Syntactic Theory: A Formal Introduction[1]

Ivan A. Sag   Thomas Wasow

Partial Draft of September, 1997
Comments Welcome

Contents

Chapter 1

Introduction

1.1 Two Conceptions of Grammar

This is a textbook about grammar. You may well wonder why a college would offer courses on grammar – a topic that is usually thought of as part of the junior high school curriculum (or even GRAMMAR school curriculum...). Well, the topic of this book is not the same thing that you probably think of as grammar.

What is taught as grammar in primary and secondary school is what linguists call 'prescriptive grammar'. It consists of admonitions not to use certain forms or constructions that are common in everyday speech. A prescriptive grammar might contain rules like:

- Be sure to never split an infinitive.
- Prepositions are bad to end sentences with.

Modern linguists think prescriptive grammar is for the most part a pointless activity. We view human language as a phenomenon amenable to scientific investigation, rather than something to be regulated by the decrees of authorities. Your seventh grade math teacher probably told you about the time the Nebraska legislature passed a bill establishing the value of π as 3, and everybody in the class no doubt laughed at their foolishness. Linguists regard prescriptive grammar as silly in just about that same way.[1]

So, if modern grammarians don't worry about split infinitives and the like, then what do they study? It turns out that human languages are amazingly complex systems, whose inner workings can be investigated in large part simply by consulting the intuitions of native speakers. Let us give you some examples from English.

Example 1: The adjectives *unlikely* and *improbable* are virtually synonymous: we talk about unlikely or improbable events or heroes, and we can paraphrase *It is improbable that Lee will be elected* by saying *It is unlikely that Lee will be elected.* This last sentence is synonymous with *Lee is unlikely to be elected.* So why does it sound so strange

[1]There are of course powerful social and economic reasons for learning the grammatical norms of educated people and the question of how these norms get established is a fascinating one. But these are primarily sociological matters, and are not directly relevant to the subject matter of this book.

to say, **Lee is improbable to be elected*? (In keeping with standard linguistic practice, we will use an asterisk to mark an expression that is not well-formed – that is, that doesn't 'sound good' to our ears).

Example 2: The sentences *They saw Pat with Chris* and *They saw Pat and Chris* are near paraphrases. But if you didn't catch the second name, it would be far more natural to ask, *Who did they see Pat with?* than it would be to ask **Who did they see Pat and?* Notice, by the way, that the only one of these two examples that sounds well-formed (or 'grammatical' in the linguist's sense) is the one that violates a standard prescriptive rule. The other sentence is so blatantly deviant that prescriptivists never thought to comment on the impossibility of such sentences. Prescriptive rules typically arise because people's language use is innovative, leading languages to change. If people never say something – like the bad example above – there's simply no point in bothering to make up a prescriptive rule to tell people not to say it.

Example 3: The two sentences *Something disgusting has slept in this bed* and *Something disgusting has happened in this bed* appear on the surface to be grammatically completely parallel. So why is it that the first has a passive counterpart, *This bed has been slept in by something disgusting*, whereas the second doesn't: **This bed has been happened in by something disgusting*?

These are the sorts of questions contemporary grammarians worry about. The first two will eventually be addressed in this text, but the third will not[2]. The point of introducing them here is to illustrate a fundamental fact that underlies all modern work in theoretical syntax[3]:

- Every normal speaker of any natural language has acquired an immensely rich and systematic body of unconscious knowledge, which can be investigated by consulting speakers' intuitive judgements.

In other words, knowing a language involves mastering an intricate system full of surprising regularities and idiosyncracies. Languages are phenomena of considerable complexity, which can be studied scientifically. That is, we can formulate general hypotheses about linguistic structure and test them against the facts of languages.

The study of grammar, on this conception, is a field in which hypothesis-testing is particularly easy: the linguist can simply ask native speakers whether the predictions regarding well-formedness of crucial sentences are correct.[4]

[2]For extensive discussion of the third question, see Postal (1986).

[3]The term 'syntax' is often used instead of 'grammar' in technical work in linguistics. While the two terms are sometimes interchangeable, 'grammar' may also be used more broadly to cover all aspects of language structure; 'syntax', in contrast, refers only to the ways in which words combine into phrases, and phrases into sentences. While the boundaries are not sharply delineated, 'syntax' contrasts with 'semantics' (the study of linguistic meaning), 'morphology' (the study of word structure), and 'phonology' (the study of the sound patterns of language) in ways that 'grammar' does not.

[4]This methodology is not without its pitfalls. Judgements of acceptability show considerable variation across speakers. Morevoer, they can be heavily influenced by context, both linguistic and non-linguistic.

NB: A '◇' before a problem indicates that it should not be skipped. The problem involves something that will either be elaborated upon or else simply incorporated into subsequent chapters.

◇**Problem 1: Judging Examples** Indicate whether each of the following examples is grammatical or ungrammatical. If it is ungrammatical, give an intuitive explanation of what is wrong with it. If it is grammatical, determine whether it:

a. is bizarre in meaning; if so, explain why,

b. contains a feature of grammar that occurs only in a particular variety of English; if so, identify the feature. Is it stigmatized in comparison to 'standard' English?

   If you are uncertain about any judgements, feel free to consult other people. Non-native speakers of English, in particular, are encouraged to check their judgements with others.

1. Kim and Sandy is looking for a new bicycle.
2. Have you the time?
3. I've never put the book.
4. The boat floated down the river sank.
5. It ain't nobody goin' to fight nobody.
6. Terry really likes they.
7. Nixon disagreed Sandy into the ocean.
8. Aren't I invited to the party?
9. They wondered what each other would do.
10. There is eager to be fifty students in this class.
11. Dana loves sushi and so don't I.
12. Strings have been pulled many times to get people into Harvard.
13. This is the kind of problem that my doctor is easy to talk to about.
14. A long list of everyone's indiscretions were published in the newspaper.
15. Which chemical did you mix the hydrogen peroxide and?
16. There seems to be many reasons for questioning the government's account.

