

Core Syntax: A Minimalist Approach

David Adger

July 16, 2002

Contents

Preface	i
1 Core Concepts	1
1.1 What is a sentence?	1
1.1.1 Utterances, propositions and sentences	1
1.1.2 Acceptability, grammaticality and stars	2
1.1.3 Form and order	5
1.2 Tacit Knowledge	6
1.3 Syntactic theories	10
1.4 Back to Sentences	13
1.5 Summary	15
2 Morphosyntactic features	17
2.1 Introduction	17
2.2 Introducing Features	17
2.3 What are features?	20
2.3.1 Feature Systems	21
2.3.2 Interface Rules	24
2.4 Motivating Features	25
2.4.1 Major Category Features	26
2.4.2 More on the content of words	28
2.4.3 Some more minor features	33
2.5 Summary	41
3 Constituency and Theta Roles	48
3.1 Introduction	48
3.2 Constituents	48
3.3 Fundamental notions	53
3.4 Determining the head	57
3.5 Predicting the head - θ -roles and selectional features	60
3.5.1 θ Roles and the θ -Criterion	61
3.5.2 Unassigned θ -roles	64
3.5.3 Selectional Features and Lexical Representation	65
3.5.4 S-selectional Features	68
3.6 Triggering Merge by feature checking	71

3.7	Summary	76
4	Representing Phrase Structure	82
4.1	Introduction	82
4.2	The Structure of Phrases	82
4.2.1	First Merge — Complements	82
4.2.2	Second Merge — Specifiers	86
4.2.3	Adjunction	87
4.2.4	The structure of the Verb Phrase	90
4.2.5	\bar{X} -Theory	91
4.3	C-command	92
4.3.1	Defining the Relation	93
4.3.2	Reflexive Binding	94
4.3.3	Negative Polarity Items	96
4.3.4	Summary	97
4.4	Ditransitives	97
4.4.1	Three argument predicates	97
4.4.2	Ternary or Binary Branching?	97
4.4.3	A Binary Alternative	100
4.4.4	Causatives and little <i>v</i>	104
4.5	Linking	109
4.6	The Architecture of the System	113
4.7	Summary	117
5	Functional Categories I — TP	124
5.1	Introduction	124
5.2	Sentences are headed	125
5.2.1	Theoretical Considerations	125
5.2.2	Empirical Evidence	125
5.2.3	Cross Linguistic Evidence	132
5.3	Tense Marking	133
5.3.1	Tense on Main Verbs	133
5.3.2	Tense on Auxiliaries	138
5.4	Head Movement	141
5.4.1	The position of Negation	141
5.4.2	Aux to T movement	142
5.4.3	V to T movement	146
5.4.4	Aux <i>in situ</i>	148
5.5	<i>Do</i> -support	149
5.5.1	Basic Properties of <i>do</i> -support	149
5.5.2	Affixal Negation	151
5.5.3	Affixal Tense	152
5.5.4	Tense Chains	155
5.6	Summary	157

6	Subjects and Objects	166
6.1	Introduction	166
6.2	Subject Movement	167
6.2.1	Quantifier Stranding	167
6.2.2	Expletive Constructions	170
6.3	Case and EPP	172
6.3.1	Subjects and EPP-features	172
6.3.2	Objects and Locality of Feature Matching	177
6.3.3	Subject Verb Agreement	180
6.3.4	Interim Summary	182
6.4	Unaccusative Subjects	182
6.4.1	<i>Ne</i> -Cliticization	184
6.4.2	Possessor Datives in Hebrew	186
6.5	Passives	188
6.6	Adverbs	190
6.7	Word order and movement	193
6.8	Summary	195
7	Functional categories II — the DP	200
7.1	The category D	200
7.1.1	Types of D	201
7.1.2	The semantics of Ds	203
7.2	Ds as Heads	205
7.2.1	Complementary Distribution of Ds	205
7.2.2	Capturing the Complementary Distribution	207
7.2.3	Implications of this Analysis	208
7.2.4	Further Questions	209
7.2.5	Possessors — Specifiers of DP	211
7.2.6	More Null Ds	212
7.3	Θ -role Assignment in DP	216
7.4	Little <i>n</i> and the Low Structure of DP	219
7.4.1	Motivation for Little <i>n</i>	220
7.4.2	Case Checking in DP	222
7.4.3	Possessors	225
7.5	AP modification in DP	226
7.6	Order and Movement in DP	230
7.7	Summary	232
8	Functional Categories III — CP	238
8.1	Complementizers	238
8.1.1	<i>that</i>	238
8.1.2	Declaratives and Interrogatives	240
8.1.3	Subject Clauses	245
8.2	Non-finite CP Complementation	249
8.2.1	PRO and Control Clauses	251

