Contents

1 Introduction

Language has been an object of interest and investigation for thousands of years. The dialogues of Plato, which form the starting point for most of Western philosophy and science almost 2,500 years ago, contain speculations about language, particularly in the dialogue *The Cratylus*. One of the most interesting things about language is that, despite what you may have been led to believe by the popular literature, linguistic capacities of the sort that humans have don't seem to be possessed by other animals. Furthermore, every human being exposed to a language as an infant will grow into a competent speaker of a human language. In other words, language appears to be a species-specific property – something possessed by all human beings and only human beings. Our language plays a central role in human thought, society and action. Therefore, in studying language, we're exploring in a fundamental way what it means to be human.

1.1 The scientific study of language

The first question we need to ask ourselves is how we want to study language. The approach that we'll be taking in this book is most strongly influenced by the American linguist Noam Chomsky, one of the most important figures in twentieth-century linguistics. We'll see constant references to his work both in the broad outlines of the theory and in the details of particular analyses.

At the most general level, Chomsky suggests that we should investigate language by using the methods of the natural sciences – formulating hypotheses, testing those hypotheses against data, making our theory as simple as possible (as long as we make the right predictions), etc. That's not to say that this approach is guaranteed to succeed. It could turn out that using the methods of the natural sciences doesn't get us anywhere, in which case we'd need to try something else. But it seems right to say that this is where we should start.

On Chomsky's view, it's just a fact about the world that people speak and understand one another, on a par with any other fact about the natural world. I expel air from my lungs which sets air molecules in motion. Those molecules

hit your eardrum and that causes thoughts to arise in your mind which roughly correspond to the content of what I said. This is just another fact about the world, like the fact that things fall down and not up. We want to try and formulate theories which explain why the world is the way it is, and not some other way.

Implicit in this approach is also the assumption that there's a right answer and there's a wrong answer, what philosophers call a 'fact of the matter'. Human beings have some way of speaking and understanding their language. We're constructing a theory about that ability, and that theory is either right or wrong. It either correctly describes the way that human beings *actually* speak and understand their language, or it doesn't. For this reason, it's important for us to be as explicit and clear as possible in constructing our theory. That way, we can more easily check to see whether our theory makes the right predictions or not.

This isn't the only possible approach to language, though. For example, the linguist Zelig Harris (who was Chomsky's PhD supervisor) believed that the study of language was more akin to the study of literature. It was a question of illuminating a subject from different points of view, with no one view being 'right' or 'better' than another. However, I'm inclined to agree with Chomsky that we should begin our investigation of language using the methods of the natural sciences, and at least see how far we can get with it.

What are we studying?

So if we're agreed that we want to approach language using the methods of the natural sciences, where should we begin? First, our ability to speak and understand our language must reside somewhere in the body. But where? The brain seems like a good bet. Certain kinds of brain injuries, such as strokes, can impair one's ability to speak and understand one's language. Similarly, if electrical activity in the brain is temporarily disrupted, this can also create language impairments. (This is commonly done during brain operations to map out which exact areas of a person's brain are used for language, in order to ensure that the surgeon minimizes the removal of tissue from those areas.) By contrast, if we look at the foot or the heart, there doesn't seem to be any injury to those areas which can impair our linguistic abilities in the same way. Therefore, at its most basic level, in doing linguistics we're studying the human brain. Whatever our linguistic capacities are and however they're organized, they're part of the human brain.

Having established that we're studying the brain, and that we want to study the brain scientifically, you might wonder why we don't start talking in this book about various kinds of complicated brain-imaging scanners. Whenever you see a science documentary about the brain, the focus is almost exclusively on CAT scans, PET scans, MRI scans, EEGs and the like. Society seems to create the impression that, in order to do 'Science' (with a capital S), you need to be wearing a white lab coat and use lots of expensive equipment. However, as we mentioned, doing science is really just about having a certain attitude and approach towards what you're studying. As long as you're

constructing and testing theories with the goal of discovering something about the natural world, you're doing science.

In linguistics, it's not that we *couldn't* do brain scans. The problem is that not really enough is understood about the brain at that level for it to be a useful thing to do at this stage. At the low, 'nitty-gritty' level, we're not 100 per cent certain what's responsible for our ability to speak and understand a language. Many people argue that neurons are relevant, but there could also be a role to play for electrical configurations, chemical activity, some combination of these things or maybe even something else that we have yet to discover. Nobody is really sure what to look at. So, while complicated experiments with brain scans are part of the popular idea of 'studying the brain', there's no point in doing them if you don't really understand what the results might mean.

If we want to study the linguistic system of the brain scientifically, and we're not going to be doing brain scans, how should we approach it? What we have to do is treat the linguistic system of the brain like a 'black box'. We'll try to put things into one side of the black box and see what comes out the other side. From careful observation about what goes in and what comes out, we can begin to theorize about what has to be inside the black box in order to explain the results that we're getting. Eventually, we'll get to the stage where we can make predictions: if X goes into the black box, then Y should come out. If Y doesn't come out then our theory is wrong somewhere, and we'll need to change it. In fact, many philosophers see this process of 'falsification' as the litmus test for science. If you're being explicit enough to know that what you're saying is false, and then you try to correct your theory, you're doing science.

