldots e'_n \rangle \langle g_1 \ldots g_k \rangle$), audience y (or if $e'_1 = e'_n$ = S, x I^+'s that F of x-self);

or

(2) (for some p) m = p and there is in G the convention that (x) (y) (u) if x and y are members of G and u is x's utterance of z in circumstances in which u is free for I-ing that p, audience y, then x, by uttering u, I's that p, audience y.

CHAPTER IV

SENTENCE MEANING

1. The Convention which Governs the Meaning of Sentences

Unlike simple signals, sentences are not governed by conventions which assign them their meanings one by one. Hence the fact that a sentence means such and such is not captured by the definition of meaning latterly presented. There exists no convention among English speakers which specifically governs the use of the string of marks (or vocables) "She brandished her clarinet like a tomahawk". For there is no mutual knowledge among English speakers that that sequence has ever before been uttered—let alone uttered with a certain meaning, nor has there been any agreement among English speakers that that sequence should mean a certain kind of thing.

Moreover, it is not only novel utterances which make it clear that sentences are not governed singly by conventions. Any English speaker is aware, at least in some rudimentary way, that the meaning of the most well-worn sentence depends, not on some special convention governing it, but on the properties of its parts and structural features, which occur with similar roles in other sentences. While simple signals are semantically noncomposite, sentences are semantically structured, or composite.

Being semantically noncomposite is quite compatible with having, in some natural sense, a perceptibly complex physical structure. The

whole fact of the semantical noncompositeness of a signal lies in there being some one convention which assigns meaning to that signal alone. Indeed, the meaning of an utterance-type is noncomposite just in case that utterance-type and that meaning are related by a convention of the kind whose form is given in the definition of meaning for simple signals.

What is it, then, for an utterance type to be composite, or to have semantical structure? Sentences have subsidiary features and parts which they share with other sentences, and the contribution of each such part or feature to the meaning of the diverse sentences in which it occurs is, in some sense, the same.(By a subsidiary feature I mean such different sentential properties as mood, tense, transformational structure, and indeed, any relevant feature of a sentence which does not consist in its having a separable part. The distinction between part and feature is not completely sharp, nor is it in itself of much theoretical importance. The meaning of a sentence is determined by the sum of its relevant parts and features.)

A not uncommon view of the matter is that the semantical compositeness of a sentence lies in its meaning's being a product of the independently constituted and identifiable meanings of its parts and features. So, on this way of looking at the matter, just as the meaning of a simple signal is determined by <u>one</u> convention, which governs it alone, the meaning of a sentence is determined by several or many conventions, each of which governs one of the sentence's semantically active parts and features. So perhaps each word and each structural feature has its own convention, the meaning of a whole sentence resulting from the joint applicability of the several conventions which apply to

its parts.

Now there is a perfectly good sense in which the meaning of a sentence is determined by the meaning of its parts and features. But as we shall see, the concept of word meaning is secondary, and derivative from the concept of sentence meaning. One might say, rather broadly, that sentences are the primary units of language, that they possess meaning by virtue of being conventional means for speakers' meaning or meaning! such and such, or for performing certain non-meaning-identifiable kinds of illocutionary act, and that the meaningfulness of words, phrases, and structural features is an abstraction from the various meanings of the sentences in which they occur. This claim is not <u>obviously</u> true, for it may appear, at first glance, to be incompatible with certain facts about language—the principal one being that words are learned one by one. The compatibility will have to be shown.

The main reason for denying that there are conventions which <u>separately</u> govern words and structural features is that the supposition of their existence is otiose both in explaining the composite nature of sentence meaning, and in stating the conditions under which a sentence is a conventional means for meaning, meaning!, or I-ing that such and such. My primary concern is to give an adequate account of sentence meaning. But that does not mean that other important semantical notions will be neglected, or will not be assimilable to my account.

By a convention, I mean a convention to perform a certain kind of act (or to refrain from performing a certain kind of act) only when

(or whenever) certain circumstances obtain. Consequently, if there are conventions which severally govern words, then to specify such a convention for a particular word w would be to specify what is required of the utterer of w in the relevant circumstances and context.

In order to appreciate the superfluousness of word conventions in the theory of meaning, it is not necessary that we examine all conceivable candidates for such conventions. What I shall do is to consider a certain fairly general form of possible appropriate conventions, and it should be clear that the considerations which are there relevant would be, mutatis mutandis, relevant to word conventions of any other proposed form.

Let us represent word-conventions like this: "Utter w in circumstances C (in linguistic context K) only when ...m...". In the final clause, 'm' stands for whatever is supposed to be the isolable word-meaning with which the convention associates w, by requiring the utterer to do something with regard to m (meaning it, or expressing it, perhaps?) in the utterance of w in the right circumstances. For want of a better candidate, we might suppose that the English convention which governs "chestnut", associates it with the property of being a chestnut, by, say, requiring one who utters it in a certain kind of context to mean by his whole utterance something or other about some chestnut, or about chestnuts, or about chestnuthood. If there were word-conventions, one might reasonably expect them to be like this.

What I shall now do is to investigate the question how conventions of this general type might figure in the determination of the meaning

of a certain class of sentences. Consider these: 'Some man is mortal', 'Some ape is mortal', 'Some man is ferocious', 'Some ape is ferocious'. I shall suppose that these four sentences are composed of only five isolable semantical elements—namely, the words 'man', 'ape', 'mortal', and 'ferocious', and the sentential framework 'some—is...'. Taking the last as a unified constituent can only help the hypothesis that there are word conventions; but, as will become obvious, the simplification could easily be dispensed with.

Now if separate word conventions are relevant to the determination of the meanings of sentences, then something of the following kind ought to be adequate for this set of four sentences.

First, we have conventions which assign meanings to the four words—namely:

1) Utter 'man' in C_1, only when ...a[a is a man]... .

2) Utter 'ape' in C_2, only when —a[a is an ape]—.

3) Utter 'mortal' in C_3, only when ***a[a is mortal]***

4) Utter 'ferocious' in C_4, only when +++a[a is ferocious]+++.

Now the convention which governs the predicative context 'Some—is...' should be such that the speaker whose utterance is in accordance with it, and with 1) to 4), means by "Some ape is mortal" that some ape is mortal, and similarly for the other sentences. There are various ways of effecting this, the following being as simple as any. Let 'w expresses F' be true of w and F if they are related by one of the conventions 1) to 4). We then have:

5) If w_1 expresses E and w_2 expresses F, then utter ⌜Some w_1 is w_2⌝

in circumstances in which the utterance is free for meaning that some E is F, only if you mean that some E is F.

Now it should be clear that 5) does all the real work in assigning meanings to these sentences; for it does not matter exactly what relation is denoted by "expresses". Provided that the right correlations are made, then, regardless of whether 1) to 4) obtain, 5) will adequately determine the meaning of each sentence.

Consider, for example, the following merely enumerative relation. For all w and F, R(w,F) if and only if either w = 'man' and F = a⌈a is a man⌉, or w = 'ape' and F = a⌈a is an ape⌉, or w = 'mortal' and F = a⌈a is mortal⌉, or w = 'ferocious' and F = a⌈a is ferocious⌉. We may now determine the meanings of the four sentences by reference to just one convention—namely,

5') If $R(w_1,E)$ and $R(w_2,F)$, then utter ⌜Some w_1 is w_2⌝ in circumstances in which the utterance is free for meaning that some E is F, only if one means that some E is F.

Conventions 1) to 4) become otiose in accounting for the semantics of this set of sentences. But that is not to say that the four words do not possess striking semantical properties as a result of their contributing to the meanings of the sentences which contain them. For if convention 5') obtains in a certain group, then, by that very fact, a certain semantical relation holds between 'ape' and a⌈a is an ape⌉, and I obviously do not mean merely the trivial relation R. We might say, for example, that w 'connotes' F just in case, for some R', R'(w,F) and 5') holds of R'. 'Connoting' is, as defined, a quite local

semantical relation. I shall return to the question of less localized kinds of word meaning in section 6.

