Stimulus, Response, Meaning

Jonathan Bennett

from *American Philosophical Quarterly* monograph no. 9 (1975), pp. 55–88.

1. Introduction

It has been thought that the meanings of some utterances might be explained or defined through their roles as responses, or through their roles as stimuli. I shall use the label 'SRM'—for 'stimulus-response meaning-theory'—to name a certain disjunctive view about this. One disjunct, speaker's SRM, says that in some natural language L there are many values of E whose meanings can be expressed in the form: *whenever any mature L-user undergoes a stimulus of kind S, he utters E by way of response*. The other disjunct, hearer's SRM, says that in some natural language L there are many values of E whose meanings can be expressed in the form: *whenever any mature L-user hears E uttered, he does something of kind A by way of response*.

Each half of SRM is in disfavor, and rightly so. Still, I believe that a suitably abstract criticism of SRM can throw light on some aspects of the concept of meaning; and that will be my task in Part I of this paper. Speaker's and hearer's SRM have been well attacked by Paul Ziff and Max Black respectively.[1] I shall follow in their footsteps; but I shall make their case against SRM as abstract as possible, arguing that it still stands, indeed more strongly than ever, when accompanied by great tolerance towards SRM on points of detail.

Throughout Part I the terms 'stimulus' and 'response' will have very thin meanings. Let S be any kind of event which an organism can undergo, the kind being definable in terms of physical and/or sensory events; and let R be any kind of movement an organism can make, the kind being defined purely physicalistically. If x's undergoing an S is very likely to result in its making an R, then for x the S's are stimuli and the R's are responses. Each R is a response to the stimulus which caused it.

So S-R behaviour-patterns are dispositions of the form: when an S occurs in x, it often or always makes an R immediately thereafter. To begin with, I shall attend only to strong stimulus-response links, where x makes an R whenever it undergoes an S. Weak links, where the occurrence of an S raises the probability of an R to something less than 1, will be discussed in §6.

I shall not need in Part I to mention different ways in which organisms can acquire given S-R patterns in their behaviour (e.g. by conditioning). More generally, I shall ignore the fact that 'stimulus' and 'response' are terms

[1] Max Black, *Language and Philosophy* (Ithaca, New York, 1949), chapters 7, 8, 9. Paul Ziff, 'A Response to "Stimulus Meaning"', *The Philosophical Review*, vol. 79 (1970), pp. 63–74.

in a substantial psychological theory. That will become relevant in Part II, when I discuss reasons for wanting to bring the notions of stimulus and response to bear somehow on questions about meaning and language, as a preliminary to examining Quine's ways of bringing them to bear. Quine's position is not SRM, as I have defined the latter; but it will be easier to assess once SRM has been dispatched.

I present SRM only as saying that the meanings of some substantial set of expressions can be described in certain ways, and not as including the further thesis that those meaning-assignments can generate an account of all the meanings in the language. Nor do I take SRM to assert that one can explain in stimulus-response terms what it is for an expression to have a meaning at all. Many stimulus-response meaning-theorists have included these further theses within the scope of their ambitions; but I can afford to let the former of them wait until §11, and I do not want to discuss the latter thesis at all. All attempts so far to explain in stimulus-response terms what it is for something to be meaningful or linguistic, such as Morris's account of 'sign' or Skinner's of 'verbal act', have in my opinion been sad failures.[1] I think they were doomed to fail, because the concept of meaning involves that of intention, and this cannot be elucidated in purely SR terms.[2] But I shall not argue for this view here. Rather, I shall assume that we know or guess that a certain tribe utter certain expressions meaningfully or linguistically, and that our concern is only to establish S-R generalizations which will let us say what those expressions mean. That is how Quine starts his application of stimulus-response to questions about meaning: 'A rabbit scurries by, the native says "Gavagai", and the linguist notes down [a] tentative translation.'[3] That the native is saying something is not questioned.

In what follows, I use 'utterance' as a vague technical term, to cover any bit of behaviour which is linguistic or involves meaning of the same general kind as language does. By the word 'expression' I mean 'type of utterance-product', where 'product' covers whatever is made in uttering, whether it be an inscription etc., or only a noise or movement. The expressions we shalt be concerned with will all be sentence-like, having meanings like those of whole sentences rather than words or phrases. Reasons for this will be given in §11.

1 Charles Morris, *Signs, Language and Behaviour* (New York, 1946), pp. 5–17. B. F. Skinner, *Verbal Behaviour* (New York, 1957), p. 14.
2 See Skinner, *op. cit.*, pp. 144–5.
3 W. V. Quine, *Word and Object* (New York, 1960), p. 29.

Part I

2. Speaker's SRM

Behaviour connects with meaning through understanding. To state in behavioral terms what expression E means, we must describe behaviour which is typical of someone who understands E, knows what E means. So speaker's SRM needs a kind of stimulus S such that

(∀x) (∀t) [(ϕx at t & x undergoes an S at t)
→ (x utters E at T+d)]

is true, where ϕ picks out enough mature members of the tribe to make it plausible to say that they are the ones who understand E. Then a tribesman's understanding E will be shown by his having a certain S-R pattern in his behaviour, namely his uttering E whenever he undergoes an S. That suggests an account of what speakers mean by E: anyone uttering E means by it that an S has just occurred in him.

Someone's undergoing an S will not usually interest his hearers much unless it is evidence for the occurrence of some kind of event outside his body. So we want to be able to say things of this form: Anyone uttering E means by it that an S* has just occurred nearby, where S* is an event-kind such that someone's undergoing an S is adequate evidence that an S* has just occurred in his environment.

A systematic relation between each linguistically relevant S and some external S*, as well as providing the tribe with interesting subject-matter, also enables them to check on one another's linguistic performances and to correct them where necessary. SRM requires this, but I presume that any viable theory of meaning would similarly emphasize uses of language which are 'realist' in the sense of having as their subject-matter events outside the speaker. In Quine's words: 'Language is socially inculcated and controlled; the inculcation and control turn strictly on the keying of sentences to shared stimulation.'[1] And, of course, anything which helps the tribal language-teacher may also help our endeavors to translate the tribe's language.

This implies that to understand the tribe's language we must relate our sensory constitution to theirs. There might be deep differences, making things hard for us; but Ziff's attempt to use this against Quine seems to reflect a misreading.[2] Quine need not limit the differences between our sensory constitutions and the tribesmen's: all he needs is that the differences be known.

Ziff writes as though even that were demanding too much. He asks rhetorically: 'Can a defect of the crystalline lens be of linguistic significance?' I submit that the right answer is 'Yes'. How could we make any intellectual assessment of creatures about whose sensory intake we were ignorant? Presumably Ziff himself would have agreed with this, once, when he wrote: 'To be able to speak and to understand English one must have. . . sensory organs capable of making certain contrasts.'[3]

Summing up, then: speaker's SRM requires situations of the following sort. Expression E is systematically related to event-kind S in the speaker's body, and through that with external-event-kind S*. E relates to S as response to stimulus, S to S* as evidence to that which is evidenced, and E to S* as expression to subject-matter or as report to event-reported. If E is always to be uttered because the

[1] W. V. Quine, *Ontological Relativity and Other Essays* (New York and London, 1968), p. 81.
[2] Ziff, 'A Response to "Stimulus Meaning"', *op. cit.*, pp. 70–71. For next quotation, see ibid., p. 71.
[3] Paul Ziff, *Semantic Analysis* (Ithaca, New York, 1960) §80.

speaker has just undergone an S, and yet to mean that an S* has just occurred, then the S-S* relationship must be generally available to the tribe: they must nearly all know that undergoing an S is adequate evidence for the occurrence of an S* in the vicinity. Also, there must not be any significant number of them who take anything other than the occurrence of an S as evidence that an S* has occurred. These two constraints give one way of expressing Quine's notion of an 'observation sentence', i.e. 'one on which all speakers of the language give the same verdict when given the same concurrent stimulation', with nobody deviating from this because he has relevant 'collateral information'.[1]

3. *Langue* and *parole*

If we had a plausible generalization to the effect that when any ϕ tribesman undergoes an S he then utters E, ought we to relinquish it just because one ϕ tribesman does not conform to it, or because various ϕ tribesmen infringe it on rare occasions? No, we ought not. In general, speaker's SRM can ignore any linguistic behaviour which reflects the speaker's suffering a lapse of memory or a slip of the tongue, or his joking or teasing, or his being linguistically eccentric or badly brought up. To use the terminology of F. de Saussure, speaker's SRM need not worry about any divergences between *langue* and *parole*.[2] I shall explain this.

Langue is the language considered as an independent structure, to which users can relate in various ways, e.g. using it well or badly. *Parole* is what linguistically goes on, complete with bad grammar, faulty memory, mispronunciation, and so on.

