

human information-processing was to establish meaningful probabilities to be attached to the different possible signals or brain-states concerned. After a flourish of 'applications of information theory' in psychology and biology which underrated the difficulty of this requirement, it has now come to be recognized that information theory has more to offer to the biologist in terms of its qualitative concepts than of its quantitative measures, though these can sometimes be useful in setting upper or lower limits to information-processing performance.

It also became clear that to avoid conceptual confusion it was not sufficient to preface the word 'information' with distinguishing adjectives such as 'selective', 'structural' and 'metrical'. Our chief terminological need was for a way of keeping the notion of information per se distinct from all measures of 'amount-of-information'. Rather than invent still more neologisms, I took to using 'information-content' (qualified as selective, structural or metrical) to denote the latter, leaving 'information' free to be used in its original everyday sense (see Chapter 6).

Meaning and Mechanism

FOREWORD

Claude Shannon's measure of selective information-content was framed explicitly to require no reference to the meaning of the information selected in response to a communication signal. The communication process was viewed as a transaction between terminals whose task was confined to the generation and reproduction of symbols. The understanding of these symbols was declared to be outwith the concern of the communication engineers, and the saying went abroad that Information Theory had no place for the concept of meaning.

The second part of this collection (see for example Chapters 5, 6 and 7) contains a number of papers aiming to narrow the gap between semantic and other aspects of the concept of Information as these were gradually hammered out. The text of the talk that follows* outlines a way of thinking about the function of an utterance which these papers use as a conceptual bridge between the mechanistic and semantic levels. By picturing an item of information as a kind of tool that operates upon the recipient's internal 'state of conditional readiness', we can conveniently define its meaning on the one hand, and its information-content (in various senses) on the other. This approach also offers a criterion of meaningfulness and meaninglessness which seems more realistic and less Procrustean than the 'verifiability' or 'falsifiability' criteria canvassed by some linguistic philosophers.

* First broadcast in January 1960 on the B.B.C. Third Programme. Reprinted (revised) in *Common Factor*, No. 5, 1968, pp. 57-65.

A human conversation depends on many processes which a scientist would call 'mechanical', in the sense that only physical categories of cause and effect are needed to describe and explain them. Puffs of air, produced by vibration of the speaker's larynx, echo around the cavities of his mouth and result in a characteristic sequence of sound waves. These travel through space and vibrate the sensitive membrane of the listener's ear, giving rise to nerve impulses, and so on. Now, until the chain of explanation reaches the nervous system, nobody minds its mechanistic flavour. True, it has made no reference to the *meaning* of what is being said; but this, we might say, would obviously be premature. Questions of meaning need not arise until we bring in the human links in the chain.

As far as it goes, I think this answer is fair enough; but as a clue to the relation between the meaning of a message and its mechanical embodiment it is far from satisfactory, on two main counts.

In the first place, it seems to imply a rather worrying and muddling sort of discontinuity. The original speaker, we suppose, means something by what he says. His utterance has meaning — at least for him. Yet in the next stages of the chain of explanation (the generation of sound waves and all the rest of it) all signs of his meaning seem to have disappeared. Discussion at this level proceeds in exactly the same terms whether the air is handling the outpourings of a genius or the jabber of a monkey. Yet finally, when the message reaches the ear of a human listener, its 'meaning' seems to pop up again from nowhere, and concerns him far more than the physical properties of the sound wave which the speaker actually produces. There are in fact two awkward discontinuities in this way of telling the story: a jump from meaningful utterance to meaningless air vibrations; and then back again to meaningful utterance. The question that seems to be raised by this way of looking at the process is *how the message recovers its meaning*. Whether or not this is a sensible question remains to be seen.

The second objection to an easy division of human communication into mechanical stages and meaningful stages is a more serious one. The process by which the sound pressure-waves titillate the eardrums is mechanical. So far, so good. But what happens next? Isn't the process by which the nerves convey these titillations to the brain a mechanical one? And what then? Isn't the next, neurophysiological stage, though still puzzling in detail, plainly a mechanical one too? Where are we to draw the line?

