Contents

CHAPTER 1

Theory in Language Study

This book is about the ways in which theory can help us to think about language. It is about how theory can get us started in looking for answers to the significant and interesting questions that thinking about language raises. It is important to remember that we are not doing something different or unusual when we use theory in this way. It does not make much sense to think about theoretical approaches to language as forming a specialist branch of linguistics, distinct from normal or non-theoretical linguistics. In fact, when it comes to doing linguistics it is almost impossible to avoid theory of one sort or another. As soon as linguists make an assumption about language, or choose a method by which to investigate it, there is a theoretical aspect to what they are doing. Later in Part I we will look at three different assumptions about language, and consider the effects these have had on how people have studied language and what conclusions about it they have drawn.

The effects of assumptions about language are not restricted to big theories of language, or to the differences between the major schools of thought in linguistics. Even apparently straightforward descriptions of language, including those discussed in the other companion books in this series, are based on theoretical commitments. For instance *Discovering Language* (Jeffries, 2006) discusses the elements that make up the phrases, clauses and sentences of English. This is not the only possible way of describing a language; it is based on some ideas that linguists have developed, and that they have found useful, to show order and pattern in what might otherwise seem chaotic and random. *Studying Language* (Clark, 2006) includes practical information on how to record and transcribe tape recordings. This too draws on some assumptions: in particular the assumption that examples of 'real life' language use is an important source of evidence in linguistics. This is an interesting idea, but it is far from being self-evidently true. We will see soon that many linguists

would simply not agree with it. In general, differences in opinion about how to go about doing linguistics, or arguments about methodology, tend to mask much bigger underlying differences. They generally result from differences over the question of what the subject matter of linguistics actually is.

So linguists start out with some assumptions about language, and these influence how they decide to investigate it and may well have an effect on what conclusions they reach. Sometimes these assumptions are made clear at the start of a book or article, but sometimes it takes a bit of effort to work out what they are. It is a good idea to get into the habit of spotting such assumptions in whatever you are reading. This will make it easier to assess the significance of the argument and the conclusions, and also help you to compare different pieces of work that have drawn on different assumptions. As a first example, here is the very beginning of *Word and Object*, a collection of essays by the American philosopher W. V. O. Quine:

> Language is a social art. In acquiring it we have to depend entirely on intersubjectively available cues as to what to say and when. Hence there is no justification for collating linguistic meanings, unless in terms of men's dispositions to respond overtly to socially observable stimulations (Quine, 1960, p. ix).

Quine is well known for his belief that language exists only as a form of behaviour: that in studying language we must look only at what people do on particular occasions. We will look at this belief and its implications in Chapter 2, but we can see his position set out clearly and explicitly in these opening lines. Quine defines language entirely in terms of social interaction. It is an art, or a skill that people employ when interacting with each other. Learning it involves picking up cues, or indications as to how to employ this skill. Quine rules out the possibility of discussing meaning in language ('collating linguistic meanings') in any way other than in relation to how people behave in certain situations. In all cases, what linguists are describing must be available for observation. Any utterance they discuss must be made overtly, and the relevant context must be observable. Quine goes on to discuss language in complex and intricate ways in *Word and Object*, but everything he says about it is based on these specific premises.

Compare Quine's opening remarks with those of another important thinker about language. Noam Chomsky has very different views from Quine on what language is, and we will consider these views in more detail in Chapters 2 and 9. Very briefly, Chomsky sees language as a formal system; a language is a set of grammatical rules that exist in people's minds and is capable of producing all the possible sentences. When we study language we have to be careful not to be distracted by all sorts of social and personal

factors that might influence how people speak but are nothing to do with that set of grammatical rules. If we look at just the first two sentences of Chomsky's first book, *Syntactic Structures*, we can identify some of these ideas:

> Syntax is the study of the principles and processes by which sentences are constructed in particular languages. Syntactic investigation of a given language has as its goal the construction of a grammar that can be viewed as a device of some sort for producing the sentences of the language under analysis (Chomsky, 1957, p. 11).

Chomsky tells us that there are 'principles and processes' that put together the sentences of a language, and that the business of syntax, which for Chomsky is the central part of linguistics, is to compile a list of these principles and processes, or a grammar of the language. However there are other views in this short extract that are not made so obvious but are equally important to understanding the ideas about language that are being put forward. Chomsky assumes that a language is essentially made up of sentences and that these sentences can be studied in their own right. This is an interesting point of view, but it is far from being obviously true and, as we will see, many linguists would disagree with it. Chomsky is further assuming that sentences exist independently of any situation or occasion of use. There is simply no mention of society or contexts, or even of the people who use a language. This is all very different from Quine's bold declaration that language exists only insofar as it used in society.