Since linguists rarely make any effort to control for such effects, not all of the data employed in the syntax literature should be accepted without question. On the other hand, many judgements are so unequivocal that they can clearly be relied on. In more delicate cases, many linguists have begun to supplement judgements with data from actual usage, by examining grammatical patterns found in written and spoken corpora. The use of multiple sources and types of evidence is always a good idea in empirical investigations. See Schütze (1996) for a detailed discussion of methodological issues surrounding the use of judgement data in syntactic research.

1.2 An Extended Example

To get a feel for the research in this paradigm of grammatical scholarship, consider the following question:

> In what linguistic environments do English speakers use reflexive pronouns (that is, forms like *herself* or *ourselves*), and where does it sound better to use a nonreflexive pronoun (like *her*, *she*, *us*, or *we*)?

Consider, first, some very basic examples:

(1) a. *We like us.

b. We like ourselves.

c. She likes her. [she ≠ her]

d. She likes herself.

e. Nobody likes us.

f. *Nobody likes ourselves.

g. *Ourselves like us.

h. *Ourselves like ourselves.

These examples suggest a generalization along the following lines:

Hypothesis I: A reflexive pronoun can appear in a sentence only if that sentence also contains a preceding expression that has the same reference (i.e. a preceding COREFERENTIAL **expression); a nonreflexive pronoun cannot appear in a sentence that contains such an expression.**

The following examples are different from the previous ones in various ways, so they provide a first test of our initial hypothesis:

(2) a. She voted for her. [she ≠ her]

b. She voted for herself.

c. We voted for her.

d. *We voted for herself.

e. *We gave us presents.

f. We gave ourselves presents.

g. *We gave presents to us.

h. We gave presents to ourselves.

i. *We gave us to the cause.

j. We gave ourselves to the cause.

k. *Nobody told us about us.

l. Nobody told us about ourselves.

m. *Nobody told ourselves about us.

n. *Nobody told ourselves about ourselves.

These examples are all covered by Hypothesis I, lending it some intial plausibility. But here are some counterexamples:

(3) a. We think that nobody likes us.

b. *We think that nobody likes ourselves.

Intuitively, the difference between these examples and the earlier ones is that the sentences in (3) contain subordinate clauses, whereas (2) and (1) contain only simple sentences.

Problem 2: Applying Hypothesis I It isn't actually the mere presence of the subordinate clauses in (3) that makes the difference. The following examples contain subordinate clauses, but are covered by Hypothesis I.

(i) We think that she voted for her. [she $\neq$ her]
(ii) We think that she voted for herself.
(iii) *We think that herself voted for her.
(iv) *We think that herself voted for herself.

a. Explain how Hypothesis I accounts for the data in (i)-(iv).
b. What is it about the subordinate clauses in (3) that makes them different from those in (i)-(iv) with respect to Hypothesis I?

So we can consider the following hypothesis:

Hypothesis II: A reflexive pronoun can appear in a clause only if that clause also contains a preceding, coreferential expression; a nonreflexive pronoun cannot appear in any clause that contains such an expression.

For sentences with only one clause (such as (1)-(2), Hypothesis II makes the same claims as Hypothesis I. But it correctly permits (3a) because *we* and *us* are in different clauses, and it rules out (3b) because *we* and *ourselves* are in different clauses.

However, Hypothesis II won't work, as stated:

(4) a. Our friends like us.

b. *Our friends like ourselves.

c. Those pictures of us offended us.

d. *Those pictures of us offended ourselves.

e. We found a letter to us in the trash.

f. *We found a letter to ourselves in the trash.

What's going on here? The acceptable examples of reflexive pronouns have typically been cases where the reflexive is functioning as an object of a verb (or the object of a preposition that goes with the verb) and the antecedent – that is, the expression it is coreferential with – is the subject or a preceding object of the same verb. If we think of a verb as denoting some sort of action or state, then the subject and objects (or prepositional objects) normally denote the participants in that action or state. These are often referred to as the ARGUMENTS of the verb. In the examples in (4), unlike many of the earlier examples, the reflexive pronouns and their antecedents are not arguments of the same verb (or, in other words, they are not COARGUMENTS). For example in (4b), *our* is just PART of the subject of the verb *like*, and hence not an argument of the verb; rather, it is *our friends* that denotes participants in the liking relation. Similarly, in (4e) the arguments of *found* are *we* and *a letter to us*; *us* only PART of an argument of *found*.

Hypothesis III: A reflexive pronoun must be an argument of a verb that has another, preceding argument with the same reference. A nonreflexive pronoun cannot appear as an argument of a verb that has a preceding coreferential argument.

Each of the examples in (4) contains two coreferential expressions (*we, us, our*, or *ourselves*), but none of them contains two coreferential expressions that are arguments of the same verb. Hypothesis III correctly rules out just those sentences in (4) in which the second of the two coreferential expressions is the reflexive pronoun *ourselves*.

Now consider the following cases:

(5) a. Vote for us!

b. *Vote for ourselves!

c. *Vote for you!

d. Vote for yourself!

In (5d), for the first time, we find a well-formed reflexives with no antecedent. If we don't want to append an ad hoc codicil to Hypothesis III[5], we will need to posit a hidden subject (namely, *you*) in imperatives sentences.

Similar arguments can be made with respect to the following sentences.

[5]For example, an extra clause that says: 'unless the sentence is imperative, in which case a second person reflexive is well-formed and a second person nonreflexive pronoun is not.'

(6) a. We appealed to him_1 to vote for him_2. [$\text{him}_1 \neq \text{him}_2$]

b. We appealed to him to vote for himself.

c. We appealed to him to vote for us.

(7) a. We appeared to him to vote for him.

b. *We appeared to him to vote for himself.

c. We appeared to him to vote for ourselves.

In (6), the pronouns indicate that *him* is functioning as the subject of *vote*, but it looks like it is the object of the preposition *to*, not an argument of *vote*. Likewise, in (7), the pronouns suggest that *we* should be analyzed as an argument of *vote*, but its position suggests that it is an argument of *appeared*. So, on the face of it, such examples are problematical for Hypothesis III, unless we posit arguments that are in some sense missing. We will return to the analysis of such cases in later chapters.