Now just as separate conventions for words are otiose, so are conventions which correspond to particular syntactic structures. Of course, we could go on a bit, introducing more conventions at the level of 5')—for example, one for 'Every...is—'. And by introducing more and more such forms, and increasing the number of words, we might generate a large number of sentences, which are governed by the set of conventions which correspond each to some sentential structure.

But such a set of sentences lacks the recursive features of a natural language, and, curiously enough, once we introduce recursiveness, our set of conventions reduces to one. There is a quite general form of convention which governs sets of sentences which might properly be called languages.

The point can be made simply, by expanding our original set of four sentences to infinity, through the introduction of two simple recursive elements—namely, disjunction and negation. So, as well as containing the four simple sentences, our new set will also contain all strings which satisfy these conditions: the result of connecting two sentences with 'or' is a sentence, and the result of prefixing a sentence with 'it is not the case that' is a sentence. Let the new sentences have their obvious meanings. Now consider the possibility that these meanings are generated by the introduction of two new conventions. We might have, for example:

6) If S is appropriately related to p, and S' is thus related

to q, utter ⌜S or S⌝ in circumstances in which the utterance is free for meaning that p or q, only if you mean that p or q.

7) Similarly for negation.

Now the obvious question is: what is the 'appropriate' relation between S and p? S may itself be a disjunction of any degree of internal complexity, containing embeddings of negations and disjunctions. If the relation in question holds solely by virtue of the obtaining of convention 6) and convention 7), then clearly those conventions are not independent of each other. For, even if we could put them jointly into an appropriate recursive form, they would no longer be conventions which could exist separately. It would, rather, be much more accurate to speak of one complex convention. On the other hand, if the 'appropriate' relation is defined independently of conventions 6) and 7), then those conventions will merely be reduplicating some of its work. For the independent relation will then associate each sentence with the right proposition, and so we should look for a single convention which is based on that relation in something like the way in which 5') is based on R.

The intended effect of convention 5') and putative conventions 6) and 7) may be achieved more simply and perspicuously as follows. We define a certain function f which maps each member of the infinite set of sentences onto its appropriate proposition. The definition presupposes no semantical concepts.

Let x be a sentence just in case x is one of the four originals or results from prefixing "it is not the case that" to a sentence, or

of placing "or" between two sentences. So:

f(x) = p if and only if

(i) for some w_1 and w_2 and some E and F, $R(w_1,E)$ and $R(w_2,F)$, and $x = \ulcorner$Some w_1 is $w_2\urcorner$, and p is the proposition that some E is F; or

(ii) for some sentence y, $x = \ulcorner$it is not the case that $y\urcorner$, and f(y) = q, and p is the proposition that not q; or

(iii) for some sentences y and z, $x = \ulcorner y$ or $z\urcorner$, and f(y) = q and f(z) = r, and p is the proposition that q or r.

On this basis, the following convention is adequate for determining the meaning of every sentence of the infinite set:

8) (x) (p) if f(x) = p utter x in circumstances in which the utterance is free for meaning that p, only if you mean that p.

The separate conventions are thus shown to be otiose. Nevertheless, the words and syntactic structures of these sentences possess significant semantical properties by virtue of convention 8)'s obtaining. Moreover, the convention suggests the form of all conventions which govern propositional languages whose sentences are unambiguous—for example, any standard first-order language. For such languages, meaning is defined as follows:

z means p in G

if, and only if,

for some f, f(z) = p, and

there is in G the convention that (x)(y)(u)(w)(q) if f(w) = q,

and x and y are members of G, and u is x's utterance of w in circumstances in which u is free for meaning that q, audience y, then x means by u that q, audience y.

Of course, the general definiendum for sentence meaning will not be 'z means p', but rather 'z means (I)m'. And since sentences may be ambiguous, there will not generally be a function which maps the sentences of a language onto their meanings. We may, rather, speak more accurately of there being, for each language, a certain recursively specifiable (abstract) relation which holds between each sentence of the language and some pair, or pairs, of the form $\langle I,m\rangle$. The following, then, defines sentence meaning.

z means (I')m in G

if, and only if,

for some R, $R(z,\langle I',m\rangle)$ and

there is in G the convention that $(x)(y)(u)(w)(p)(I)(F)(e_1)\ldots(e_n)$ if x and y are members of G and i) if $R(z,\langle I,\langle F,\langle e_1\ldots e_n\rangle\rangle\rangle)$, and if u is x's utterance of z in circumstances in which there are specifications $e'_1\ldots e'_n$ of $e_1\ldots e_n$, and a sequence s, such that u is free for I-ing $P(F,\langle e'_1\ldots e'_n\rangle,x,y,u)$ of s, audience y (or, if $e'_1 = e'_n = S$, u is free for I^+-ing that F of x-self), then there are specifications $e'_1\ldots e'_n$ of $e_1\ldots e_n$, and complexes $g_1\ldots g_k$ such that x super-refers to $g_1\ldots g_k$ qua $e'_1\ldots e'_n$, and x I's(or I^+'s) $P(F,\langle e'_1\ldots e'_n\rangle, x,y,u)$ of $F(\langle e'_1\ldots e'_n\rangle, \langle g_1\ldots g_k\rangle)$, audience y, (or if $e'_1 = e'_n = S$, x I's that F of x-self), and ii) if $R(z,\langle I,p\rangle)$, and if u is x's utterance of z in circumstances

in which u is free for I-ing that p, audience y, then x by uttering u I's that p, audience y.

This definition allows for a sentence's being ambiguous; z may mean (I)m in G, and also mean (I')m' in G. For, as I pointed out in the first chapter, the utterance of, say, a two-ways ambiguous sentence will typically be such that, either some explanation overrides the hypothesis that the sentence is being uttered with one of its two conventional meanings, or the hypothesis that it is being uttered with a certain one of its conventional meanings overrides the hypothesis that it is being uttered with the other, or perhaps, both are meant. In each case, the convention is satisfied. But if the signal is free for being uttered with both meanings, and the audience cannot infer which is meant, then the utterance is ambiguous, and violates the convention.

If the compositeness of the meaning of a sentence does not consist in there being separate conventions governing sentential parts and features, in what does it consist? We wish to say of a disjunction 'S or S'' that it has a certain semantical structure. Suppose that we are dealing with a purely propositional language whose sentences are unambiguous. There is then a function f, which is governed by the relevant kind of convention, and there are propositions p and t such that f('S or S'') = the proposition that p or t. Now, the fact that 'or' has the semantical force of disjunction in this sentence is merely the fact that, for all sentences x and y of the same language, and for

all propositions q and r, if f(x) = q and f(y) = r then f(⌜x or y⌝) = the proposition that q or r. The semantical structure of a sentence consists in the totality of such facts about its parts and features. I shall return to this matter in section 6.

2. A Problem Concerning the Nature of Meanings

This definition of sentence meaning presupposes that for each language there is a relation which associates each sentence with those instances of '⟨I,m⟩' which are its meanings in that language. Now there are legitimate questions concerning both the availability of such relations and the nature of those abstract entities which are their relata. Let us, to clarify the issues, confine our attention for the moment to propositional languages all of whose sentences have unambiguous indicative meaning.

For any such language, there exists, according to my definition, a function which maps each of its sentences onto the proposition which is its meaning. An interesting language of this kind will contain infinitely many sentences, with infinitely diverse propositions as their meanings, any first order language with identity being an example. But, if the language is to be humanly learnable, this infinite member of associations of sentences with propositions must be finitely representable—that is, the relevant function must be recursively definable. Where f is such a function, if f(x) = p, then it must be possible to show, in a finite number of steps, that this is so.

If, for example, the language contained only a certain finite set

of atomic sentences, as well as all sentences constructible from them using truth-functional connectives, then, except for the assignment of propositions to atomic sentences, all of the conditions which define the relevant function would be like this: if f(S) = p and f(S') = q, then f(⌜S and S'⌝) = the proposition that p and q. This seems unproblematical until one asks: what does the expression "the proposition that p and q" denote?

That the question is real may be seen in this way. Given only locutions of the form 'the proposition that —', and other vague semi-technical cousins of it, like 'the property (relation) of being —', we have no regular or exact way of describing how propositions might systematically be generated from other propositions. How, for example, might we introduce existential quantification?