One way of explaining this is to say that Chomsky's ideas are 'reductionist'. That is, he tackles an apparently unmanageable task, that of explaining everything about human language, by concentrating only on one aspect of the problem. He reduces the task to a discussion of simpler and therefore potentially more manageable features. For Chomsky, explaining language does not involve taking on the full range of factors that might affect how people speak. It is the much more restricted, although still immensely complex, task of describing the grammatical rules that produce sentences. In order to recognise this belief, you have to read his words carefully and think about what is left unsaid as well as what is made explicit.

Many linguists have criticised Chomsky for his reductionism, and we will look at some of their arguments in more detail later in Part I. One branch of linguistics that has taken a broader approach, arguing that explaining language is not just about grammar, is sociolinguistics. Linguists in this branch claim that any explanation of language has to take account of the social situations in which language occurs and the social factors that affect how speakers relate to each other. In sociolinguistics too we can find some assumptions that influence the process and outcomes of the study of lan-

guage. In order to understand this we will consider Bernard Spolsky's *Sociolinguistics* (1998). This is a very different type of book, with a different purpose, from Chomsky's *Syntactic Structures*. It is a textbook, designed to summarise and explain other people's work, rather than to put forward a new set of ideas. And of course it is concerned with the ideas of sociolinguistics rather than those of syntactic theory. In fact, in the opening pages of his book Spolsky makes a point of drawing a distinction between Chomsky's assumptions about language and those of sociolinguistics:

> Sociolinguistics is the field that studies the relation between language and society, between the users of language and the social structures in which the users of language live. It is a field of study that assumes that human society is made up of many related patterns and behaviours, some of which are linguistic (Spolsky, 1998, p. 3).

Here too there are both explicit and less obvious claims about language and linguistics. Spolsky offers a clear account of the main ideas in sociolinguistics. Speakers of languages live in societies, so in order fully to understand language linguists must take into account the structures of society. Language is a form of behaviour, and as such is comparable with a number of other forms of behaviour. But the assumptions do not end there, and in fact we can detect some reductionism even in this statement of the aims of sociolinguistics. In order to do sociolinguistics at all it is necessary to accept that language is in some way a distinct, identifiable type of behaviour, one that can be isolated from all other forms of behaviour. Otherwise sociolinguists would not be able to say that some types of behaviour are linguistic and others are not. Furthermore the very idea that people can be described as 'users of language' implies that language exists separately from its users or from any instance of linguistic behaviour. These may well seem like reasonable assumptions to make, but again they are far from accepted by everyone. Above all it is important to remember that they are assumptions rather than obvious facts, and that the way that sociolinguists work and the conclusions they draw may well be influenced by them.

Even those approaches to language study that are sometimes said to be theory-free, or at least based entirely on observable facts rather than speculation, are necessarily involved in some founding assumptions about language and some reductions. Corpus linguistics is a case in point. We will consider this approach to language study in Chapter 2. It involves looking at large quantities of written or spoken language, recorded in electronic form on computers, in order to find out about various features of the language under investigation. Sticking for the time being with opening lines, here is the one from a textbook on corpus linguistics: 'Corpus linguistics is perhaps best

described for the moment in simple terms as the study of language based on examples of "real life" language use' (McEnery and Wilson, 1996, p. 1). This is a clear and apparently straight-forward definition, but it is not free from assumptions and reductions. First of all there is the reference to 'real life' examples. This suggests that text recorded on a computer some time after the original event of writing or speaking is authentic, or true to life. It may well be a necessary assumption to make for corpus linguistics to get off the ground, or for corpus linguists to be able to say anything interesting about language. But it is not obviously or uncontroversially true. It might be possible to claim that once corpus linguists take an example away from the context in which it was originally produced they are altering it in significant ways. In effect they are relying on the assumption that one can extract language from a situation, or reduce a complex action to a purely linguistic element, and still be left with something authentic to study. There are further assumptions implicit in the phrase 'language use', which we have already considered in relation to sociolinguistics. Again the authors of this textbook are assuming that language has the potential to be 'used' in different circumstances, hence that it in some sense exists independently of those individual instances of use.

It may well be the case that linguistics is always, necessarily, reductive; it is just the choice of specific reductions, or perhaps the scale of reductionism, that varies. In order for language study to be possible at all, it is necessary to abstract away from the coincidental features of language use and concentrate on what is essential to language itself. This leaves open the question of which features are coincidental and which are essential, and language theorists have made widely different decisions on this issue.