1.2 The competence/performance distinction

In this book, we'll be focusing on a particular aspect of the linguistic system, namely syntax. That is, we'll be focusing on sentence structure and 'grammar'. What sort of experiments should we run in order to investigate our grammatical capacities? How do we investigate this grammatical black box? As Chomsky has noted on many occasions, there seems to be a distinction between what people *know* and what they *do*. Put another way, what we actually say during the course of ordinary conversation isn't always an accurate reflection of our linguistic capacities. We can make mistakes, such as slips of the tongue, for example. Chomsky refers to this distinction as the *competence/performance* distinction. Therefore, we don't want to focus our attention on what comes out of our mouths. We want to focus on the underlying system.

Studying behaviour

In the early twentieth century, though, it was common for linguists to study 'linguistic behaviour'. They would record as many speakers and as much natural conversation as they could, and then analyse these recordings and construct a theory based on what was found. This approach would seem to be perfectly in keeping with our goal of investigating language using the methods of the

natural sciences. If we were studying, say, dominance hierarchies in ape groups, we would probably spend a lot of time observing and recording their behaviour in different situations and then constructing a theory from that data.

However, there a number of reasons why that approach isn't the best way to investigate the syntactic capacities of human beings. When you're actually producing speech in real time, there are all sorts of factors which come into play – how tired you are, whether you're paying attention, etc. As a result, *actual* speech is full of 'mistakes', like false starts, slips of the tongue and times when you start a sentence and then don't know how to finish it. I have myself on occasion produced sentences like (1) in conversation:

(1) This is the film that you said that was good.

Now despite having produced this sentence in conversation, I would admit, if it were pointed out to me, that the sentence in (1) is *not* part of my language. Intuitively, I know there's something wrong with (1), and, more specifically, the problem occurs in the phrase 'said that was good'. In order for the sentence to be well-formed, the *that* must be deleted, creating (2). (This is true for most dialects of British and American English, although not all.)

(2) This is the film that you said ( ) was good.

Thus, there is a competence/performance distinction. What I *do* with my language doesn't always accurately reflect what I *know* about my language because of the additional burdens involved in actually producing speech in real time.

This means that we could run into problems if we were to simply record a large amount of natural speech. It would presumably contain 'errors' of the sort we've just seen. But we would have no principled way of deciding what was an error and what wasn't without reference to the speaker's *knowledge* of what is an error and what isn't. In real life, you know whether a person has been interrupted, or broken off a sentence and started a new one, but that's because you're filtering the data through your own knowledge of the language system.

Another problem is that there are sentences which are perfectly grammatical, but which come up rarely, if at all, in ordinary conversation. In illustrating this point, Gregg (1989) discusses the following sentence:

(3) Mary ate an apple, and Sue a pear.

He says of this sentence (p. 22): 'I do not produce [sentences of this type] because they sound incredibly affected, and I have enough problems of that sort already'. Nonetheless, if you're a native speaker of English, you know that (3) forms a part of your language. We would want our theory to explain why English speakers accept (3). However, it's not clear that we'd be able to do that if we just go out into the world and record people, since it might not come up very often in our sample. In fact, from a purely statistical point of view, sentences like this probably occur even less frequently than 'errors'.

Related to some of the points we've been making above, you also know that certain strings of words do *not* form a part of your language. When we were discussing 'errors' just a moment ago, I mentioned that (1) was ungrammatical in my dialect of English. However, this aspect of your linguistic knowledge isn't something that you'd discover if you were to just record people. People make mistakes, but that's different from getting systematic evidence that certain strings of words are ungrammatical.

For all of these reasons, it seems as though simply recording speakers has problems as a method for discovering how the black box that constitutes our linguistic capacities works.

An alternative to studying behaviour: studying knowledge

Recall above the distinction between *competence* and *performance* which Chomsky draws our attention to. If studying performance (linguistic behaviour) has the problems that we've seen, perhaps there's a way that we can focus on competence (linguistic knowledge) directly; that is to say, focusing on the cognitive system in the brain that *underlies* our linguistic behaviour, rather than looking at the behaviour itself (see Figure 1.1).

