1

PROGRAMME FOR A SYNTHESIS

§1. Identifying language

Someone visited the planet Margo, and found that its inhabitants have a language. How could such a discovery be made? We can usually tell at a glance, or at a hearing, whether language is being used by humans; because our species has a physically recognizable sort of behaviour ('vocalizing'), in which vocal cords, tongue and lips collaborate to produce variegated sequences of sounds which we use for linguistic purposes and, as adults, for little else. That is why one can recognize something as a conversation, say, without understanding it. But this basis for recognizing a performance as linguistic is not illuminating: it infers 'They are using a language' from 'They are vocalizing' on the grounds that *most human vocalizing is linguistic*; and this must rest on some other basis for classifying performances as linguistic. What is it?

Here is one answer: 'I just directly know that my own vocal performances are linguistic – there is no basis on which I recognize them as such. Then I can classify some other people's vocalizings as linguistic because they involve my language: I know that you are using language, for instance, because I understand what you say. So I can learn that certain humans vocalize only when using language; others could similarly learn this about other groups; and by assembling all the testimonies we learn that throughout the species most vocal behaviour is linguistic.' That answer, though, teaches us nothing about language, and merely rubs our noses in it. Of course when I use a language I know that I do, but that still leaves me unclear about what this knowledge amounts to, that is, unclear about my concept of language. And of course when you speak I understand; but my proper confidence about this does not greatly help me to understand what understanding is.

Nor does any of this help to explain how we could discover that the Margolese have a language. No one is bilingual as between Margolese and any human language, and the Margolese do not communicate vocally;[1] so how could we learn that they have a language?

[1] See Kurt Vonnegut, Jr, *Breakfast of Champions* (New York, 1973), ch. 5.

PROGRAMME FOR A SYNTHESIS

The Margolese are just a dramatic device. The fact is that we can, in fairly agreed ways, sort out behaviour into linguistic and non-linguistic – not just actual behaviour but also many possible sorts which do not actually occur. I want to know what we go by when we divide performances into the linguistic ones and the others. 'So you plan to explore the concept of language, or the meaning of "language". There may be something to be discovered about this, but is it worth the effort?'

Yes, because we might find that the concept of language is governed by a set of rules which are orderly enough to be manageable, yet complex enough to be interesting. We do have some such rich, disciplined conceptual structures, and they are worth studying in their own right.

Yes, also, because the inquiry is directed not only at 'language' but also at *language* – not just the word but the phenomenon. One's own first language is so familiar and all-encompassing that one cannot easily command a clear view of it; it is hard to resolve the blur of detail into a sharply focussed picture which highlights just those features which mainly serve to make language important to us. One way of developing such a picture is to discover which features of our language are relevant to our agreeing that it is a language.

This, incidentally, could also motivate a careful consideration of whether to count as 'languages' any of the natural communication systems of known non-human animals, or how to classify the supposed languages which have been taught to captive chimpanzees in recent years. So far from being a merely verbal matter, something which 'depends on the definition of language rather than on the observation of what animals do' and can be contemptuously left to 'those who have a taste for such questions',[2] this may involve a serious consideration of aspects of the human condition. When we have the raw behavioural facts about a given chimpanzee, the question of whether we should credit her with having a language forces us to compare her achievements with our own, thus coming to see ourselves more clearly. The endeavour to get straight about chimpanzees, bees, dolphins and the like – which I have reluctantly decided to omit from this book – can spring from the belief that the proper study of mankind is man.

Sometimes the question whether a given behavioural system should be counted as a language is a matter for decision or stipulation rather

[2] Beatrice T. Gardner and R. Allen Gardner, 'Two Way Communication with an Infant Chimpanzee', in A. M. Schrier and F. Stollnitz (eds.), *Behavior of Non-Human Primates*, vol. 5 (New York and London, 1971), p. 181.

than for conceptual discovery; but even then discovery is involved, for one cannot stipulate reasonably unless one has knowledge. The question of whether to count a given system as a language ought to depend upon how it resembles paradigm languages, and that cannot be known without conceptual investigation.

