3. Do we *teach* children to speak?

3.1 Some popular ideas

What's difficult to understand about a child acquiring language? Isn't it self-evident? Doesn't everyone know how it happens?

The usual belief goes something like this: one day the mother's leaning over the end of the cot, talking to the baby and perhaps trying to get it to say its first word. 'Say Mum-Mum,' she says, or words to that effect, and after a few gurgles and splutters the baby manages to say something that sounds a bit like 'Mum-Mum'.

Perhaps the baby has just opened its lips and made a natural sound; but no matter. The mother will seize on this performance with delight, smile very warmly and repeat the little game. Her pleasure is infectious, and the baby enjoys the game as much as she does. It smiles, and maybe even laughs.

Pretty soon, the baby is very good at responding with 'Mum-Mum', so the mother introduces another word – perhaps even 'Dad-Dad'. After a number of trials, the baby manages to say something pretty close, and is rewarded with more beaming smiles.

The baby's vocabulary is built up word by word, until it has quite a repertoire of words it can say: probably the names of brothers, sisters and animals, as well as the names of various objects and activities.

Everyone knows that this is so, don't they? And the next part of this procedure is that the mother teaches the baby phrases, one by one, and then sentences, and then longer utterances, until the baby has learnt the language. So what's the problem?

The problem is that that's almost certainly *not* the way the baby acquires language. Well, it possibly *is* the way the baby learns those first dozen or so words, but it can't be the way most of it proceeds, as we'll see a little later.

3.2 The sources of these beliefs

If the account given in the preceding section isn't correct, why do so many people believe it? First of all, because it seems plausible. Everyone has seen mothers go through those steps with a baby, and the baby does, after all, learn to talk. But in spite of this apparent plausibility, it's worth repeating that there's good evidence that those procedures don't teach the baby the language, though they might get it started.

The second reason people believe that this is the way children learn their language is that in the past, it was more or less what scientists believed. It became the standard account, and to most scientists concerned with these questions (i.e. linguists and psychologists) it seemed self-evident, too.

It will help our aims if we have a detailed look at the set of beliefs that those linguists and psychologists had, a set which was collectively known as behaviourism. Although I'll maintain that it was a mistaken set of beliefs, we can learn a lot by seeing why it was mistaken.

The journey we're about to undertake will perhaps seem surprising to the reader. In order to discuss the different views on how children acquire their first language, we're going to look at different attitudes to what happens in the brain/mind when people speak. It may seem that we're straying far from what we set out to do; but I think you'll come to see that all this is relevant, and that it will lead us to considerable illumination about how the different points of view came about, and what different claims their proponents make.

In the preceding paragraph I referred to the brain/mind, and I'd now better discuss what the difference is. When linguists, psychologists and other cognitive scientists think about the physical aspects of the brain, they call it that: the brain. But when they concern themselves with its psychological aspects, they tend to call it the mind. The relationship between the mind and the brain has been a continuing problem at least since Descartes. For a very interesting discussion of this matter, see Chomsky (1995a, 1996b), which is reprinted in Chomsky (1997).

Linguists tend to talk of the mind rather than the brain, because the things that linguists study are abstract, and it isn't clear how they relate to the physical side of the brain. But Chomsky (1980: 31) claims, 'We may think of the study of mental faculties as actually being a study of the body – specifically the brain – conducted at a certain level of abstraction.' As you can see, there's still room for a little mystery.

3.3 Early behaviourism

3.3.1 Pavlov's research

Ivan Pavlov, a Russian physiologist, showed as early as 1906 that certain behaviour by animals could be described in mechanistic terms and could therefore be predicted and controlled.

Pavlov's first step was to attach some tubes to the saliva ducts in the tongue of a dog, and measure how much saliva came from the tongue when the dog had a plate of meat put in front of it.