♢**Problem 3: Reciprocals** English has a 'reciprocal' expression *each other* (think of it as a single word for present purposes) which behaves in some ways like reflexive pronouns. For example, a direct object *each other* must refer to the subject, and a subject *each other* cannot refer to the direct object:

1. They like each other.
2. *Each other like(s) them.

a. Construct examples parallel to those in (1)–(3) to test whether the basic behavior of reciprocals is similar to that of reflexives.

b. Construct examples parallel to those in (5)–(7) to test whether the behavior of reciprocals is similar to that of reflexives in imperative sentences and in sentences containing *appeal* and *appear*.

c. Are there any constraints that the reciprocal imposes on its antecedent that reflexives don't impose? [Hint: what change to (1d) and (6b) did you have to make in order to construct the corresponding well-formed reciprocal sentence?]

d. Consider the following contrast:

- They lost each other's books.
- *They lost themselves' books.

Discuss how such examples bear on the applicability of Hypothesis III to reciprocals.

You can see that things get quite complex quite fast, requiring abstract notions like 'coreference', being 'arguments of the same verb', and allowing arguments to be missing from the sentence but 'understood', for purposes of the rules for pronoun type. And we've only scratched the surface of this problem. For example, all the versions of the rules we have come up with so far predict that nonreflexive forms of a pronoun should appear only in positions where their reflexive counterparts are impossible. But this is not quite true, as the following examples illustrate:

(8) a. We wrapped the blankets around us.

b. We wrapped the blankets around ourselves.

c. We admired the pictures of us in the album.

d. We admired the pictures of ourselves in the album.

It should be evident by now that formulating precise rules characterizing where English speakers use reflexive pronouns and where they use nonreflexive pronouns will be a difficult task. It is one we will return to in Chapter 7. The point of discussing it here was to make the following points:

1. Normal language use involves mastery of a complex system, which is not directly accessible to consciousness.
2. Speakers' tacit knowledge of language can be studied by formulating hypotheses and testing their predictions against intuitive judgements of well-formedness.
3. The theoretical machinery required for a viable grammatical analysis may be quite abstract.

1.3 Remarks on the History of the Study of Grammar

This conception of grammar is quite a recent development. Until about 1800, almost all linguistics was primarily prescriptive. Traditional grammar (going back hundreds, even thousands of years – to ancient India and ancient Greece) was developed largely in response to the inevitable changing of language, which is always (even today) seen by most people as deterioration. Prescriptive grammars have always been attempts to codify the 'correct' way of talking. Hence, they have concentrated on relatively peripheral aspects of language structure. On the other hand, they have also provided many useful concepts for the sort of grammar we'll be doing. For example, our notion of parts of speech, as well as the most familiar examples (such as noun and verb) came from the Greeks.

A critical turning point in the history of linguistics took place at the end of the eighteenth century. It was discovered at that time that there was an historical connection among most of the languages of Europe, as well as Sanskrit and other languages of India (plus some languages in between).[6] This led to a tremendous flowering of historical linguistics, reconstructing the family tree of the Indo-European languages by comparing the modern

[6] The discovery is often attributed to Sir William Jones who announced such a relationship in a 1786 address, but others had noted affinities among these languages before him.

languages with each other and with older texts. Most of this effort concerned itself with the correspondences between individual words and the sounds within those words. But syntactic comparison and reconstruction are also possible, and some was initiated during this period.

In the early twentieth century, many linguists, following the lead of the Swiss scholar, Ferdinand de Saussure, turned their attention from historical (or 'diachronic'[7]) studies to 'synchronic'[8] analysis of languages – that is, to the characterization of languages at a given point in time. The attention to synchronic studies encouraged the investigation of languages without writing systems, which are much harder to study diachronically, since there is no record of their earlier forms.

In the United States, these developments led linguists to pay far more attention to the indigenous languages of the Americas. Beginning with the work of the anthropological linguist Franz Boas, American linguistics for the first half of the twentieth century was very much concerned with the immense diversity of languages. The Indo-European languages, which were the focus of most nineteenth century linguistic research, constitute only a tiny fraction of the approximately five thousand known languages. In broadening this perspective, American linguists put great stress on developing ways of describing languages that would not forcibly impose the structure of a familiar language (such as Latin or English) on something very different. Most, though by no means all, of this work emphasized the differences among languages. Some linguists, notably Edward Sapir and Benjamin Lee Whorf, talked about how looking at language could provide insights into how people think. They tended to emphasize alleged differences among the thought-patterns of speakers of different languages. For our purposes, their most important claim is that studying the structure of language can give insight into human cognitive processes. This is an idea that has wide currency today, and, as we shall see below, it constitutes one of the most interesting motivations for studying syntax.

In the period around World War II, a number of things happened to set the stage for a revolutionary change in the study of syntax. One was that great advances in mathematical logic had provided formal tools that seemed well suited to be applied to natural languages. A related development was the invention of the computer. Though early computers were unbelievably slow and expensive by today's standards, some people immediately saw their potential for natural language applications, such as machine translation or voice typewriters. Indeed, one of the very first useful applications of a computer was language related, namely, its use by the British – under the leadership of the great logician, Alan Turing – to break German codes and decipher their secret messages during the Second World War.

A third relevant development around mid-century was the decline of behaviorism in the social sciences. Like many other disciplines, linguistics in America at that time had been dominated by behaviorist thinking. That is, it was considered unscientific to posit mental entities or states to account for human behaviors. Rather, everything was supposed to be described in terms of correlations between stimuli and responses. Abstract models of what might be going on inside people's minds were taboo. Around 1950, some psychologists began to question these methodological restrictions and to show how they made it impossible to explain certain kinds of facts. This set the stage for a serious rethinking of the goals and

[7]From the Greek: *dia* 'across' plus *chronos* 'time'

[8]*syn* 'same, together' plus *chronos*.

methods of linguistic research.