§2. The emphasis on behaviour

I stress linguistic *behaviour* because when we consider extra-terrestrials or animals or one another, outward behaviour such as the making of movements or noises is what we must go by. It may some day become possible to base claims about language on facts about neural events inside its users. Such events, however, are no part of our actual basis for distinguishing the linguistic from the non-linguistic, and indeed we know little about them. They may eventually come to dominate our picture of what language is, but I want to see what picture can be drawn with the materials now available.

Although I shall ignore neural events, I shall attend to mental ones. (Or, if physicalism is true, I shall attend to neural events in their guise as – or considered in their role as – mental events.) But statements about minds are based upon facts about behaviour, and I shall never introduce any mentalistic concept without first displaying its behavioural credentials, saying what sorts of physical behaviour would entitle us to apply it. Throughout the present work, indeed, I shall be interested in mentalistic concepts primarily as aids to intellectually organizing behavioural data, treating mental items as theoretical entities which are postulated by a certain kind of theory to explain behaviour. There is more than that to mental concepts, no doubt; but the 'more', though significant in some parts of the philosophy of mind, seems to be safely negligible in philosophy of language of the kind I am doing.

I therefore give short shrift to the private or inner uses of language, for instance as an aid to ordering one's thoughts. I have to, because I do not know how to handle such phenomena except on the basis of an understanding of the outer uses of language.

Some writers deplore any such behavioural emphasis in the philosophy of language. Chomsky, notably, has attacked those students of language who attend to linguistic behaviour at the expense of what 'underlies' it, as when he writes against Wittgenstein:

If we interpret [Wittgenstein] as merely circumscribing the task of the philosopher, limiting it to a 'purely descriptive method', to descriptions which 'are not hints of

explanations', then . . . he is proposing . . . to concentrate on evidence, . . . putting aside the question of what the evidence is evidence for. The traditional answer to this question was that the observed phenomena constitute evidence for an under- lying mental reality; it would not have surprised any traditional theorist of language or mind that evidence falls into unilluminating networks of family resemblances. It remains to establish the fact that there is some point in restricting one's activities to arrangement of data which are no longer regarded as evidence for the construction of a theory of language or a theory of mind.[3]

There are two targets here. One is Wittgenstein's reluctance to specu- late about the mind, and the other is his refusal to look for illuminating general theories of any sort at all. I sympathize with Wittgenstein's anti-mentalism in the philosophy of language; as I have explained, I shall use mentalistic concepts only under a tight behavioural rein. But I side with Chomsky on the other matter. Wittgenstein has an anti- theoretic streak which invites Chomsky's jibe about 'restricting one's activities to arrangement of data'. The 'networks of family resemblances' into which Wittgenstein sometimes tries to arrange data can indeed be un-illuminating, as Chomsky says, and the passage where Wittgenstein deploys his opinions about them is pertinent to my main theme:

Someone might object against me: 'You take the easy way out! You talk about all sorts of [linguistic activities], but have nowhere said what the essence . . . of language is: what is common to all these activities, and what makes them into language or parts of language . . .' And this is true. – Instead of producing some- thing common to all that we call language, I am saying that these phenomena have no one thing in common which makes us use the same word for all, – but that they are *related* to one another in many different ways. And it is because of this relationship, or these relationships, that we call them all 'language'.[4]

Since Wittgenstein gives no coherent evidence that 'these phenomena have no one thing in common which makes us use the same word for all', this is just dogmatic defeatism. And when, elsewhere, he empha- sizes 'the particular case' as an antidote to 'the craving for generality',[5] he seems to be hostile to something which is of the essence of the life of reason. None of that, however, is implied by the policy of rooting one's theory of language in facts about behaviour. I shall try through my heavily behavioural theory to establish general results of a certain kind – results which will let me, one might say, 'produce something

[3] N. Chomsky, 'Some Empirical Assumptions in Modern Philosophy of Language', in S. Morgenbesser *et al.* (eds.), *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel* (New York, 1969), pp. 280–1.
[4] L. Wittgenstein, *Philosophical Investigations* (Oxford, 1953), §65.
[5] L. Wittgenstein, *The Blue Book* (Oxford, 1964), p. 18.