In the next chapter we will look at three different theoretical approaches to the study of language, and consider what has motivated each of them and what implications they have. Each of these approaches is based on a particular view of what linguists study. This might at first glance seem like a very odd statement; surely what linguists study is language? But as we will see, the question of what language is, or how we can best define the subject matter of linguistics, is far from being a straightforward matter. Each of the three approaches suggests a specific answer to the question 'What is language?'

This question is different from another important question that looks deceptively similar, namely 'What is *a* language?' Attempting to answer this second question makes linguists think about different possible ways of defining any particular language such as English, Russian or Urdu. It means deciding what makes the difference between saying that two people speak different languages and saying that they speak different dialects of the same language. And it involves explaining why we might want to say that although a language changes over time it remains fundamentally the same language. But

many linguists see the question 'What is language?' as a more fundamental problem, and claim that it is necessary to make some decisions about this before going on to any of the more specific issues.

Before we turn to some of the answers that have been suggested to this question, it is worth thinking about the general issue of why they are important in the first place, or of why theory is necessary or even appropriate in the study of language. It is not hard to imagine the arguments that people who want to oppose theory in language study might put. They might argue that language is a natural phenomenon, and that it is therefore not appropriate to try to explain it by means of some abstract, artificial theory. They might point out that language is far more complex than any neat theoretical description could explain, because it is decided on a day-to-day or even minute-to-minute basis by the needs and wishes of individual speakers who use it in everyday situations. Or they might complain that you can prove any theory you care to come up with, so long as you are prepared to select your examples carefully to match your theory. If we think about each of these arguments in turn we will find some interesting things to say in reply to the critics of theory, and in the process discover something about the importance of theory in language study.

We can answer the first argument without having to deny that language is a natural phenomenon. Indeed the subject matter of linguistics is often described as being 'natural language' to distinguish it from invented, artificial languages such as those used in logic and in computing. Knowing and using a language seems to be a natural part of what it is to be a human being, and human languages are not invented or designed. But to agree that something is a natural phenomenon is not the same as admitting that it should not be explained by a theory. Indeed many theories in other areas of study suggest explanations of natural phenomena. The theory of evolution offers an account of the origin of the vast variety of life forms on the planet, and an explanation of the natural fact that species apparently change over time. The theory of gravity explains why objects fall to the ground when dropped, how planets orbit and why we do not fall off the Earth. And the theory of molecular structure gives us a way of discussing the actions and interactions of the different substances around us. In each case the theoretical explanation does not imply that the subject matter under investigation must also be theoretical or artificial. The point of the theory is to say more than could be said by a simple description of the relevant facts, and in each case the theory has been viewed as successful because it has had something interesting to say about those facts. Of course these three examples are all taken from the physical sciences: from biology, physics and chemistry respectively. Some linguists claim that linguistics too is a science, and that linguistic theory explains the phenomena of language just as, say, physical theory explains the phenomena

of matter, energy and time. Not everyone agrees that linguistics can be classified quite so straightforwardly as a science, but linguists are generally happy to compare linguistic theories to scientific theories, and to see scientific theories as providing a good basis for what a theory should be like. As we will see in the rest of Part I, linguists are not always in agreement about what makes a theory truly scientific.

We can also use this relationship between linguistics and scientific theories in our response to the second criticism. In the sciences the complexity of a subject matter is no reason to abandon the search for a theory to explain it. On the contrary, it is often precisely because a natural phenomenon appears to be extremely complex that scientists look for a theoretical explanation. It is not that they want to 'explain away' the apparent complexity, or to pretend that things are much neater and more straightforward than they really are in order to make life easier. Rather scientists realise that it is only by focusing on certain features of a situation that they can say something constructive and systematic about what is going on. A physicist describing the path of a leaf falling from a tree could not possibly hope to give an exact description of all the relevant factors, including air resistance and wind direction, and the interaction of the shape and structure of the leaf with these. But it is still worth using the theory of gravity to explain something about the movement of the leaf, rather than just giving up because explaining everything is too complicated. A similar argument can be used in relation to language study. What any individual says at any particular moment might well be decided by a hugely complex set of factors, including the personality of the speaker, what mood he or she is in, what – if anything – has been said immediately before. It would be unrealistic to expect to be able to devise a theory that took all these different factors into consideration. It would also be undesirable; a theory that managed to account for everything would necessarily be so broad and so general that it would end up telling us nothing very interesting at all. But it is still possible to say something about the language the speaker is using, even if this forms only one part of the complex situation that holds when the language is used.