Nested within that distinction are two smaller ones—between dialect and idiolect, and between competence and performance. What Saussure called *langue* is the language of some tribe (a dialect), and this can be distinguished from the language of any individual tribesman (his idiolect). But someone's idiolect—the language he has, knows, understands—is a matter of his linguistic competence, and this is distinct from the actual linguistic performance that he puts up, which can involve misuses of his own idiolect. So we have the dialect/ idiolect or many/one distinction, and at the level of idiolect we have the competence/performance distinction. The distinction between *langue* and *parole* is best seen as straddling these two other distinctions: it is the distinction between competence of many (*langue*) and performance of one (*parole*). I propose that we allow speaker's SRM to pretend that there is no important work to be done by the distinction between *langue* and *parole*. Let us take speaker's SRM to be addressed to an idealized linguistic situation which contains no memory-failure, no slips of the tongue, no lying, and in general no phenomena which would give work to the distinction between competence and performance, or that between dialect and idiolect.

This procedure is justified by its results: I hope to show that it is interesting to see what difficulties SRM encounters even when aided by an idealization of the sort described. But the pretense that *parole* perfectly mirrors *langue* can also be justified in a different way, as follows.

We have to approach competence through performance, and dialect through idiolect; and so—combining the two points—all our data about *langue* must consist in facts about *parole*, i.e. the actual linguistic performances of individuals. (These may include reports on their 'linguistic intuitions'; but there must be other data as well, if only so that we can know which performances *are* reports on linguistic intuitions.)

[1] W. V. Quine, *Ontological Relativity*, op. cit., pp. 86–87.

[2] F. de Saussure, *Course in General Linguistics* (New York, 1959).

Now, the early stages of studying a language would be impossible if *parole* diverged too far from *langue*—e.g. if all the speakers under study were liars with bad memories and a poor grasp of the grammar of their language. One needs, to start with, linguistic performances which reflect the facts about the language itself faithfully enough to support some grammatical and other hypotheses which approximate to the truth. Given the beginnings of a set of such hypotheses, one can then start to cope with performances which seem to conflict with them, explaining such conflicts in terms of stupidity or forgetfulness or clumsiness with syntax—or any of the matters which serve to divorce *parole* from *langue*. (Analogously, one might explore the principles of missile-trajectories in a vacuum by studying missiles in the atmosphere, ignoring air-resistance to start with, and then introducing it later so as to make one's principles agree better with one's observations.) But the procedure will be difficult if all those performance-bending factors have to be taken into account from the outset—as though one were to approach the physics of missile-trajectories in a vacuum by experimenting with feathers in the atmosphere. Fortunately, the gap between *parole* and *langue* cannot ordinarily be unmanageably wide, simply because the users of a given language have to understand one another. Someone whose performances were like feathers in the wind would defeat not only translators but also his fellow-tribesmen.

Now, someone may be engaged not in studying any actual language, but in trying to understand what main elements are involved in the study of a language. He may be concerned, as a philosopher, to command a clear over-all picture of what a language is, what knowledge of a language is, and how the concept of behaviour relates to both of these. In that context, it might be legitimate initially to neglect the *langue*/*parole* distinction, pretending that hypotheses about *langue* are just hypotheses about the linguistic behaviour of individual tribesmen. All the complications which divorce *parole* from *langue* can be introduced later, when their introduction will not cloud the picture.

It is in that spirit that I think it is reasonable to let speaker's SRM pretend that we are dealing with an idealized linguistic situation in which *parole* perfectly reflects *langue*. A similar idealization underlies some other approaches to meaning, including Quine's. N. Chomsky says:

> As this is defined [by Quine], the stimulus meaning of a word varies widely with the level of attention, set, gullibility, mood, visual acuity, cortical lesions, etc., while the meaning and reference of a term are independent of these factors.[1]

The factors in question, however, all pertain to the gap between *langue* and *parole*; and Quine's procedure is best seen as involving an idealization according to which that gap is closed.

Whether Chomsky has ever said anything against this sort of idealization, I am not sure. He does endorse 'the classical Saussurian assumption of the logical priority of the study of *langue*',[2] but I do not know what it is for a 'study' to have 'logical priority'. Certainly, nothing I have said conflicts with Chomsky's views about how his interests should be pursued, but I am less sure about some of his animadversions against others' ways of pursuing *their* interests.

[1] Noam Chomsky, 'Current Issues in Linguistic Theory', in *The Structure of Language*, ed. by J. A. Fodor and J. J. Katz (Englewood Cliffs, New Jersey, 1964), pp. 80–81. See also Noam Chomsky, 'Quine's Empirical Assumptions' in *Words and Objections*, ed. by Donald Davidson and Jaakko Hintikka (Dordrecht, 1969), pp.64–65.

[2] Noam Chomsky, 'Current Issues in Linguistic Theory', *op. cit.*, p. 52.

4. Why speaker's SRM fails

Ziff says: 'Viable regularities of the form "If α is the case, then a speaker does β" are simply not to be found.[1] If this were true, even under the idealization I have been defending, then speaker's SRM would be doomed. In fact, Ziff's statement is not quite true as it stands, but we shall see in §12 that the counter-examples to it do not help speaker's SRM. In the meantime, I shall adopt the more cautious claim that there are no viable regularities of the form 'If α is the case, then a speaker does β', where doing β is uttering something which means that α is the case. I shall express this by saying that there are no true S $\rightarrow$ E generalizations, i.e. ones of the form 'Whenever a speaker undergoes an S he utters E' where E means that an S* event has just occurred, this being a kind of event which is systematically related to the occurrence of an S event in the speaker.

Has anyone ever thought that there are S $\rightarrow$ E generalizations? I believe so, although SRM theorists usually stop short of saying so openly. When they do say so, the results are unattractive, as here for example:

> If you want a person to utter the word 'chair', one of the best ways is to let him see an unusual chair. This stimulus will evoke a tact [sc. descriptive comment], because he knows that in the past comments about unusual things have usually led to approval and generalized reinforcement.[2]

Although not many avowals are as open as this, there is much in the SRM literature which shows that SRM theorists have not appreciated how very far short of the truth every nonvacuous S $\rightarrow$ E generalization falls.

The reason why they fall short is simple. If an S $\rightarrow$ E generalization were true, there would be an expression E which every competent speaker would utter whenever he was in a situation of having just undergone an S or just observed an S*, which is to say that he *would* utter E whenever he *could truthfully* utter E. But this is not even approximately right as a requirement for a competent or E-standing speaker, for silence never convicts one of semantic error or ignorance.[3]

SRM is not helped to cope with this difficulty by its freedom to idealize the linguistic situation, pretending that *parole* mirrors *langue*. That idealization, however generously construed, does not imply that every tribesman always utters whatever he could truthfully utter at that time. Taciturnity is not one of the factors that drives *parole* apart from *langue*. Anyway, the abolition of the right to silence, which would be needed to rescue SRM from the present difficulty, is not just unwarrantable but impossible for any language which is not poverty stricken. A potential speaker's stimuli at a given time may entitle him to utter indefinitely many different expressions; he cannot utter them all; and so the S $\rightarrow$ E generalizations concerning the ones he does not utter will all be falsified by his silence about them.[4]

Although there are no S $\rightarrow$ E generalizations, there are—at least in the idealized linguistic situation—many E $\rightarrow$ S generalizations. The reasons for this are instructive. In any language there are values of S and E such that: if a speaker undergoes an S he can truthfully utter E, or it is permissible for him to utter E, or E is thereby verified. (What speaker's SRM tries to do is to capture in a purely behavioral formulation that notion of an utterance's being rendered true

[1] P. Ziff, 'A Response to "Stimulus Meaning"', *op. cit.*, p. 73. See also his *Semantic Analysis*, *op. cit.*, §§46, 54.
[2] George A. Miller, *Language and Communication* (New York. 1951), p. 166.
[3] Thus John R. Searle, *Speech Acts* (Cambridge. 1969), p. 190n.; and David K. Lewis, *Convention: a Philosophical Study* (Cambridge, 1969). p. 160.
[4] Thus P. Ziff, 'A Response to "Stimulus Meaning"', *op. cit.*, p. 67.

or—in the very broadest sense—permissible.) Furthermore, we can plausibly suppose that in any given language there are observation sentences, that is, expressions E such that the only role of E is to report that an S has just been undergone or an S* just observed. If this is right, then we have truths of the form: **(1)** *It is permissible to utter E if and only if an S has just occurred*; from which it follows trivially that **(2)** *If it is permissible to utter E, then an S has just occurred*. Now, we can get from that to something in a strictly behavioral idiom, not using 'true' or 'permissible' etc., by means of the idealization whereby *langue* is identified with *parole*. That idealization pretends that speakers never lie, do not suffer from lapses of memory or slips of the tongue, and in general never speak impermissibly—for in this context impermissible speech is just speech in which something false is said. That amounts to pretending that **(3)** *If E is uttered, then it is permissible to utter E*; and from **(2)** and **(3)** together it follows that **(4)** *If E is uttered, then an S has occurred*; which is just to say that a certain E $\rightarrow$ S generalization holds.