8.2.2	ECM clauses	252
8.2.3	C in Non-finite Clauses	253
8.2.4	Null case and PRO	255
8.2.5	Cross-clausal case checking and ECM	256
8.2.6	Raising	260
8.2.7	Object Control	266
8.2.8	Summary and further implications	269
8.3	Adjunct Clauses	270
8.4	Cross-Linguistic Implications	271
8.5	Summary	275
9	Wh Movement	282
9.1	Introduction	282
9.2	Elements that undergo wh-movement	285
9.2.1	Wh-expressions	285
9.2.2	Interpretation of wh-questions	288
9.3	Mechanics of wh-movement	289
9.3.1	Checking wh-features	289
9.3.2	Refining the system	293
9.3.3	Other wh-movements	295
9.3.4	Embedded wh-questions	297
9.4	Subject wh-questions	297
9.5	Long distance wh-movement	300
9.6	Wh-in-situ	302
9.7	Superiority	304
9.8	Cross-linguistic variation in wh-constructions	305
9.9	Summary	307
10	Locality	313
10.1	Introduction	313
10.2	Empirical evidence for local movement	314
10.2.1	Floating quantifiers	314
10.2.2	Complementizer agreement again	316
10.2.3	Successive inversion phenomena	318
10.3	Ensuring successive wh-movement	320
10.3.1	Locality in feature matching	320
10.3.2	Solving an old Puzzle	321
10.4	Island Phenomena	323
10.4.1	Wh-Islands	323
10.4.2	DP islands	325
10.4.3	Subject Islands	327
10.4.4	Adjunct Islands	332
10.5	Summary	332

Preface

Core syntax is a somewhat cheeky name for a textbook. There are many ‘core’ aspects of syntax, depending on the viewpoint you take. In this book, I have tried to highlight three aspects. Many working syntacticians will disagree that these are *the* core areas, but I hope that everyone will agree that they at least provide a consistent perspective.

Firstly, I have tried to emphasize that an important part of doing syntax is constructing a consistent theoretical system within a broad research agenda. The research framework that I have adopted is that of the Minimalist Program, a perspective that has emerged from the work of Noam Chomsky. In the book, I develop certain theoretical ideas which stem from the concerns articulated by Chomsky, and I use these ideas to build a system. I have tried to make the system as consistent (both conceptually and notationally) as I can, and I have tried to emphasize the interactions between theory, analysis and data, and how developments in any one of these areas affects the others.

The second way in which the material of the book could be considered ‘core’ is that I have tried to cover a broad range of phenomena which form a (somewhat nebulous) area of empirical importance as far as the recent history of the subject is concerned. These phenomena have also been chosen because they can be used to motivate or to challenge the theoretical ideas that are being developed. I hope that by doing this, students will both get an overview of some important areas of empirical concern, and an appreciation of how syntactic arguments are developed.

The third aspect which I have concentrated on is methodological, and it has to do with the importance of rigour in helping to clarify ideas, and the concomitant importance of formalization as a skill which will help students to think through syntactic questions. The strictly derivational approach I have adopted, and the explicit presentation of how derivations are constructed, should help students to learn to be comfortable with some of the formal aspects of syntactic theory.

I hope that the book’s emphasis on these three aims will mean that it will outlive the details of current theoretical fashion. The particulars of the theory will certainly change, but the kinds of syntactic argumentation deployed in the field, and the importance of rigour and consistency will surely remain, as will, I think, the kinds of questions raised by the research agenda adopted here.

Aside from these three aims, I have also tried to tease out how a theory based on invariant syntactic principles interacting with the parameterization of featural properties of functional heads, really does provide at least one answer to the poverty of stimulus questions raised by Chomsky. The sections in the chapters which concentrate on this aspect of syntax also try to raise the further question about how parameterization itself can be restricted. This

highlights a number of open questions, as does the approach to island phenomena sketched in the final chapters. I hope that at least some students will be encouraged to try to find out answers to these and related questions themselves.

Because the Minimalist Programme is a research agenda rather than a particular theory, and because one of my main aims in this book is to introduce students to how internally consistent syntactic systems are constructed, I have melded argumentation and analyses from the literature with ideas that have developed from the particular theoretical choices I have made. For example, the checking theory I have adopted allows variation in the strength of the values of features, rather than in the features themselves, with effects on numerous analyses and on the encoding of parametric variation; the approach to verbal inflection I propose is different from others on the market (although I do discuss the alternatives); the analyses I sketch for passivization, for subject *wh*-questions and for verb-second, for example, are not derived directly from the primary literature, but rather arise from the particular version of Minimalism developed here. This was the best way I could see of achieving my three aims of theoretical consistency, a reasonable coverage of core phenomena, and a fair amount of analytical rigour. I hope that readers knowledgeable about the current Minimalist literature will forgive any liberties I have taken.

Finally, some thank-yous: firstly to three anonymous OUP reviewers, two of whom later became onymous as Bob Borsley and Liliane Haegeman. All three reviewers went well beyond the call of duty, and provided the book with much needed direction, and I'm hugely grateful to them, as I am to my editor at OUP, John Davey, for advice and encouragement. I'd also like to thank Bernadette Plunkett and Gillian Ramchand, both of whom used earlier versions of the book in their teaching, and who provided me with detailed comments about how it was structured. These two join a band of trusty linguistic old faithfuls who have been on call for me whenever I needed to check out that an idea wasn't completely insane. Others in this band are Daniel Harbour, Ian Roberts, George Tsoulas, and the Tequila Cabal. Final thanks are to Anson, *sine quo nihil*.