The position is complicated by our enormous ignorance of what goes on in the central nervous system. Some people would take advantage of this to suggest that somewhere in the brain the laws of mechanical causation, on which we have been relying, may give place to 'something else' — associated with the action of mind — and this could be the stage at which (as they would say) 'meaning is given' to the buzz of physical brain activity.

For reasons that I won't go into now, I must say I find this suggestion unattractive. To begin with, it seems entirely arbitrary and without scientific foundation, and relies only on our ignorance of brain organization. But my chief objection is that it is an unnecessary way of seeking to safeguard the meaningfulness of human communication — a way that commends itself only if we insist on regarding meaningful and mechanical processes as mutually exclusive, so that to describe something as 'mechanical' implies automatically that it is 'meaningless'. If, as I hope to show, this opposition of 'meaningful' and 'mechanical' is false, the momentum of the whole debate between 'mechanists' and 'anti-mechanists' disappears.

A FUNCTIONAL APPROACH

Suppose we now make a fresh start on the problem of meaning by asking what difference it makes to you when you receive and understand the meaning of a message. For our

purpose it will be sufficient if we confine ourselves to indicative sentences;* suppose, for example, someone tells you 'it's raining'. What happens? You may be immersed in a book, and may not feel inclined even to grunt an acknowledgement. But this does not mean that your understanding of the message has had no effect on you. If a sudden call comes for you to go out of doors, for example, you may now be ready to reach for umbrella or mac. If someone comes in, you are likely to ask whether he got wet; and so on. What has been affected by your understanding of the message is not necessarily what you do — as some behaviourists have suggested — but rather what you would be *ready* to do *if* given (relevant) circumstances arose. It is quite possible that relevant circumstances may never arise, so that a naively behaviouristic approach would reveal no sign that you had understood the message. It is not your behaviour, but rather your state of *conditional readiness* for behaviour, which betokens the meaning (to you) of the message you heard.

At this stage you may wonder whether we are not merely replacing one problem by another — this idea of 'states of conditional readiness' may sound every bit as vague and elusive as the notion of meaning itself. The all-important point, however, is that we can talk about 'states of readiness' in relation either to human beings or to mechanical systems — or indeed to human beings as mechanical systems. The notion offers us a conceptual bridge between the two ways of talking whose relationship has puzzled us.

Let me try to make it more definite by way of an example. Think of a railway signal-box controlling a large shunting-yard. At any given moment, the configuration of levers in the box defines what the yard is *ready* to do to any waggon that happens to come along. There may in fact be no waggons moving; there may be some tracks on which no waggons will

* For a discussion of non-indicative meaning in these terms see Chapters 4 to 8 and Chapter 11.

move for years; but this is no obstacle to a definition of the total state of conditional readiness of the yard, as betokened by the total configuration of lever-settings, which determines what would happen *if* any given circumstance arose. A change of a lever that controls a disused siding may cause no visible change in the activity of the yard; but it makes a perfectly definite change in its state of conditional readiness.

Now we do not know at present how the 'signal-box' in our heads controls our own conditional readiness to act and react in all possible circumstances; but it seems reasonable in our present state of knowledge to suppose that there is a brain mechanism that does this. I don't mean of course that the brain is organized on the simple principles of railway signalling; there is vastly more subtlety in its blend of spontaneous and controlled activity, which allows us to think of it as determining only the probabilities of different actions in given circumstances, rather than anything as rigid as the motions of railway waggons. What is more, the brain differs from a shunting-yard in being a self-guiding system; and what we are chiefly concerned with is its state of readiness for goal-directed, adaptive activity: activity with a purpose. It is your conditional readiness for this kind of activity that is modified and moulded according to your understanding of the information you receive.