After the dog had been trained to expect the meat to be presented in this way, Pavlov added a new element to the procedure: he not only presented the plate of meat to the dog, but also rang a bell at the same time. He estimated that after this procedure had been followed a number of times, the dog would associate the sound of the bell with the presentation of the meat, and that it would salivate even if the meat wasn't presented, as long as the bell was rung. He then showed experimentally that this was indeed the case. The dog had been 'conditioned' to respond to the bell, and the bell was called the 'conditioned stimulus'.

What Pavlov had shown in this remarkable experiment was that one stimulus could be substituted for another, so as to elicit the same response. By using conditioning, it was possible, at least in part, to predict and control animal behaviour.

3.3.2 J. B. Watson

In 1913 J. B. Watson, an American psychologist, launched a new model of psychology and called it 'behaviourism', or, more accurately, since Watson was an American, 'behaviorism' (with no 'u'). It attempted to make psychology into a rigorous science – every bit as rigorous as the established sciences like physics and chemistry.

3.3.3 The behaviourist conception of science

The success of the 'hard' sciences seemed to be in part due to the fact that they followed 'scientific method', a set of procedures which were supposed to be used by all scientists. The only kind of evidence that was scientifically respectable was evidence that could be publicly inspected,

was open to scrutiny by other scientists, and was therefore considered to be objective.

Scientific method was adopted wholeheartedly by behaviourist psychologists, who were (quite naturally) intent on making their science as respectable as the very successful physical sciences. They must have become all the more convinced that they were on the right track when they saw the spectacular advances in physics during the 1920s and 1930s.

But as a matter of fact, so-called 'scientific method' wasn't consistently followed in the hard sciences; see Bohm (1983) and Collins and Pinch (1993). At best, scientific method seems to have been an ideal which was constantly breached, even by those who professed to believe in it. It was not a method for making scientific breakthroughs, but a method for writing up the research later (and not always even that).

Just as religious converts are often more zealous than the other adherents of the faith, many psychologists became more 'pure' about scientific method than most scientists. They insist, for instance, more than physicists usually do, that all accounts of the research must be written in an 'impersonal' style, in which passive sentences are favoured over active ones, and it's forbidden to use the pronoun 'I'.

The myth is that the 'scientific method' style is more 'objective'. But that style doesn't guarantee objectivity, and, conversely, the writer can be just as objective without it.

These beliefs about the nature of science raise an important issue for us in our quest to find out how children acquire their first language. If you believe that the only scientific evidence is what can be publicly observed and inspected, it follows that feelings, thoughts, mental pictures, ideas, meanings and a host of other things that are said to take place in the brain are beyond the scope of science. They simply can't be studied scientifically.

To be sure, under such a view, the physical aspects of the brain are not beyond the scope of science. Operations on the brain can be observed, and the neurological structure of the brain studied. The trouble is, no matter how you open a brain and inspect it, you won't find anything like thoughts or mental pictures inside. They may exist, said the behaviourists (though most of them were pretty sceptical about that), but science can't study them, because there's no objective way of examining them.

It followed that any scientific account of how children learnt their first language couldn't make any reference to anything that was going

on inside the baby's head, such as thoughts or ideas. It would have to be a description entirely 'from the outside'. That's why this approach was called 'behaviourism', because only behaviour could be studied scientifically. Psychology, in fact, was (and often still is) defined as the study of behaviour.

3.3.4 Are human beings just complex machines?

Along with the view described in the preceding section went another: that human beings are basically complex machines. For thousands of years there have been two competing views about how much control human beings have over their own actions. One view is that we're free to decide what actions we'll take. People who accept that are said to believe in 'free will'. The other view is that we have no control at all over our actions, for everything we do is the result of the forces acting on us in our environment. We may think we have freedom to choose which way we'll act, but that's an illusion. People who accept this are said to believe in 'determinism', because they believe that everything is determined by forces acting on the individual.

Suppose someone is tempted to steal some money. According to the determinist view, if the forces acting to make the person steal it are stronger than the ones acting to prevent them from doing so, then they will steal the money; otherwise they won't.