Chapter 1

Core Concepts

1.1 What is a sentence?

1.1.1 Utterances, propositions and sentences

Certain utterances of human languages seem to have a special status, in that they express what you might call “complete thoughts”. So if I were to say “It’s rather hot, today”, any native English speaker will interpret me as conveying a message which is somehow complete. Note that *you* could say exactly the same thing: “It’s rather hot, today”, which will also convey a complete message. However, if I do this on Sunday, and you do it on Monday, the message communicated is different. Linguists say that the **proposition** expressed is different.

A proposition is that aspect of the meaning of a sentence which allows us to say “Yes, that’s true” or “No, that’s false”. It describes a state of affairs that holds in the world, and its correspondance with that state of affairs allows us to attribute truth or falsity to the proposition. There are other aspects of sentence meaning which we will address later in the book, but propositional meaning will be the most relevant for us here.

Note that, even though we have expressed different propositions, we have both used exactly the same linguistic form. We have both said the same **sentence**. This little scenario gives us some grasp on a core idea in syntax, the idea of a sentence as an abstraction over **utterances** which have the same form. Linguistic form is not important to a proposition. The same proposition is conveyed by the English sentence (1), the French sentence (2), and the Scottish Gaelic sentence (3), even though these sentences are in different languages:

- (1) John saw Stephan
- (2) Jean a vu Stephan
- (3) Chunnaic Iain Stephan

So a proposition is that aspect of the meaning of a sentence which says something about a state of affairs, and an utterance is an actual use of a sentence. How do we define sentence itself then?

Take any act of linguistic communication, an utterance of (1) by me to you, for example. Somehow you glean a proposition from my utterance of (1). How do you do this? The

common-sense answer is that it's because we both know the same language. Focussing in more precisely on the question of how we define sentences, it appears that there is something about my knowledge of English which is shared with your knowledge of English, and that this includes how to form sentences of English, and how the proposition expressed by a sentence depends on its form. Clearly the form is important, since if you were to utter some sequence of sounds that did not form a sentence of English, then I'd have a much more difficult task in understanding what proposition you were trying to convey.

Part of our shared knowledge of language, then, allows us to construct sentences, which we can then utter. Again, the idea of a sentence is more abstract than the idea of an utterance (which is something that you can hear, record, feed into a computer as sound waves etc). A sentence itself is something which can't be recorded, heard or electronically manipulated, only **uses** of sentences can. These stretches of words that you are now reading, delimited by capital letters and full stops and interspersed with other markings, are uses of sentences. The sentences themselves are defined by the knowledge of English that I put to use in writing them, and that you use in comprehending them. Although it sounds counter-intuitive, what you see on this page are technically utterances, in that they have an external, physical manifestation. Sentences, on the other hand, are internal, mental entities, which have an abstract form.

1.1.2 Acceptability, grammaticality and stars

The form of sentences is what will mainly concern us in this book. The example I went through above showed that the same sentence form can be used to express different propositions on different occasions of utterance. We can also express the same proposition using different sentence forms. So most people would agree that essentially the same message is conveyed by both the sentences below:

- (4) That monkey is eating the banana.
- (5) The banana is being eaten by that monkey.

And most people would find it difficult to say what message is conveyed by (6) and would reject it as a sentence of English:

- (6) By is eaten monkey banana that the being.

Why is (6) not an **acceptable** sentence of English? Well obviously the **order** of the words matters. (6) uses exactly the same words as (5) but the order in which the words come in (6) is somehow not an order which English allows. Acceptability, however, doesn't just depend upon the order of words. (7) is just as unacceptable as (6), but the order of the words seems to be fine (it's just like the order in (4)), it's the *form* of the words that makes the sentence unacceptable, the word after *is* should have the form *eating*, not *ate*:

- (7) That monkey is ate the banana

As well as clearly unacceptable sentences like (7) and (6), we also find sentences which seem odd in some way.

(8) The amoeba coughed and then it fainted.

The acceptability of this sentence depends upon the context that it is uttered in. (8) is unacceptable to most speakers if someone just says it out of the blue. However, one can easily imagine contexts where (8) is quite fine: in a fantasy novel for example, or as a description about what's going on in a children's cartoon. The form of (8) does not appear to be a determining factor in our judgements of its acceptability, it is rather the proposition that is expressed which we baulk at. The acceptability of (8), then, appears to be dependent on our view of the world. In most people's view of the world, amoebas don't cough or faint, and so we judge (8) as unacceptable because it conflicts with our expectations about what words can mean, and about how the world works.

There are also cases of unacceptable sentences which appear to be semantically plausible, but which seem to be awkward. An example of such a sentence is (9):

(9) I looked the number which you picked out at random by using a needle and a phone-book up.

This sentence becomes more acceptable if the word *up* is placed immediately after the word *look*:

(10) I looked up the number which you picked out at random by using a needle and a phonebook.