In the early 1950s, a young man named Noam Chomsky entered the field of linguistics. In the late 50s, he published three things whose impact on the study of syntax was revolutionary. One was a set of mathematical results, establishing the foundations of what is now called 'formal language theory'. These results have been seminal in theoretical computer science, and are crucial underpinnings for computational work on natural language. The second was a little book called *Syntactic Structures* that presented a new formalism for grammatical description and analyzed a substantial fragment of English in terms of the new formalism. The third was a review of B.F. Skinner's (1957) book *Verbal Behavior*. Skinner was one of the most influential psychologists of the time and an extreme behaviorist. Chomsky's review was scathing and devastating. In many people's minds, it marks the end of behaviorism's dominance in American social science.

Since about 1960, Chomsky has been the dominant figure in linguistics. As it happens, the 1960s were a period of unprecedented growth in American academia. Most linguistics departments in the United States were established in the period between 1960 and 1980. This helped solidify Chomsky's dominant position.

One of the central tenets of the Chomskyan approach to syntax, known as 'generative grammar', has already been introduced. It is that hypotheses about linguistic structure should be made precise enough to be testable. A second, somewhat more controversial one is that the object of study should be the unconscious knowledge underlying ordinary language use. A third fundamental claim of Chomsky's concerns the biological basis of human linguistic abilities. We will return to this claim in the next section.

Within these general guidelines there is room for many different theories of grammar. Over the past forty years, generative grammarians have explored a wide variety of choices of formalism and theoretical vocabulary. We present a brief summary of these in the next chapter, to help situate the approach presented here within a broader intellectual landscape.

1.4 Why Study Syntax?

Students in syntax courses often ask what the point of such classes is: why should one study syntax?

Of course one has to distinguish this question from a closely related one: why DO people study syntax? The answer to that question is perhaps simpler: exploring the structure of language is an intellectually challenging and, for many people, intrinsically fascinating activity. It is like working on a gigantic puzzle – one so large that it could occupy many lifetimes. Thus, as in any scientific discipline, many researchers are simply captivated by the complex mysteries presented by the data themselves – in this case a seemingly endless, diverse array of languages past, present and future.

This is, of course, similar to the reasons scholars in any scientific field pursue their research: natural curiousity and fascination with some domain of study. Basic research is not typically driven by the possibility of applications. Although looking for science that will be useful in the short term might be the best strategy for someone seeking personal fortune, it wouldn't be the best strategy for a society looking for long term benefit from the scientific research it supports. Basic scientific investigation has proven over the centuries to have long-term payoffs, even when the applications were not evident at the time the research was carried out. For example, work in logic and the foundations of mathematics in the first decades of the twentieth century laid the theoretical foundations for the development of the digital computer, but the scholars who did this work were not concerned with its possible applications. Likewise, we don't believe there is any need for linguistic research to be justified on the basis of its foreseeable uses. Nonetheless, we will mention three interrelated but distinct reasons that one might give for the value of studying the syntax of human languages.

1.4.1 A Window on the Structure of the Mind

One intellectually important rationale for the study of syntax has been offered by Chomsky. In essence, it is that language – and particularly, its grammatical organization – can provide an especially clear window on the structure of the human mind[9].

Chomsky says that the most remarkable fact about human language is the discrepancy between its apparent complexity and the ease with which children acquire it. The structure of any natural language is far more complicated than those of artificial languages or of even the most sophisticated mathematical systems. Yet learning computer languages or mathematics requires intensive instruction (and many students still never master them), whereas every normal child learns at least one natural language merely through exposure. This amazing fact cries out for explanation[10].

Chomsky's proposed explanation is that most of the complexity of languages does not have to be learned, because we are born knowing about it. That is, our brains are 'hardwired' to learn languages of certain types.

More generally, Chomsky has argued that the human mind is highly modular. That is, we have special-purpose 'mental organs' that are designed to do particular sorts of tasks in particular sorts of ways. The language organ (which, in Chomsky's view, has several largely autonomous sub-modules) is of particular interest because language is such a pervasive and unique part of human nature. All people have language, and (he claims) no other species is capable of learning anything much like human language. Hence, in studying the structure of human languages, we are investigating a central aspect of human nature.

This idea has drawn enormous attention not only from linguists but also from people outside of linguistics, especially psychologists and philosophers. Scholars in these fields have been highly divided about Chomsky's innateness claims. Many cognitive psychologists see Chomsky's work as a model for how mental faculties should be studied, while others argue that the mind (or brain) should be regarded as a general-purpose thinking device, without specialized modules. In philosophy, Chomsky provoked much comment by claiming that his work constitutes a modern version of Descartes' doctrine of innate ideas.

Chomsky's innateness thesis and the interdisciplinary dialogue it stimulated were major factors in the birth of the new interdisciplinary field of 'cognitive science' in the 1970s. (An

---

[9]See Katz and Postal (1991) for arguments against the dominant Chomskyan conception of linguistics as essentially concerned with psychological facts.

[10]Chomsky was certainly not the first person to remark on the extraordinary facility with which children learn language, but, by giving it a central place in his work, he has focused attention on it.

even more important factor was the rapid evolution of computers, with the concomitant growth of artificial intelligence and the idea that the computer could be used as a model of the mind). Chomsky and his disciples have been major contributors to cognitive science in the subsequent decades.

One theoretical consequence of Chomsky's innateness claim is that all languages must share most of their structure. This is because all children learn the languages spoken around them, irrespective of where their ancestors came from. Hence, the innate knowledge that Chomsky claims makes language acquisition possible must be common to all human beings. If this knowledge also determines most aspects of grammatical structure, as Chomsky says it does, then all languages must be essentially alike. This is a very strong universal claim.

In fact, Chomsky tends to use the term 'Universal Grammar' as synonymous with the innate endowment that permits language acquisition. A great deal of the syntactic research since the late 1960s has been concerned with identifying linguistic universals, especially those that could plausibly be claimed to reflect innate mental structures operative in language acquisition. As we proceed to develop the grammar in this text, we will ask which aspects of our grammar are peculiar to English and which might plausibly be considered universal.

If Chomsky is right about the innateness of the language faculty, then it has a number of practical consequences, especially in fields like language instruction and therapy for language disorders. For example, since there is evidence that people's innate ability to learn languages is far more powerful very early in life (specifically, before puberty) than later, it seems most sensible that elementary education should have a heavy emphasis on language, and that foreign language instruction should NOT be left until secondary school, as it is in most American schools today.