The final criticism suggested above, the one about the possibility of proving just about any theory by choosing your examples, is repeated quite frequently and is potentially a serious one. It would indeed be inappropriate to make grand claims for the success of a theory that explained a small selection of examples but could not be applied beyond that limited set. But in fact any linguistic theory that is treated as interesting or significant must be able to cope with examples beyond those it was initially designed to explain. We will consider in Section 1.2 how a theory might be tested in relation to those further examples, and what would determine whether it could cope with them successfully.

As we have seen, the study of language must necessarily include some element of theory, so it can be misleading to talk about theoretical approaches to language as though they were in some way specialised or unusual. As soon as linguists start describing something as complex as language, they have to make decisions about what sort of things it contains, and about which features are similar and which are different. These decisions are imposed on the language rather than just being there in the evidence. The linguists are using some sort of theory of language. However theory in language study also makes it possible to go beyond these types of description to consider more general questions about language, and it is with these that we will mainly be concerned in this book.

When it comes to these more general questions, theory is usually a much more central and prominent part of the study. This is not to say that the theory takes over, or becomes the main focus of study in its own right. But it does mean that the theory allows us to focus on some of the features of language that we cannot so obviously observe and describe. It allows us to consider, for instance, the nature of the relationship between words and the things in the world they describe, or between language and thought. This is part explains why the issue of the appropriate type of evidence, or of what linguists should study, is so significant. If we want to know what speech sounds there are in a certain language, or whether subjects come before verbs or the other way round, we can collect some examples and examine them to find an answer. But there is no obvious evidence that we can collect and examine in search of an answer to the bigger questions. In these cases, theory allows us to go beyond a list of examples to consider, for instance, what makes language unique, and how we use it to make sense of the world around us and relate to each other.

1.1 Types of data

We have seen that different approaches to linguistics can start from different answers to the question 'What is language?' To put this another way, they start from different points of view about the nature of the data in linguistics. 'Data' is the name given to the set of information, or observable facts, in any particular area of study. Strictly speaking any one such piece of information is called a datum, and data is the plural of this word, describing several separate pieces of information, or even the full set of all relevant pieces of information. Nowadays, however, the word data tends to be treated grammatically as though it were singular. You will come across sentences such as 'my data is selected from a range of different sources' at least as frequently as the more strictly correct 'My data are selected from a range of different sources'. To fit in with this practice we will use 'data' as a singular noun.

What counts as relevant data of course varies from subject to subject. The data for an investigation in geography, for instance, might be information about the population in different cities or the rainfall in different countries. Relevant data in economics might include information about prices, incomes and savings. So there is no doubt that the data for linguistic study comes from language, but we are still left with the question of what type of data language has to offer, or of what language actually consists of. When we were considering the ways in which theories become involved whenever we think about language, we compared Noam Chomsky's ideas about what linguists should study with those put forward by sociolinguists. These two sets of ideas include very different takes on what language is, and as a direct consequence the two styles of linguistics use different types of data.

Chomsky belongs to what is known as the 'mentalist' tradition in linguistics. We will look at mentalist approaches to linguistics in more detail in Chapter 2, when we consider the view that language is a state of mind. As we will see, the set of grammatical rules that for Chomsky define a language exist in the minds of people who speak that language. Therefore according to Chomsky linguists are studying something about the human mind. Everything they need to know about a language is contained in the mind of someone who speaks that language. So the best way to study a language is to ask speakers questions about what they know, for instance by suggesting various possible sentences and asking them to judge whether these conform to the rules of their grammar. Not only do linguists not need to look elsewhere for their data, it could actually be misleading for them to do so. If they look at examples of people using language, there is a danger that they might end up concentrating on things to do with context or interpersonal relationships that have nothing to do with grammatical rules.

Because the data for mentalist linguistics comes from asking people to think about what they know about a language, or to introspect their own knowledge, it is sometimes described as introspective data. Mentalist linguists argue that it is perfectly acceptable for them to do their own introspection, and to use their own knowledge of the language as their data. Language data can be provided by any speaker of the language. Mentalist linguists do, however, attach particular significance to the notion of a 'native speaker': someone who has learnt the language by growing up with it and acquiring it as a child, rather than learning it later in life. It is only from native speakers that they can be sure of getting reliable data on a particular language. For this reason mentalist linguists studying a language other than their own need to find one or more native speakers of that language to tell them about the language, or to act as their informants.