Suppose the propositional function F is associated in the right way with the open sentence S. What, then, is the meaning of ⌜(Ev)S⌝? The simple answer that it is the proposition that something is F is inadequate for two reasons. First, S may have free variables other than v, in which case the meaning of ⌜(Ev)S⌝ should be another propositional function. So we ought to have some apparatus for describing the mapping of an n-adic propositional function onto an (n-1)-adic propositional function by the operation which corresponds to existential quantification. But notice that merely introducing some notation which purports to do this, and declaring with a wave of the hand that it does, achieves nothing. In order to understand the relevant operations we need a more precise idea of what propositions and their associates are.

Secondly, even in the case in which S is associated with a monadic propositional function, or property, so that it seems we can say that the meaning of ⌜(Ev)S⌝ is the proposition that something is F, there is still a difficulty. Our rule says something like this, presumably: a) if f'(S) = F, then f(⌜(Ev)S⌝) = the proposition that something is F. Suppose that we have: b) f'(S') = the property of being chubby. How do we get from a) and b) to the desired: c) f(⌜(Ev)S'⌝) = the proposition that something is chubby? I know of no generally valid form of inference which does the trick, given just this unanalyzed material.

No clarification results from the following kind of answer: where f'(S) is F, the value of f for ⌜(Ev)S⌝ is the proposition or propositional function which is <u>the existential quantification of F</u>. What is that item? The locution does not provide us with that independent specification of the value of the function which is needed to explain just what the function is.

As we shall see, once we have an account of what propositions and propositional functions are, the general nature of the functions which map sentences onto them becomes quite describable. The difficulty of specifying the actual relevant functions varies considerably with the language involved. To do it for a first-order language would be trivial. But the relevant relation from the sentences of English to the right instances of '⟨I,m⟩' is, of course, quite complex. And there is no reason to suppose that it is capable of elegant statement in completely general clauses. The semantics of English differs from

the laws of physics in being the result of more or less haphazard human activity and invention.

In a propositional language, the meaning of a sentence is what the conventional utterer means by it. In other words, we have this equivalence: z means p if, and only if, any conventional utterer of z means thereby that p. Hence, propositions, insofar as they are the meanings of sentences, are also the items to which propositional attitudes relate us. Consequently, there is another way of posing the question as to what propositions are—namely, what kinds of things are the denotata of the that-clauses which follow propositional attitude verbs? This is a problem about the semantics of a particular part of English, and so, to this extent, the general theory of meaning requires, as a preliminary, that we clarify a certain bit of the structure of English.

The supposition that that-clauses have denotata—are "denoting phrases"—may seem peculiar. Nevertheless it has been a central assumption in modern logical theory. Frege and Church, for example, take that-clauses as denoting senses; and writers of a more nominalist persuasion, like Quine, take them as denoting their embedded sentences or sentence-tokens.

The issues here are not completely clear, but they center around the need for attributing logical form to sentences. Davidson has recently put forward a view of logical form which clarifies matters considerably. Although Frege, Carnap and (until recently) Quine have

not formulated the problem as Davidson does, one might take Davidson's suggestion as explicating at least partially their intuitive notion of logical form.

Instead of speaking of the meanings of sentences, Davidson, fearing that no systematic sense can be made of meanings, speaks of the truth-conditions of sentences, in Tarski's sense. Let us continue to take into account only sentences which are propositional, that is, sentences which, having no indexicals or similar open referential devices, have fixed truth-values. Whatever else the speaker of English might be said to know about his language, he knows the truth-conditions of its indicative sentences. Moreover, we must suppose that what the speaker knows with regard to truth-conditions is finitely representable—otherwise English is not a learnable language. Now, knowing the truth conditions of a sentence S is knowing something of this form: S is true if and only if —. But to state consistent finite conditions which imply all of the relevant instances of this form—and thus to represent finitely the general knowledge of truth-conditions which the speaker is assumed to possess—presupposes that there are just a finite number of devices, some of them recursive, which determine the truth-conditions of sentences. We have some idea as to what such devices are, namely, predication, quantification, conjunction, negation, and so on.

Now giving the logical form of a sentence, suggests Davidson, is virtually the same as showing how its truth conditions depend upon the application of some combination of those devices which are adequate

for stating truth-conditions for the whole language to which it belongs. Consequently, in logical theory since Frege, there has been a central assumption, which, until Davidson, was not very clear—namely, that to account for the logical form of any class of sentences—for example, those which contain propositional attitude verbs—requires showing how their truth conditions can be generated from the appropriate finite stock of logical devices.

As we all know, the embedded sentences in propositional attitude contexts cannot be treated truth-functionally, nor as the operanda of quantificational operations. There are an infinite number of such sentences, and so we cannot treat the whole verb-phrase (including the that-clause) as merely a predicate of the subject, without breaking further into its structure. For that would make the truth-conditions of those sentences not finitely representable. So, failing truth-functional and quantificational devices, the only remaining candidate in classical logical theory is treating the main verb of propositional attitude sentences—e.g. 'believes'—as a relational expression, and to treat that-clauses as having denotata, on the assumption that we can finitely represent that further relation which obtains between any member of the infinite set of that-clauses and its denotatum. 'Jones believes that p' has—on a certain level—the same logical form as 'Jones kicked the dog'. They differ in the kind of denotational devices used to specify the items to which Jones is being said to be related.

If we accept such a denotational account of the logical role of that-clauses, we can approach the question as to what propositions are

as subject of this dual constraint—namely, that propositions are both the meanings of sentences and the denotations of that-clauses. Notice that there will be an intimate connection between the nature of the function which maps sentences onto their meanings, and the nature of the denotational function which maps the that-clauses of (say) English onto their denotata. The problems are virtually equivalent, since (as Frege first pointed out) the denotatum, or reference, of a that-clause, is the meaning, or sense, of its embedded sentence.

One must be careful though, to keep a proper view of the different places of the two functions in our current concerns; perhaps the matter may be looked at like this. We are attempting to give a general account of what sentence meaning is, for <u>all</u> languages. This account, though is itself presented in a particular language—a mixture of English and certain bits of first-order and set-theoretical notation. Let us call this language our metalanguage, even though there is no one object language which is relevant to our current concerns. We are purporting to quantify over certain entities which are the relata of propositional verbs ('to mean that')—that is, the denotations of that-clauses in our metalanguage. But this assumption about that-clauses is justified only if it is true of our metalanguage itself that there is a certain function which is recursively specifiable and maps that-clauses onto their denotata. But it is obvious that at least as a kind of <u>formal</u> problem, this is equivalent to the question whether there is a function which maps the sentences of English (taken as the object language of a theory of meaning) onto their meanings, and

virtually equivalent to that question for other languages.

One should keep in mind that the notation used in Chapters II and III for propositional function abstraction—'$a_1 \ldots a_n[F(a_1 \ldots a_n)]$'—is rather like a symbolized equivalent of that-clauses for propositional functions. Although its use should be intuitively clear enough there, the current question is virtually equivalent to the question of what the semantics of that notation is to be—that is, how to specify finitely the denotata of all such expressions. What follows will, in effect, justify the previous intuitive use of that notation.

3. Propositions and Intensions

What, then, are propositions? Let us begin with a well-worn clue, confining the discussion still to indicative propositional sentences. To know the meaning of such a sentence is to know the circumstances under which it is true. So it might seem that we should identify the proposition expressed by a sentence with the truth-conditions of that sentence, provided that we can get a reified enough notion of the latter. We must be careful with the sense of "truth-conditions" which is here appropriate. For it is clearly not the case that merely to specify conditions under which a sentence is in fact true is sufficient to specify the sentence's meaning. For not every true instance of

S is true if and only if —

gives the meaning of S. As it happens, 'Snow is white' is true if and only if grass is green.

In knowing the truth-conditions of a sentence, one knows what the world would be like, given its truth, regardless of how else things might stand—that is, one knows how to tell (so to speak) worlds in which the sentence is true from worlds in which it is false. On account of this, it has become common in some quarters to identify propositions (truth-conditions) with functions from possible worlds to truth-values. For example, the proposition that snow is white is the function which maps a possible (state of the) world w onto T just in case snow is white in w, and otherwise onto F. This function might then be considered as a candidate for the role of denotatum of the clause 'that snow is white', and for the role of meaning of the sentence "Snow is white" and of all sentences which are synonymous with it. The upshot would be curious; on this view, the denotation function for that-clauses maps them into a certain class of functions, and the meaning-function for each language maps its (propositional) sentences into that class.