ENERGY VERSUS FORM

But with these cautions in mind, I do suggest that we can think metaphorically of the receipt of a message as causing a change in the 'lever-settings' of the brain, and so in the human receiver's state of readiness, in much the same sense as in the case of our shunting-yard. This will allow us to relate the concepts of meaning and mechanism in a very natural way.

On the one hand, we could say that all that is required to make the brain-levers change is the purely mechanical

energy of the message, without reference to its meaning. In one sense this is true. On the other hand, changing the physical energy, say by doubling the loudness, of the message you hear, would not normally make much difference to its effect on your brain-levers; whereas a change in its meaning, even if it left the total energy unaltered, could make all the difference in the world.

The mechanical *energy* of a message must be sufficient to do the mechanical job that eventually resets the brain-levers; but the selective job, of determining which levers shall move, depends on the *form* of the message, and on the state of your brain before you hear it. This is where the meaning of the message comes in. As long as we think only of what actually happens, we may be able to make do with explanations solely in terms of physical energy. It isn't until we consider the range of other states of readiness, that *might have been selected but weren't*, that the notion of meaning comes into its own. A change in meaning implies a different selection from the range of states of readiness. A meaningless message is one that makes no selection from the range. An ambiguous message is one that could make more than one selection.

A WORKING DEFINITION OF MEANING

And so we could go on; but these examples are perhaps sufficient to lead us to a working definition of meaning, in this context. It looks as if the meaning of a message can be defined very simply as its selective function on the range of the recipient's states of conditional readiness for goal-directed activity; so that the meaning of a message to you is its selective function on the range of your states of conditional readiness.

Defined in this way, meaning is clearly a relationship between message and recipient rather than a unique property of the message alone. Thus the conundrum with which we began — the puzzle of the missing meaning — appears to be a spurious one. For the original speaker, the meaning of

what he says is the selective function he wants it to perform on the listener's range of states of readiness. This is distinct — and may, as we shall see, be quite different — from the *effective* meaning to the listener, which is of course the selective function actually performed; and both of these may differ from the *conventional* meaning, which is the selective function calculated for a 'standard recipient'. In between the speaking and the hearing — at the stages of mechanical transmission — no comparable selective operations occur, so that questions of meaning do not normally arise; but at any stage we can still define the intended meaning of the message as its intended selective function, and so on. No queer discontinuity occurs along the way.

The jabber of monkeys is thus distinguishable even at the mechanical level from the outpouring of the genius, if we ask the right mechanical question — about its ultimate selective function. Traced far enough along the chain, the effects of the signals emitted by the genius should — or at least could in principle — modify your bodily state of readiness in a manner intended by him; whereas nothing comparable could be said for the sound of monkeys — unless indeed you had some exceptional acquaintance with their vocal habits!

MESSAGES AS 'KEYS'

Let me now digress a moment to introduce one further feature into our analogy of the signal-box. On the older one-track railway lines, the engine-driver carries a key, which must be inserted and turned in a special signalling lock at the station before his train can be cleared into the next section of the lines. His signal, we may say, is key-operated. Imagine now a complete signal-box, working on the same principle. Insert a key of a given shape into the box, and you make a certain selection from the range of possible configurations of the signal-levers. Insert another, and the selection you make is different.

What I am leading up to is the idea that the brain, considered as a signal-box, is also 'key-operated'. The physical embodiment of a message, which eventually acts on the brain, may be likened to the key that sets up a certain configuration of levers in the signal-box we have been talking about. Its selective function, like that of a key, depends both on its shape and on the arrangement of levers it meets.