That explains why Ziff is right when he says: 'One can. . . hope to formulate regularities of the form "If a speaker does β, then generally α is the case".'[1] (The word 'generally' there leaves room for all the ways *parole* can misrepresent *langue*.) In a nutshell: we can usefully idealize a language in ways which imply that it exhibits regularities of the form E $\rightarrow$ S, but there is no remotely defensible idealization under which a language exhibits significant regularities of the form S $\rightarrow$ E. This is because we can pretend that whatever is uttered is permissible, but not that whatever is permissible is uttered.

So SRM's fatal defect is that it points the arrow towards the utterance, whereas it should be pointed from the utterance towards the world.

That double claim about the arrow's direction can also be explained in a quite different way. The essential fact is that *utterances are for hearers*. (A solitary man's language can work for him only if he can somehow function both as speaker and as hearer. For example, he may, as 'speaker' at one time, address his later self as 'hearer', relying on writing or on memory of his own linguistic performances.) The primary purpose of a linguistic act is to confront some hearer with an utterance he can use; but to use an utterance is to get from it to something else (the world); and so the hearer needs an arrow pointing from the utterance towards something else. An arrow pointing towards the utterance might help with the minority activity of *predicting and controlling* linguistic behaviour,[2] but in the majority pursuit of *understanding* what is said in linguistic behaviour it is worthless.

5. Amplifying the stimulus

Perhaps we can still point the arrow towards the utterance if we point it from a strengthened stimulus which, in effect, gives the speaker a reason to utter E rather than something else or nothing at all. Schematically, as well as knowing that he could correctly utter E, the speaker must know that he has a hearer whom he wants to hear E being uttered.

Leonard Bloomfield points the arrow towards the utterance: 'Every member of the social group must upon suitable occasion utter the proper speech-sounds. . . . Every child that is born in a group acquires these habits of speech.'[3]

[1] Ibid., p. 73.
[2] See B. F. Skinner, *Verbal Behaviour*, *op. cit.*, p. 3.
[3] Leonard Bloomfield, *Language* (London, 1935), p. 29.

That is true only if the account of an occasion's 'suitability' includes all the needed motivating material; and that, be it remembered, must register itself upon the speaker as a stimulus. Analyzing Jill's request to Jack for an apple, Bloomfield says:

> She was hungry; that is, some of her muscles were contracting etc. Perhaps she was also thirsty. . . The light-waves reflected from the red apple struck her eyes: She saw Jack by her side. Her past dealings with Jack now enter the picture; let us suppose that they consisted in some ordinary relation, like that of brother and sister. . . All these events which precede Jill's speech and concern her, we call the speaker's stimulus.[1]

By these standards of hospitality, the 'speaker's stimulus' can easily accommodate causally sufficient conditions for the utterance. But Jill said something unambiguous meaning 'Get me that apple', and SRM must associate that with a single kind of stimulus S such that the occurrence of an S in a mature tribesman is necessary and sufficient for him then to utter what Jill uttered. Can it be done? A single stimulus-kind may be pretty complex, with a generous allowance of conjunctions and disjunctions; but the sky is not the limit.

Let us grant that every utterance of 'Get me that apple' involves stimuli of kinds H and A, correlated respectively with being hungry and with seeing an apple. Jill's set of motivating data must include as well as H and A, sensory evidence that she has a companion who will get her the apple if (but only if) he hears her utter 'Get me that apple'. This is where 'her past dealings with Jack' come in: she recognizes him as her loving brother who will oblige if. . . etc. But the stimuli underlying that recognition—even if they are described in terms of their relations to past stimuli Jill has undergone—cannot have a place in a general account of the conditions for the utterance of 'Get me that apple'. For often the intended hearer will not relate to the speaker as Jack does to Jill: the possible relationships are legion, and could not possibly be caught in a single generalization.

But the situation of SRM is even worse than that suggests. Let us pretend that all Jill's compatriots are equally helpful to everyone, so that they are equally suitable hearers of 'Get me that apple', whoever says it; let us postulate a stimulus-kind T, correlated with observing a compatriot within earshot; and let us suppose, even more wildly, that no member of the tribe ever gets himself an apple if he can get someone else bring it to him. Then we can say, as the basic fact about 'Get me that apple', that tribesmen utter it when and only when they have just undergone stimuli of the kinds H, A, and T.

Yet even that lavish helping of assumptions and pretences does not rescue speaker's SRM, for it has relied on two special features of Bloomfield's example—features which we cannot suppose, even as an idealization, to be present in some favored, basic, central class of linguistic acts.

The two features are combined in this: Bloomfield's example is a request for food. **(1)** Because it is a request, a suitable intended hearer must be disposed to be helpful; but a suitable hearer for 'There is an apple', say, might be someone who is not helpful but hungry. What stimulus-kind goes with perceiving someone who is hungry? Even the lax standards I have lately adopted will not let me solve the problem by pretending that everyone is always hungry! **(2)** Food answers to a need (hunger) which perhaps manifests itself to the subject through a single stimulus-kind; but other

[1] Ibid., p. 23; see also p. 74.

needs—e.g. for a stick or a pebble—do not at all correlate with stimulus-kinds in this way. As I noted earlier, a 'single stimulus-kind' may be a disjunctively complex affair, but not so complex as to have a disjunct for every kind of stimulus that might betoken a need for a stick! To offer that as part of the common factor in all utterings of 'Hand me that stick' would be playing with words.

Combining the two points: compare 'Give me that apple' with 'There is a stick'. Each involves a stimulus indicating the presence of an apple or a stick. Then the differences start. The apple-request needs stimuli betokening the speaker's hunger and the presence of a helpful hearer: there was no problem about hunger, and we faked a treatment of the hearer as well. What the stick-remark needs, however, is a stimulus-kind corresponding to perceiving someone who needs a stick. And if we stop assuming that speakers always want to satisfy hearer's needs, the required stimulus-kind is even more forlornly impossible; for then it has to correspond to perceiving someone whom one wants to know of the presence of a stick—because one likes him and he needs a stick, or one hates him and thinks he will harm himself with the stick, or. . . . The desiderated stimulus-kind has by now vanished over the horizon. I shall return to one aspect of this in §9.

6. Weakening the link

Can we rescue SRM by weakening the stimulus-response link, so that someone's undergoing an S has only to make him more likely to utter E? Now that the strong-link version of SRM is dead, the weak-link version, which I dropped in §1, ought to be examined.

Skinner makes heavy use of 'probability', and indeed it occurs in one of his few gestures towards SRM. This is when he explains a word's naming a given object (for a given speaker) as that speaker's being more likely to utter the word in the object's presence than in its absence—an explanation which Chomsky has no trouble in demolishing.[1] Chomsky also shows that Skinner is confused about how a response's 'probability' relates to its 'strength', but that is another Skinnerian special feature which we can ignore. Again, Chomsky has often remarked that most sentences have almost zero probability of utterance.[2] But that is consistent with some language's containing a substantial set of sentences—perhaps a few hundred of them—each of which is uttered pretty often. Speaker's SRM aims to express the meaning of each member of that set in terms of the kind of stimulus that raises the probability of the sentence's being uttered, and the question of whether this could be done remains to be discussed.

There will be no hope of success unless, for each sentence in the favored set, there is a uniquely correlated stimulus-kind S whose occurrence makes it correct or permissible to utter E—let us say for short that S *verifies* E. If E means 'That is a stick', then S is the stimulus-kind that goes with perceiving a stick. What SRM in its weak-link variant has to do is to capture the content of 'S verifies E' in a statement about how the occurrence of S in a tribesman makes him more likely to utter E. But what statement? There are just three *prima facie* possibilities.

(1) *Whenever a tribesman undergoes an S, he thereby becomes more likely to utter E.* But if someone undergoes an S, and this is what verifies E, the probability of his then

[1] B. F. Skinner, *Verbal Behaviour*, *op. cit.*, p. 115. Noam Chomsky, Review of Skinner's *Verbal Behaviour*, in *The Structure of Language*, *op. cit.*, pp. 553–554. For Chomsky's next point, see ibid., p. 555.

[2] Noam Chomsky, *Syntactic Structures* (The Hague, 1957), pp. 15–16; 'Quine's Empirical Assumptions', *op. cit.*, pp. 57–58.

uttering E may not rise above zero. For example, his seeing a stick will not make him more likely to utter 'There is a stick' if he thinks his hearer is looking for something to hit him with. To develop this point would be to repeat §5.

(2) *There are circumstances C such that: whenever C obtains and a tribesman undergoes an S, he thereby becomes more likely to utter E.* Although this is true for the S that verifies E, it also holds for other stimulus-kinds as well; and so it will not do. If E is 'That is a stick', the formula holds good if S is the stimulus-kind that goes with seeing a stick, because C can cover everything else—the helpful audience, the need for a stick, and so on. But for that same E the formula equally holds good if S is the stimulus kind that goes with having an itchy back, because we can again let C cover whatever else is needed the helpful audience, the presence of a stick, and so on. So the formula lets us attribute to E the meanings of 'That is a stick' and 'My back itches', and of course endless other meanings as well; not by altering our account of E and the conditions of its utterance, but just by re-applying formula (2) to it in different ways. This could be repeated for any expression whatsoever. So SRM based on formula (2) would imply that every sentence is multiply ambiguous; and that refutes it.