1.4.2 A Window on the Mind's Activity

If you stop and think about it, it's really quite amazing that people succeed in communicating using English (or any other natural language, for that matter). English seems to have a number of design properties that get in the way of efficient and accurate communication of the kind that routinely takes place.

First, it is massively ambiguous. Individual words, for example, often have not just one, but a number of meanings, as illustrated in (9).

(9) a. Leslie used a *pen*. ('a writing implement')

b. We put the pigs in a *pen*. ('a fenced enclosure')

c. They should *pen* the letter quickly. ('to write')

d. The judge sent them to the *pen* for a decade. ('a penitentiary')

(10) a. The cheetah will *run* down the hill. ('to move fast')

b. The president will *run*. ('to be a political candidate')

c. The car won't *run*. ('to function properly')

d. This trail should *run* over the hill. ('to lead')

e. This dye will *run*. ('to dissolve and spread')

f. This room will run $200 or more. ('to cost').

g. She can *run* an accelerator. ('to operate')

h. They will *run* the risk. ('to incur')

i. These stockings will *run*. ('to tear')

j. There is a *run* in that stocking. ('a tear')

k. We need another *run* to win. ('a score in baseball')

l. Fats won with a *run* of 20. ('a sequence of successful shots in a game of pool')

To make matters worse, many sentences are ambiguous not because they contain ambiguous words, but rather because the words they contain can be related to one another in more than one way, as illustrated in (11).

(11) a. Lee saw the student with a telescope.

b. I forgot how good beer tastes.

(11a) can be interpreted as providing information about which student Lee saw (the one with a telescope) or about what instrument Lee used (the telescope) to see the student. Similarly, (11b) can convey either that the speaker forgot how GOOD beer (as opposed to BAD or MEDIOCRE beer) tastes, or else that the speaker forgot that beer (in general) tastes good. These differences are often talked about in terms of which element a word like *with* or *good* is modifying (the verb or the noun).

Lexical and modificational ambiguity interact to produce a bewildering array of (often comical) ambiguities, like these:

(12) a. Visiting relatives can be boring.

b. If only Superman would stop flying planes!

c. That's a new car dealership.

d. I know you like the back of my hand. (Think!)

e. An earthquake in Roumania moved buildings as far away as Moscow and Rome.

f. The German shepherd turned on its master.

g. I saw that gas can explode.

h. Max is on the phone now.

i. The only thing capable of consuming this food has four legs and flies.

j. I saw her duck.

Problem 4: Ambiguity Give a brief description of each ambiguity illustrated in (12). We will return to many of these examples – or closely related ones – later in the book.

This is not the end of the worrisome design properties of human language. Many words are used to refer to different things on different occasions of utterance. Pronouns like *them*, *(s)he*, *this*, and *that* pick out different referents almost every time they are used. Even seemingly determinate pronouns like *we* don't pin down exactly which set of people the speaker is referring to (compare *We have two kids/a city council/a lieutenant governer/50 states/oxygen-based life here.*). Moreover, although certain proper names like *Sally Ride* or *Sandra Day O'Connor* might reliably pick out the same person almost every time they are used, most conversations are full of uses of names like *Chris*, *Pat*, *Leslie*, *Sandy*, *Bo*, etc. that vary wildly in their reference, depending on who's talking to whom and what they're talking about.

Add to this the observation that some expressions seem to make reference to 'covert elements' that don't exactly correspond to any one word. So expressions like *in charge* and *afterwards* make reference to missing elements of some kind – bits of the meaning that have to be supplied from context. Otherwise, discourses like the following wouldn't make sense, or would at best be incomplete.

(13) a. I'm creating a committee. Kim – you're in charge. [in charge of what? – the committee]

b. Lights go out at ten. There will be no talking afterwards. [after what? – after ten]

The way something is said can also have a significant effect on the meaning. A rising intonation, for example, on a one word utterance like *Coffee?* would very naturally convey 'Do you want some coffee?'. Alternatively, it might be used to convey that 'coffee' is being offered as a tentative answer to some question (say, 'What was Columbia's former number one cash crop?'). Or even, in the right context, the same utterance might be used to confirm that a given liquid was in fact coffee. Intonational meaning can be vivified in striking ways.

Finally, note that communication using language leaves a great deal unsaid. So if I say to you *Can you give me a hand here?*, I'm not just requesting information about your abilities, I'm asking you to help me out. This is the unmistakable communicative intent, but it wasn't said. Other examples of such inference are similar, but perhaps more subtle. A famous example[11] is the letter of recommendation saying that the candidate in question has outstanding penmanship (and saying nothing more than that!).

Summing all this up, what we have just seen is that the messages conveyed by utterances of sentences are multiply ambiguous, vague, and uncertain. Yet somehow, in spite of this,

[11]This example is one of many due to the late H. Paul Grice, the philosopher whose work forms the starting point for much work in linguistics on problems of PRAGMATICS, how people 'read between the lines' in natural conversation; see Grice (1989).

those of us who know the language are able to use it to trasmit messages to one another with considerable accuracy – far more accuracy than the language itself would seem to permit. Those readers who have any experience with computer programming or with mathematical logic will appreciate this dilemma instantly. The very idea of designing a programming language or a logical language whose predicates are ambiguous or whose variables are left without assigned values is unthinkable. No computer can process linguistic expressions unless it 'knows' precisely what the expressions mean and what to do with them.

The fact of the matter is that human language users are able to do something that modern science doesn't understand well enough to replicate via computer. Somehow, people are able to use nonlinguistic information in such a way that they never even see most of the unwanted intepretations of words, phrases and sentences. Consider again the various senses of the word *pen*. The 'writing implement' sense is more common –i.e. more frequent in the language you've been exposed to (unless you're a farmer or a prisoner) – and so there is an inherent bias toward that sense. You can think of this in terms of 'weighting' or 'degrees of activation' of word senses. In a context where farm animals are being discussed, though, the weights shift – the senses more closely associated with the subject matter of the discourse become stronger in this case. As people direct their attention to and through a given dialogue, these sense preferences can fluctuate considerably. The human sense selection capability is incredibly robust, yet we have only minimal understanding of the cognitive mechanisms that are at work. How exactly does context facilitate lexical sense lookup?