Another way of making this kind of sentence more acceptable is by making the distance between *look* and *up* shorter:

(11) I looked the number up.

The unacceptability of (9) might have an explanation in terms of how we process the sentence. In an intuitive sense, the words *look* and *up* are closely associated with each other. In (10), they are pronounced together, while in (11) they are separated by a small number of words. In (9), however, there are fifteen words between *look* and *up*. One hypothesis to pursue would be that, on processing the sentence, the hearer accesses the word *look* and expects the word *up* to appear next. However, the hearer has to wait quite some time for the expected word to appear, and presumably process a fairly complex structure in the meantime. We might, therefore, put the unacceptability of (9) down to the fact that the connection between *look* and *up* is difficult to make because there is too much other processing going on.

This kind of explanation for unacceptability is known as a **parsing** explanation. It assumes that the problem with the sentence is not that it does not conform to the rules of the language, but rather that human beings processing (or parsing) the sentence have a hard time assigning the right structure to it. In the most extreme cases we might conclude that the sentence can't be parsed at all—it is **unparsable**.

Notice that this kind of explanation does not appear to be available for examples like (7). We cannot make the sentence any better by reducing the distance between the relevant words. The problem appears to be that the language simply requires a particular relationship to hold between these words, and in (7) it simply doesn't hold. This kind of explanation for the unacceptability of (7) is known as a **grammaticality** explanation, and sentences like

(7) are termed **ungrammatical**.

Syntacticians express the claim that a particular sequence of words is not a grammatical sentence of the language under discussion, by marking that sequence with a star, like this:

(12) *By is eaten monkey banana that the being

(13) *The monkey is ate the banana

Sometimes a sequence of words is called a **string**. Putting a star at the start of a string is a claim that it isn't a grammatical sentence of the language in question.

Acceptability is a judgement about the status of a string as part of a language, and as such it's an intuition that speakers of the language have. The claim that a sentence is difficult to parse, or that it is ungrammatical, is an attempt to explain its (un)acceptability. As linguists, we cannot know in advance whether a string of words is unacceptable because it is difficult to parse, because it is ungrammatical, or because of some other factor.

The acceptability of a sentence will also often depend upon whether the sentence expresses the intended meaning. Here the judgement doesn't state whether the string is acceptable or not, but rather whether the string is assigned the meaning that is specified. So sometimes you might see the claim that a sequence of words is "star (*) under the intended interpretation". What this means is that the sentence is unacceptable as an expression of a particular proposition. This use of the star is most common when contrasting the meanings of sentences. For example, (14) has the meaning that there are some people in the garden, and not the meaning that it's generally a property of people that they're in the garden, while (15) works in the opposite way, it means that, as a rule, people are stupid. It cannot mean that there exist some stupid people. It is therefore said to be starred under this interpretation.

(14) People are in the garden

(15) People are stupid

What we have here is a case where the form of the sentence is fine, and the sentence is perfectly grammatical, but the interpretation of the sentence is not the one that one might expect, given the interpretation of other similar sentences.

Here is another example of the same kind of thing

(16) How did Julie ask if Jenny left?

This sentence can be a question about the way that Julie asked something (loudly, rudely, etc.), and it is perfectly fine under this interpretation. However, the discussion of this example might make it clear that the intended interpretation is a question about the way that Jenny left (angrily, jauntily etc.). This kind of interpretation is perfectly possible for a sentence like (17):

(17) How did Julie say that Jenny left?

However, this is *not* a possible meaning of (16). In that case, we would say that (16) is starred under this interpretation.

Why do linguists use the star to express the idea that a sentence doesn't have the meaning that might be expected, as well as the idea that a string isn't actually a sentence at all? The

reason is that the propositions expressed by sentences are not just dependent on the words in a sentence, but also on the ways that those words are put together. It's our knowledge of language that allows us to put together words into sentences, and something about that knowledge simply stops us from putting the words in (17) together in the same way as those in (16). So speakers of English can't assign a structure to (16) that will give us the same interpretation as the structure for (17).

Remember that we assign a star to a sentence if we think that the explanation for its unacceptability is that it does not conform to the requirements of the grammar of the language under discussion. Put more directly, we assume that speakers can't assign a structure to the particular string in question at all. In the example we have just discussed, exactly the same situation obtains: speakers can assign two structures to (17) but only one to (16). The absence of one of the structures can be marked with a star, as long as we refer in the surrounding text to the intended interpretation, so that we know which structure is relevant.

1.1.3 Form and order

Let's go back to more obvious facts about the form of a sentence. We can see that the order of words is also relevant to the message (proposition) conveyed by the sentence, so (18) does not mean the same as (4), even though it consists of exactly the same words:

(18) That banana is eating the monkey.

As was mentioned above, there also appear to be certain relationships between the forms of the words in sentences. So in (18), the first and third words have the form they do because of the form of the second word. If we change the form of the second word, then the forms of the first and third words have to change too:

(19) Those monkeys are eating the banana.

The same relationship holds in (19), and we can't, for example, swap *are* for *is*, or vice versa:

(20) *Those monkey are eating the banana.