Such functions from worlds to truth-values belong to a more general class of functions which may be called <u>intensions</u>. For, the parallel treatment of open sentences (and sentences containing indexicals, and any simple or complex predicative expressions) would identify their meanings with functions from possible worlds to certain sets—in the case of monadic open sentences, to the set of items which satisfy a given condition in that world, and in the case of n-adic open sentences, to the set of n-tuples of items whose members, in the right order, satisfy a given condition in that world. For example, on this view ab⌈a is the brother of b⌉ is a function which maps a world w onto

the set of all pairs ⟨x,y⟩ such that x is the brother of y in w. To know the meaning of "He is his brother" does not require knowing actual fraternal connections, but rather knowing under what circumstances, regardless of how else the world is, that sentence would hold of a given pair.

The identification of propositions with intensions cannot, however, be seriously maintained, at least not if propositions are to be the denotata of that-clauses, and the meanings of sentences. Consider these:

(a) Nicholas believes that some swans are black.

(b) Nicholas believes that some swans are black and the reals are non-denumerable.

It would seem that (1a) might be true although (1b) is false. But the intensions which correspond to the two that-clauses are one and the same. For if a function maps w onto T just in case some swans are black in w then that function is identical with the function which maps w onto T just in case some swans are black and the reals are non-denumerable in w, since the latter condition holds in all possible worlds.

One could be tendentious here and claim that the position of that-clauses is, in such sentences, both referential and non-extensional. Such a defense of the identification of propositions and intensions should, however, be taken seriously only if there were some positive account of the non-extensionality of such occurrences. Notice that these occurrences cannot be assimilated to those discussed in Chapter

II, since the latter occur within that-clauses, and it is that fact about them which enables us to account for their non-extensionality.

So, on certain natural assumptions, the non-equivalence of (1a) and (1b) rules out identifying the proposition denoted by a that-clause with the intension which corresponds to its truth-conditions. Similar considerations show that the meanings of sentences cannot be identified with their corresponding intensions. For, all logically equivalent sentences would then have the same meaning, since the relevant intension for each is the function which maps all worlds onto T. It would be implied that, for example, these are synonymous:

(2a) 2+3 = 5, and

(2b) The reals are not denumerable.

And, finally, for quite parallel reasons, the identification would imply this false sentence:

(3) In uttering the English sentence "Snow is white", the conventional utterer means thereby that snow is white and the reals are not denumerable.

Carnap[1] isolated the problem by introducing the concept of intensional isomorphism. For the contexts ⌜N believes that S⌝ and ⌜N believes that S'⌝ to be equivalent, not only must the embedded sentences S and S' be logically equivalent—that is, have the same associated intensions—but also, S and S' must be intensionally isomorphic. Now two sentences are intensionally isomorphic only if

---

[1] Meaning and Necessity, pp. 56-59.

they are constructed from logically equivalent elements by the same logical operations. It is a rather difficult question just what degree of isomorphism must exist between two sentences for us to count them as synonymous, or between two that-clauses for us to count them as being co-referential, but, as we shall see, we need no exact account of this.

Clearly, the notion of intensional isomorphism does not, in itself, solve our problem, for we still need to show what kinds of items are the denotata of that-clauses. Nevertheless, we now have a suggestive restriction on the candidates for such items. They must be such that ⌜that S⌝ and ⌜that S'⌝ denote the same such item if and only if S and S' are intensionally isomorphic. This immediately suggests that propositions themselves have structures which mirror the semantical structure of the clauses which name them and the sentences which have them as their meanings.

4. Intensional structures

Consider again the kind of function which maps the sentences of a language onto the intensions which embody their truth-conditions. Let us take a simple language, which contains a finite number of atomic sentences $S_1 \ldots S_n$ and all truth functions of them, expressed by the connectives '&' (conjunction) and '-' (negation).

The overall function for this language—call it 'f'—would work like this. In assigning an intension to '$-((S_1\&S_2)\&-S_3)$' f takes into account, so to speak, what its value for the constituent '$(S_1\&S_2)\&-S_3$' is.

If the latter is q, then the intension of the whole is p, which is related to q thus: for every w, if q(w) = F, then p(w) = T, and vice-versa. So the function f maps a whole sentence onto an intension on the basis of two facts: (a) what the main connective of the sentence is, and (b) what the intensions of the chief constituents are.

But it is important to notice that there is also a direct, generalizable, functional relation between q and p themselves, which is stateable without reference to anything linguistic. For we might say that p is the value of the negation function for q as argument. N(q) = p iff for each (w), q(w) = T iff p(w) = F, and q(w) = F iff p(w) = T.

On this basis, there is an alternative way of achieving what f does, but in a slightly more roundabout fashion. Suppose the function f' mapped our sentence '$-((S_1 \& S_2) \& -S_3)$', not directly onto an intension, but onto the structured entity $\langle N,q \rangle$ where q is, again, the intension of '$(S_1 \& S_2) \& -S_3$'. Although f' does not immediately specify the truth-conditions of our sentence, it nevertheless specifies something which itself uniquely determines those truth conditions. The important point is that f' takes into account the same facts as f—namely, what the main connective is and what the chief constituent's intension is—and, moreover, in the course of indirectly yielding what f does, introduces a bit of articulated structure into the entity which it assigns to the sentence. The function f', of course, does not go far enough to provide what we need to solve the problem of intensional isomorphism. But the further step is simple.

Let us say that the "conjunction function", K, is that function which maps pairs of intensions onto intensions like this:

K(r,s) = p iff (w) if r(w) = T and s(w) = T, then p(w) = T;

otherwise p(w) = F.

Using this function and the negation function—which I shall call structural functions—we may now represent the overall intensional structure of the original sentence as being:

$\langle N,\langle K,\{\langle K,\{p_1,p_2\}\rangle, \langle N,p_3\rangle\}\rangle\rangle$

where p_i is the intension of the atomic sentence S_i. I shall call this kind of hierarchically arranged item an 'intensional structure' (or merely 'structure'). Its unwieldiness is only apparent, for the whole class of such entities is trivially definable. The only intensions which actually occur at any level are those which are most embedded. Nevertheless to each constituent of a structure there corresponds an intension which is recoverable by applying its most embedded structural functions to its basic constituents, and then, in turn, the next most embedded functions to the result of the previous step, until the intension of the relevant constituent is reached.

The general class is defined recursively as follows.

x is an intensional structure if, and only if,

(i) x is an intension, or

(ii) there are structures $y_1 \ldots y_n$, and an n-ary structural function Z, such that $x = \langle Z,\langle y_1 \ldots y_n\rangle\rangle$; or

(iii) there are structures $y_1 \ldots y_n$, and an n-ary structural function Z, such that $x = \langle Z,\{y_1 \ldots y_n\}\rangle$.

There are three features of the definition which may require explanation. First, the structural function, and hence also the second member of the structure, may be more than binary, although all of the plausible candidates for English are either binary (conjunction, causal connection) or singulary (existential quantification, specific kinds of adverbial modification). But obviously there could be languages with duodeciary connectives. Secondly, the second members of the structures admitted by (ii) are ordered, while those admitted by (iii) are not. This feature of the definition looks forward to the use to which structures are put as meanings. The point is this: for any sentences S and S', ⌜S and S'⌝ is, on at least one reading, not only logically equivalent to, but also synonymous with ⌜S' and S⌝ (or so it seems to me). Hence the meanings of its constituents should not be ordered in the meaning of the whole. This is to be contrasted with ⌜If S then S'⌝, whose non-equivalence to ⌜If S' then S⌝ explains the ordering of the second member of (ii). Thirdly, notice that the definition does not specify what kinds of intensions are introduced by the first clause. There may be many structures which are not candidates for being the meanings of sentences, but that matters little.