In other cases, it's hard to explain disambiguation so easily in terms of affinity to the discourse domain. Consider this contrast:

(14) a. They found the book on the table.

b. They found the book on the atom.

The preposition *on* modifies the verb in (14a) and the noun in (14b), yet it seems that nothing short of rather complex reasoning about the relative size of objects would enable someone to choose which meaning (i.e. which modification) made sense. And we do this kind of thing very fast, as you can see from (15).

(15) After finding the book on the atom, Sandy went into class, confident that there would be no further obstacles to getting that term paper done.

When you read this sentence, there's no strong feeling that you were 'garden pathed', i.e. derailed by an incorrect interpretation midsentence. The decision about how to construe *on the atom* is made well before the words *class* or *confident* are even encountered.

When we process language, we integrate encyclopedic knowledge, plausibility information, frequency biases, discourse information, and perhaps more. Although we don't yet know exactly how we do it, it's clear that we do it very quickly and reasonably accurately. Trying to model this integration is probably the most important research task facing the study of language in the coming millenium.

And syntax plays a crucial role in all this. It provides constraints on how sentences can or cannot be construed. So the discourse context may provide a bias for the 'fenced enclosure' sense of *pen*, but it is the syntactic context that determines whether *pen* occurs as a noun or a verb. Syntax is also of particular importance to the development of language processing models because it is a domain of knowledge that can be characterized more precisely perhaps than some of the other kinds of knowledge that are involved.

When we understand how language processing works, we will probably also understand quite a bit more about how cognitive processes work in general. This in turn will no doubt enable us to develop better ways of teaching language. We should also be better able to help people who have communicative impairments (and more general cognitive disorders) of various kinds. The study of human language processing is an important subarea of the study of human cognition, and it is one that can benefit immensely from precise characterization of linguistic knowledge of the sort that syntacticians seek to provide.

1.4.3 Natural Language Technologies

Grammar has more utilitarian applications, as well. One of the most promising areas for applying syntactic research is in the development of useful and robust natural language technologies. What do we mean by the phrase 'natural language technologies'? Roughly, what we have in mind is any sort of computer application that involves natural languages like English, Japanese, or Swahili in essential ways. These include devices that translate from one language into another (or perhaps more realistically, that provide translation assistance to someone with less than perfect command of a language), that understand spoken language (to varying degrees), that automatically retrieve information from large bodies of text stored on-line, or that help the disabled to communicate.

There is one application that rather obviously must incorporate a great deal of grammatical information, namely grammar checkers for word processing. Most modern word processing systems include a grammar checking facility, along with their spell-checkers. These tend to focus on the concerns of prescriptive grammar, which may be appropriate for the sorts of documents they are generally used on, but often leads to spurious 'corrections'. Moreover, they typically depend on superficial pattern matching for finding things likely to be grammatical errors, rather than employing in depth grammatical analysis. In short, grammar checkers can benefit from incorporating the results of research in syntax.

Other applications in which grammatical knowledge is clearly essential include those in which well-formed natural language output must be generated by the computer. For example, reliable software for translating one language into another must incorporate some representation of the grammar of the target language. If it did not, it would either produce ill-formed output, or it would be limited to some fixed repertoire of sentence templates.

Even where usable natural language technologies can be developed that are not informed by grammatical research, it is often the case that they can be made more robust by including a principled syntactic component. For example, a project at Stanford University's Center for the Study of Language and Information is developing software to reduce the number of keystrokes needed to input text. This has many potential uses, including facilitating the use of computers by individuals with motor disabilities or temporary impairments such as carpal tunnel syndrome. It is clear that knowledge of the grammar of English can help in predicting what words are likely to come next at an arbitrary point in a sentence. Software that makes such predictions and offers the user a set of choices for the next word or the remainder of an entire sentence – each of which can be inserted with a single keystroke – can be of great value in a wide variety of situations. Word prediction can likewise facilitate the disambiguation of noisy signals in continuous speech recognition and handwriting recognition.

But it's not obvious that all types of natural language technologies need to be sensitive to grammatical information. Say, for example, you were designing a system to extract information from an on-line database by typing in English questions (rather than requiring use of a special database query language, as is the case with most existing database systems). Some computer scientists have argued that full grammatical analysis of the queries is not necessary. Instead, they claim, all that is needed is a program that can extract the essential semantic information out of the queries. Many grammatical details don't seem necessary in order to understand the queries, so it has been argued that they can be ignored for the purpose of this application. Even here, however, a strong case can be made for the value of including a syntactic component in the software.

To see why, imagine that we are dealing with a database in a law office, containing information about the firm's past and present cases, including records of witnesses' testimony. It turns out that, without careful attention to certain details of English grammar, there are questions we might want to ask of this database that could be misanalyzed, and hence answered incorrectly.

Let's start with our old friend, the rules for reflexive and nonreflexive pronouns. Since formal database query languages don't make any such distinction, one might think that it wouldn't be necessary for an English interface to do so either. But suppose a user asked one of the following questions:

(16) a. Which witnesses work with defendants who supervise them?

b. Which witnesses work with defendants who supervise themselves?

Obviously, these two questions will have different answers, so an English language 'front end' that didn't incoporate some rules for distinguishing reflexive and nonreflexive pronouns would sometimes give wrong answers.

In fact, it isn't enough to tell reflexive from nonreflexive pronouns: a database system would need to be able to tell different reflexive pronouns apart. The next two sentences, for example, are identical except for the plurality of the reflexive pronouns:

(17) a. List all witnesses for the defendant who represented himself.

b. List all witnesses for the defendant who represented themselves.

Again, the appropriate answers would be different. So a system that didn't pay attention to whether pronouns are singular or plural couldn't be trusted.