It will now be evident that we have sidestepped our second conundrum (about the point in the nervous system at which incoming physical stimuli acquire meaning). From our present point of view, it makes no more sense than to ask about the point inside the key-operated signal-box at which a key 'acquires selective power'. There is, no doubt, a point inside the box at which the 'selective power' of the key is exercised; but of course it had this power all along. Similarly, there is, no doubt, a stage in the central nervous process at which the selective function of a message is exercised, and your brain is set up in conditional readiness to match the state of affairs that the message betokens. But we should, I think, rightly decline to regard this as the acquisition of meaning by a hitherto meaningless incoming pattern; for anyone who knew enough of the state and mechanism of your brain (*ex hypothesi*) would have been able in principle to tell that the pattern possessed that meaning for you, even if in fact you never received it. If you know enough about a signal-box, you can determine the 'selective power' that a particular key has for it without ever inserting the key.

STRUCTURAL VERSUS FUNCTIONAL CRITERIA

The analogy of a signal-box and key is of course desperately incomplete; for the brain responds to messages in a far more lively way than signal-levers to a key, and its reactions may even include a change in 'code' such that the same key can have different functions on two successive occasions. But the illustration will serve to introduce a topic of recent

debate among philosophical linguists and logicians, who are professionally concerned to find an objective basis for talk about meaning. Proponents of one school of thought, whom for short we may call the 'structuralists', concentrate on the structure of what is said, and try to analyse meaning by breaking sentences and words down into their logical components. An extreme form of this approach, advocated at one time by Bertrand Russell,³⁰ went by the name of 'logical atomism': the idea that the meaning of a complex utterance could be broken down into a conjunction of elementary or 'atomic' sentences, each so simple that it could not be further analysed. Though applicable in some scientific contexts, this notion runs into many difficulties with ordinary language. Advocates of the other school, whom we may term the 'operationalists', look for meaning rather in the pattern of the use that is made of words. This approach is generally associated with the name of Wittgenstein⁴⁰ whose later years were spent in vigorous repudiation of the atomistic view of language earlier expounded in his famous *Tractatus Logico-Philosophicus*.

Each party claims an insight into the nature of meaning that the other lacks; and although an outsider here must needs tread delicately, I believe that the mechanistic analogies we have been considering may help us to see not only the justice of these rival contentions, but also the possibility of a unified approach that finds room for both. For a similar argument could be imagined between people on the one hand who wanted to define the 'selective power' of a key by looking at its shape, and those on the other who preferred to watch and see how it was used.

The first, the structuralists, could be quite successful if only all signal-boxes were of a fixed pattern; but if it turned out that some boxes changed their internal arrangement with time, for example, so that the same key had a different selective power at different times, their structural approach would naturally lose prestige. The second, the operationalists, would be in their element if keys behaved in a way that could not be

predicted from their structure; but if in fact most keys of a certain shape were found consistently to make a selection predictable from that shape, their operational method would seem very roundabout and inefficient and lacking in insight.

As a parable of the situation in semantics this is a little unfair, but it will serve our present purpose. In each case, I think, trouble arises through taking an experimental *criterion* as an objective *definition*. Both the structure of a key and the pattern of its use are necessary experimental pointers to its 'selective power'; but neither is suitable as an overall definition of the notion. Similarly, both the structure of a message and its pattern of use would seem to be necessary experimental pointers to its 'meaning'; but we only breed confusion if we try to turn either of them into a definition of what they betoken.

Does our mechanistic analysis offer any clues, then, to a synthesis between these rival lines of thought? Let us recapitulate briefly. We have been thinking of human communication in terms of the mechanical pattern of cause-and-effect that we have assumed to embody it from start to finish. We have considered the human recipient's brain as a physical system, with at any time a certain (statistical) 'state of conditional readiness', roughly analogous to what we called the 'state of conditional readiness' of a railway shunting-yard. The object of communication is to select some particular conditional readiness in the recipient from the range of states that are possible. The intended meaning of the communication is then definable as the selective function that it is intended to exercise on the range of possible states. Its effective meaning is the selective function that it actually performs.