4

common to all that we call language'. Chomsky does not try to do this. In one place he implies an account of what 'language' is, an account which is quite wrong if offered as analytic of the concept of language; but in a later version the account is addressed to what 'human language' is, which makes it look like empirical description rather than conceptual analysis.[6]

Chomsky has also opposed the behavioural emphasis in the philosophy of language because it focusses on how language is used rather than on how it is structured. Let us examine this antithesis between use and structure.

Any human language is certainly a complex, orderly structure. Its complexity shows in the richness of what can be expressed in it, and its orderliness in the fact that everyone can produce brand-new sentences which others immediately understand, a feat which is miraculous unless we all have at our command some general rules or principles which govern our language. Actually to state the rules is evidently a difficult task; there is no natural language for which linguists have found a complete syntactic and semantic theory embodying all the facts about the grammatical structures in the language and how they contribute to the meanings of its sentences. Indeed, the experts even disagree about what the broad shape of such a theory would be (see §76 below). That hints at the size and complexity of this intellectual area; but it does not indicate the power of the work that has been done in it, under Chomsky's acknowledged leadership, in the past twenty years.

As well as facts about how a language is structured, there are facts about how it is used – the circumstances under which various sentences may be uttered, the purposes of the utterers, the effects on the hearers, and so on – and my main subject-matter lies in this area. Chomsky disapproves. 'If one wants to find out something significant about the nature of language', he writes, 'I think it is important to look not at its uses, which may be almost any imaginable, but, rather, at its structure.' It is futile to examine the uses of language, he thinks, because these 'may be almost any imaginable'. No doubt they may; but perhaps we can bring order into the chaos by mastering some central and basic kinds of language-use, and then with their aid elucidating others. Even if many uses are never covered by the theory, the latter may still have lit up the centre of our problem-area.

For a start, I agree with those philosophers of language who give

[6] For those two references, and for every remaining quotation from Chomsky in this section, see Chomsky (1967), pp. 73–4, and Chomsky (1968), pp. 61–2.

prominence to the use of language in communication between an utterer and an audience – I shall often say 'speaker' and 'hearer', just for euphony and not implying a restriction to vocal behaviour or to auditory intake of it. Furthermore, it seems likely that communication is primarily a matter of a speaker's seeking either to inform a hearer of something or to enjoin some action upon him. (Strictly speaking, 'enjoining' involves *authoritatively* trying to get someone to act; but I shall always use it to cover not just commanding but also requesting, advising and the like; and analogously with 'injunction'. There seems to be no word which is just right for the job, except perhaps 'impera-tive' – but that lacks a corresponding verb, and also has the disadvant-age that it suggests a grammatical form rather than a use of language.) I suggest that we may get a handhold on the uses of language by starting with its uses as a means for information and injunction. Chomsky disagrees:

> There is little reason to believe that the primary use of language is to modify behavior or modify thought. Language can be used for all kinds of other purposes. It can be used to inform or to mislead, to clarify one's thoughts or to show how clever one is; or, in fact, it can be used for play in a very general sense . . . I am not using language any less, if I do not care whether I convince anyone or change anyone's behavior or change his thoughts.

Granted that language has many uses other than to inform and enjoin, why does Chomsky deny that these uses of it are primary? He does not say, and does not even explain what he means by 'primary'.

Perhaps information and injunction are not the most frequent uses of language on our planet, but why should we care about that merely statistical matter? The following seems more important. There could be a community which used language to inform and/or enjoin but for noth-ing else; but it is doubtful if one could describe in detail a community which used language for 'all kinds of purposes' but never to inform or enjoin. If that is right, then those two uses of language are 'primary' in one sense which is clearly relevant to the philosophy of language; and I conjecture that no other language-use has that kind of primacy.

That informing and/or enjoining *are* indispensable to a working communal language is an unobvious claim which needs defending – perhaps by arguing that a language must be used in those ways if it is to be closely enough linked with the world for its users to be able to learn it. But I shall not pursue this, for I do not actually need to claim primacy for informing and enjoining. It is enough for my purposes if

those two uses of language make a good starting point for a certain kind of philosophical inquiry. I am sure they do, and I shall try to demonstrate this by writing a book which goes well because of its well-chosen starting point.