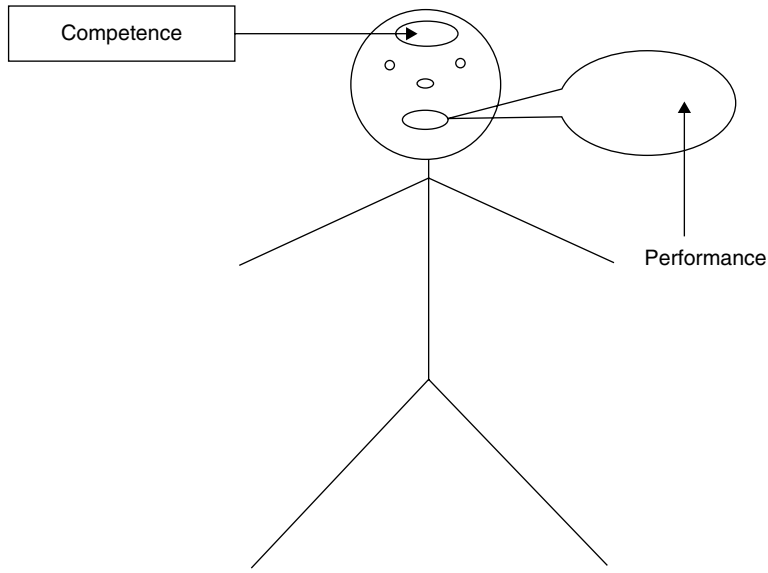

Figure 1.1 The object of study in linguistics

To do this, the standard experimental task is to solicit judgements from native speakers about the sentences of their language. Part of what you know about your language is that certain strings of words are part of that language and certain other strings are not. For example, the following string of words forms part of my language, a fairly standard dialect of American English:

(4) I haven't read the book yet, but I will tomorrow.

However, the words in a different order create a sentence which is not a part of my language:

(5) *I yet the read book haven't, but will tomorrow I.

For a native speaker of many dialects of British English, on the other hand, neither (4) nor (5) forms a part of their language. They are both ungrammatical. The grammatical version of (4) in their dialects is (6), which to me sounds very strange.

(6) I haven't read the book yet, but I will do tomorrow.

By feeding test sentences to native speakers of a language, and asking for their intuitive judgements/reactions, we can begin to understand what's inside the black box. We can begin to understand how their linguistic system is set up.

Consulting native-speaker intuitions in this way would seem to avoid, or at least minimize, the problems that we saw above when we considered simply recording speakers. There is still the problem of error due to interfering factors (which we'll discuss below). However, the burdens associated with what you're asking them to do are at least reduced. If you just have to provide your intuitive judgements, that's easier than having a 'real' conversation, with all of the distractions that that entails. You also are able to test both less frequently occurring structures and structures which you expect to be ungrammatical. Although people may not use sentences like (3) in ordinary conversation, you can still get judgements regarding them.

Grammaticality vs acceptability

As we mentioned just a moment ago, the grammaticality judgement task, which focuses on linguistic competence rather than linguistic performance, would seem to minimize errors due to interfering factors. However, it's important to see that, even when you're using grammaticality judgements, you must still consider whether the data is 'reliable'. There are things which can influence native-speaker intuitions which are not due to the linguistic system proper, but which are intimately related, including affectedness (which we mentioned above in discussing example (3)) and stylistics. For this reason, in addition to a competence/performance distinction, we also have a distinction between grammaticality and acceptability.

For example, sentences which themselves have sentences as their subjects can often be somewhat awkward stylistically.

(7) ?Once [that John was leaving] became clear, we abandoned the project.

In (7), the underlined adverbial clause contains a sentence *that John was leaving* which is itself in the subject position of the sentence *that John was leaving became clear*. This can result in a sentence which doesn't flow and which is

difficult to understand. To make the sentence easier to understand, the sentential subject is often put at the end and a dummy pronoun *it* placed in subject position:

(8) Once *it* became clear *that John was leaving*, we abandoned the project.

Unlike the somewhat awkward (7), (8) is perfect. However, despite the fact that (7) and (8) differ in *acceptability*, they are both *grammatical*, in the sense that our grammar (theory) should generate them. Our judgements in (7) are being influenced by extra-grammatical factors. In this case, stylistics and difficulty of processing.

This point is also illustrated by a story that Professor Jill Carrier, one of my old teachers, related to me about some fieldwork that she was doing with a native speaker of Tagalog, a language of the Philippines. She had a list of sentences that she was checking with the informant, and they were all confirmed as grammatical in Tagalog as expected. However, there was one particular sentence where the informant said 'Oh, no. Nobody would ever say that. That's impossible'. This was deeply puzzling, as from a structural point of view it seemed to be identical to various other sentences which were all reported to be perfectly fine. After much thought and further questioning, she discovered that the problem was that the example sentence happened to involve children swimming in the ocean. The reason that the informant had objected to the sentence was because the ocean was full of sharks and no children would ever be swimming in the ocean. It was too dangerous. Clearly, in that instance, it wasn't that the informant's *grammar* didn't generate the sentence in question. Instead, non-linguistic factors – knowledge of the real world – were influencing his judgement. What we are trying to do is isolate purely linguistic capacities.

Thus, whenever we're discussing grammaticality judgements or consulting with informants, we need to make sure that we keep this distinction between acceptability and grammaticality in mind. There may be plenty of sentences that our mental grammar generates. There are some, like (3), which are said quite infrequently, but there are also sentences which we would *never* say for one reason or another. If we want to correctly characterize how the black box of our linguistic capacities works, we want to be sure that our theory captures *all* of the sentences that it allows.

1.3 Some simple approaches to grammaticality judgements

In the previous sections, we've argued that we should focus in on a particular aspect of our linguistic knowledge: grammaticality judgements. You know instantly and without prior instruction that certain sentences of your language are grammatical and that certain other ones are ungrammatical. But how do you know what you know? Where do these grammaticality judgements come from?