This view has a certain plausibility about it. The machines that we use in everyday life, such as washing machines and vacuum cleaners, act according to what buttons we push. We are the forces manipulating them: they start if we press a certain button, and they stop if we press a certain button (which may be the same one). In the case of a washing machine, all its actions after starting are governed by the settings we make before we start it. It can't do anything that it hasn't been designed and programmed to do. Perhaps human beings are like that, too, only much, much more complex. At least, that's what a determinist thinks, and that's what behaviourists thought, too.

Perhaps you find this idea, that we could all be just complicated machines, a rather repulsive one. But before you rush to embrace the opposite point of view, that human beings are free to choose what they do, you should realize that this view is not without its difficulties, either.

What does it mean, to say that we're free to choose? Does it, for instance, mean that our actions are not caused? If so, they must be random, and although this might be a kind of freedom, it isn't a very satisfying one. In what sense, then, are we free?

The discussion so far may seem rather surprising. You opened this book, presumably, wanting to find out how children were able to learn to use their first language, and here we are, discussing whether human beings have free will or not.

The reason is that the use of language is something that is very close to the essence of being human. Our answer to the question of how children acquire their first language will depend on the nature of human beings. If we're just complicated machines, then one kind of answer will suffice; but if we're more than that, if we have some kind of freedom in choosing what we want to say, then a rather different kind of explanation will be necessary.

It isn't altogether clear to what extent the brain is just a complex machine, and to what extent human thinking is somehow 'free'. Although I believe that the behaviourists were wrong in trying to describe people as complex machines, I don't think they were silly for making the claim, in the era in which they first made it.

3.3.5 The work of Leonard Bloomfield

Leonard Bloomfield, a linguist, published his ideas on language in his book *Language* (1933). It took a behaviourist stance.[1] He points out (p. 32) that human conduct, including speech, has great variability, and that this has led to the development of two opposing theories about it.

Bloomfield describes the first theory as follows:

> The mentalistic theory, which is by far the older, and still prevails both in the popular view and among men of science, supposes that the variability of human conduct is due to the interference of some non-physical factor ... that is present in every human being. This spirit, according to the mentalistic view, is entirely different from material things, and accordingly follows some other kind of causation or perhaps none at all.

Bloomfield and other behaviourists were at that time causing a swing away from that view to the second kind of theory, which was known as the *materialist* or *mechanistic* theory. This claims that speech is so

complex only because the human body is a very complex system. Bloomfield goes on (p. 33):

> We could foretell a person's actions (for instance, whether a certain stimulus would lead him to speak, and, if so, the exact words he will utter), only if we knew the exact structure of his body at the moment, or, what comes to the same thing, if we knew the exact make-up of his organism at some early stage – say at birth or before – and then had a record of every change in that organism, including every stimulus that had ever affected the organism.

Now, Bloomfield was quite aware that the matter was complicated. Even if we know a lot about a speaker and about the stimuli to which he is subject, he wrote, 'we usually cannot predict whether he will speak or what he will say' (p. 32). Nevertheless, he believed that it was possible *in principle* to predict people's actions, including what they say.[2]

We'll be returning later to the conflict between the two theories cited by Bloomfield. In fact, it will play a continuing part in the story that is to be told, though changing as it goes along.

3.4 The work of B. F. Skinner

The work on conditioning done by Pavlov was later to be of considerable interest to B. F. Skinner (1904–90), an American who was one of the most famous psychologists of his day. Skinner too carried out some very impressive conditioning experiments with rats and pigeons.

Pavlov's famous experiment had simply involved a physiological conditioning, in which the change in response took place within the body of the dog. But Skinner wanted to see whether the external behaviour of animals could be changed by conditioning.

He developed the so-called Skinner box, which was set up in such a way that if a rat in the box pressed a lever, a food pellet would be released. The rat did it the first time more or less accidentally, but soon learnt to press the lever at will. More complicated conditions for releasing the food pellet were then introduced; for instance, the rat might have to press a light switch as well as the lever before it would be rewarded.