The picture is rather different in sociolinguistics. Linguists who are interested in the relationship between language and social context necessarily

have to find out about how language is used in different contexts. Relying on their own intuitive knowledge, or asking informants what they know about their language, will not be enough because it will not produce any information about how different contexts affect how people speak. Sociolinguists, and those working in any of the other branches of linguistics that study how language is used in social situations, go out and collect their data. Generally this involves recording conversations on tape recorders or video cameras, and often transcribing these recordings into a written version. Sociolinguists also sometimes use questionnaires to gather their data, in the tradition of work in social sciences such as sociology.

Differences between mentalist linguists and those working in traditions such as sociolinguistics sometimes lead to disagreements, and even heated debate, about what is the 'right' type of data to use in linguistics. But ultimately this is not a debate that can be settled. Different types of data are suitable for different types of study. As long as there are different ideas about the answer to the question 'What is language?' there will be linguists working with different types of data. And the wide variety of possible answers to that question is one of the things that make linguistics such a fascinating subject.

1.2 Types of theory

Nearly all present-day linguists are keen to explain what they do as 'descriptive', as distinct from 'prescriptive', linguistics. They want to distance themselves from any notion that they are trying to lay down rules for how people ought to behave when using language, or to label some forms of language as correct and others as incorrect. This is what goes on in prescriptive linguistics, such as that found in some old-fashioned grammar books. A prescriptive linguist might tell you that it is incorrect to end a sentence in English with a preposition: *That is the man about whom I told you* is correct, but *That is the man I told you about* includes a grammatical mistake. A descriptive linguist, on the other hand, will be more interested in observing what speakers of English actually do. They might note that the second version is much more likely to occur in spoken and written English than the first version or comment that it will appear more natural to most native speakers of English. Indeed descriptive linguists believe that the rule about not ending sentences with prepositions is a more or less arbitrary one that has very little to do with English itself but has been imposed on the language by prescriptive linguists because of some facts about Latin grammar.

Prescriptive linguistics is generally now seen as marginal and old-fashioned. However there are some current approaches to linguistics that make claims about how people ought to use language rather than describing how

they do use it. For instance feminist linguists and advocates of political correctness argue that the pronouns *he* and *his* ought not to be used as generics in sentences such as *Every student must hand his essay in on time*, because such uses exclude women or present the male as the norm. This might be seen as a form of modern-day prescriptivism, and indeed adherents of political correctness are quite happy to describe it as a justifiably prescriptivist movement.

Theoretical approaches to language generally belong to the descriptive tradition in linguistics. Whatever assumptions about language they start out with, their purpose is to find out something about language as it actually is, not as a purist would like it to be. As we have seen, linguists use theory to say something more about language than they could from just an accumulation of data. However not everyone who uses theory to go beyond the data and consider wider questions uses the same methods. That is, not everyone makes their theory and their data work together in the same way. In order to understand this we need to look a little at how scientists have thought about the function and operation of their theories. This is because ideas about theory that are most closely associated with the natural sciences have also had a significant impact on the ways in which language has been, and still is, studied.

Imagine that you are a scientist interested in finding out about some aspect of the world: say about types of flower. You might well decide to start your study by going out and looking at some flowers and finding out what you can about them. You will probably decide quite quickly that you will need to work with theory to some extent if you are going to get anywhere. Without theory all you can do is accumulate data: perhaps flowers that you have collected, or descriptions of flowers that you have observed. If someone asks you what you have found out so far, all you can do is show them this collection of data. You findings are just as complicated, and just as unstructured, as the data itself.

A good first step might be to try to make some general statements about your data. Perhaps you have collected a number of white daisies, but no daisies of any other colour. Instead of just listing your data (*Here's a white daisy . . . here's a white daisy . . . here's a white daisy*) it would be more useful and interesting to make a general statement, such as *All the daisies I have collected are white*. This is better than the simple list of data but it still does not tell you much about your area of study. It tells you about the particular examples of daisies you happen to have come across, not about daisies in general. Things get a lot more interesting if you try going beyond your immediate data and saying something such as *All daisies are white*. In doing this you are making a general statement about how the world is, based on some specific pieces of evidence. To put it another way, you are making a prediction; you are predicting that if you go on collecting data you will find many more instances of white daisies but no instances of daisies of any other colour.