(3) The remaining possibility is to say, for each E and S, *what* the circumstances are under which an S raises the probability of E's being uttered. Since **(1)** 'Under any. . . ' is too strong, and **(2)** 'Under some. . . ' is too weak, we must try for a middle-strength 'Under just these circumstances . . . '. For a given E and the S which verifies it, we may be able disjunctively to specify a set of circumstances C such that: when and only when C obtains, the occurrence of an S raises the probability of E's being uttered. And that statement—in which C is a constant, not a variable—may hold true for E in relation to the S that verifies it and not in relation to any other kind of stimulus. But to implement this proposal we need some way of selecting the right C for a given E, and that requires independent knowledge of what E means. To see that this is so, consider how proposal **(3)** could be expressed in a general formula. It would have to take the form: 'For each E_i identify the S_i such that in circumstances. . . the occurrence of an S_i raises the probability of E's being uttered.' But the blank would have to be filled by a reference to circumstances 'whose absence reduces the probability of E_i's utterance to zero, *though they do not themselves verify* E_i'. That last clause, which gets the relevant material properly apportioned between C and S, has an ineliminable reference to verification (or meaning), and so it does not conform to SRM's program for expressing facts about meanings in a purely behavioral idiom. Indeed, if the program were built upon that sort of statement it would be not just non-behavioral but downright trivial.

So each probabilistic version of SRM is either much too strong, or much too weak, or entirely nugatory.

7. Hearer's SRM

We must consider utterances not just as responses but also as stimuli. As well as my seeing flames and shouting 'Fire!', there is your hearing my shout and running for water. This brings us to hearer's SRM, which aims to express E's meaning as a fact about what anyone who understands E will do upon hearing it. What hearer's SRM needs are values of E for which this holds:

$$(\forall x)\,(\forall t)\,[\,[(\phi x \text{ at } t) \,\&\, (x \text{ hears E at } t)] \rightarrow (x \text{ does A at } t+d)]$$

where ϕ picks out enough mature tribesmen to make it plausible to say that they are the ones who understand E. The obstacles to finding generalizations of this form are significantly different from those which speaker's SRM

encountered.

Speaker's SRM claims **(i)** that for certain values of E there are general truths of the form 'S *verifies* E', where S is a kind of stimulus, and proposes **(ii)** to express these in purely behavioral terms. Subject to certain idealizations, **(i)** is true; but speaker's SRM came to grief over its proposal **(ii)**. Now, hearer's SRM claims **(i)** that there are general truths of the form 'E *dictates* A', where A is a kind of action, and proposes **(ii)** to express these in purely behavioral terms. In this case, if **(i)** were true there might not be too much trouble over **(ii)**. Given that E dictates A (which means that doing A could count as 'acting on' E), it is not obviously absurd to say that in a suitably idealized situation every understanding hearer of E will do A.

At least it is not as absurd as what speaker's SRM must say to make good its proposal **(ii)**, namely that every competent speaker who undergoes an S will utter E. There are clear reasons for this difference. Speaker's SRM requires an injunction of the form: 'At every moment, utter every O sentence which is true at that moment', where O sentences are the ones to which speaker's SRM applies. But if the language is fairly rich in O sentences, this injunction cannot be obeyed. The world is the trigger, and it cannot be prevented from offering stimuli; so the only way to make it possible to obey the above injunction would be to impoverish the language. In contrast with this, the trigger for the hearer's response is always an utterance, and there is no special difficulty about reducing the number of utterances. To create a situation where every utterance can be acted upon, we need not impoverish the language, but have merely to restrain our use of it. So even if we have a rich repertoire of sentences of the sort to which hearer's SRM is to be applied (call them O* sentences), we may be able to obey the injunction 'At every moment act upon every O* sentence which you have just heard.' Furthermore, there is a contrast in respect not just of possibility but also of desirability. There are reasons that might sanely be given for trying to bring it about that every uttered O* sentence is acted upon, and so that might be a feature of an ideal linguistic situation; but what could possibly make it seem desirable that every true O sentence should be uttered?

So hearer's SRM may be able to make a case for its claim that hearers do whatever is dictated by what they hear. Where it runs into trouble is over its claim about expressions dictating kinds of action. A given utterance may dictate a certain action to a given hearer, in the sense that the performance of that action constituted, for that hearer at that time, acting upon what he heard, or doing what was dictated or indicated by what he heard. But hearer's SRM is about expressions, i.e. types of utterance-product; and my formula 'E dictates A' is really short-hand for 'Any utterance of E dictates an action of kind A to any understanding hearer'. There are no true generalizations of this kind for indicative values of E (imperatives will be discussed in §9). I now proceed to explain why.

8. Why hearer's SRM fails

A particular uttering of E may dictate a certain action to a particular hearer. What does this 'dictating' depend upon?

It might depend upon special knowledge on the hearer's part—knowledge not shared by all other hearers. Let us set that aside, just as we earlier set aside S-S* relationships which are not common knowledge throughout the tribe.

What action is dictated by the utterance of E must depend upon the hearer's 'circumstances'—meaning all the facts about his physical capacities, mental aptitudes, environment, and so on. What action is dictated to a hearer by 'Here comes a tiger' will depend whether he is young and armed or old and

unarmed, on whether there is a climbable tree within reach, and so on. There is no escape from this dependence on circumstances, for any utterance; and so any generalization about what E dictates is vulnerable to what I shall call circumstantial variation.

Still, we may be able to accommodate circumstantial variation within the 'single kind of action' which E is said to dictate. That is, we may be able to say something of the form 'E dictates A' by using an 'A' which has the form 'doing B if one is F, doing C if one is G, doing D if one is H, . . . ' and so on, through each relevantly different set of circumstances (F, G, H, . . .) and each associated kind of action (B, C, D, . . .). This procedure, which I have borrowed from Lewis, is not objectionable in principle.[1]

But we cannot even discuss how it might work out in practice, because we have slid around a much larger obstacle that lies in the path of hearer's SRM. This is the following fact: If E is indicative, then its dictating the action A to a given hearer depends partly upon his 'value-system'—by which I mean his likes and dislikes, wants, approvals, ideals, revulsions and so forth.[2] This introduces a fresh range of variations—call them value-variations—over and above those due to differences in circumstances. For example, 'Here comes a goat' may dictate different actions to you and to me just because I like goat's milk and you want a goatskin rug. I suggested that we absorb circumstantial variation by taking each sentence to dictate a 'kind of action' specified by a disjunction in which each disjunct is a circumstance-action pair; but now we need a disjunction of triples, each comprising a value-system and a set of circumstances and an action. Another way of putting it: we need to list the relevant circumstance-action pairs for one value-system, and then for another,. . . and so on through every relevantly different value-system.

This cannot be done. Suppose that we have a disjunctive account of the 'kind of action' that is always dictated by E for any hearer who has a certain value-system—mine, say. Then we can say: 'For anyone who has Bennett's value-system, hearing E dictates that he should do B if he is F, do C if he is G, . . . ' and so on. But we still could not specify all relevantly different value-systems just by ringing changes on mine; for a value-system might differ from mine in ways which demand the mention of circumstances which are not mentioned at all in stating how my value-system relates to E. Let E mean 'Here comes a tiger!', and suppose that the relevant facts about my value-system are that I put my own safety before everything else, other people's safety next, the pleasure of watching wild animals next; I do not like eating tigers; and I do not kill things for fun. (The oversimplified nature of this does not affect my present point.) Now, someone whose value-system differs from mine need not like killing for fun, or dislike watching wild animals, etc. The differences may rather be of this sort: he likes seeing tigers chase white men but not black men; he enjoys steaks from female tigers but not from male ones; he worships teeth taken from tigers killed during rainstorms; and so on. No limit can be set in advance to what value-variations there can be, and so the required list of value-circumstance-action triples cannot be constructed.

Lacking a list, we can fall back on a general characterization which trivializes the whole endeavor: what a hearing of E dictates to any hearer is the action which is suitable for someone with his value-system, in his circumstances, who has just heard E.

[1] D. K. Lewis, *Convention*, *op. cit.*, p. 131.

[2] Ibid., pp. 160–161.

The literature of hearer's SRM does lapse into triviality in just this way. Morris, for example, says that a 'sign' is whatever 'causes a disposition in some organism to respond under certain conditions'.[1] I agree with Black that, according to this, 'Any stimulus which has some causal influence upon subsequent behaviour will. . . have to count as a sign.'[2] But there is also a different point, namely Morris's failure to specify the content of any one of these meaning-constituting 'dispositions'. He explains what dispositions are, and what it is for a sign to change its meaning by changing in respect of what dispositions it sets up in hearers. But he seems not to notice that he cannot state, even sketchily and with the help of copious idealizations and simplifications, the behavioral disposition which a given expression causes in any hearer. His account of what a disposition is does perhaps cover circumstantial variations: hearing E can be said to cause in hearers a disposition to do B if they are F, to do C if they are G, and so on; but the problem of value-variation is just overlooked.