(21) *That monkeys is eating the banana.

Native speakers of English generally judge such sentences not to be part of their language, and therefore as ungrammatical. The relationship between the second word in these sentences, and the other words which seem to take their lead from it, is known as **agreement**, and we will come back to it later on in the book.

Word order and agreement are fairly obvious properties of sentences. When the right word order, or the right agreement relationships are not expressed in a sentence, native speakers know this, and judge the sentence as unacceptable. These phenomena are not obviously susceptible to an explanation in terms of parsing, and we hypothesize that the problem is a purely structural one.

The meanings of sentences are also things which native speakers have strong intuitions about, as we saw above. Most obviously, native speaker of English intuitively know that (4) and (18) differ in meaning, and that (22) has two meanings (i.e., it is **ambiguous**):

(22) The monkey wrote a poem on its favourite banana

One meaning of (22) involves the banana being written about, while the other meaning has the surface of the banana actually being written on.

All of these phenomena, and this is only a tiny taster, arise because sentences of English, and all other languages are not just built out of words strung one after the other in an utterance or on a page, but rather, as mentioned already, because they have a **structure** which specifies certain relationships as holding between the fragments of language that the sentence is made out of. One of the major tasks of linguists who are interested in syntax is to discover what the basic building blocks of language are and what relationships hold between them. The general term that is used to refer to these relationships is the **syntactic structure** of the sentence.

1.2 Tacit Knowledge

When I used the word **knowledge** in the last section, I meant knowledge in the particular technical sense that's used in linguistics: knowledge of language isn't like knowledge of history. My knowledge of history is extremely **partial, forgettable, explicit** and **learned**. I know that William Wallace was a Scottish Highlander that fought against the English some centuries ago. I'm not sure when he was born, or when he met his (rather grizzly) death - so my knowledge about him is partial. My knowledge of English structures is at the other end of the spectrum. Although you might introduce a new word to me, it's rare that you'll be able to introduce to me a new way of building up a sentence of English (unless that structure is special because of the dialect you speak, or because it used to be a structure of English but is no longer). Similarly, I am sure that, at some time, I knew the date of William Wallace's death, but I've forgotten it. There's no sense in which I can be said to forget how to construct sentences of English, unless I suffer a serious brain trauma — something which certainly isn't necessary to make me forget historical dates.

Perhaps the most important difference between knowledge of language and other types of knowledge is that the latter are easily made explicit: I can tell you that Scotland merged its Parliament with England in 1707, and then opened a new devolved parliament in 1999. However, trying to explain how sentences are constructed is something which isn't accessible to my conscious mind: all I can do is use various types of scientific method to try to uncover how it works. Even simple statements like “a sentence is constructed from an initial subject followed by a predicate” immediately throws us into the realm of technical notions. What's a subject? What's a predicate? Is this the only way to build up a sentence, and if not, then what are the other things involved?

Yet we all possess a highly intricate system that allows us to determine whether certain utterances correspond to sentences of our native language. You know, for example, if you are a native English speaker, that I can utter (23), and that this corresponds to a sentence of English, whereas if I utter (24), although you might be able to glean some message from it, it simply doesn't correspond to any sentence of English:

(23) Anson is incredibly difficult to please.

(24) *Anson is incredibly difficult to be pleased.

How do you know this? Are you able to articulate why (24) is unacceptable? In fact, it's just not obvious why (24) should be so bad. The kind of knowledge of language which allows you to make this judgement is not accessible to you — the technical term for this is that it is **tacit knowledge**.

Finally, knowledge of language is acquired in a different way to other types of knowledge. Most children have a command over the core structures of their language by the time they are three. The following sentences were produced by a child at the age of three years and two months (see Pinker 1994, Chapter 7):

(25) Can I put my head in the mailbox so the mailman can know where I are and put me in the mailbox?

(26) Can I keep the screwdriver just like a carpenter keep the screwdriver?

These sentences contain errors in agreement ((25) contains *I are* rather than *I am* and (26) contains *a carpenter keep* rather than *a carpenter keeps*), but display a complex structure involving combining sentences and operating on their parts (you will meet many of the syntactic operations that are involved in constructing these sentences later on in the book). Even more remarkable are experimental studies which show that children at this young age seem to have flawless access to aspects of the structure of sentences which do not seem to come from the data that they are exposed to.

One classic example comes from the syntactic rule which forms yes/no questions in English. Yes/no questions look as follows:

(27) Has Jenny eaten a cake?

(28) Will Anson come to the party?

Certainly, such questions appear in the data that young children are exposed to. It might seem a simple task to work out the rules for forming such a question. The examples above are clearly related to the following sentences:

(29) Jenny has eaten a cake.

(30) Anson will come to the party.

But notice that when a child hears a yes/no question like those above, that child might be able to dream up any number of rules to construct them from the related sentences:

- (31)
1. swap the first two words around.
 2. swap the first verbal element with the first noun phrase.
 3. put the verbal element that follows the subject in front of it.