Skinner did other notable conditioning experiments, too. In one famous one, he taught pigeons to walk in the shape of a figure eight, simply by rewarding them with food pellets every time they happened to walk in the right direction to achieve part of this feat. If a pigeon just

by accident walked in a brief arc that conformed to a part of a circle, it was rewarded. Then, bit by bit, Skinner extended the proportion of the circle that would be rewarded, and gradually the pigeons learnt to walk around in a circle, in order to achieve the reward of the food. Finally, Skinner taught them to reverse direction and walk the other way in a circle, so as to complete the figure eight.

This method of rewarding 'correct' behaviour was called *reinforcement*. Skinner's experiments were very successful in showing that certain behaviour in animals could be developed and controlled by conditioning and reinforcement.

Next, Skinner wanted to apply his techniques of conditioning and reinforcement to human behaviour. In everyday talk, we say that people have been 'conditioned', and that's why they act in a particular way. This popular usage will give you a rough idea of the technical meaning of the term, though psychologists use it in a rather more precise way.

Skinner held the two most prominent beliefs of behaviourists, already described: (a) that behaviour must be described entirely in terms of what is observable, without any reference to what goes on in the head; and (b) that human beings are just complex machines.

It's quite possible that some of our habitual behaviour can be accounted for in terms of conditioning and reinforcement; for example, much of our behaviour while driving a car may be of this kind. But Skinner then made a very important leap. He claimed that conditioning could also account for human language learning and language behaviour.

His account of the nature of language and language acquisition was presented in his book *Verbal Behavior* (Skinner, 1957). The book had been in preparation for some twenty years before that.

Skinner felt that his work on animals could be 'extended to human behavior without serious modification' (p. 3). And the behaviour to which it could be so easily extended included language behaviour. 'The basic processes and relations which give verbal behavior its special characteristics are now fairly well understood,' Skinner wrote. 'Much of the experimental work responsible for this advance has been carried out on other species, but the results have proved to be surprisingly free of species restrictions.' Thus Skinner hoped to be able to describe the whole of human behaviour, including language behaviour, in terms of stimulus, response, conditioning and reinforcement.

Suppose someone uttered the word 'chair' in the presence of such a piece of furniture. Skinner would say that certain properties in that

piece of furniture were acting as stimuli to elicit that particular verbal response. If it had been a different piece of furniture, say a bed, then the rather different properties of that piece of furniture would have caused a different response: the word 'bed'. Thus, what we say is always a response to some stimulus (or stimuli) in the environment. All behaviour can therefore be explained in terms of the 'building bricks' of *stimulus* and *response*.

I happen to think that Skinner's view of language behaviour is wrong, but it seemed reasonable for him to put it forward at that stage in history. The idea that very complex happenings can be built by multiplying some very simple ones is certainly plausible. This becomes evident when we consider the way in which all matter is built up of atomic particles, or the way in which computers are programmed in very complex ways, simply by using the fact that current in a circuit can be either flowing or not flowing – on or off. These two states can be represented by 1 and 0 respectively, in very intricate arrangements.

Another example of stimulus and response that Skinner mentions is that of a student who's learning to pick the composer of an unfamiliar piece of music, or to name the artist who painted an unfamiliar picture, or the school to which the artist belongs. Certain properties in the music might lead the student to respond with 'Mozart' in the one case or with 'Dutch' in the other. These responses are then reinforced by the community with 'right', or punished with 'wrong'.

Still another example is that of the child who's doing arithmetic and is praised and thus reinforced if (s)he gives the response 'four' to the question 'two and two?'. We will be looking at some crucial criticisms of Skinner's proposal in Chapter 5, but it's perhaps worthwhile commenting here that this account is inadequate for the general case. What about the response 'ninety-two' in response to the stimulus 'thirty-five plus fifty-seven'? Many people are capable of giving this response, perhaps after a hesitation of a second or two. It is surely absurd to claim that such a response is made because of previous reinforcement in giving that answer.

Or, for that matter, what about the answer 'six hundred and twelve' in response to the question 'two hundred and seventy-four plus three hundred and thirty-eight'? Surely anyone who can give this answer engages in some mental computation, but Skinner can't allow such an explanation, since for him anything that goes on in the brain is outside the scope of scientific explanation. So he persists in considering only the externals of behaviour.