It is not overlooked by Stevenson, who has a theory of meaning—a 'general theory of signs'—which has much in common with hearer's SRM.[3] Because of the problems of circumstantial variation, he comes to rest here : 'We must be content, then, to say that descriptive meaning is the disposition of a sign to affect cognition.' Because his primary commitment is not to a behavioral account of meaning (so that he is not a typical SRM theorist), he can afford to leave open the question of whether 'cognition' can in its turn be analyzed in terms of behavioral dispositions.

9. Speakers and hearers

Most SRM theorists attend primarily to the speaker or primarily to the hearer. They sometimes imply that they accept both the disjuncts in SRM, but I can find no careful relating of the two to one another.

Bloomfield mentions both sides of the story, and then leans towards the speaker's:

> In the causal sequence
>
> > speaker's situation → speech → hearer's response
>
> the speaker's situation, as the earlier term, will usually present a simpler aspect than the hearer's response: therefore we usually discuss and define meanings in terms of a speaker's stimulus.[4]

I am baffled by the implication that what is earlier is usually simpler. Still, one might neglect the hearer just to reduce the amount on one's plate. Thus Stevenson chooses to 'consider meaning-situations entirely from the point of view of a hearer, neglecting any. . . speaker', but only, he says, so that things will be 'simplified'.[5] Someone who drops one side just for that reason presumably thinks that substantially the same results will be reached whichever side one tackles. That seems to be Bloomfield's view here:

> We have defined the meaning of a linguistic form as the situation in which the speaker utters it and the response which it calls forth in the hearer. The speaker's situation and the hearer's response are closely co-ordinated, thanks to the circumstance that every one of us learns to act indifferently as a speaker

[1] C. Morris, *Signs, Language and Behaviour*, *op. cit.*, p. 10.
[2] M. Black, *Language and Philosophy*, *op. cit.*, p. 176.
[3] Charles E. Stevenson, *Ethics and Language* (New Haven, 1944), p. 37. For next quotation see ibid., p. 67.
[4] L. Bloomfield, *Language*, op. cit., p. 138.
[5] C. Stevenson, *Ethics and Language*, *op. cit.*, pp. 53–54.

or as a hearer.[1]

I do not see how that circumstance could explain that co-ordination. Anyway, I now proceed to show that there is a deep, structural reason why the co-ordination does not exist.

Hearer's SRM came to grief over value-variation. This affects indicative or fact-stating sentences, but not imperatives. A hearer's value-system affects whether he obeys a given command, but will not much affect what he counts as obedience to it. Understanding the command involves knowing which actions would constitute obedience to it and which would not; and values do not come into it. So, given an idealization according to which imperatives are always obeyed, the meaning of any imperative E could be expressed in the form 'Anyone who hears E does A', where A is some fairly unitary action-kind—fragmented by circumstantial variation, perhaps, but not by value-variation. (Here and elsewhere I ignore problems about hearing versus overhearing, and so on, trusting that the reader is also prepared to neglect the negligible.) So hearer's SRM does not collapse over imperatives as it does over indicatives. That helps to explain Stevenson's penchant for hearer's SRM; for he was concerned with moral judgments and in general with utterances having 'emotive meaning', and these are like imperatives in the relevant respect.

With the difference between imperatives and indicatives in mind, let us look back at the speaker's side of the linguistic transaction. When I spoke of stimuli which 'verify' an utterance, I implied that the utterance is not imperative, for imperatives cannot be verified because they cannot be true. Still, there is something we can say which covers imperatives as well as indicatives. We can say that something (a stimulus, or more liberally a speaker's situation) makes a given utterance suitable or appropriate or fitting—with truth being the kind of fittingness that indicatives have. Now, the question of whether an indicative utterance fits the speaker's situation does not depend upon one's value-system. 'The door is shut' fits all situations where the door is shut, and no others; and values are irrelevant. But whether a given imperative fits the speaker's situation may depend upon his value-system. Not always, because an imperative must be judged unfitting in a situation where it cannot be obeyed—e.g. 'Shut the door!' where the door is already shut. But value-systems play a part over a wide area. For example: I rightly judge that my situation makes 'Shut the door!' a good or suitable thing to say, whereas you in that situation would rightly regard it as quite unsuitable, counter-indicated; because you like traffic-noises and I dislike them.

The two preceding paragraphs yield a theorem. For indicatives, value-systems help to relate utterance to subsequent action, but not to relate situation to utterance; while for imperatives the case is precisely reversed—value-systems help to relate situation to utterance, but not to relate utterance to subsequent action. (As already noted, values affect obedience but not what counts as obedience.)

For neither kind of sentence, therefore, can the facts of the form 'S makes E fitting' be mapped onto those of the form 'E dictates A'. Indicatives have a value-complication on the right-hand side but not the. left, while imperatives have one on the left but not on the right. Quite apart from the separate difficulties encountered by speaker's and by hearer's SRM, therefore, the supposed correlation between them *cannot* exist.

[1] L. Bloomfield, *loc. cit.*

Part II

10. Why SR looks good for meaning-theory

The concepts of stimulus and response belong to a body of experiment and theory in psychology, and SRM is one attempt to apply that theory to questions about meaning. The attempt failed, for rather abstract reasons, but that does not imply that stimulus-response could not be made helpful to meaning-theory in some other way.

The desire to use 'stimulus' and 'response' in one's theory of meaning is not a mere perversity. There are reasons for it. These will have force on1y for those who want an account of meaning which (a) gives primacy to cases where a speaker describes the external world on the basis of his experiences, and (b) stresses the notion of linguistic behaviour. I shall not discuss those emphases here, beyond remarking that I favor them and agree with Quine that they are inter-connected: 'The externalizing of empiricism begets. . . a behaviorizing of meaning.'[1]

To appreciate the attractions of stimulus-response as an instrument for explaining linguistic behaviour, one must attend to the phenomenon of *conditioning*. This is a process whereby an organism may be caused to have a certain S-R pattern in its behaviour, i.e. to be disposed whenever it undergoes an S to make an R. Described as abstractly as possible, it goes as follows. Initially, an S's occurring in organism x does not affect the probability of x's then making an R. But there may be a kind K of event which is reinforcing with respect to x, S, and R, in the following sense. If x makes an R just after an S occurs in it, and a K occurs just after that, then the next S that occurs in x will raise the probability of x's then making an R; and the oftener the S-R-K sequence occurs, the greater the amount by which every subsequent S raises the probability of an R shortly thereafter, up to the limit where every S in x is followed by an R.

For example, a dog hears 'Sit!' (S), and just happens to sit (R), and is then given sugar (K); so the next time it hears 'Sit!' this raises slightly the probability that it will then sit; and the oftener this S-R-K sequence occurs in the dog's history, the nearer it will come to sitting whenever it hears 'Sit!'

Conditioning could be the work of non-sentient nature, but sometimes there is a person who applies the K reinforcements to the organism because he wants the organism to become disposed to make an R whenever it undergoes an S. It is arguable that in such a case the organism is *taught* to make an R whenever it undergoes an S, which suggests that S-R theory might yield a theory about language-teaching. It is evidently true that each person's linguistic behaviour results partly from social pressures and interventions, i.e. that 'language skills' are 'socially mediated';[2] and the claim that utterances are conditioned responses to stimuli constitutes a theory about the nature of the social mediation, i.e. an account of how older tribesmen bring younger tribesmen linguistically into line. This account, furthermore, is not altogether implausible. Insofar as language-learning involves rewards and punishments, inducements and counter-inducements, chance, experiment, rote, and so on, some provision is made for all these in the theory of conditioning.

At the same time, even if a given S-R behavioral pattern is acquired through deliberate conditioning, behaviour in

[1] W. V. Quine, 'Philosophical Progress in Language Theory', *Metaphilosophy*, vol. 1 (1970), p. 8.

[2] Roger W. Brown and Don E. Dulaney, 'A Stimulus-Response Analysis of Language and Meaning' in *Language, Thought, and Culture*, ed. by Paul Henle (Ann Arbor, 1965), p. 73.

accordance with it can still be seen as fully causally explicable. In organism x, the occurrence of an S causes the making of an R, even if it is through conditioning (teaching) that x has come to be so constituted that an S will cause it to make an R. This has the same simple, familiar logic as boiling something to make it brittle, i.e. causing it to be such that its being struck will cause it to break.

If linguistic behaviour consists in conditioned responses to stimuli, then we have a way of representing it as both biological and social, both natural and artificial, both caused and learned; and these two aspects will not be in tension with one another, but will both flow smoothly from a single, unified theory of linguistic behaviour.