It turns out that the right rule is the third one. If the child were to adopt the first rule, then he/she should produce sentences like the following, where the (a) examples give the plain sentence types and the (b) examples give the result of applying the rules:

(32) a. The man has eaten the cake.

- b. *Man the has eaten a cake
- (33) a. The woman who is singing is happy.
b. *Woman the who is singing is happy

If the child were to adopt the second rule, then we would expect the following results. Don't worry about the notion of **noun phrase** for the moment. It simply means that the first two words are grouped together:

- (34) a. The man has eaten the cake.
b. Has the man eaten the cake?
- (35) a. The woman who is singing is happy
b. *Is the woman who singing is happy

However, Crain and Nakayama (1987) conducted a series of experiments and showed that the errors expected under these hypotheses never occurred. This was true for children as young as three years and two months. All the children appear to use the third, and apparently most complicated rule, and produce the correct results:

- (36) a. The man has eaten the cake.
b. Has the man eaten the cake?
- (37) a. The woman who is singing is happy
b. Is the woman who is singing happy?

This kind of argument is known as a **Poverty of the Stimulus** argument. It attempts to show that children are in possession of kinds of knowledge about their language which simply isn't plausibly in the input data that they hear.

Even more thought provoking is that fact that, at this point in their development, children are immensely ignorant of the kinds of knowledge that become so important in later life for passing exams, earning money and preparing food. Yet they can put together highly complicated sentences, with only small deviations from what adults do. The system of knowledge they have is developed to such an extent that it far outstrips any other aspect of their cognitive development which is not obviously instinctual (vision, locomotion, eye-contact, using emotional responses to control their environment). But this knowledge hasn't been taught to them by anyone. In fact, the idea of teaching one year old children that sentences are constructed out of subjects and predicates, or that they involve putting together nouns, verbs and other things in particular ways, is absurd. In many senses, this tacit knowledge is far more like instinct than it is like other kinds of knowledge.

So why, you might ask, do we call it knowledge at all? The answer really has to do with a philosophical discussion that goes back to Plato in the fourth century BCE, who asked the question of how we can know so much with so little input (Plato's main area of discussion was geometry rather than language), a problem which Chomsky has christened **Plato's Problem**. Plato proposed, as a solution, that there are certain abstract forms which we grasp because we are human beings, and therefore our souls exist in the realm of these abstract forms. Knowing something corresponds to our souls grasping these abstract forms. Because the problem is so similar to the problem of how human beings grasp the abstract

forms associated with language, we still talk about this as a **problem of knowledge**.

Philosophers often find linguists' use of the term knowledge for the relationship we bear to our language problematic because of the common sense use of the word knowledge as 'being in possession of true facts', or 'true justified belief'. The linguists' response to this is that knowledge is as reasonable word as any, but if philosophers want to maintain some special status for the word knowledge, then we'll use something else (for example, **cognize** has been suggested). The linguists' use of the term knowledge (and more especially tacit knowledge) should be taken as a technical term in linguistics.

Another common term for tacit knowledge of language is **grammar**. This is a word to be careful of because it is used in two technical ways: (i) a grammar is a speaker's tacit knowledge of their language; (ii) a grammar is a linguist's explicit theory of a speaker's tacit knowledge of their language. I will be careful in this book to distinguish the two, but it is as well to be sensitive to this distinction, since the two uses of 'grammar' are widespread. Chomsky tends to refer to grammar in the first sense as **I-language**, where I stands for Internal/Individual. So a speaker's I-language is his or her tacit knowledge of their native language and a linguist has a theory of this I-language.

We don't know a great deal about how I-language itself is manifested physically or biologically, although it's clear that at least some aspects of it are localised in particular positions in the brain. There are two areas of the brain's left hemisphere which appear to be particularly implicated in our linguistic abilities: **Broca's area** and **Wernicke's area**. We know that these areas are important because sometimes they are affected by some pathology or trauma, and particular linguistic deficits arise. For example, people with a trauma to Broca's area have selective problems in dealing with complex grammatical constructions, although their pronunciation may be perfectly normal.

Some aspects of I-language are apparently determined by genetic factors. Researchers have studied families which appear to have a predisposition to **Specific Language Impairment** (SLI). SLI is a term for linguistic impairments which cannot be easily tied down to more general cognitive difficulties. Myrna Gopnik and her colleagues have argued over a number of years that the SLIs of one particular family can only be properly explained by genetic factors, and moreover that these genetic factors affect language specifically. Studies of identical and non-identical twins appear to back up this idea, and recently, researchers in genetics have identified the mutation of a particular gene in affected individuals (see the reference section for details).

Other aspects of I-language are clearly determined by environmental factors. A child brought up in a linguistic community will acquire the language or languages of that community, no matter what that child's genetic inheritance. The environment provided by the community provides the linguistic data which the child needs and which she/he uses, in conjunction with what is provided innately, to acquire her/his I-language.