As I mentioned, to each such structure there corresponds what we might call its associated intension, which is defined recursively in the following way.

q is the associated intension of the structure x

('Int(x) = q') if and only if,

(i) x is identical with some intension, and

$q = x$; or

(ii) there are structures $y_1 \ldots y_n$ and a structural function Z such that $x = \langle Z, \langle y_1 \ldots y_n \rangle\rangle$, or $x = \langle Z, \{ y_1 \ldots y_n \} \rangle$ and $q = Z(Int(y_1) \ldots Int(y_n))$

That is, q is the value of the function Z when it is applied to the associated intensions of $y_1 \ldots y_n$.

The notion of an associated intension is important for the simple reason that it yields, for the right structures, the truth-conditions and satisfaction conditions of the sentences to which those structures correspond. Now, since I shall identify meanings with structures, it is an agreeable fact that these structures, or meanings, trivially determine truth-conditions and satisfaction conditions.

It must be clear that we have, in intensional structures, entities whose internal articulation is capable of mirroring the semantical structures of sentences. Hence we may say that, if we can identify them with the meanings of sentences, the problem of intensional isomorphism will be solved—that is, that a) there would be available a recursively specifiable denotation function from that-clauses to structures which will have this effect: for all S and S' there is a structure x which is the denotation of both ⌜that S⌝ and ⌜that S'⌝ if and only if S and S' are, in the required sense, intensionally isomorphic; and for each propositional language with unambiguous sentences there would exist a function from its sentences to structures, such that, for all S and S', S and S' are both mapped onto the same

structure just in case they are intensionally isomorphic.

Suppose y_1 is the meaning of 'snow is white', and y_2 is the meaning of 'the reals are not denumerable'. Without knowing precisely what those entities are, we may safely declare "Snow is white" and "Snow is white and the reals are not denumerable" to be different in meaning, even though they are logically equivalent, since y_1 is not identical with $\langle K,\{y_1,y_2\}\rangle$.

It might be useful, at this point, to have an idea of what kinds of functions might map sentences onto the structures which are their meanings. Let us take the rudimentary language which has $S_1 \ldots S_n$ as atomic sentences, together with all sentences constructible from them by conjunction and negation. Suppose that our overall function f has already mapped each atomic sentence onto a structure. (Whether the meaning of an atomic sentence should in some cases be simply identified with an intension, or should be given a more articulated structure, depends on factors which are discussed in a general way below.) The definition of the function f is then completed in these two clauses:

a) If $f(S) = x$ and $f(S') = y$, then $f(\ulcorner S\&S' \urcorner) = \langle K,\{x,y\}\rangle$; and

b) If $f(S) = x$, then $f(\ulcorner -S \urcorner) = \langle N,x\rangle$.

Notice that all of the real definitional work, so to speak, would go into the definition of the structural functions themselves. They have nothing specifically linguistic in their definition, and so the question naturally arises whether there are such things as semantical universals—that is, entities of this kind which figure in the semantics of all languages.

Let us keep in mind the various functions in our repertory. They are: (a) intensions, which are functions from worlds to extensions; (b) structural functions, which are functions from intensions to intensions; (c) the earlier mentioned overall functions which map each sentence of a language onto some intension (these are no longer of any concern); (d) the 'associated intension of' function, which goes from structures to intensions; and (e) the semantical functions, which map sentences onto the structures which are their meanings.

The question now arises: what constraints are there on the assignment of structures to sentences as their meanings, and to that-clauses as their denotata? The problem is, largely, what to count as the basic constituents of sentences and that-clauses—that is, those elements for which logical equivalence is to be the same as intensional isomorphism.

Consider this proposal. S and S' are intensionally isomorphic provided that their corresponding atomic, unquantified, constituents (sentences, open sentences), are logically equivalent and that they are (in the obvious sense) isomorphic with regard to the way in which their connectives, operators, and so on, build up the whole. Now this suggestion provides neither a necessary nor a sufficient condition of intensional isomorphism. It is not sufficient since more than logical equivalence is required for sameness of meaning at the level of what might plausibly be taken as atomic sentences—consider, for example, '2+2 = 4' and '7+5 = 12'. Nor is the condition necessary, at least if it is true that, for example, certain connectives are strictly speaking

definable in terms of others. Some have wanted to claim, for example, that ⌜If S,S⌝ means the same as ⌜Either not S, or S⌝. Even if this were not allowed, the condition would still be too strong, for it would rule out the possibility of introducing connectives as abbreviations for combinations of connectives.

The issues here are complex. On the one hand it seems natural to suppose that some phrases are synonymous with words—'bachelor' with 'unmarried man'. (But not all cases of phrases being logically equivalent to words are cases of synonymy.) With 'bachelor' and 'unmarried man', we could secure the right equivalence of meaning by having our theory of English assign to 'bachelor' a complex structure as its meaning—namely, the one which gets assigned to 'unmarried man'; and to 'man' and 'unmarried' we might be able to assign not complex structures, but simple intensions—in these cases functions from worlds to the sets of items which satisfy the relevant conditions in those worlds.

But what criteria are there for deciding (a) whether to assign a complex structure or a simple intension to a word, and (b) whether to assign just a simple structural function (say, implication) to a sentential framework like "If...,—", or to assign to it a complex structural framework? It might seem that we have to provide answers to these questions before we may legitimately use the notion of intensional structure in the definition of meaning.

Fortunately, however, the definition of sentence meaning does not require a full treatment of the notion. Consider once again the definition of meaning for languages with propositional, unambiguous, sentences.

Roughly speaking, a sentence of such a language means p just in case its conventional utterer means thereby that p. Nothing in the definition presupposes having an exact account of intensional isomorphism. One might say that the hard problem concerning the notion arises with respect to the semantics of our metalanguage, that is, with respect to the construction of the right denotation function for that-clauses. That function—f'—must satisfy the following condition of adequacy: for all sentences of our metalanguage S and S', and for all propositional attitude verbs V and person-descriptions N

$$f'(\ulcorner \text{that } S \urcorner) = f'(\ulcorner \text{that } S' \urcorner)$$

just in case this is true:

$$\ulcorner N\ V\text{'s that } S \text{ iff } N\ V\text{'s that } S' \urcorner .$$

But if we can be assured that the right such function exists for our metalanguage, then our definition of meaning itself provides the sole constraint on the function which maps the sentences of a language onto those intensional structures which are their meanings. That is, the real definition of intensional isomorphism lies in the construction of the denotation function for that-clauses (or, more precisely, the denotation relation, since that-clauses may be ambiguous), subject to the adequacy condition just given. And so the exact solution depends on a theory of English.

For non-propositional languages with ambiguous sentences—any natural language, presumably—the expected parallel obtains. In the specification of a sentence's meaning as $(M)F(e_1 \ldots e_n)$, the n-ary

propositional function F is to be identified, not with an n-ary intension, but rather with a certain intensional structure. All one can say about such a structure in general, given the information merely that it is an n-ary propositional function, is that, if it has the form $\langle Z, \langle y_1 \ldots y_k \rangle \rangle$, then $Z(Int(y_1) \ldots Int(y_k))$ is an n-ary intension.

Furthermore, the constituents of referential qualifiers themselves must, at least at certain levels of complexity, be identified with structures; that is, the notions "property" and "relation" used to define qualifiers in chapter 3 should be cashed in terms of structures.

If this all seems unmanageably complex, then that is a reflection on the structure of language, and on those intuitions of ours which require intensional isomorphism. The meaning of a sentence—if it is identified with either $\langle I, p \rangle$ (p being a structure) or with $\langle I, \langle F, \langle e_1 \ldots e_n \rangle \rangle \rangle$ (F being a structure and $e_1 \ldots e_n$ being made up of structures)—will be exactly as complicated as we normally take the structure of the sentence to be. The only alternative, it would appear, is a quite different kind of theory of meaning. But I doubt that any alternative theory exists which both takes into account the failure of purely extensional notions to capture our intuitions about meanings, and explains the assignment of meanings to infinitely many sentences on a finitary basis.