I take my stand, then, on a prediction that a certain kind of inquiry can yield useful results. Chomsky thinks that it cannot. The remarks I have quoted are part of his attack on the view that human languages can be fruitfully compared with systems of animal communication. Directly on that topic, he says:

> If we rise to the level of abstraction at which human language and animal communication systems fall together, then we find plenty of other things incorporated under the same generalizations which no one would have regarded as being continuous with language ... Consider ... the properties of purposiveness, of having syntactic organization, and of being propositional, in a sense, informative. [And in the light of these] consider the common gestures one uses in helping someone park a car. When you indicate to him by the distance between your two hands how far he is from the car behind, your actions are purposive, integrated, and propositional. But, it is unlikely that any significant purpose would be served by studying such gestures and human language within the same framework.

That is, in effect, a prediction. I hope to falsify it in this book, and to show that the questions which Chomsky here sweeps aside as unworthy of study are profitable ones.

We can treat language-use without assuming anything about grammatical structure. Indeed, much of my subsequent inquiry into language-use will concern a communication-system – perhaps it is not a 'language', properly so called – whose utterances lack grammatical structure, but whose uses are significantly like some central uses of language.

What about the converse? Chomsky pictures the investigation of linguistic structure as an independent activity, owing nothing to assumptions about the uses of language. Certainly, one can investigate grammatical structure without explicitly mentioning uses; but in §75 below I shall argue that the investigation of structure does subtly presuppose something about meanings and thus about language-use.

§3. Meaning-nominalism

A study of communication must focus upon what speakers mean by what they say, and that requires us to become clear about how 'what he meant by x' relates to 'what x means'. Consider first the ludicrous

suggestion that no one has ever known what any part of our language really means. Don't say that in that case there would be no such thing as English, for that evades the question: is it conceivable that you and Milton and Henry James and all the rest of us have always been slightly wrong about the meaning of 'don't' and of 'say' and of 'that' and of 'in' . . . and so on through the English lexicon, and also about the real force of every grammatical construction in English? Granted that we know *what we mean by those items*, might we not be wrong about *what those items mean*? Obviously not; and that must be because what an expression means in language *L* is logically connected with what the users of *L* generally mean by it.

But which way does the connection run? Should we take as basic the notion of 'what *x* means' and try to elucidate 'what *U* meant by *x*' in terms of it, or should we reverse that order?

Here is one possible answer:

> We should elucidate 'What he meant was . . .' in terms of 'The meaning of what he uttered was . . .'. When the two fail to coincide, that is always because of some wrong belief the speaker has about the meaning of what he uttered. For instance, when he said, 'The episode just before the cadenza is climatic' he meant that the passage provided a climax, because that is what he thought his sentence meant. So 'what an expression means' is the basic notion, and 'what a speaker means by an expression' can be derived from it with the aid of the concept of belief.

That approach has been favoured by many writers, and one can see why. What a speaker can mean by an utterance is usually restricted to what he thinks it means in the language, and one is tempted to explain that by simply identifying 'what he means by his utterance' with 'what he thinks his utterance means'. But other explanations are possible, including one highly persuasive one which takes 'what he means by *x*' as basic and yet does abundant justice to the fact that ordinarily one can mean by an expression only what one thinks it means.[7]

What can be said in positive support of the view that 'what *U* means by *x*' is more basic than 'what *x* means'? Technical, grammatical arguments have been given for this,[8] but one simpler consideration is the fact that no one has yet succeeded in elucidating 'what *x* means' without appealing to 'what *U* means by *x*'. Davidson has tried, and I shall explain in §79 below why I think he fails; but even he does not base

---

[7] The explanation is Keith Donnellan's, but the fullest presentation of it is in Bennett (1973), pp. 166–8.
[8] Dennis W. Stampe, 'Toward a Grammar of Meaning', *Philosophical Review*, vol. 77 (1968). Stampe's arguments are discussed in Steven Davis, 'Meaning and the Transformational Stew', *Foundations of Language*, vol. 6 (1970).

3. MEANING-NOMINALISM

'what U means by x' upon 'what x means', for he wishes to dispense with the former locution altogether. So the evidence is quite strong that 'what x means' is not the more fundamental of the two.