I-language, then, is the component of the mind of members of the human species which allows us to link together meanings (which are in part propositional in nature) with forms (sequences of sounds, strings of letters, collocations of body signs in sign languages). The nature of I-language seems to be that it is highly creative, in the sense that it can link an infinite array of meanings to an infinite number of sequences of signs. Because of this, human beings can communicate with each other in ways that no other species seems able to do, and our ability to communicate is at least partially responsible for the astonishing evolutionary

success of our species. Without I-language, that is, without an internalised syntax, we would be unable to communicate fluently, because we would be unable to externalize our messages except in the crudest ways. Syntax is, then, key to human achievement at a fundamental level. The project of modern linguistics is to investigate and try to understand this crucial phenomenon within the human mind, both as a goal in itself, and as part of a broader attempt to map our cognitive abilities.

1.3 Syntactic theories

We have so far come to two conclusions: (i) that a sentence is an abstract mental object, and (ii) that whether a sentence is part of a particular speaker's native language or not is something which that speaker tacitly knows. We are now in a position to make a first stab at saying what *syntactic* knowledge is: syntactic knowledge (or just syntax) is that part of an I-language that is concerned with the ways that sentences are put together out of basic parts. A particular syntactic theory is a linguist's theory of this, while syntactic theory in general is what linguists interested in this question are engaged in.

What is a theory? In the broadest terms, a theory is a statement of **general principles** or laws which can be used to explain some phenomena. We can have theories which apply in many different domains: theories of poetics, of social interactions and structures, of biology, of the physical world. There is no *a priori* way of deciding what constitutes an appropriate domain for a theory; that is, we do not know in advance of investigation whether a theory developed to explain big things (like the motion of planets) will also explain very small things (like the motion of sub-atomic particles). We do not even really know, in advance of study, whether there is a difference between, say, energy and matter—these concepts, which seem so different from a common sense perspective, may turn out, after theoretical investigation, to be better understood as the same.

Apart from the domain of a theory (that is, the phenomena which it purports to explain), we have to consider what kinds of statements we can make as part of the theory. What are the statements about, and what possible relations do we allow to hold in these statements? In economic theory, the statements might be about goods and services and the values attached to them; in social theory, the statements might be about the categories that group individuals together or force them apart; in Newtonian physical theory, the statements might be about mass, force, acceleration and other properties of physical bodies. Notice, once again, that we don't know what the statements are about in advance of investigation. What we do is develop a body of theoretical knowledge, and then explore this theory to determine what its consequences are. A successful theory engenders new and interesting questions, and provides accounts of phenomena in general terms. Part of providing an account of a theoretical domain is giving explanations of phenomena which have already been observed. If the theory is to be really successful, though, it should extend beyond the observed phenomena and make predictions about novel situations. These predictions should be able to be tested, and it should be possible to challenge the theory as it stands if the prediction is **falsified**.

Some of the most successful theories we have are theories of the natural world: of physics and chemistry and biology. As human beings we have what you might call a **folk understanding** of the natural world, given to us by our biology and by our experience. We know

that when it gets dark, it tends to get colder, that someone pushed off a cliff will generally fall, and that if you step onto some water you will generally not float on it. We have also developed scientific understanding of these phenomena involving concepts such as energy, gravitation and surface tension. Over the years, the theories of these concepts have become more and more remote from our folk understanding, requiring the development of highly abstract mathematical systems. Yet at the same time, these theories have become more successful: they make accurate predictions about what will happen in certain situations; they allow us to develop technologies to interact with the natural world in ways that our ancestors would have marvelled at; and they bring to light phenomena which were previously unobserved and help us to ask new and interesting questions about these.

One of the characteristics of theories of physics, genetics, biochemistry etc is that they have a **formal** aspect to them. By formal, we mean that they are developed as systems which have (reasonably) well-defined mathematical properties, so that the statements of the theory are explicit and the concepts well-defined. The importance of defining concepts explicitly is that it is then harder to misuse them. For example, if we have a definition of, say, a particular gene type which we use to explain some phenomena, and then we try to use the same gene type to explain another phenomena, we can check whether we really mean to use the same thing by investigating what the consequences of applying the same definition is in both cases. A formal approach is extremely useful, then, as a means of ensuring that our theory is really doing what we want it to do.

The approach to syntax that we will take in this book, and which is taken more generally in generative grammar, assumes that certain aspects of human psychology are similar to phenomena of the natural world, and that linguistic structure is one of those aspects. This approach is motivated by some of the phenomena we discussed in the last section: the physical structure of the brain and the genetic structure of human beings both appear to be implicated in linguistic structure. Given this, it might be possible to isolate those aspects of linguistic structure which depend on genetic and neural bases. Since those aspects of structure will have a physical base, it would be unsurprising if the kinds of theoretical approach which have been so successful in dealing with the natural world were helpful in understanding language.

Related to this perspective is a broader issue: one of the most interesting and successful endeavours in the field of human knowledge in general has been the general search for “mathematical patterns in nature” (Freidin and Vergneaud 2001, page 647 refer to this research agenda as the **Galilean approach** after the famous Italian scientist (1564-1642)). On the assumption, motivated above, that at least some structural aspects of human language are part of nature, we can ask the question: what are the mathematical patterns in language?