5. Frege's Theory Incorporated

Frege proposed that the denotatum of a that-clause is identical with the sense, or meaning, of its embedded sentence. So, where f is the function which maps the sentences of a certain restricted part of

English, which contains only unambiguous propositional sentences, onto their meanings, Frege's theory is the same as the recently suggested equivalence: the denotatum of ⌜that S⌝ is f(S).

The theory has the following further implications, as Frege pointed out. The <u>meaning</u> of a sentence of the form:

⌜A believes that S⌝

is determined by the meanings of its parts, among which is ⌜that S⌝. Now the meaning of ⌜that S⌝ is distinct from its denotatum (this is general Fregean doctrine), and hence distinct from the meaning of S. But somehow it contributes to the meaning of the whole, that is, to f(⌜A believes that S⌝). But the latter may, in its turn, be embedded in

⌜A' believes that A believes that S⌝

and once more we must distinguish between the meaning of ⌜that A believes that S⌝ and the denotatum of ⌜that A believes that S⌝, the latter being the same as f(⌜A believes that S⌝).

So, at each level of embedding we must specify two things if we are to incorporate the Fregean theory into the currently proposed theory of meaning.

(1) What is the meaning, or sense, of ⌜that S⌝ and how is it determined by the meaning, or sense, of S?; and

(2) How does the meaning of ⌜that S⌝ (partially) determine the meaning (and the intension, or truth-conditions) of ⌜A believes that S⌝?

The answers to 1) and 2) must be given in such a way as to be applicable regardless of how many embeddings there are in a that-clause.

Let us begin with sentences of the form ⌜A believes that S⌝ where S does not itself contain any propositional attitude constructions, but rather comes from a (perhaps infinite) part of English for which the meaning function f has already been defined. So we have a denotation function for all such clauses ⌜that S⌝, since the denotatum of ⌜that S⌝ = f(S) — a certain intensional structure.

But what is the meaning, or sense, of ⌜that S⌝? At this stage, I should point out that when I speak of the "meaning" of a certain sentential part which is not itself a sentence (⌜that S⌝ is a singular term), I am not implying that the expression has sentence-meaning in the sense defined in section 1 of this chapter. Rather, I am supposing that, in the quite abstract definition of the overall meaning function f, which maps sentences onto intensional structures, we will associate, in a subsidiary way, certain intensional structures with sentential parts. The correctness of any such assignment is determined entirely by whether, as a result, f effectively maps sentences onto the right intensional structures. So we may speak of "meanings" in this broad abstract sense.

As a step in determining the meaning of ⌜that S⌝, we might first determine what intension should be assigned to it, for, the meaning of ⌜that S⌝, whatever it is, should be such that it determines the intuitively correct intension. Now, in general, the intension of an expression is a function from a possible world to the extension of that expression in that world.

But the extension of a singular term is its denotatum. So, where

⌜that S⌝ = f(S) = p, the <u>intension</u> of ⌜that S⌝ is the function q such that, for all w, q(w) = p. That is, the intension of a that clause maps every world onto some one meaning (intensional structure). Notice that (a) intensions must somehow capture the conditions which determine, for any state of the world, what the relevant expression applies to (or under what circumstances the relevant sentence is true); and (b) the denotatum of a that-clause should remain the same abstract entity regardless of how the world is. One might say that (English being taken as fixed) the denotatum of ⌜that S⌝ is the same in all possible states of the world, while this is not true of ⌜the first man on the moon⌝. The proposed intension for ⌜that-S⌝ meets both these conditions in a simple fashion.

So, with regard to the clause ⌜that S⌝, we have these corresponding items:

(1) the intension of the embedded sentence S;

(2) the meaning, or sense, of S (f(S)). This is a certain intensional structure, and will, at least normally, be distinct from (1);

(3) the denotatum of ⌜that S⌝. This is identical with (2);

(4) the intension of ⌜that S⌝. If f(S) = p, then the intension of ⌜that S⌝ is the function which maps all worlds onto p. This is quite distinct from (1), (2), and (3);

(5) the meaning of ⌜that S⌝. This I shall maintain is identical with (4).

In determining the nature of the meaning of ⌜that S⌝, the important

consideration is of a kind with which we are already familiar—namely, what is necessary for the intensional isomorphism of the sentences in which it occurs. What I shall now argue is this. If ⌜that S⌝ and ⌜that S'⌝ have the same intensions, then

(a) ⌜A believes that S⌝, and

(b) ⌜A believes that S'⌝

are intensionally isomorphic—that is, have the same meaning. And if that is the case, there is no impediment to identifying the meanings, or intensional structures of ⌜that S⌝ and ⌜that S'⌝ with their intensions. As I mentioned in the last section, we must expect this identity at least sometimes; what may be surprising, in this case, is that it should happen with regard to such complex expressions as that-clauses.

The argument goes like this:

(i) If ⌜that S⌝ and ⌜that S'⌝ have the same intensions,

then S and S' are intensionally isomorphic.

(ii) If S and S' are intensionally isomorphic, then

(c) ⌜A' believes that A believes that S⌝, and

(d) ⌜A' believes that A believes that S'⌝,

are logically equivalent.

(iii) If c) and d) are logically equivalent, then

(a) ⌜A believes that S⌝, and

(b) ⌜A believes that S'⌝

are intensionally isomorphic.

(iv) Therefore, if ⌜that S⌝ and ⌜that S'⌝ have the same intensions,

(a) and (b) are intensionally isomorphic.

(All we have claimed hitherto is that, if ⌜that S⌝ and ⌜that S'⌝ have the same denotata, (a) and (b) are logically equivalent.)

Premisses (i) and (iii) are trivial; (ii) is slightly problematic. First, as to (i). If some one function p is the intension of both ⌜that S⌝ and ⌜that S'⌝, then, for all w, p(w) = f(S) and p(w) = f(S'). Therefore, f(S) = f(S'); and so S and S' are isomorphic. Next, as to (iii). If (c) and (d) are logically equivalent, the denotata of their that-clauses are the same, hence their embedded sentences (a) and (b) are intensionally isomorphic.

So what must be shown is that (ii) if S and S' are intensionally isomorphic, then (c) and (d) are logically equivalent. As far as I can see, the prima-facie counterexamples to the equivalence of (c) and (d), on this condition, all depend on confusing, speaking roughly, beliefs that someone believes that a certain thing (certain things) is (are) F, on the one hand, with beliefs that someone believes that the expression "F" describes a certain thing (certain things). For example, it might be maintained that these are not logically equivalent:

(e) Inge believes that Ivan believes that all bishops are in their seventies;

(f) Inge believes that Ivan believes that all bishops are septuagenarians.

For (it might be claimed) (e) would be true with (f) false, if Inge, knowing that Ivan is unlettered, sincerely asserts the embedded sentence of (e) while sincerely denying the embedded sentence of (f).

It seems, however, that Inge (though sincere, etc.) is misrepresenting what is in fact her belief when she (verbally) denies the embedded sentence of (f). What Inge intends to be denying, given her belief as reported in (e), is 'Ivan believes that "septuagenarian" describes all bishops'. Notice that one cannot, under pain of a vicious regress, deny a general distinction between

(h) A' believes that A believes that--- is F, and

(j) A' believes that A believes that "F" describes ---.

It sometimes happens, however, that the embedded clause of (h) is asserted or denied as if it had the meaning of the embedded clause of (j).

Now if (ii) can be maintained, we can then assert that if ⌜that S⌝, and ⌜that S'⌝ have the same intensions, (a) and (b) are intensionally isomorphic and, therefore, we may identify the meaning of ⌜that S⌝ with its intension—namely, the function which maps all worlds onto f(S). On the other hand, if (ii) is not true, some further kind of structured entity will be needed to serve as the meaning of a that-clause. And then the following account of the semantical structure of propositional attitude sentences would have to be modified. But the basic pattern should remain.

So, the crucial questions concerning Frege's theory have a fairly direct answer. The meaning, or sense, of ⌜that S⌝ is the intension just mentioned; and it is determined by the sense of S quite directly, since it is the function which maps all worlds onto f(S). Consequently, if we can recursively specify f(S) for any S,

then we can recursively specify the meaning of ⌜that S⌝ for any S.