The foregoing basis for interest in S-R as a tool in the philosophy of language has involved only the view that linguistic behaviour consists in conditioned responses to stimuli (and it could be amplified by similar attractions in the idea that someone who acts upon what he hears said is also responding to stimuli). But there is a further service that S-R might render, if we could go a step further and say that the meaning of E is somehow determined by the kinds of stimulus to which E is a proper or permissible response.

It arises from a problem which has been classically stated by K. R. Popper:

> If the statements of science are not to be accepted dogmatically, we must be able to justify them. If we demand justification by reasoned argument, in the logical sense, then we are committed to the view that *statements can be justified only by statements*. The demand that all statements are to be logically justified... is therefore bound to lead to an infinite regress. Now, if we wish to avoid the danger of dogmatism as well as an infinite regress, then it seems as if we could have recourse only to psychologism, i.e. the doctrine that statements can be justified not only by statements but also by perceptual experience.[1]

Faced with this choice—dogmatism, infinite regress, psychologism—Popper opts for a dogmatism which he finds harmless because it is a matter of practice but not principle, *and* an infinite regress which he does not mind because it is a matter of principle and not practice:

> The basic statements at which we stop... have admittedly the character of dogmas, but only in so far as we may desist from justifying them further by further arguments (or by further tests). But this kind of dogmatism is innocuous since, should the need arise, these statements can easily be tested further. I admit that this... makes the chain of deduction in principle infinite. But this 'infinite regress' is also innocuous since in our theory there is no question of trying to prove any statements by means of it. (p. 105.)

Popper rejects 'psychologism' altogether, allowing statements to be justified only by statements and never by experience. His position is not that experiences *cannot justify but can falsify* statements. Rather he is saying that experiences *cannot justify or falsify* statements or have any bearing on their reasonableness, tenability, or anything of that sort:

> The decision to accept a basic statement... is causally connected with our perceptual experiences. But we do not attempt to justify basic statements by these experiences. Experiences can motivate a decision and hence an acceptance or rejection of a statement, but a basic statement cannot be justified by them—no more than by thumping the table. (still p. 105)

[1] Karl R. Popper, *The Logic of Scientific Discovery* (London 1959) pp.93–94.

Granted, experiences can act causally on (motivate) linguistic behaviour and thus on statement-acceptance. But it seems perverse to deny that they also have a justificatory or accusatory relationship to statements—as Popper himself implies, indeed, in his remark about 'justifying them. . . by further tests'.

Still, one can sympathize with a philosopher's making that apparently perverse denial. It is natural to assume that anyone's linguistic behaviour can be causally explained, at least in part, and that the explanation will involve his experience. In this limited sense, then, one assumes that there is a causal relationship between experiences and statements. But if there is also a justificatory one, then we have a hard problem about how the two relationships are related to one another—a problem about how causes can leave room for verifications.

Now, suppose that we have some way of explaining meanings in terms of utterances' roles as responses. To keep things simple, suppose that we have, specifically, the doctrine I call speaker's SRM. If the mark of someone's understanding E is his uttering it whenever he undergoes an S, then S is the kind of experience that verifies E or makes it correct or permissible to utter E. But in a suitably conditioned speaker, according to SRM, the occurrence of an S also causes the uttering of E. So what SRM offers is a *prima facie* solution to the problem of how experiences can relate to statements both causally and by way of verification and falsification. Since SRM is dead, this does not matter much; but it is important that the same kind of solution to the problem might be offered by any theory in which the meanings of expressions were stated in terms of their roles as responses to stimuli.

The vitality of the problem can be seen from the way it arises in the context of Quine's epistemology.[1] Quine represents a body of accepted sentences as a globe immersed in experience; and pressures from the outside are transmitted inwards, the distance and route of transmission depending upon factors which Quine has explained. That much is fairly well understood. But how do pressures pass through the surface of the globe? When one starts thinking about them from inside the globe, one approaches the surface asymptotically, always interposing one more sentence between oneself and the non-linguistic world. (Popper would probably replace 'justify by further tests' by 'justify by further reports of tests'.) Approaching the surface from the outside, on the other hand, one tends to reach it with purely causal pressures: this experience will—as a matter of biology, not logic—dispose him to behave thus and so with regard to that sentence. From the inside we do not reach experience, and from the outside we give experience an executive but not a judicial role. Our need might be met by an SR theory of meaning, for such a theory could represent beliefs or statement-acceptances or linguistic acts as being justified (not just caused) by experiences (not just statements about them). SRM might have rendered this service if it had not fallen at the first hurdle.

Although Quine is no adherent of SRM, he does use the notions of stimulus and response in the philosophy of language. There is some evidence that he values those notions partly because they may solve the problem presented in this section.

[1] W. V. Quine, 'Two Dogmas of Empiricism' in *From a Logical Point of View* (Cambridge, Mass., 1953).

11. Two dangers

I mentioned in §1 that some writers have wanted to base an account of all the meanings in a language on an initial S-R account of the meanings of some favored set of expressions. This second stage of the program would need care and ingenuity, especially in view of the fact that most sentences have meanings which obviously could not be directly described in S-R terms because they are obviously not semantically related to the circumstances in which they are uttered. As Chomsky remarks, 'A minute amount of verbal behaviour, outside the nursery, consists of such remarks as *This is red* and *There is a man*.'[1] He notes that an S-R theorist might say that any utterance may be a response to a purely private stimulus—an S with no associated S*. But that claim could not help to preserve the view (which Chomsky is not considering) that the meaning of what is uttered is defined by the kind of stimulus to which it is a response; for any stimulus which is to define or determine a meaning must be (systematically related to something which is) accessible to the casual observer or language-teacher.

So even if we had an S-R account of the meanings of a core of sentences, it is not obvious how this could be parlayed into an account of all the meanings in a language. Bloomfield tries to acknowledge the problem:

> People very often utter a word like *apple* when no apple at all is present. We may call this *displaced speech*. The frequency and importance of displaced speech is obvious. We recall the infant 'asking for' his doll. . . . In other ways, too, we utter linguistic forms when the typical stimulus is absent. . . . Lying, irony, jesting, poetry, narrative fiction. . . The displaced uses of speech are derived in fairly uniform ways from its primary value, and require no special discussion.[2]

Bloomfield must have under-estimated the 'frequency and importance' of the phenomenon he here declines to discuss. The contrast of 'displaced' with 'primary' suggests a peripheral use of language, as does the choice of such examples as lying, irony and jokes, rather than physics, history, and gossip.

As for the phrase 'derived in fairly uniform ways'—that seems to be mere bluster. Bloomfield implies that his 'second stage' problem, of deriving a general account of meanings from his initial S-R account of some meanings, has been solved or at least is clearly soluble. That is just false.

Quine's application of stimulus-response to meaning-theory is not open to criticism on this score. He envisages giving an S-R account (to be described in §12) of the meanings of a core of 'observation sentences', and then using these as a check on 'analytical hypotheses' through which meanings are to be assigned to all other sentences. There is no suggestion that the notions of stimulus and response can be directly applied to every sentence, nor does Quine share Bloomfield's naive optimism about how the S-R core of sentences relates to all the others. On the contrary, Quine defends a kind of scepticism about the very notion of meaning except in application to 'observation sentences', maintaining that an inevitable 'indeterminacy' besets any attempt to assign meanings to all the sentences in a language.[3] His basis for this includes an insight into how very complex and unstraightforward are the explanatory routes from an account of the meanings of observation sentences through

1 N. Chomsky in *The Structure of Language*, ed. by Fodor and Katz, *op. cit.*, pp. 570–571; see Jonathan Bennett, *Rationality* (London, 1964), pp. 83–84.
2 L. Bloomfield, *Language*, *op. cit.*, pp. 141–142.
3 W. V. Quine, *Word and Object*, *op. cit.*, §§15–16.

to any account of the other meanings in the language. This 'indeterminacy' thesis, whether or not it is true, at least shows how free Quine's work is from one kind of inadequacy which is common in S-R theorizing about meaning.

In the quoted passage, Bloomfield sets the problem up wrongly, in a manner which involves the second 'danger' I want to discuss in this section. Specifically, he focuses the problem on words rather than sentences. I have taken SRM to be addressed to expressions which are like sentences in that the uttering of them is saying something. The troubles that hearer's SRM encounters if it starts with words rather than sentences are well analyzed by Black.[1] But speaker's SRM also collapses instantly unless it starts with something like sentences.

I am sure that any attempt to bring stimulus-response theory to bear upon questions of meaning would have to be addressed in the first instance to sentences rather than to words. But I do not offer that as a criticism—quite the contrary. Most philosophers of language now agree that sentence-meaning is primary, in the sense that to explain what it is for a word to have meaning one must refer to the word's contribution to the meanings of sentences containing it. Any S-R meaning-theory, I conjecture, must further contend that sentence-meaning could come first in the order of teaching, learning, translating—so that the translation of a language might begin by assigning meanings to whole sentences in it. I believe that this further contention is true. Of course we usually understand sentences through understanding their constituent words, which explains 'the fact—which is of the essence of language—that we can understand new sentences which we have never heard before'.[2] But there is no conflict here, for the order of understanding could be: some sentences, then all words plus syntax, then the remaining infinity of sentences.