The interesting thing is that what Skinner presented as his theory of how children learn language was basically the same as the popular theory described earlier in the chapter – only, of course, Skinner's description was more complex and technical. In fact, the reason that the popular theory is so widely held is possibly because Skinner's work had such a strong influence. But it may also be that the popular theory has been around a lot longer than Skinner's work; I somehow suspect it has.

3.5 Why can't that be the way children acquire language?

Why am I so confident that that couldn't be the way children acquire their native language? Because it implies that the method that the mother uses to teach the baby is simply that of providing a model utterance for it to imitate. We would then have to assume that the baby stores up each utterance in its memory and on later occasions produces one of the utterances from its store, just when it's needed. That's what I call the 'human tape recorder' theory of language acquisition. Well, what's wrong with it?

The answer to that question was expressed very well by George Miller, an American psycholinguist, many years ago (1970: 82–3).

> If you interrupt a speaker at some randomly chosen instant, there will be, on the average, about ten words that form grammatical and meaningful continuations.[3] Often only one word is admissible and sometimes there are thousands, but on the average it works out to about ten. (If you think this estimate too low, I will not object: larger estimates strengthen the argument.) A simple English sentence can easily run to a length of twenty words, so elementary arithmetic tells us that there must be at least 10^{20} such sentences that a person who knows English must know how to deal with. It would take 100,000,000,000 centuries (one thousand times the estimated age of the earth) to utter all the admissible twenty-word sentences of English. Thus, the probability that you might have heard any particular twenty-word sentence before is negligible. Unless it is a cliché, every sentence must come to you as a novel combination of morphemes.[4] Yet you can interpret it at once if you know the English language. With these facts in mind, it is impossible to argue that we learn to understand sentences from teachers who have pronounced each one and explained what it meant.

Note that George Miller was talking only of the twenty-word sentences of English. There's good reason to think that if we take all sentences into account, there are an indefinitely large number of sentences and potential sentences in the English language.

Mathematicians have made us familiar with the fact that there's an infinite set of numbers. If anyone told us the highest number they could think of, we could always exceed it by simply adding 1 to it; so there's no end to how many numbers there are. What I'm claiming is that something similar applies to sentences: there's no end to how many there are. If you haven't met that claim before, you might find it incredible, but it's nevertheless true.

It's easy to demonstrate trivially that it must be true. Consider a sentence such as *There were two grains of sand in the box*. It's possible to replace the word *two* by the word(s) for any number up to infinity.[5]

While granting that that gives an infinite number of sentences, you may nevertheless feel that's a pretty uninteresting kind of infinitude. Agreed. But there are several other pieces of evidence that might convince you that there are an infinite number of potential sentences in the English language (or any other language).

Try a simple exercise. Go to the nearest large library – maybe a national or state library, or a university library. Choose a section of the library at random. Point your finger at a random shelf, and then run your hand along the shelf and stop it at a random book. Take the book down and open it at a random page, then run your finger down the page and stop it at a random sentence. And then see if you can find EXACTLY THE SAME SENTENCE anywhere else in the library. You may well spend the rest of your life at the task.

Of course, I know you won't do that, but I think even if you do the exercise for a little while, you'll soon become convinced that there's a great multitude of different sentences – many more than you would have imagined. Remember, it has to be exactly the same sentence, not just one that resembles the first one you put your finger on.

Let's come at the matter another way. Those of you who are native speakers of English (and many who aren't) will be familiar with the rigmarole that's taught to children, which begins with the words 'This is the house that Jack built.' As you know, the formula keeps expanding, bit by bit, until eventually it becomes 'This is the cock that crowed at dawn and woke the priest all shaven and shorn that married the man all tattered and torn that loved the maiden all forlorn that milked the cow with the crumpled horn that tossed the dog that worried the cat

that chased the rat that ate the malt that lay in the house that Jack built.'