Now how does the meaning of ⌜that S⌝ contribute to the meaning of the sentences in which it occurs? Let us take as an example the class of sentences ⌜Someone believes that S⌝.

If f(S) = q, then the meaning of ⌜that S⌝ is p, which is the intension which maps all worlds onto the proposition (intensional structure) q. Suppose that B is the structural function which maps an intension like p onto the intension r (corresponding to ⌜x believes that S⌝), where r maps each world onto the set of persons in that world who believe the corresponding proposition q. At this stage, then, we have: the meaning of ⌜x believes that S⌝ is $\langle B, \{p\} \rangle$, and it follows that its associated intension (satisfaction conditions) is r.

The rest is simple: let Z be the structural function corresponding to 'someone', its definition being straightforward. The meaning of ⌜Someone believes that S⌝ is, on the above assignments, $\langle Z, \langle B, \{p\} \rangle \rangle$ and its associated intension, or truth-conditions, is the function which maps each world onto T just in case someone believes q (i.e., f(S)) in that world.

This sentence may in its turn be embedded in a further propositional attitude context. Since we have effectively determined the meaning of our newly embedded sentence, we can effectively determine the meaning of ⌜that someone believes that S⌝, and then, in the manner just exhibited, we may determine the meaning of the new sentence. And so on, indefinitely.

A word about the structural function B. In chapter 2, I spoke

of propositional belief as being a relation between a person and a proposition. But in that case, we should treat the verb 'believes' as having a relation as its meaning (and hence a certain intensional structure as its meaning) whereas on the current treatment, it is being placed on a par with operators like negation and quantification since it is assigned a structural function. The issues here are quite complex, and ultimately the matter rests on what the simplest theory for this part of English would be. But, suppose we had assigned to 'believes' a certain relation B*. We then have to determine the meaning of ⌜x believes that S⌝. So we would have a certain structural function A which, speaking very roughly, maps the meanings of a verb and its direct object onto the meaning of the verb-phrase; for the kind of case we are dealing with, A is straightforwardly definable. So, then, the meaning of ⌜x believes that S⌝ is $\langle A, \langle B^*, p \rangle \rangle$. The previous B combines the effect of A and B*.

So far I have dealt only with cases in which the embedded sentence in a that-clause is fully propositional. Certainly Frege held that this covered all the cases. There is, nevertheless, a direct extension of the essential Fregean theory to that-clauses which are not propositional. I argued in chapter 2 that that kind of case often occurs.

Let the sentence S mean $F(e_1 \ldots e_n)$. If S were embedded in ⌜x believes that S⌝, the denotatum of the that-clause would be, not a proposition, but a certain propositional function (which is an intensional structure). We cannot always, in such a case, identify the denotatum of ⌜that S⌝ straightforwardly with the propositional

function F. That happens, to return the claims of chapter 2, only when all the referential positions in S remain referential and fully <u>extensional</u> in ⌜x believes that S⌝.

However, to keep matters simple, let us just deal with that case here. The non-extensional kind of case may, in fact, be dealt with in a complicated but specifiable way. On the extensional interpretation, the overall meaning relation for English (to drop the pretense that sentences are unambiguous) will relate S to F—the n-ary propositional function which is (the relevant aspect of) its meaning. Now the meaning of ⌜that S⌝ is the intension which maps each world onto F; call that intension 'p'. So, then, the propositional function which is (the relevant aspect of) the meaning of ⌜x believes that S⌝, is $\langle B,\{p\}\rangle$, which is n+1-ary. Roughly speaking, its associated intension is satisfied by an n+1-tuple $\langle y,x_1 \ldots x_n\rangle$ just in case y believes that F holds of $\langle x_1 \ldots x_n\rangle$. Keeping the same case, consider ⌜N believes that S⌝, where N is a referring expression whose associated qualifier is e'. In the full use of 'means', which takes into account referential qualifiers, that sentence means $\langle B,\{p\}\rangle$ $(e',e_1 \ldots e_n)$. The referring positions in the embedded sentence thus remain (on this extensional construal) fully referential.

Hence, although Frege's explicit theory of that-clauses required them to be fully propositional, the central element in the theory has a direct extension, on which we may deal with that-clauses which contain normally referential occurrences of singular terms.

6. Word Meaning and Logical Syntax

In this concluding section I shall make tentative suggestions regarding the connection between the account I have given of sentence meaning, and the notions of word meaning and logical form. Earlier I maintained that the meaningfulness of sentential parts and features is derivative drom the meaningfulness of the sentences in which they occur. That is, the semantical properties of the words, phrases and syntactic structures of the language of G all flow from the existence in G of a convention governing some overall sentence-meaning relation or function.

It was necessary to develop a systematic and slightly artificial style of attributing meaning to sentences before we could provide an adequate account of sentence meaning. The situation with regard to word-meaning is as much in need of regimentation, with the added complication that there will be no one straightforward locution for assigning meanings to words. Philosophers and logicians have developed a vocabulary of a kind, but it provides at best a fragmentary taxonomy of word meanings: predicates "express" properties, names "denote", and a certain symbol is the existential quantifier. But how, for example, might we assign to adverbs or to words like 'because' their specific semantical properties? We could strictly vindicate the claim that word-meaning is derivative from sentence-meaning, only if an appropriate set of definienda had been developed and we were actually to provide definitions in terms of the convention which governs sentence meaning. Now, in fact, which word-meaning concepts will turn out to be useful in the description of natural languages

will depend upon investigations of their structure. What I shall do now is to indicate what might be some of the relevant considerations concerning their definability.

Let us describe a simple language L, and consider some relations between its sentential parts and certain extra-linguistic items like properties and qualifiers. The atomic sentences of L consist of some one of the expressions 'P_1'...'P_{25}', followed by some one of the expressions 'a_1'...'a_{16}'. L also has recursive conjunction and negation. Now, the meaning of the atomic sentences is specified like this: The meaning function for the group whose language this is maps each atomic sentence $\ulcorner P_i a_j \urcorner$ onto $\langle F_i e_j \rangle$, where each P_i has its own property F_i associated in this way with it and each a_j has its own qualifier e_j associated with it. Each atomic sentence then, for some F and e, means F(e).

Now if one came across a group in which these 25x16 (= 400) utterance types (the "atomic sentences") each had the meaning just described, one would immediately associate each P and each a with the corresponding properties and qualifiers, as if the latter were the meanings of the former "words". One might say that the P's were "predicates", and that each expressed a certain property F, and that the a's were "descriptions" and that each connoted a certain qualifier e. And it may seem possible to define the notions "expresses" and "connotes", as genuine word-meaning concepts, entirely on the basis of the contribution which the separate sentence parts abstractly make

to the meaning of their sentences—and to define these notions, moreover, in a relatively language independent way. So we have:

A. w <u>expresses</u> F in G iff

for some R, f, g and h,

(1) R is governed by the sentence-meaning convention in G

(2) (F')(e)(w')(w")(z) if w' belongs to category h and w" belongs to category g, and f(w') = F and f(w") = e, and z = ⌜w'w"⌝, then R(z,⟨F,e⟩); and

(3) w belongs to h, and f(w) = F.

A quite parallel definition might be constructed for 'w <u>connotes</u> e in G'. And on the basis of these, we might define a <u>predicate</u> as a sentential part which expresses something, and a <u>description</u> as one which connotes something.

Now, the notions "expresses" and "connotes", as defined, hold of each pair P_i, F_i, and a_j, e_j, respectively, and so it may appear that they capture, to some extent, what it is for sentential parts to have their own semantical properties in that full sense in which we regard words in English as having their own meanings. Of course, the notions are a bit local, since they require of a predicate and description that they occur as discrete linear parts, in left and right positions respectively. But that feature of the definition is easily eliminable (as we shall see).

Unfortunately, this kind of specification of abstract correlations between the features of utterance types and certain elements in the meanings of those utterance types is insufficient to constitute the

former as having their own semantical properties. For, it is consistent with what we know of G (though improbable) that the relevant correspondences between word parts, and properties or qualifiers, is quite, as it were, accidental.