On this point SRM theorists have usually done badly.[3] Speaker's SRM has not visibly crumbled in their hands, as a result, only because they have conflated words with unstructured sentences. Osgood, for example, bases his whole theory of language on 'the acquisition of object labels':

> The adult holds up. . . a doll within the child's field of vision and says 'doll'; the child (ideally) makes an imitative approximation. . . and reaches for the object; the parent. . . immediately applies. . . rewards. In theoretical terms the visual cues from the object are being conditioned to the correct vocal response. The final step in labeling. . . is making the correct vocal response to the object as a stimulus without needing an adult model.[4]

Parts of Osgood's discussion show that this child is supposed to be acquiring 'doll', a *word to be used in* 'This is a doll' and other sentences;[5] but the passage is tolerable only if the child is seen as acquiring 'Doll!', a *sentence which means the same as* 'This is a doll'. Similar examples abound in the literature.[6]

---

1 M. Black, *Language and Philosophy*, *op. cit.*, pp. 174–175, 205–206.

2 Michael Dummett, 'Nominalism', *The Philosophical Review* vol. 65 (1956) p. 492.

3 See C. Morris, *Signs, Language and Behaviour*, *op. cit.*, pp. 15–27 (re 'sign of food'); C. K. Ogden and I. A. Richards, *The Meaning of Meaning* (London 1923), pp. 14 ff., 404; Skinner, *Verbal Behavior*, *op. cit.*, pp. 19–21.

4 Charles E. Osgood, *Method and Theory in Experimental Psychology* (New York, 1953), pp. 688–689.

5 C. E. Osgood, *op. cit.*, pp. 690 ff. (re 'hammer').

6 L. Bloomfield, *Language*, *op. cit.*, pp. 139–140; R. W. Brown and D. E . Dulaney, 'A Stimulus-Response Analysis of Language and Meaning', *op. cit.*, pp. 67–68; Roger Brown, *Words and Things* (New York, 1958), pp. 7–9; Theodore Thass-Thienemann, *Symbolic Behaviour* (New York, 1968), pp. 17–20.

Quine is again in the clear. He carefully distinguishes unstructured sentences from words, starting with the former and introducing the latter through 'analytical hypotheses'.[1] Yet Katz, speaking of 'stimulus-response theories of meaning such as. . . Quine's, says: 'The basic idea underlying each of these conceptions of meaning would limit a semantic component to the task of trying to state the meanings of words individually.'[2] By this remark Katz achieves something which is rare even in the grey world of philosophical controversy—a perfect injustice, a precise reversal of the truth.

12. Quine: Answers as responses

In my final three sections I present and comment upon Quine's use of the notions of stimulus and response in meaning-theory. Quine has generously helped me with these sections, but no doubt he would still disagree with them.

It is worth repeating that the basic relationship between S and E that a theory of meaning has to capture is: *the occurrence of an S is what makes it correct or permissible to utter E*. Quine offers to express this in behavioral terms, without using 'correct' or the like, as follows: *If a tribesman is asked 'E?', he answers affirmatively if an S has just occurred and negatively otherwise*. (To ask him 'E?' is to put E to him interrogatively.) This has the general stimulus-response form—it says that in such and such conditions the tribesman *will* utter so and so. But it does not try to define E's meaning through its role as a response, but rather through its role as part of a stimulus; and *a fortiori* it does not imply that anyone's understanding E is manifested in his uttering it whenever he could truthfully do so; and so it does not collapse as speaker's SRM does.

Quine's 'behavioral approximation' to 'the notion of confirmatory and disconfirmatory experiences' (p. 64) is not in fact purely behavioral. We have to recognize tribal assent and dissent, as Quine remarks (pp. 29–30), and also tribal interrogatives—which could be tricky because it could depend not on intonation but purely on grammatical transformations. Also, will the interrogated tribesman answer truthfully? We can fairly pretend that he will not lie or stay silent out of malice or sulkiness.[3] But suppose he sees that the linguist's questions cannot be asked for the usual purpose of gathering information about cabbages and rabbits etc., but does not guess what their unusual purpose is? In that case, the questions will not be 'pretty sure to elicit' answers (Quine, p. 17.), and it is not clear that legitimate idealizations can take care of this possibility. (Quine approaches this last point when he says: 'The native may dissent from "Gavagai" in plain sight of the rabbit's ears, because the rabbit is in no position for shooting: he has misjudged the linguist's motive for asking "Gavagai?".' (p. 39) But this native has surely misjudged a meaning rather than a motive.)

So Quine's query device works for a given language only if assent and dissent and interrogatives are identified early, and speakers submit cooperatively to interrogation by linguists. So the device might fail.

There is, however, a more important limitation on the query device. I contend that we can easily imagine a language L such that: L resembles actual natural languages in all its main aspects, and raises all the basic problems actual languages raise; but the query device is absolutely inapplicable to L because L has no means for asking and

[1] W. V. Quine, *Word and Object*, *op. cit.*, pp. 69–72.
[2] Jerrold J. Katz, *The Philosophy of Language* (New York and London, 1966), pp. 313–314.
[3] See p. Ziff. 'A Response to 'Stimulus Meaning', *op. cit.*, p. 71.

answering 'whether' questions—there just is not any way of putting expressions to the L-speakers interrogatively.

I submit, then, that Quine's query device cannot solve any problems about language as such, for example, the one about how experience can bear upon language semantically as well as causally. This may make the query device inadequate for the work Quine wants it to do. Consider, for example, his account of what an 'observation sentence' is. Informally, an observation sentence is one whose truth-value can be read off from the experiential content of the speaker's situation without special background knowledge. Quine offers to 'clean up' this notion by defining it through the query device (pp. 42–43). This is fine; but one hopes that the notion of an observation sentence can also be 'rendered respectable' without relying on 'whether' questions. That notion is involved in explaining and defending Quine's thesis about the indeterminacy of radical translation of *any* language; and so some limitation on Quine's intentions seems to be implied by the fact that there could be natural languages to which the query device could not be applied.

This point can be generalized. Quine can adduce the answering of 'Whether' questions as refuting Ziff's claim that there are no viable regularities of the form 'If α is the case, then a speaker does β'. (From now on I shall use 'S $\rightarrow$ E generalization' to cover any report of such a regularity. The special constraints I put on the phrase in §4 are now dropped.) Indeed, one can generate examples of such regularities at will, by describing values of S which combine physical torture with demands that E be uttered: for many values of E, such means will work almost infallibly for anyone who understands what is being demanded. The vital point about such examples is not that in them S involves torture, but that in them S involves language, in the form of a demand or request that something be done. I conjecture that we cannot make trouble for Ziff's claim except through examples having that special feature, namely an S which involves language; and that any use of language which served this purpose would be theoretically dispensable in the way 'whether' questions are, so that a language might easily lack the means for it while still resembling natural languages in all its main features. If that is right, then Ziff's claim is as good as true from the standpoint of the most general problems in the philosophy of language.

13. Quine: Learning and theories

The view that linguistic performances are conditioned responses to stimuli yields a plausible account of some early episodes in language-learning. This is clearly part of the service that Quine expects stimulus-response to render. He says:

> Words mean only as their use in sentences is conditioned to sensory stimuli, verbal and otherwise. Any realistic theory of evidence must be inseparable from the psychology of stimulus and response, applied to sentences (p. 17).

And also:

> Conditioned response does retain a key role in language-learning. It is the entering wedge to any particular lexicon, for it is how we learn... simple observation sentences... by ostension. Learning by ostension is learning by simple induction, and the mechanism of such learning is conditioning.[1]

This use of the notions of stimulus and response must be a sheer addition to their use in the query device; for

[1] W. V. Quine, 'Comment on Chomsky' in *Language and Philosophy*, ed. Sidney Hook (New York, 1969), pp. 96–97.

very young children are not conditioned to answer 'whether' questions—'Prompted assent is no game for such small children.'[1] But then what *does* go on, in the way of setting up dispositions to respond to stimuli, in the early stages of language-learning? Quine says:

> The child's early learning of a verbal response depends on society's reinforcement of the response in association with the stimulations that merit the response, from society's point of view, and society's discouragement of it otherwise. (p. 82)

In a similar vein, he describes the upshot of conditioning thus: 'In future the approach of the mother's face succeeds as a stimulus for further utterances of "Mama"' (p. 81); and also 'The baby learns. . . to say his word when appropriately irritated and not otherwise.' (p. 92) The idea seems to be that the child is so conditioned that whenever an S occurs it utters E, or at least becomes much more likely to utter E. Neither of these is any good as a basis for describing meanings; but then Quine is not here presenting a theory of meaning, but simply applying stimulus-response notions to early language-learning.