To start with, let's take some first guesses at what might be going on. Just to be methodical, we'll start with some simple possibilities and work through to more complicated ones. Hopefully, you'll be convinced that these various simplified accounts aren't plausible.

The 'parrot' theory

The first, most naïve account of how you come to have grammaticality judgements would be that you judge a sentence grammatical if you've heard it before, and judge it ungrammatical if you haven't heard it before. We could call this the 'parrot theory'. This is pretty obviously wrong, I think. One of the most striking features of language is its unbounded capacity for innovation. It's quite likely that a large number of the sentences in this paragraph, not to mention the book so far, will be sentences that you've personally never come across before (unless you've heard me lecture). However, none of these sentences strikes you as strange or unfamiliar. I can even make up a sentence which I'm 100 per cent sure is new in the history of the English language (at least at the time of writing):

(9) Although Huxley's *The Doors of Perception* was never intended to be taken internally, there were many people in the late 1960s who failed to heed the warnings of prominent scholars and critics.

I'd bet a lot of money that (9) has never been produced ever in the entire history of the English language, yet you see it as a well-formed sentence of English without any effort whatsoever. I take that as fairly conclusive proof that your grammaticality judgements aren't based on simply whether you've heard the sentence before. Otherwise, you ought to judge (9) as ungrammatical, and you don't.

Other seemingly plausible suggestions

It also appears as though various other seemingly simple and plausible ideas don't play a role in our grammaticality judgements. For example, judgements about grammaticality are totally independent of what, if anything, the sentence means. Here is a famous example sentence from Chomsky (1957):

(10) Colorless green ideas sleep furiously.

Personally, I don't have the first clue what this sentence means. Ideas don't normally have colours, but even if they could, they certainly couldn't be colourless and green simultaneously. I'm also hard pressed to understand what sleeping furiously might entail. However, one thing I do know is that (10) is different from (11), where the words are in the reverse order:

(11) *Furiously sleep ideas green colorless.

The difference is that (10), despite its bizarre non-meaning, is syntactically well-formed. (11), on the other hand, is not syntactically well-formed.

The reverse situation can be found as well – that is, sentences which *do* make sense, but which we instantly judge to be ungrammatical. Consider (12):

(12) *The child seems sleeping.

Unlike (10) or (11), we have no problem assigning a meaning to the sentence in (12). If a foreigner had produced it, we would have no difficulty understanding what he or she intended to convey. However, we all know instantly that (12) is not a grammatical sentence of English, even though its meaning is clear.

In a similar way, ease of use doesn't seem to have any implications for grammaticality. There are plenty of sentences which are grammatically well-formed but which are not easy to use for one reason or another. Consider tongue-twisters, for example:

(13) She sells seashells by the seashore.

The fact that they're difficult to use (difficult to pronounce, in this case) is precisely why they're of interest. But crucially, they're not ungrammatical.

Sentences can also be difficult to use because they're difficult to parse. So-called 'garden path' sentences are ones in which your parser is 'tricked' into assuming that a sentence has a structure which needs to be revised when you get to the end of the sentence. A famous example is:

(14) The horse raced past the barn fell.

The effect when reading a garden-path sentence is usually lessened, but they can be quite startling when you hear one spoken to you. That's because when you're parsing the spoken sentence in real time, your parser is usually fooled into thinking that the past participle *raced* is in fact the main verb of the sentence. You don't realize you've made a mistake until you hear *fell*, and you're suddenly left with a word that you don't know what to do with. You've then frantically got to go back and figure out where you went wrong, which is harder to do when the sentence is spoken as opposed to read. You then realize that *raced* is supposed to be a past participle rather than the main verb. It's the horse *which was raced past the barn*. In spite of all this though, when the intended reading is pointed out to you, you realize that the sentence is in fact grammatical.

Analogy

A slightly more sophisticated version of the 'heard it before' theory is less obviously stupid. It could be that you judge a sentence as grammatical if it's 'similar to' or 'analogous to' a sentence that you've heard before or that you know is grammatical. Call this the 'analogy' theory. The essential problem is that, as

an explanation of your linguistic abilities, the whole concept just looks wrong. Consider the following pair of sentences:

(15) a. It is likely that John is here.
b. It is probable that John is here.

Likely and *probable* mean almost exactly the same thing, and so it's perhaps not surprising that they can both appear in the 'it is X that John is here' frame, as we see in (15). However, consider (16):

(16) a. John is likely to be here.
b. *John is probable to be here.

If analogy or similarity were the main factor in explaining why we have the grammaticality judgements that we have, it's not clear at all how or why (16b) is ungrammatical. *Likely* and *probable* are synonyms, they behave identically in (15), yet you know instantly and without prior instruction that (15b) is grammatical and (16b) is ungrammatical. Why does analogy fail here?

An even more telling example, I think, can be seen in the sentences that we were looking at earlier in the context of the competence/performance distinction above, having to do with the presence or absence of *that*. In virtually all situations, *that* can be either present or omitted with no change in grammaticality:

(17) a. I know *that* Mary is here.
b. I know Mary is here.

This optionality seems to remain when there is a question, rather than a statement:

(18) a. Who knows that Mary is here?
b. Who knows Mary is here?