Even if it is not, however, the other may not be fundamental either. Perhaps each must be explained through the other, so that together they make a tight little conceptual circle, as I shall argue that the concepts of intention and belief do. The state of the literature used to tempt one to draw that conclusion; but the situation changed when H. P. Grice showed how to give a clear sense to 'By uttering x, U meant that P' without implying anything about any language – not that x meant that P in any language, or that U thought it did, or that U thought that his hearers thought it did, or . . . etc.[9] Grice's theory analyses the concept of meaning which is involved in language as distinct from the one involved in 'Those clouds mean rain', but the analysis itself presupposes nothing about language. Indeed, it treats of what an individual means by an utterance on a particular occasion, without reference to what anybody did, does, will or would mean on any other occasion.

So now we can have a meaning-nominalist account of language, that is, one which treats as basic the individual instance of meaning, by one speaker at one time, and gives a derivative status to every kind of general statement about meanings – what the speaker usually means by x, what speakers generally mean by x, what x means in the language, and so on. Grice's work is controversial, and some of its critics object to that aspect of it which I have been praising, namely its meaning-nominalism.[10] I find their criticisms unconvincing; but I choose to answer them only indirectly, by trying to carry out a thorough meaning-nominalist programme.

Some years ago I started an inquiry in the philosophy of language by asking how a certain animal communication-system fell short of counting as a language.[11] Although that work may still be worth reading, I dislike its central question – What mainly marks us off from the brutes? – and my answer to it was wrong (see §§31–3 below). My book's gravest fault, however, was its failure to profit from Grice's insights about meaning. I could not avail myself of Grice, I now see, because I did not begin in a nominalist way. Instead of attending to isolated cases of communication, I started with a communication-

[9] Grice (1957).
[10] See Paul Ziff, 'On H. P. Grice's Account of Meaning', *Analysis*, vol. 28 (1967–8); N. L. Wilson, 'Grice on Meaning: the Ultimate Counter-example', *Nous*, vol. 4 (1970).
[11] Bennett (1964).

9

system which is defined by certain *general* behaviour-patterns of honey-bees. This time I shall do better.

Even if our Gricean foundation is secure, can we find the materials to complete the structure? We are to start with 'what U meant by x on that occasion' and to elucidate 'what x means in L' in terms of that; but how? It seems obvious that the notion of 'what x means in L' is a covertly general one: what x means in language L is connected with what L-users *generally* do or would mean by x. But we cannot simply equate 'what x means in L' with 'what L-users generally mean by x': mere generality is too weak a foundation for the concept of meaning-in-L. For one thing, it lets us say that someone is using x in an unusual way, but not that he is using it wrongly; yet we do have the notion of linguistic malpractice, and a theory of language should account for it. It might be replied that the notion of misuse of language deserves no respect – that it merely reflects our deplorable tendency to condemn unusual kinds of behaviour. I cannot disprove this; but it is *prima facie* implausible, and we ought not to accept it without looking for an account of 'what x means' which will support the notion of wrong uses of language.

There is also a tougher objection to the equation of 'what x means' with 'what people generally mean by x'. A statement of the form 'Whenever anyone in this community utters x he means by it that P', which expresses mere generality, could be true by sheer coincidence: it does not connect one speaker's meaning that P by x with any other speaker's doing so. But it seems clear that we need something non-coincidental. If x is to mean that P in L, then your using it to mean that P must somehow be connected with my doing so. And the connection cannot be merely this: we both use it to mean that P because in our language it means that P. For that reverses the desired explanatory order.

But if mere generality is not enough, what shall we add to it? Some have said that language involves not merely regularities or general practices but rules or conventions. Early in this century, I think, most philosophers of language put their money on 'convention'; but no one explained clearly what a convention is, and suspicions grew that the concept of convention, if we did once become clear about it, would prove inapt for our purposes after all.[12] At about the same time, 'rule' came much into favour, perhaps under the influence of Wittgenstein's

---

[12] See the closing paragraphs of W. V. Quine, 'Truth by Convention', reprinted in H. Feigl and W. Sellars (eds.), *Readings in Philosophical Analysis* (New York, 1949); and William P. Alston, *Philosophy of Language* (Englewood Cliffs, N.J., 1964), pp. 56–8.