Suppose that, in G', the sound sequence 'Oh ha!' means "the deer approaches" and 'Oh ho!' means "the boar approaches". It would clearly be jumping to conclusions to assume that 'oh' means "approaches" (as it were). The coincidence in sound may be irrelevant to the coincidence in the common part of the two meanings. (They may be simple signals, each with a different history.) Nevertheless, by the definition of "expresses", it happens that 'oh' expresses the property of approaching.

The chances of the correspondence being accidental with the four hundred "atomic sentences" of G are slim; nevertheless it is logically possible. The mere largeness of scope of the coincidence is not logically sufficient to establish word meaning.

On the other hand, as I urged earlier, there are no separate conventions which assign meanings to words apart from the overall convention which assigns meanings to sentences. But it would now seem that the mere obtaining of the latter kind of convention is not sufficient to ensure the existence of the separate semantical properties of words—not even when the right correspondences exist on a sufficiently large scale. I shall now suggest the general area in which a reasonable account of word meaning might be found.

Although 'expresses' is not sufficient to capture even one kind of word meaning, it and similar concepts play a crucial role in the

determination of what word meaning is. Suppose that members of G learned the communication system of G described above, by gradually learning (in a rough way, to be sure) that, whatever the total system in G is, certain P's express certain F's in G and certain a's connote certain e's in G. Suppose that it is a psychological fact that they learn the communication practices of their group by means of learning such separate facts as that certain items (some P's and a's) express or connote in the abstract sense defined: in that case, those relationships between sentence constituents and sentence-meaning-elements take on a psychological feature which they do not otherwise have. The point is that the psychological character of these relations turns entirely on how the members of the group learn, and maintain competence in employing, the overall meaning-relation in the group as that is governed by the convention for sentence meaning in the group. Clearly, these remarks are intended only as programmatic preliminaries to a proper account of word meaning.

One important kind of worry concerning the view that the semantical properties of words are derivative from sentence-meaning conventions is that children appear to learn words before they can have any idea of the semantical role of words in sentences. But this important fact is not inconsistent with my general claim. Children do indeed learn that there are striking associations between certain sound-types and certain thing- or event-types. In a sense, one might say that the child knows that "dog" denotes dogs (means dogginess). But a much

fairer and less intellectualist view of infantile achievement is that what the child learns is that there is some important relation or other between vocables and things, this learning being reinforced in various ways. Only when the child begins to learn the use of words in communication does he learn that the striking association of words with types of things is one of meaning. But we, to whom the important association between word and type of thing is intrinsically one of meaning, most naturally report that correlation which the child learns as being meaning. A slightly parallel case: we say that the child learns who various persons' fathers are, but we may not thereby be implying that the child has even the most minimal biological information.

I shall conclude with some rough indications of what logical form and logical syntax are. The semantical relations between words and their meanings named by 'express' and 'connote' (assuming the incorporation of those psychological features of the kind mentioned) are dependent on that very local kind of sentence structure in which predicates occur to the left of descriptions. But, of course, there are other possible arrangements of words which might have the force of predication. Moreover, there is no reason why predicates and descriptions should themselves be discrete linear sentential parts; it would be more appropriate to speak with full generality of sentential features, as well as parts, as being "constituents" of sentences.

If we generalize the semantical property of "expressing" something,

we should try to impose as few physical constraints as possible on the sentential constituents which possess that property. Sentential constituents constitute the sentence, but not always as physical parts. To generalize over all such constituting relations, we may speak of syntactic functions, which map sentential constituents—linear parts or other kinds of features—onto (intermediate sentential constituents, and finally onto) whole sentences. The most obvious such syntactic function is left-right concatenation, which maps x and y onto ⌜xy⌝. So, if we let 'V' range over the class of all possible syntactic functions, we may generalize the definition of 'expresses', as follows:

w expresses F in G iff

For some R, f, g, h and V

(i) The meaning convention in G governs R; and

(ii) (F')(e)(w')(w")(z) if w' belongs to the set of features or parts g, and w" belongs to the set of features or parts h, and f(w') = F' and f(w") = e andz = V(w',w"), then R(z,⟨F',e⟩); and

(iii) w belongs to g, and f(w) = F.

A parallel definition of "connotes" is constructible.

In terms of these defined notions, we may define the categories Pred and Des. A sentential constituent is Pred just in case it expresses something. Notice that the definition of 'expresses' allows this to be true of complex constituents (depending on 'f'). Similarly for 'Des'.

Just as we have, in this abstract way, isolated the semantical contributions of sentential constituents, so we may define semantical properties for the syntactic functions themselves. For example:

V is <u>predicative</u> in G for (constituent-classes) g and h if, and only if

For some R and f

(i) R is governed by the meaning convention in G;

(ii) (F)(e)(w)(w')(z) if w belongs to g, and w' belongs to h and f(w) = F and f(w') = e and z = V(w,w'), then R(z,⟨F,e⟩).

This gives V a significant semantical property only if (a) the classes of constituents g and h for which V is predicative is large enough not to be trivial (for example, the classes of P's and a's in L); and (b) the typical member of G's knowledge of the predicativeness of V plays the right kind of role (of the kind left unspecified) in his learning and control of the meaning-relation in G.

On these assumptions, the notion allows us to describe the logical form of a certain class of sentences in a way which is radically independent of the features of any particular language. We may say that an utterance-type z is <u>Pred (Des)</u> in G just in case there is a g, h, x, y and V such that V(x,y) = z, and x is <u>Pred</u> and belongs to g, y is <u>Des</u> and belongs to h, and V is predicative for g and h in G. It logically follows from a sentence's being <u>Pred (Des)</u> in G that, for some F and e, it means <u>F(e)</u> in G.

We could, in fairly obvious ways, define a vocabulary for describing sentences as being <u>conjunctions</u> of constituents whose

logical form is also there described. And I have no doubt that, with sufficient ingenuity, an adequate logico-syntactic vocabulary could be developed for the description of all English sentences at a level which accords with our intuitions. Such a vocabulary would describe a sentence (a) as being constituted by parts and features each of which are associated by the overall meaning function with such and such aspects of meaning, and (b) as having those parts assembled by syntactic functions which are described only by what role they play in contributing to the meaning of the sentential parts or sentences which are their values. This logico-syntactical vocabulary would be defined entirely in terms of the general kind of convention which governs sentence meaning, together with, of course, a stock of logical and set-theoretical devices.

Here is a guess as to what a logico-syntactical vocabulary would, in part, contain. First, there would be a vocabulary for assigning meanings to word and to certain features -- in these terms the lexicon of a language could be stated; and, secondly, a set of terms for word and feature categories; and, thirdly, a vocabulary for describing the semantical contributions of syntactic functions. In this last class would occur descriptions of the semantical contribution of transformations, which might be viewed as functions which map sentences, as described at a certain high syntactical level, onto sentences.

The differences between formal and logical syntax are radical. In the definition of the overall meaning relation of the kind governed by the sentence meaning convention of a group, sentences are described entirely

in terms of their formal syntactical features, these being ultimately purely physical. Word and feature categories are defined entirely by enumeration since, in natural languages, if not in formal languages, there are no striking physical resemblances which give unity to word categories. Certain functional relations which take on great semantical significance have bleak definitions. Consider the tedium of defining the function which maps count nouns onto their plurals -- 'fish' onto 'fish' and 'goose' onto 'geese'. The real unity and character of grammatical categories and functions is semantical, and derivative from the convention which gives the sentences of a language their meanings.

WORKS REFERRED TO

Carnap, Rudolf. Meaning and Necessity, Chicago; University, 2nd ed., 1956

Grice, H.P. "Meaning", The Philosophical Review, LXVI, 1957, pp. 377-388

_____. "Utterer's Meaning and Intentions", The Philosophical Review, LXXVIII, 1969, pp. 147-177.

_____. "Utterer's Meaning, Sentence Meaning, and Word Meaning", Foundations of Language, 4, 1968, pp. 225-242.

Quine, W.V. "Quantifiers and Propositional Attitudes", The Ways of Paradox, New York, 1966, pp. 183-194.

Schiffer, S.R. Meaning, Oxford University Press, 1972 (forthcoming).

Strawson, P.F. "Intention and Convention in Speech Acts", The Philosophical Review, LXXIII, 1964, pp. 439-460.