(19) a. Who do you think that Mary saw?
b. Who do you think Mary saw?

However, there is one exception to this otherwise entirely regular pattern. When the *subject* of the embedded clause is questioned, the presence of the complementizer *that* is no longer optional. It is obligatorily absent. (Again, this is true for *almost* all dialects of British and American English.)

(20) a. *Who did you say that left the room?
b. Who did you say left the room?

Immediately, and without any prior instruction, you know that there's something wrong with (20a), even if you can't put your finger on what it is.

This is totally unexpected if something like analogy or similarity is relevant for understanding our intuitions about grammaticality. What we have is a paradigm which is absolutely and utterly regular, except for the one exception in (20a). If analogy were relevant, then (20a) should be grammatical, *on analogy* with (17a), (18a) and (19a). This really makes it look as though analogy or similarity isn't relevant for grammaticality.

Conclusion

What seems to be going on is that our intuitions about grammaticality stem from the fact that the brain contains a system for analysing sentences. When presented with a sentence of English, it's analysed by the cognitive system that you possess, providing you with a judgement about its acceptability. You have internalized a 'grammar' of your language.

What we're really doing in syntax, then, is something that you won't find in any other subject. In exploring this mental computational system and how it works, what we're doing is, in a sense, nothing more than telling you things that you already know. We are attempting to develop a theory about what's already in your head. However, that doesn't mean that you can put down the book and run away. Although you do, in an important sense, already know everything that we'll be talking about, the problem is that you don't *know* that you know it. It's knowledge, but unconscious knowledge. Syntactic theory is about making this knowledge explicit.

1.4 Language acquisition and universal grammar

We've come to a pretty radical conclusion. Part of your linguistic abilities consists of having judgements about whether sentences are grammatical in your language or not. We concluded that this was because your brain contains a cognitive system which analyses (or 'generates') sentences. In this section, we're going to ask a simple question, which turns out to have a complicated answer: how did this cognitive system get there? How did you acquire this ability to judge the grammaticality of sentences? The answer to this acquisition question will turn out to have a huge influence on the way we conduct our inquiry into syntax.

With respect to this issue, the linguist Noam Chomsky has observed what he calls 'the logical problem of language acquisition'. In a nutshell, what adults come to know about their language goes far beyond anything they were actually exposed to as children. Examples (15)–(20) are prime examples. *Likely* and *probable* are about as close synonyms as you can get. Yet you know instantly that *John is likely to be here* is grammatical but *John is probable to be here* isn't. However possible it might be that *one particular* speaker's experience enabled him or her to figure these facts out, it can't possibly be that *everyone's* experience was enough. Yet this is something that *every* speaker of English knows.

Here's another example, which we'll be looking at in the next chapter. (21) is a perfectly well-formed sentence of English:

(21) Jack met the student from England and I met the one from France.

The element *one* here seems to be filling in for *student*. (21) means *Jack met the student from England and I met the student from France*. However, if we change the prepositional phrase, things change radically:

(22) *Jack met the student of physics and I met the one of chemistry.

For some reason, although *one* can substitute for *student* in (21), it can't in (22).

So the question is: how did you come to know these facts? It certainly couldn't have come from any explicit instruction by teachers. It's unlikely that any of them would have noticed these facts. It's equally unlikely that the data you were exposed to when you were acquiring your native language was itself detailed enough for you to be able to figure out these facts. And, even more to the point, these facts are again something that *every* speaker of English knows. No matter how likely it is that you somehow managed to get exactly the right data, it's unlikely that *every* speaker of English got exactly the right data. Our experiences as children were all radically different. It looks like what we have is a case of knowledge being manufactured from nothing, not just in you, but in every speaker of English. How is this possible? This, as Chomsky sees it, is the logical problem of language acquisition.

Chomsky has a logical yet radical suggestion in order to solve this problem: if we have knowledge that couldn't have come from the data we were exposed to, it must have come from somewhere else. Chomsky proposes that this 'somewhere else' is an inborn language acquisition device specific to the human species. (This is what we mean by 'species-specific'.) Just as it is genetically predetermined that a human being will grow arms and not wings, it is genetically predetermined that a human being will acquire a human language. Children come pre-equipped to the task of language acquisition with a device that is designed to take certain kinds of acoustic disturbances ('sounds') as its input, and as its output creates a cognitive system for understanding and using language. If we assume that being exposed to a language acts as a trigger for the unfolding of a genetic program, on a par with the development of limbs and organs, this would seem to solve the questions about language that we've just raised. You end up with a linguistic system that goes well beyond the particular input data you were exposed to, in precisely the same way that organ development 'goes beyond' the particular 'nutritional inputs' (i.e. food) that you were exposed to as a child.

Chomsky's hypothesis is that this 'language acquisition device', more commonly known as Universal Grammar (UG), has two central parts to it which aid in the task of language acquisition: *principles* and *parameters*. In fact, the general approach to syntax that we'll be arguing for in this book gets

its name from these parts of UG. It's called the 'Principles and Parameters' approach.