The method of relating theory and data that we have just been considering is known as the inductive method, or as working by induction. An inductive investigation of any phenomenon starts with, and is driven by, the data. It involves gathering relevant examples of the subject matter under investigation, and then looking closely at this data to discern any general trends or patterns. As well as being a valid scientific method, induction is the method by which we generally make sense of the world around us. In many areas of life we are not constantly surprised by events because we are constantly making predictions about what is likely to happen next. These predictions are produced by generalising from individual pieces of relevant past experience. If we see thick black clouds when we are going out, we are likely to take an umbrella with us. This is not because we know for certain that it will rain, but because we have had a number of past experiences of black clouds being followed by rain. Our past experience gives us a strong enough motive to agree with the statement *Black clouds mean rain*, and certainly a strong enough motive to take an umbrella to avoid getting wet. All this seems no more than common sense. We would not get anywhere in life if we stopped to ponder about the nature and certainty of our knowledge all the time. The prediction that black clouds mean rain is not definite. The many instances of black clouds followed by rain we have experienced in the past do not guarantee that this particular black cloud will be followed by rain on this occasion. But the past connection makes our expectation of rain highly probable, and this is good enough for us.

Some theorists have questioned whether this type of probabilistic knowledge is good enough for science. Take for instance the claim that *The sun will rise in the east tomorrow*. As far as ordinary life is concerned we are nearly certain of the truth of this prediction. There have been so many instances of the sun rising in the east that it seems certain that this will also be the case tomorrow. We have positive knowledge of numerous previous instances of the sun rising in the east, and from these we draw a probabilistic but in practice near-certain prediction. But speaking in strictly scientific terms we do not know that tomorrow the sun will rise in the east. However strong our positive knowledge of past events, it is possible that the sun will not rise at all tomorrow, or will surprise us by rising in the west instead. In a similar although rather less dramatic fashion, however many white daisies we collect our claim that *All daisies are white* remains probabilistic rather than definite. Even a very large number of white daisies does not rule out the possibility that somewhere out there is a red daisy.

There is a further problem for induction as a scientific method, beyond the issue of the probabilistic nature of inductive statements. This problem has to do with the nature of the data that counts as support for an inductive statement, and it takes the form of a philosophical paradox. That is, the nature of

induction seems to lead inevitably to a particular conclusion that is bizarre in itself, and contrary to what we expect. This particular paradox was pointed out by Carl Hempel, a mid-twentieth-century philosopher who began to doubt the validity of the inductive method.

Hempel's example was *All ravens are black*, a claim that we might reasonably reach by induction, after observing lots of examples of black ravens and no examples of ravens of any other colour. Hempel pointed out that, logically speaking, the claim that *All ravens are black* is equivalent to the claim that *All non-black things are non-ravens*. That is, if all ravens are black then anything that is not black cannot be a raven; it must be something else. If *All ravens are black* is true then *All non-black things are non-ravens* must also be true. One consequence of this logical equivalence is that if we find some evidence that supports the claim that *All non-black things are non-ravens*, then this piece of evidence must also support the equivalent claim that *All ravens are black*. Our inductive claim is supported every time we see a black raven. But it is also supported every time we see something that is not a raven and is not black. A white horse, a red apple, a purple tree, a blue cow: an observation of any of these things provides evidence in support of the claim that *All ravens are black*. Here is the paradox. It seems utterly bizarre, and contrary to all our expectations of the 'common sense' nature of induction, that a blue cow should count as evidence in support of the claim that A*ll ravens are black*. But this unacceptable conclusion seems to be a necessary consequence of the nature of induction itself.

Positive knowledge based on generalisations from individual observations was highly valued in the philosophy of science in the middle of the twentieth century, particularly within the school of thought known as logical positivism. In Chapter 4 we will look at the consequences of this approach for claims about meaning in language. As we will see, these claims proved very controversial, and in response to them some wholly different ideas were put forward as to how best to think about language. In the philosophy of science, too, thinkers began to question induction as a method. This was in part because of the paradox about relevant evidence that we have just considered, which seemed to call into question the validity of the inductive method itself. It was also because of worries about how far scientists could get by generalising from particular observations. One person who was particularly associated with the reaction against the inductive method of logical positivism was Karl Popper. The alternative to induction, closely associated with the work of Popper, is nowadays generally known as deduction, or the deductive method of science.