You'll notice that this is all one very long sentence which has constantly been extended by adding another clause. The question now is, how many more clauses could we add to the rigmarole and still have a grammatical sentence of English? The answer is surely that we could add an indefinitely large number without becoming ungrammatical. We would, of course, become boring after a while, or we would stop because we had run out of breath, or because all the seconds in our lifetime had elapsed. But notice that none of these things would render the sentence ungrammatical. In fact each of those events has nothing to do with grammar.

Since we could have stopped the sentence at the end of any of the clauses, each of these potential stopping-points marks the end of a potential sentence. And since the sentence can be infinitely long, there's an infinite number of potential sentences implied.

To repeat George Miller's point, it is impossible to argue that we learn to understand sentences from teachers who have pronounced each one and explained what it meant.

Let's play just one more game. Get a piece of paper and write down the longest sentence you know. However long a sentence you write, I will then suggest to you that you actually know a longer one, because if I begin with the words 'Jack said that . . .' and then add on your long sentence, you'll almost certainly agree that that's also a sentence of English.

Suppose you've already thought of that, and have started your sentence with 'Jack said . . .' Then I'll suggest to you that if I begin with the words 'Mary claimed that . . .', then carry on with 'Jack said that . . .' and complete your very long sentence, you'll almost certainly agree that that too is a possible sentence of English.

To cut a long story short, I'll always be able to add another clause to the beginning so as to make a longer sentence than the one you suggested. It's similar to always being able to add one number, however high someone counts. Again, the sentence might become boring, but it wouldn't become ungrammatical.

There's another reason why I don't believe children learn the language by imitation from their parents and others. When they make so-called 'mistakes', children often say things that no adult would say, and which are highly unlikely to be failed imitations of what adults would say. I'm thinking of a case where a child says *I goed* instead of *I*

went. *I goed* sounds nothing like *I went*, which is what an adult would say, so it can hardly be a failed imitation. But there's more to it than that. The form that the child produces, *goed*, is exactly what we might predict the past tense of *go* would be, if the language were regular, and all past tenses were formed on the analogy of *walk – walked*; *love – loved*; *start – started*, etc. These three are not phonetically identical in their past-tense endings, but the endings are all forms of the same morpheme; that is, despite their phonetic differences, the words have the same grammatical significance. This means that the child is acting as though (s)he's constructed a rule to which many of the verbs in the language conform, but has over-generalized it so that it's wrongly applied to a verb which is an exception to that rule.

What I've been criticizing in this chapter is behaviourism as a theory which can explain normal language acquisition. I'm told that speech pathologists have applied uses of behaviourism which are said to be helpful in the treatment of both children and adults. That may be an area where conditioning does have a role to play. I'm a great believer in the dictum 'If it works, use it.' Nothing beats that.

Although I've given some reasons above for being suspicious of Skinner's claims about normal language and language acquisition, I haven't yet presented the main criticisms of his theory. This is because Skinner's most trenchant and cogent critic is Noam Chomsky, whose work will be described in Chapters 5 and 6.

Notes

1. Neil Smith points out to me that Bloomfield was actually a behaviourist for only part of his career and that his early mentalist and his late descriptive work ignored his own precepts.
2. I am indebted to Neil Smith for his remark (in a private communication) that 'It's ironic that quantum indeterminacy [Heisenberg] saw the light of day in the hard sciences in 1927.'
3. For example, if you interrupt a speaker after he has said 'The other night, I wanted . . .', it's possible that the utterance could have continued with the words *a drink*; *to visit a friend*; *Sue to see me receiving a prize*, etc. In each case the continuation would be both grammatical and meaningful.

4. Morphemes are the smallest meaningful parts of words. So, *unusual* consists of two morphemes, *un-* and *usual*; *hotel* consists of only one morpheme; and *playgrounds* consists of three, *play*, *ground* and *-s* (signifying plurality).
5. This assumes a box of infinite size, as well, but imaginary boxes will do just as well as real ones for the purpose of judging whether the sentence is grammatically possible.