SEMANTIC UNITS AND CRITERIA OF TRUTHFULNESS

Now the big difference between the brain and a signal-box, of course, comes in here; for, as we noted, the effect of an

input to the brain is liable to depend in a complicated way on what has gone before it; and it is only in special cases, usually after long training, that one could rely on getting exactly the same effect from the same input at different times. The moral — and it is an important one — is that what we may call the 'semantic units' of communication — corresponding to the keys we have been talking about — are often made up of a whole group of words or even of sentences acting together, rather than single words acting in isolation. The selective job is done in easy stages, as it were, and the final sharpening-up of the state of readiness may sometimes have to wait a long time for some of the words necessary to complete it. (This is particularly true, of course, in poetic and religious forms of speech — but that's another story.)

It might be argued on these lines that the concept of semantic unit must be extended to include the whole linguistic experience of the listener; but this in most cases would be unrealistic. Our point is simply that in many quite ordinary cases the unit of communication can and does extend far beyond the boundaries of a grammatically complete and apparently unambiguous statement, so that it becomes unrealistic to attach labels such as 'true' or 'false' to the particular statement in isolation. For truthful communication, what matters is whether the present message, coupled with what has gone before it, is calculated to select a state of readiness to match the actual state of affairs. If it is not, then the communication is untruthful, regardless of whether we would be inclined to accept its component sentences as true if we took them in isolation. It is easy to see from a mechanistic angle how even a physical change in tone of voice might alter the state of readiness selected from an appropriate one to an inappropriate one. I can deceive you by telling you the truth, if I do so in such a way that I know you will not believe me. More insidiously, of course, I can do so by simply selecting 'true information' which I know will give you a false impression in the absence of other data. In short, by insisting

that criteria of truthfulness should be applied to semantic units rather than to statements, we demand an even higher standard of honesty in human discourse than is achieved by simple propositional logic.

In the end, then, we find ourselves agreeing with both the 'structuralists' and the 'operationalists'. Meaning is indeed inseparable from use, for it is a relation between message and recipient which may differ from one recipient to another. Thus far our sympathies must be with the operationalists. On the other hand, provided that we consider semantic units instead of words, we can agree with the 'structuralists' in believing that the structure is a major determinant of selective function. Different though they are in detailed contents, our brains embody basic principles of organization that have to be reckoned with — and in that sense reflected — in any linguistic system. Indeed I would go so far as to hope that when — if ever — the outlines of the 'shunting-yard' of our brains become a little clearer, one of the tools that will help us further to unravel its structure may well be the analysis of human language.

POSTSCRIPT

A standard source of the line of thought we have called 'operationalist' is Wittgenstein's Philosophical Investigations.⁴⁰ A more 'structuralist' approach is taken by Ullman in The Principles of Semantics.³⁸ Contemporary trends may be followed in: The Structure of Language, edited by J. A. Fodor and J. J. Katz.⁵

A recent general review entitled 'Opinions about Language' is given in: What is Language? — A New Approach to Linguistic Description by R. M. W. Dixon.³

CHAPTER 4

What Makes a Question?

In our search for a conceptual bridge between mechanism and meaning, we have so far considered only indicative sentences. The third of these introductory talks tackles the problem of defining the meaning of interrogative utterances in a mechanistic context. The notion of 'conditional readiness' is here illustrated in terms of a switchboard (of a rather indeterministic sort) instead of a signal-box; but the approach is essentially the same as in Chapter 3.*

Asking questions, like walking, is something we learn to do from infancy — as every parent knows all too well. It comes so naturally, in fact, that to ask how we do it may seem more than a little queer. And yet, when you think about it, there is something rather remarkable about this ability of ours to coax or wring information out of one another just by uttering a few words of our own; and when you come to study communication between human beings from a scientific point of view, as the new science of 'information theory' tries to do, it is not at all obvious what distinguishes a question from all the other noises a man can make, or how to describe the process by which questions produce the effects they do. This is the problem I want to discuss — the problem of describing scientifically how questions work.

* Broadcast on B.B.C. Third Programme and first published in *The Listener*, May 5th, 1960.