An introduction to principles and parameters

The idea is that the *principles* of Universal Grammar account for the respects in which all human languages are the same. Most of the topics that we'll be discussing in this book concern principles of UG, and we'll be drawing data from English and many other languages in order to argue that principles of UG are at work.

Parameters, on the other hand, are intended to capture certain ways in which languages differ. Now, obviously, one way that languages differ is that they have different words for different concepts. So dog is *dog* in English, *hund* in German, *perro* in Spanish, etc. These are differences that simply have to be memorized by the child, and it's only when the child is exposed to the language that he or she will know what the words are going to be.

However, there are other ways in which languages differ. For example, the languages of the world seem to divide up into *null subject languages* and *non-null subject languages*. English is a non-null subject language. That is to say, English sentences must *always* have a subject, even when that subject appears to contribute nothing to the meaning of the sentence. Consider (23):

(23) a. It is raining.
b. It seems that John is here.

The subject of the sentences in (23) is *it*, but this pronoun is totally meaningless. You're not pointing to something when you say (23a) or (23b). The pronoun doesn't refer to anything. However, because English requires a pronounced subject at all times, these 'dummy' pronouns must be inserted. By contrast, in a null subject language like Spanish, these dummy pronouns are omitted. The subject is simply not pronounced.

(24) a. Llueve.
rains
'It's raining.'
b. Parece que Juan está aquí.
seems that Juan is here
'It seems that Juan is here.'

The hypothesis is that UG contains a parameter called the Null Subject Parameter. Therefore, in some sense, the child already knows from birth that the language that he or she will be exposed to will either be a null subject language or a non-null subject language. Once the child begins to hear data, it is a relatively simple matter to set the parameter one way or the other. A parameter is like a question. Setting the parameter is a matter of answering that question in light of the evidence from the particular language.

1.5 Description and explanation: the case of structure-dependent rules

One important benefit that our postulation of an inborn Universal Grammar also gives us is the ability to go beyond merely describing individual languages. It puts us in a position to *explain* why particular languages are the way they are. For example, one of the things which seems to characterize all languages that have been studied so far is that they form yes/no questions (and, in fact, all questions) in a *structure-dependent* fashion.

By way of illustrating the structure-dependence of yes/no questions in English, consider the following examples:

(25) a. The agent has left the building.
b. Has the agent left the building?

In (25a) we have a declarative statement. In (25b), we have the corresponding yes/no question; that is to say, a question which asks for a yes or no answer. Imagine you're a Martian scientist trying to discover what the rule is for English yes/no question formation. (Since you're a Martian, you're not familiar with English or any other human language.) On the basis of the pair in (25), a simple rule that looks like it accounts for the data is (26):

(26) Move the third word to the front of the sentence.

This rule not only accounts for the pair in (25), but also an infinite number of other sentences, including pairs like those in (27) and (28):

(27) a. The man is here.
b. Is the man here?

(28) a. My friend will go to the store.
b. Will my friend go to the store?

The rule in (26) does not require analysing the sentence in terms of its structure. You just go along the words one by one until you hit the third word, then you move that word to the front of the sentence. It is therefore a *structure-independent* rule.

Interestingly, although the structure-independent rule in (26) is very simple to apply, you quickly discover that it isn't right. Consider the (declarative) sentence in (29):

(29) The man from the bank will remember nothing.

If we were to apply the structure-independent rule in (26), we would scan along from left to right until we encountered the third word, which in this case is *from*. We move *from* to the front of the sentence, and we get (30):

(30) *From the man the bank will remember nothing?

This is clearly not a grammatical sentence of English. If we want to make the declarative sentence in (29) into a question which asks for a yes or no answer, we need to move *will* to the front of the sentence:

(31) Will the man from the bank remember nothing?

So clearly our rule for yes/no question formation in (26) isn't right. But what is?

One thing to notice about (30) is that our rule in (26), which just counted words from left to right, ended up moving a preposition (*from*) to the front of the sentence. However, we know that in order to form a yes/no question, what we need to do is move an *auxiliary verb* to the front of the sentence. (Let's put aside the cases where we have no auxiliary, in which case you'd have to insert the 'dummy' auxiliary *do*.)

(32) a. *Should* John have gone to the store?
b. *Will* the book arrive on time?
c. *Has* Mary defeated her opponents?

It's possible that we can save a structure-independent approach by modifying our rule in (26) to something like (33):

(33) Move the first auxiliary verb to the front of the sentence.

That rule *does* correctly get us (31), and all of the other cases that we've seen so far, but it's still structure-independent. It doesn't require that we know anything about the *structure* of the sentence. We just scan from left to right until we hit the first auxiliary verb, and then move it to the front of the sentence.

However, (33) isn't the English rule either. Consider (34):

(34) The man who has written the book will be followed.

Using that rule, we'd predict that the yes/no question corresponding to (34) should be (35):

(35) *Has the man who written the book will be followed?

If this isn't the right rule either, then what is it?

Since you're a speaker of English, a human language, and not really a Martian scientist, it'll probably seem obvious what the 'right' answer is. But it's important to see how unexpected that answer is. The 'right' answer is something like (36):

(36) Move the first auxiliary verb *in the predicate* to the front of the sentence.