We could summarise the deductive method, in contrast to the inductive method, by saying that it is theory-driven rather than data-driven. The theory is the focus and the starting point, with the data playing a supporting

but vital role as a means of testing the theory. Scientists who want to explain some aspect of the natural world deductively start with a general claim, or hypothesis, about how things are. Popper himself was rather vague about where the all-important hypothesis comes from. He seems to have seen it as resulting from inspiration, rather as if the scientist were a creative artist:

> The question how it happens that a new idea occurs to a man – whether it is a musical theme, a dramatic conflict, or a scientific theory – may be of great interest to empirical psychology; but it is irrelevant to the logical analysis of scientific knowledge (Popper, 2002, p. 7).

In reality the beginnings of a hypothesis are unlikely to be completely detached from the data. Some initial observation, or a very small set of data, may suggest a general explanation that is developed into a theory. If this theory explains the initial observations, then so far it is adequate and useful. But to be a successful deductive theory it must do more than that: it must stand up to the test of further data. In the deductive method the data does not function to provide repeated confirmations of the theory. In fact the theory may be such that it cannot be directly supported by any positive data. The theory may be a hypothesis about the underlying, unobservable causes of observable phenomena.

What is important to a deductive hypothesis, indeed what makes it a valid scientific statement, is not that it can be supported by data but rather that it can be falsified. That is, the hypothesis must make specific predictions that can be tested against data and can in principle be proved wrong. This may at first sight seem like a very strange claim to make. To be of any use a theory must be capable of being proved wrong. But for supporters of the deductive method the only theories that can be taken seriously are those for which it is possible to state clearly what would count as a counterexample. The theory tells us something about how the world is, and makes predictions about what we will find in the world as a result. The relevant data is then used to test the validity of these predictions, and therefore of the theory itself. Scientists who use the deductive method are not looking for lots of data to support their theory. They are on the look-out for any one piece of data that will falsify their theory. If we deductively test our claim that *All daisies are white* we are not interested in accumulating instances of white daisies. Rather we are interested in the possibility of discovering a daisy that is not white. A single daisy that is red or blue or yellow will act as a counterexample to the predictions made by our theory, and will falsify it.

When a falsifying counterexample to a theory is identified, scientists have two options. They may decide that the theory is not going to work; the new data includes serious counterexamples, or examples that go completely

against what the theory predicts and there is no option but to abandon the theory. Or they may decide that the theory can be retained in basically its original form but that it needs to be modified in some way to explain the new as well as the original data. This leads scientists to refine their theory, and yet more data may cause them to refine it further; the theory is in a constant state of modification and improvement. Importantly, there is no need for an unattainable hope of proving the theory. If scientists collect some new data and it fits the theory, that may boost their confidence in the theory but it can never prove that it is correct; it is always possible that the next bit of data may present problems. What is important is that the theory should be falsifiable. That is, the scientists should know with confidence what sort of data would lead them to acknowledge that there was a problem with the theory, and to modify or abandon it. If it is not possible to imagine any type of data that would falsify the theory, then the theory itself is of no value because it can never be put to the test by means of data. To be described as 'unfalsifiable' is probably the worst possible insult for a deductive theory.

On the face of it, adopting a deductive approach might seem to be a very precarious undertaking. Deductive scientists commit themselves to a theory, based perhaps on little more than a hunch, and then live in the constant expectation that the next piece of data they come across might falsify it. Induction might produce probabilistic statements, but the claims of deduction are always contingent. As Popper forcefully claimed, 'every scientific statement must remain *tentative for ever*' (Popper, 2002, p. 280, emphasis in original). But those who support the deductive method argue that it is only by taking these risks that they are able to say interesting things about the world. Using the inductive method can take you beyond a mere list of the data, but not much beyond it. You can generalise over your observations and make predictions about similar observations in the future. But with the deductive method you can go beyond the evidence immediately available to you and speculate about things for which you have no direct evidence, perhaps things for which no direct evidence is possible. As long as your theory makes definite predictions, and therefore can be falsified, you will be making a valid claim about how the world is.

Deduction may have become established as part of the reaction against the problems of induction, but it did not replace induction and become the only legitimate scientific method. And deduction is not without problems of its own. For instance some people have argued that neither Popper nor anyone else has ever successfully defined what counts as falsification. Popper certainly seemed to imply that any one observation that provided a counterexample to a prediction made by a theory would be enough to falsify that theory. But some thinkers have questioned whether that is really a satisfactory state of affairs; should a whole theory really be modified or disposed of

in response to one little piece of data? And even if we accept that a theory can be falsified in this way, we might still want to claim that some theories are interesting and merit careful consideration even after they have been falsified. They may still offer interesting insights or helpful ways of looking at the world. To a strict deductionist, once a theory is falsified it must be abandoned in favour of a modified version or a different theory altogether. Some thinkers claim that it might be worth hanging on to the old theory, even if this means keeping it alongside a new or different one.