He may well be right. It seems plausible to suppose that the language-learning child goes through phases in which it does conform to S $\rightarrow$ E patterns, perhaps with a somewhat weakened arrow. It may even be necessary to one's eventual grasp of the language that such phases be passed through. But such phases are at best only stepping-stones, even if essential ones, to linguistic competence. To regard them as constituting even a primitive form of linguistic competence would lead back into all the troubles of speaker's SRM. Quine says:

> The operant act may be the random babbling of something like 'Mama' at some moment when, by coincidence, the mother's face is looming. The mother, pleased at being named, rewards this random act, and so in future the approach of the mother's face succeeds as a stimulus for further utterances of 'Mama'. The child has learned an occasion sentence. (p. 81)

I suggest that the last sentence is misleading. The child as described has merely come to conform to a strong or weak S $\rightarrow$ E pattern: to say that that is learning an occasion sentence is to imply that there are forms of linguistic competence which are adequately describable in S $\rightarrow$ E terms: and that is false. That is a marginal quibble, however. My main concern—commenting rather than quarrelling—is to separate the use of stimulus-response in the query device from its use in describing early language-learning.

Quine also uses the notions of stimulus, conditioning, etc. in discussing general theories. This is quite distinct from the query device and from language-learning, and Quine's procedures here seem to be more open to criticism. He says:

> The power of a non-verbal stimulus to elicit a given sentence commonly depends on earlier associations of sentences with sentences. . . Someone mixes the contents of two test tubes, observes a green tint, and says 'There was copper in it'. Here the sentence is elicited by a non-verbal stimulus, but the stimulus depends for its efficacy upon an earlier network of associations of words with words; viz., one's learning of chemical theory (pp. 10–11).

What can Quine mean, of a purely behavioral sort, by 'association of sentences with sentences'? Well, a little later he says this:

[1] W. V. Quine, *Word and Object*, op. cit., p. 81.

> In the series of sentence-to-sentence associations ultimately linking 'The stuff has gone green' with 'There was copper in it', all steps but the last are evidently unspoken. Some may be sketchily but inaudibly spoken, but more are just skipped as the theory becomes second nature. Such skipping... seems a basically humdrum affair: a transitivity of conditioning (p. 12).

So when there is no 'skipping' the relevant sentences *are* uttered. It seems that my being conditioned to accept the theoretical sentence of the form $(\forall x)(\phi x \rightarrow \psi x)$ is my being conditioned, whenever I utter or hear a sentence of the form ϕa, to *utter* the corresponding sentence of the form ψa (though sometimes I may 'skip' this because of 'transitivity of conditioning'). Quine seems here to have committed himself to S → E patterns of a quite indefensible sort.

No doubt that was not his intention; but that implies that he has used expressions like 'sentence to sentence conditioning' with too little explanation. He cannot be giving 'conditioning' and the rest their full theoretical load, as technical terms in animal psychology; yet he does not explain how their load is to be lightened.

14. Reversing the arrow

Properly to assess Quine's uses of stimulus and response, we must consider one more emendation of speaker's SRM. I showed that we could not rescue it by amplifying the stimulus or weakening the link, but I did not discuss the third obvious remedy—namely reversing the direction of the arrow. If there are true E → S generalizations, at least in the idealized linguistic situation, why should we not build an account of meanings upon *them*? It seems plausible to suppose that the meanings of some expressions could be captured in statements of the form 'If a speaker utters E, he has just undergone an S [or observed an S*]'.

Such a reversed-arrow variant on speaker's SRM would no longer assign causes to utterances; and so it would not single-handed support the view that utterances are biological events with their place in the causal flow. This may be regrettable, but it is hardly fatal.

A more important point arises if we view arrow-reversal from the standpoint of stimulus-response theory generally. In place of the standard S → R, perhaps with a weakened arrow, we are now going to have to make room for generalizations of the form (¬S → ¬R), if an S does not occur, the organism does not make an R. This suggests a study of conditioned *non-responses*, which would be a very odd inquiry! Conditioned *inhibitions* are something else again: they are a legitimate subject of inquiry and a possible object of psycho-technology. But an inhibition is not fully described by the form 'If an S does not occur, then x does not make an R', for if it were then death would be the possession of all inhibitions. The concept of inhibition has, rather, this shape: 'In seeking food the rats often make R's, sometimes when an S has not just occurred; but after they have been conditioned *they continue their food-seeking activities as before except that* they never make an R unless an S has just occurred.' Here the generalization of the form (¬S → ¬R) expresses an inhibition—a conditioned non-response which is not automatically satisfied by their never doing R—because it occurs as a qualification of a statement about something that the rats *do* do.

All this re-applies when we try to express meanings through (¬S → ¬E) generalizations. These need to be embedded in a positive account of the ongoing linguistic activity which the generalizations serve to qualify, or else accompanied by a candid admission that such an account is being presupposed but not provided. The latter alternative would involve using the reversed arrow in statements of the

form 'The tribe use E meaningfully or linguistically, their use of it being subject to an inhibition expressed by ($\neg S \rightarrow \neg E$).' That, however, would make clear that the largest and most interesting part of the story was not being attempted. I implied in §1 that I was going to consider the use of S-R in saying what certain expressions mean but not in explaining what it is for something to have a meaning or to be linguistic, but that limitation looks more serious now than it did before the arrow was reversed. It was plausible to suppose that the $S \rightarrow E$ generalizations were sufficient to establish that the tribe had *something like* linguistic behaviour; but $E \rightarrow S$ generalizations, on their own, come nowhere near establishing even that much. Any one of them could be true because the tribe never utters anything;[1] and although we can stipulate that each generalization must be often instantiated within range of a potential hearer, this still falls a long way short of guaranteeing that the tribe's behaviour will look anything like the sensible use of a language.

All this throws an interesting light on a certain passage of Quine's. A good deal of his discussion of language can be seen as an attempt to show how the logical relationship between experience and language is a special kind of causal relationship between experience and linguistic behaviour.[2] Within this, there is one notably persuasive episode, which describes experience as (logically) refuting a theory and as (causally) inhibiting something:

> Prediction is in effect the conjectural anticipation of further sensory evidence for a foregone conclusion. When a prediction comes out wrong, what we have is a divergent and troublesome sensory stimulation that tends to inhibit that once foregone conclusion, and so to extinguish the sentence-to-sentence conditionings that led to the prediction. Thus it is that theories wither when their predictions fail (p. 18).

This sounds promising because in it the arrow is reversed: the topic is falsification, not verification; inhibitions rather than responses. But what is being inhibited? When a theory meets up with a recalcitrant experience, what (according to Quine) is one thereby stopped from doing? The passage hints at two different answers. **(1)** What is inhibited is a predicting, an uttering: had the recalcitrant experience not occurred, one would have uttered something which its occurrence causes one not to utter. But that presupposes that the ($\neg S \rightarrow \neg E$) generalizations are embedded in, and serve to qualify or restrict, generalizations of the form $S \rightarrow E$; which takes us back to the fatal flaw of speaker's SRM. This is an echo of something I noted in §13, namely the suggestion that Quine equates theory-acceptance with a disposition whenever one encounters a certain sentence to utter a certain other sentence. **(2)** Alternatively, what is inhibited or extinguished is a belief. This reading is suggested by the expressions 'anticipation' and 'foregone conclusion'. I believe that an adequate theory of language must give a central role to the concepts of belief and of intention; and that an important task confronting behavioral meaning-theory is to give a behavioral analysis of these. (Not an S-R analysis. I think that the concepts of intention and belief get work to do precisely when the grip of stimulus-response statements is broken.[3] But that is a topic for another day.) Since Quine offers no analysis of these concepts, he cannot rest weight on either of them without thereby compromising the behavioral nature of his enterprise.

1 D. K. Lewis, *Convention*, *op. cit.*, p. 178.
2 W. V. Quine, *Word and Object*, *op. cit.*, §§3–5.
3 J. Bennett, *Rationality*, *op. cit.*, pp. 36–37, 43–45.

It is a matter for congratulation that Quine does not take the standard S-R short-cuts to the concepts of belief and intention. But he is on record as thinking that they cannot be handled in unitary behavioral analyses of any kind: he regards them as part of a 'dramatic idiom' which has no place in a scientific account of behaviour.[1] I suggest that that helps to explain Quine's various attempts—discussed in §§12–13 above—to overload the concepts of stimulus and response. The trouble is that those concepts are being made to bear a load which properly belongs to concepts which Quine will not allow himself in the philosophy of language because he thinks that they cannot be represented as cleanly behavioral—namely the concepts of intention and belief.[2]

[1] W. V. Quine, *Word and Object*, *op. cit.*, p. 219; *Ontological Relativity*, *op. cit.*, pp. 146–147.

[2] I have been helped in writing this paper by my colleagues S. C. Coval, Howard Jackson, Edwin Levy, Alan Loveland, and Thomas E. Patton. If the paper now has a coherent shape, this is due to the constructive intervention of Donald G. Brown.