You don't just move along and pick the third word or the first auxiliary in a *structure-independent* way. The rule that you've internalized requires treating the sentence not just as a string of words coming one after the other. Instead, you treat the sentence as if it has structure. The English rule for yes/no question formation is structure-dependent. In this case, at the very least you have to divide the sentence up into a subject and a predicate:

(37) a. Subject: The man who has seen the book
b. Predicate: will be followed

Having identified the predicate, you can then find the first auxiliary in the predicate and move it to the front of the sentence. This rule will also account for the other sentences that we've seen:

(38) a. Subject: The agent
b. Predicate: has left the building

(39) a. Subject: The man
b. Predicate: is here

(40) a. Subject: My friend
b. Predicate: will go to the store

(41) a. Subject: The man from the bank
b. Predicate: will remember nothing

So it seems that even a simple thing like turning a declarative sentence into a yes/no question requires an analysis that is, from a computational point of view, somewhat complicated. You don't just scan across and pick the first auxiliary or the tenth word or something. You need to analyse the sentence into a subject and a predicate in order to find the main verb phrase.

Now the reason for this long discussion about structure-dependent vs structure-independent rules is to make a point about UG. If we look at all the languages of the world that have so far been studied, the number of languages which use structure-independent rules is exactly none. There seems to be no language in the world which uses structure-independent rules for anything. No language, for example, forms questions by moving the third word to the front of the sentence or by taking the words in the declarative and reading them in reverse order.

As we mentioned above, if we were just looking at language after language in a descriptive way, we'd have no explanation for this fact. It would just be a coincidence. We'd look at English and say, 'Wow, no structure-independent rules'. We'd then move on to Swahili and say, 'Wow, no structure-independent rules'. And so on and so forth. On the other hand, if we as humans possess a genetically determined UG, then we can explain why we find what we find. There's something about UG which doesn't allow for structure-independent rules.

We can even go further than this, however. If we're right to say that UG doesn't allow for structure-independent rules, and that's what accounts for the fact that no language has them, then we make a further prediction that we can check. Assuming that the language acquisition process consists of language data interacting with UG, then we should predict that no child ever hypothesizes a structure-independent rule during the course of language acquisition. If UG is always there in the background, helping the child process the input data, then the child should never use structure-independent rules at any stage. UG doesn't allow them.

Interestingly, this prediction also seems to be verified. No child is ever exposed to sentence pairs like (27) and (28) above and hypothesizes a computationally simple structure-independent rule like (26) or (33). They always and automatically make a structure-dependent choice. It might be the wrong structure-dependent choice, but it's still a structure-dependent choice. Thus, it would seem that they never hypothesize a structure-independent rule for anything because it's never an option. Children come pre-equipped to the task of language learning knowing that all grammatical rules have to be structure-dependent.

Other sources of evidence

There's some interesting corroborating evidence as well from studies of an autistic patient, who's called Christopher in the literature. Christopher has some fairly severe cognitive deficits and has to be looked after in an institution, but, like some autistic patients, he also has a talent. That is, there is one area where he seems to perform much better than ordinary people. Christopher's particular talent is languages. He enjoys learning new ones and does so effortlessly and at a tremendously rapid pace. He can translate back and forth between them perfectly, getting all the grammatical endings right, etc. If UG is implicated in Christopher's talent, then we would expect to see that reflected in the way he acquires his languages.

Structure-dependency is actually one of the things that Christopher has been tested on, and the results seem to support the existence of UG. Christopher and a control group of normal people were presented with two languages to learn. One was Berber, a North African language, which has incredibly complex rules and constructions. The other was a made-up language which crucially contained *structure-independent* rules. Although Christopher mastered Berber without difficulty just as he had other languages, he struggled with the made-up language, and was never able to master it. The control group struggled with Berber, and also struggled initially with the made-up language, although they eventually did better than Christopher on the made-up language.

This would seem to provide strong corroborating evidence for the claim that UG exists and that it prohibits structure-independent rules. By hypothesis, Christopher's considerable linguistic resources come from UG. His general problem-solving abilities are not very good. By contrast, the control group may not have the access to UG that Christopher has (being adults), but they do have good general problem-solving skills. This would explain why the control group had difficulties with Berber which Christopher did not, and why they had more

success with the made-up language. Since the made-up language contained rules which were prohibited by UG, the control group could switch gears and treat it as though it were a puzzle, calling on their general problem-solving abilities. Since Christopher has very little resources in this area, he was unable to succeed with it.

Description and explanation

So, ultimately, while we're interested in individual languages, we're also interested in discovering what we can about UG, the initial state of the language faculty in the mind/brain. As a result, we have a constant flow of ideas back and forth between individual languages and UG. In order to account for some particular fact about English, we might hypothesize that UG is implicated. But a hypothesis about the make-up of UG makes a prediction about other languages. We therefore then need to check those other languages to see if the prediction is verified. This process is illustrated in Figure 1.2.