A case in point is Newton's theory of gravity. Newton's laws were accepted in science for a long time because they were able to explain a lot about how the world, and indeed the universe, worked. There was evidence, however, that might have been seen as falsifying the laws. In particular the orbit of Mercury did not behave as they predicted. But to abandon the laws altogether for this reason would have left science with no account at all of the movement of bodies. So the useful but obviously flawed account was kept, together with the evidence that technically falsified it. This data was eventually explained by Einstein's general theory of relativity. Newton's theory is so accurate and explanatory, however, that it is still used today in calculations for space travel. Even though it has been falsified it is still useful.

Theorists in all fields generally accept that they are unlikely to come up with a definitive account of their subject matter that will be proved to be correct and will be universally accepted. Indeed it makes little sense to talk of proving any type of theory. If you take the inductive approach you can find as much data as you like to support your theory, but it will always remain probabilistic. As we have seen, if you are deductive in your method the issue of proof simply does not arise. A valid theory is one that has not been falsified – yet. What makes a theory of language or of anything else interesting is the light it can shed on a particular area of inquiry: the particular way of looking at the matter that it suggests. Karl Popper argued that 'Theories are nets cast to catch what we call "the world"; to rationalize, to explain, and to master it' (Popper, 2002, pp. 37–8).

Both deduction and induction have proved to be important tools for language study. When we look at some of the ideas of Noam Chomsky in more detail in Chapter 2 we will see how the deductive method has informed his studies and his conclusions. Chomsky once defended his method against the accusation that it was too speculative and not properly grounded in observation by arguing that serious hypotheses, about language or anything else, will always go beyond the available evidence; 'if they did not, they would be without interest' (Chomsky, 1969, p. 66). For Chomsky it is the big, risky claims made possible by deduction that are the most interesting, not the generalisations of induction. Induction too has its advocates in linguistics. It is used in those branches of linguistics that are based on the analysis of large

amounts of language data. We briefly considered this approach in relation to corpus linguistics earlier in this chapter. Induction is also used – in fact is a guiding principle – in conversation analysis. This is based on careful attention to recorded examples of genuine conversations, and is concerned with the structures and patterns found in even the most casual of conversations. Here is a passage from an introduction to conversation analysis (CA) that is particularly concerned with the issue of methodology:

> CA views the empirical conduct of speakers as the central resource out of which analysis must develop. Furthermore, what is said provides not only the data underlying analysis, but also the evidence for hypotheses and conclusions: it is participants' conduct itself that must provide evidence for the presence of units, existence of patterns, and formation of rules. To this end, CA searches for recurrent patterns, distributions, and forms of organization in large corpora of talk (Schiffrin, 1994, p. 236).

For a conversation analyst, then, the raw data of recorded and transcribed conversation is not used simply as a test for theories of conversation. It is actually the driving force behind any legitimate analysis. If the analyst is interested in identifying patterns and rules, these must be built up from the evidence of the data. The analyst must proceed by making generalisations about the relevant data. For that reason 'large corpora' or huge amounts of data are necessary. It will not be possible just to try out a series of isolated examples to check whether a particular theory is falsified. On the contrary, the more data there is, the stronger will be the analyst's conviction in the conclusions drawn from it. This is clearly recognisable as the inductive scientific method.

Both induction and deduction are still used, and actively defended, as methods in language study. As with so many of the issues and debates we are considering here, it is not a case of one being necessarily right and the other wrong, or of deciding conclusively which is the better. Both methods have things to recommend them and can offer interesting insights into language. The important thing is to be aware of what type of method any individual linguist is using, and to remember the limitations of this method and its possible implications for the resultant account of language. The choices of method that thinkers make are generally related to what they consider to be the most important aspects of the subject under study. So linguists who believe it is most important to study how people generally use language in different situations will be drawn to the inductive method. Those interested in explaining language in terms of unseen but powerful mental capacities will be drawn to the deductive method. So the choice of method often depends on individual theorists' views of what language essentially is. We will consider this issue in the next chapter.

1.3 Further reading

See Clark (2006, ch. 1) for a discussion of issues that are relevant to the question 'What is a language?', which, as we have seen, should be distinguished from the question 'What is language?'; for a description of the processes and methodologies of sociolinguistics, which we have considered in relation to different types of data; and for practical advice on how to use these in your own research. In Chapter 2 of that volume there is a discussion of the methodology of conversation analysis, which we have considered in relation to the inductive method.

Index