Theoretical linguistics, as we will tackle it here, attempts to build theories of linguistic structure which posit various abstract concepts and determine the relations that hold between them. The consequences of these theories are then explored, and the concepts and relations in the theory thereby undergo development. Part of the general idea of developing theories is to keep them as simple as possible. We try to posit as few concepts and relations as we can get away with (an idea that goes back to the mediaeval English philosopher William of Ockham (1280-1347) and which is sometimes dubbed **Ockham’s Razor**). Of course the effects of the theory might be extremely complex, but the base of the theory is kept as simple as possible. This injunction, is, in part, a methodological one: it is easier to work with fewer

concepts than with more.

Syntacticians, then, are interested in positing general, explicit, falsifiable and simple theories about syntactic structure. Part of these theories is a statement of what the basic units of syntax are (technically these are known as syntactic **primes**, or syntactic **formatives**), and in what ways these units can be related to each other (that is, the **syntactic relations**). With a theory in place, a syntactician will develop an **analysis** of some set of data. The analysis relates various aspects of the data to concepts in the theory: it states what constitute the syntactic primes within the body of data, and investigates how the syntactic relations between those primes are established. Many different analyses may be possible of a particular set of data within a theory, although the theory will constrain the space of possible analyses. A good analysis may have consequences for the theory: it may show that this theory successfully extends to a new domain of data or it may show that the theory has to be modified in some way.

Syntacticians usually want their theories and the related analyses to be as **descriptively adequate** as possible: this means that the structures that the analysis assigns to sentences have to be the kinds of structures that will interface naturally with other aspects of the theory of language: they should allow semantic theories to derive the right meaning for the sentences; psycholinguistic theories to predict how the sentences are processed; theories of historical change to predict how languages change on the basis of these structures, and so on. Descriptive adequacy is a much more stringent condition on the theory than mere **observational adequacy**, which simply says “get the right words in the right order”. Descriptive adequacy requires the theory to assign the right structures to the sentences in question.

There is another level of adequacy, **explanatory adequacy**, which is concerned with capturing the commonalties that all languages share but, at the same time, allowing only those languages which are actually possible human languages. An explanatorily adequate theory will be able to explain how, given a particular set of input data, a child fixes on the I-language that they do, rather than on some other I-language which is compatible with the data but makes the wrong predictions.

The aspects of I-language which are common to all of us are known as **Universal Grammar (UG)**, and a theory of Universal Grammar will state the commonalties that hold across all possible languages, (often called **Principles** of UG) and in what ways individual I-languages may diverge from these commonalties (known as the Parameters of Variation of UG - or just **Parameters**). For example, UG may require that all languages have agreement relations between a verb and its subject, but some languages may not pronounce this relationship, or it may obtain also between a verb and its object.

Now, although there are only a few parameters, the interaction of these with each other may lead to fairly divergent structures in the sentences which particular I-languages allow. However, if you think of a child who is born with UG, she or he does not need much exposure to particular sentence types in order to be able to infer just how the parameters for the language they are acquiring need to be specified, because there are only a few parameters, so there are only a few possible analyses. This means that a theory structured in this way will stand a chance of explaining exactly why children acquire language so quickly, and on the basis of data which does not seem to contain enough information to develop the grammar *sui generis* (see references to section 1.2).

1.4 Back to Sentences

One of the core properties of I-language is that it characterises a potentially infinite number of sentences, and it is this that leads to language having such an amazing capacity to express abstract thoughts and concepts. Many of the sentences we use in our everyday life, we've never spoken before, and we seem to have no trouble creating new sentences to describe new situations. We appear to be able to use our I language to relate an infinite variety of thoughts to an infinite number of sentences. How can this be?

Perhaps human beings simply have this huge number of sentences stored in their brains, and I-languages are just lists of these. This is impossible. Imagine you were to start with a sentence like:

(38) Paul had an affair

and then add to it

(39) Paul had two affairs

(40) Paul had three affairs . . .

(41) Paul had eighty eight billion sixty three million forty-four thousand nine hundred at twenty three affairs

(42) Paul had eighty eight billion sixty three million forty-four thousand nine hundred at twenty four affairs . . .

as you can readily see, Paul's promiscuity is no match for the number of sentences we can have to describe it!

There are many ways to show that the human brain just isn't big enough to store all of the potential sentences. Consider (43):

(43) Anson thought Julie had fainted

Notice that we can take this sentence, and **embed** it inside another sentence, to give:

(44) You said that Anson thought that Julie had fainted

This process does not end here. We can now embed (44) inside another sentence:

(45) David wrote that you said that Anson thought that Julie had fainted

Of course we can keep this going, in principle, forever. In practice, one or the other of us will keel over from exhaustion, boredom or death. But the point remains that there isn't an upper limit to the number of sentences that we can in principle produce or understand. Each time we perform the embedding operation, we have a new sentence. Since we can continue the embedding operation indefinitely (at least in principle), there are an infinite number of such sentences.

Our ability to do these things *in principle* comes from our I-language. This is distinguished from our actual performance, which is what happens every time we use our knowledge of language on a particular occasion to produce an utterance, a written form, or a series of