Even features of English grammar that seem really useless – things that appear to be entirely redundant – can be needed for the analysis of some sentences that might well be used in a human-computer interaction. Consider, for example, English subject-verb agreement (a topic we will return to in some detail in chapters 3–5). Since subjects are marked as singular or plural – *the dog* vs. *the dogs* – marking verbs for the same thing – *barks* vs. *bark* – seems to add nothing. We would have little trouble understanding someone who always left subject agreement off of verbs. In fact, English doesn't even mark past-tense verbs (other than forms of *be*) for subject agreement. But we don't miss agreement in the past tense, because it is semantically redundant. One might conjecture, therefore, that an English database querying system might be able simply to ignore agreement.

It turns out, however, that, once again, examples can be constructed in which agreement marking on the verb is the only indicator of a crucial semantic distinction. This is the case with the following pair:

(18) a. List associates of each witness who speaks Spanish.

b. List associates of each witness who speak Spanish.

In the first sentence, it is the witnesses in question who are the Spanish-speakers; in the second, it is their associates. These will, in general, not lead to the same answer.

Such examples could be multiplied, but these should be enough to make the point: Building truly robust natural language technologies – that is, software that will allow you to interact with your computer in YOUR language, rather than having to learn ITS language – requires careful and detailed analysis of grammatical structure and how it influences meaning. Short-cuts that rely on semantic heuristics, guesses, or simple pattern-matching will inevitably make mistakes. Computational linguists have learned through painful experience that it is unwise to ignore some type of locution on the grounds that 'nobody will ever use that construction'.

In short, the study of grammar, while motivated primarily simply by intellectual curiousity, has some fairly obvious uses, even in the relatively short term.

1.5 Conclusion

In this chapter, we have drawn an important distinction between PRESCRIPTIVE and DESCRIPTIVE grammar. In addition, we have provided an illustration of the kind of syntactic puzzles we will focus on later in the text. Finally, we provided an overview of some of the reasons people have found the study of syntax to be of inherent interest or useful. In the next chapter, we look at some very simple formal models that might be proposed for the grammars of natural languages and discuss some of their shortcomings. We then briefly survey some of the theories of grammar that have been developed within the generative tradition.

1.6 Further Reading

An entertaining and knowledgeable exposition of modern linguistics and its implications is provided by Pinker (1994). A somewhat more scholarly survey with a slightly different focus is presented by Jackendoff (1994). For discussion of prescriptive grammar, see Nunberg (1983) and chapter 12 of Pinker's book (an edited version of which was published in *The New Republic*, January 31, 1994). For an overview of linguistic science in the 19th century, see Pedersen (1959). A succinct survey of the history of linguistics is provided by Robins (1967).

Among Chomsky's many writings on the implications of language acquisition for the study of the mind, we would especially recommend Chomsky (1959) and Chomsky (1972); a more recent, but much more difficult work is Chomsky (1986a). There have been few recent attempts at surveying work in (human or machine) sentence processing. J. A. Fodor, Bever, and Garrett (1974) was a comprehensive review of early psycholinguistic work within the Chomskyan paradigm, but it is now quite dated. Garrett (1990) and J. D. Fodor (1995) are more recent, but much more limited in scope.

Chapter 2

Some Theories of Grammar

2.1 Introduction

Among the key points in the previous chapter were the following:

- Language is rule-governed.
- The rules aren't the ones we were taught in school.
- Much of our linguistic knowledge is unconscious, so we have to get at it indirectly; one way of doing this is by consulting intuitions of what sounds good.

Our central objectives in this text are to:

- work on developing a set of rules that will deal with most of English (i.e. we ideally want a grammar that can tell us for any arbitrary string of English words whether or not it is a well-formed sentence);
- consider how the grammar of English differs from the grammar of other languages (or how the grammar of standard American English differs from those of other varieties of English);
- look at the kinds of rules we need, and consider what (if anything) this tells us about human linguistic abilities.

In developing the informal rules for reflexive and nonreflexive pronouns in Chapter 1, we assumed that we already knew a lot about the structure of the sentences we were looking at – that is, we talked about subjects, objects, clauses, etc. In fact, a fully worked out theory of reflexive and nonreflexive pronouns is going to require that many other aspects of syntactic theory get worked out first. We begin this grammar development process in the present chapter.

We will consider several candidates for theories of English grammar. We begin with approaches that are quite simple-minded, quickly dismissing them. We spend some time on a formalism known as 'context-free grammar', which serves as a starting point for most modern theories of syntax. We end with a brief overview of some of the most important schools of thought within the paradigm of generative grammar, situating the approach developed in this text with respect to some alternatives.

2.2 Two Simplistic Syntactic Theories

2.2.1 Lists as Grammars

The simplest imaginable syntactic theory would be that a grammar consists of a list of all the well-formed sentences in the language. The most obvious problem with such a proposal is that the list would have to be too long. There is no fixed finite bound on the length of English sentences, as can be seen from the following sequence:

(1) Some sentences go on and on.
Some sentences go on and on and on.
Some sentences go on and on and on and on.
Some sentences go on and on and on and on and on.
...

Every sentence in this sequence is acceptable English. Since there is no bound on their size, it follows that the number of sentences in the list must be infinite. Hence, there are infinitely many sentences of English. Since human brains are finite, they cannot store infinite lists. Consequently, there must be some more compact way of encoding the grammatical knowledge that speakers of English possess.

Moreover, there are generalizations about the structure of English that an adequate grammar should express. For example, suppose the list in (1) were replaced by one in which every other sentence reversed the order of the words *some* and *sentences*:

(2) Some sentences go on and on.

*Sentences some go on and on.

*Some sentences go on and on and on.

Sentences some go on and on and on.

Some sentences go on and on and on and on.

*Sentences some go on and on and on and on.

*Some sentences go on and on and on and on and on.

Sentences some go on and on and on and on and on. ...

Of course, the sentences with the word 'sentences' before the word 'some' are not well-formed English. Moreover, no natural language exhibits patterns of that sort – in this case, having word-order depend on whether the length of the sentence is divisible by 4. A syntactic theory that sheds light on human linguistic abilities ought to explain why such patterns do not occur in human languages. But a theory that said grammars consisted of lists of sentences could not do that. If grammars were just lists, then there would be no patterns that would be excluded – and none that would be expected, either.