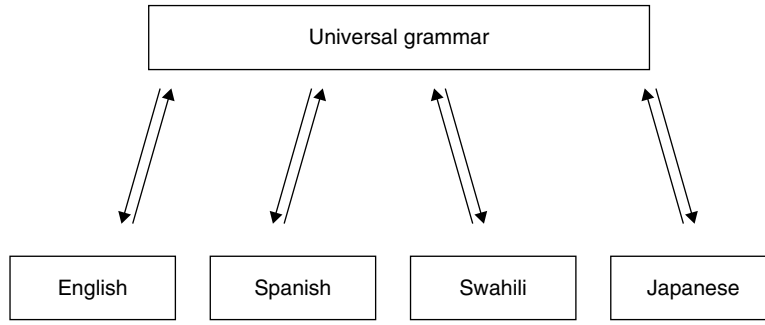

Figure 1.2 The relationship between languages and Universal Grammar

In this way, we can make a distinction between two kinds of arguments: *conceptual* and *empirical*. Roughly speaking, when we're making an argument that goes *from* UG *to* a particular language, we're dealing with a conceptual argument. Our theory about how UG works says X, and so therefore we expect that we'll find X when we look at individual languages. When we're going *from* individual languages *to* UG, then we're making an *empirical* argument. We've got some fact X about a language, and so we need to work that fact X into our theory about UG.

We'll try as much as possible to make arguments of both types. Rather like a pair of scissors, we'll work best when we come from both directions at once. In that way, we can be as secure about our conclusions as possible.

1.6 Introduction to the book

Now that we've discussed some general things about the study of language, I want to turn to a specific introduction to the subsequent chapters. Each chapter covers a specific topic, from sentence structure to transformations and beyond.

Several of the later chapters, however, return at a more advanced level to issues raised in earlier chapters. This way, you can get an introduction to a topic in an earlier chapter, and then come back in a later chapter to more recent proposals. If you work all the way through to the end of the book, you should be well prepared, with a little bit of coaching, to take the next step – reading actual articles in the linguistics literature.

In the individual chapters, one of the things you'll find is that there are in-text exercises scattered throughout. There are two kinds of these. The first kind is labelled 'Exercise'. These are just like the exercises at the end of the chapter, but placed right after the relevant point of the discussion in the chapter. These are there to give you a chance to practise using the theory while the relevant concepts are fresh in your mind; answers to these can be found at the end of the chapter. The second kind of exercise is labelled 'Before Continuing'. These exercises will usually come at a point where we've encountered a problem with the theory. The reader is encouraged to stop for a minute before reading on and give some thought to how to solve the problem or address the question. The immediately following discussion will then provide an answer to the question. Just like anything else, the easiest way to get comfortable with the theory is to practise it. Pretty soon the whole thing will be second nature. As I mentioned, the end of each chapter also contains 'Further Exercises' for you to practise the things that you've learned.

Each chapter also has a 'Further Reading' section. In many cases, the references will be to articles that we've actually discussed in the chapter. Sometimes in the text I will have omitted complexities that aren't relevant to the argument at hand, or I have translated an argument from a slightly different framework into the one we're using. Where this occurs I will mention it in the 'Further Reading'. One of the things that's slightly different about this book is that each chapter will conclude with a section entitled 'Open Issue'. These sections are designed to take something from the chapter a step further or in a controversial direction. I thought it was important to include things like this to show you that we're not dealing with a theory that's fixed in stone, with all problems solved and nothing new to think about. They're meant to be starting points for your own thinking and to provide suggestions for research projects. You might try to find evidence from a different language for an Open Issue proposal or perhaps look for counter-examples and discuss whether or not they represent real problems.

Summary

In this chapter, I've tried to sketch out some of the basic assumptions which underlie the more technical parts of the rest of the book. As a native speaker of a language, one aspect of your linguistic abilities is the ability to judge whether a given string of words forms a part of the language you speak or not. In other words, you have judgements about grammaticality and ungrammaticality. This ability stems from the grammar of your language which you unconsciously 'constructed', based on the language acquisition data that you were exposed to

in childhood. However, what your grammar contains seems to go well beyond the data you were exposed to. Additionally, speakers of a language all converge more or less on the same grammar despite the wide variation in particular data they happened to have been exposed to. For these reasons, it's hypothesized that there is a language acquisition device, or Universal Grammar, which is part of our common genetic endowment as humans. In trying to develop a theory about 'English', by which we mean the internalized grammar of a speaker of English, we also need to construct a theory of Universal Grammar.

FURTHER READING

Probably the most accessible work by Chomsky himself on these general topics is Chomsky (1988). It's a transcription of a series of lectures given at the University of Managua for a general audience; that is, for people who didn't know anything particularly about linguistics. One caveat, though: since the audience were native Spanish speakers, the examples he picks, though not complex, are taken from Spanish rather than English. There are translations of everything, though, and it's actually good practice for getting used to non-English examples. Most standard Chomskyian linguistics textbooks don't devote much discussion to general, 'nature of linguistics' questions, but a notable exception is Haegeman (1994). Chapter 1 of that book has some very good discussions of general issues and the setting out of the Chomskyian linguistic programme. For people interested in reading more about the autistic patient, Christopher, and what he can tell us about the Chomskyian approach to language, you can check out Smith and Tsimpli (1995).

Index