This form of argument – that a certain type of grammar fails to 'capture a linguistically significant generalization' is very common in generative grammar. It takes for granted the idea that language is 'rule governed', i.e. that language is a combinatoric system whose operations are 'out there' to be discovered by empirical investigation. If a particular characterization of the way a language works leads to redundancy and complications, it's assumed to be the wrong characterization of the grammar of that language. We will see this kind of argumentation again, in connection with more plausible proposals than those we have looked at so far.

2.2.2 Regular Expressions

A natural first step towards allowing grammars to capture generalizations is to classify words into what are often called 'parts of speech' or 'grammatical categories'. There are large numbers of words that behave very similarly syntactically. For example, the words *apple, book, color*, and *dog* all can appear in roughly the same contexts, such as the following:

(3) a. That __ surprised me.

b. I noticed the __ .

c. They were interested in his __ .

d. This is my favorite __ .

Moreover, they all have plural forms that can be constructed in similar ways (orthographically, simply by adding an *-s*).

Traditionally, the vocabulary is sorted into nouns, verbs, etc. based on loose semantic characterizations (e.g. 'a noun is a word that refers to a person, place, or thing'). While there is undoubtedly a grain of insight at the heart of such definitions, we can make use of this division into grammatical categories without committing ourselves to any semantic basis for them. For our purposes, it is sufficient that there are classes of words which may occur grammatically in the same environments. Our theory of grammar can capture their common behavior by formulating patterns or rules in terms of categories, not individual words.

Someone might, then, propose as a grammar of English that we have a list of patterns, stated in terms of grammatical categories, together with a lexicon – that is, a list of words and their categories. For example, the patterns could include (among many others):

(4) a. ARTICLE NOUN VERB

b. ARTICLE NOUN VERB ARTICLE NOUN

And the lexicon could include (likewise, among many others):

(5) a. Articles: a, the

b. Nouns: cat, dog

c. Verbs: attacked, scratched

This mini-grammar licenses forty well-formed English sentences, and captures a few generalizations. However, a grammar that consists of lists of patterns still suffers from the first drawback of the theory of grammars as lists: it can only account for a finite number of sentences, but natural languages are infinite. For example, such a grammar will still be incapable of dealing with all of the sentences in the infinite sequence illustrated in (1).

We can enhance our theory of grammar to permit infinite numbers of sentences by introducing abbreviatory devices. In particular, the problem associated with (2) can be handled using what is known as the 'Kleene star' (after the logician Stephen Kleene). Notated as a superscripted asterisk, the Kleene star is interpreted to mean that the expression it is attached to can be repeated any finite number of times (including zero). Thus, the examples in (2) could be abbreviated as follows:

(6) Some sentences go on and on [and on]*.

A closely related notation is a superscripted plus-sign (called Kleene-plus), meaning one or more occurrences of the expression it is attached to. Hence, another way of expressing the same pattern would be:

(7) Some sentences go on [and on]$^+$.

We shall employ these, as well as two other common abbreviatory devices. The first is simply to put parentheses around material that is optional For example, the two sentence patterns in (4) could be collapsed into: ARTICLE NOUN VERB (ARTICLE NOUN). The second abbreviatory device is a vertical bar, which is used to separate alternatives[1]. For example, if we wished to expand the mini-grammar in (4) to include sentences like *The dog looked angry*, we could add the pattern ARTICLE NOUN VERB ADJECTIVE and collapse it with the previous patterns as: ARTICLE NOUN VERB (ARTICLE NOUN)|ADJECTIVE. Of course, we would also have to add the verb *looked* and the adjective *angry* to the lexicon[2].

Patterns making use of the devices just described – Kleene star, Kleene-plus, parentheses for optionality, and the vertical bar for alternatives – are known as 'regular expressions'[3]. A great deal is known about what sorts of patterns can and cannot be represented with regular expressions (see, for example, Hopcroft and Ullman (1979: chaps. 2 and 3)), and a number of scholars have argued that natural languages in fact exhibit patterns that are beyond the descriptive capacity of regular expressions (see Bar-Hillel and Shamir (1960: secs. 5 and 6)). The most convincing arguments for employing a grammatical formalism richer than regular expressions, however, have to do with the need to capture generalizations.

In (5), the string ARTICLE NOUN occurs twice, once before the verb and once after it. Notice that there are other options possible in both of these positions:

(8) a. Dogs chase cats.

---

[1]This is the notation standardly used in computer science and in the study of mathematical properties of grammatical systems. Descriptive linguists tend to use curly brackets to annotate alternatives.

[2]This extension of the grammar would license some unacceptables strings, e.g., **The cat scratched angry*. This sort of overgeneration is always a danger when extending a grammar, as we will see in subsequent chapters.

[3]Readers who use computers with the UNIX operating system may be familiar with the command 'grep'. This stands for 'Get Regular Expression and Print'.

b. A large dog chased a small cat.

c. A dog with brown spots chased a cat with no tail.

Moreover, these are not the only positions in which the same strings can occur:

(9) a. Some people yell at (the) (noisy) dogs (in my neighborhood).

b. Some people consider (the) (noisy) dogs (in my neighborhood) dangerous.

Even with the abbreviatory devices available in regular expressions, the same lengthy string of symbols – something like (ARTICLE) (ADJECTIVE) NOUN (PREPOSITION ARTICLE NOUN) – will have to appear over and over again in the patterns that constitute the grammar. Moreover, the recurring patterns are in fact considerably more complicated than those illustrated so far. Strings of other forms, such as *the noisy annoying dogs*, *the dogs that live in my neighborhood*, or *Rover, Fido, and Lassie* can all occur in just the same positions. It would clearly simplify the grammar if we could give this apparently infinite set of strings a name and say that any string from the set can appear in certain positions in a sentence.

Furthermore, as we have already seen, an adequate theory of syntax must somehow come to grips with the fact that a given string of words can sometimes be put together in more than one way. If there is no more to grammar than lists of recurring patterns, where these are defined in terms of parts of speech, then there is no apparent way to talk about the ambiguity of sentences like those in (10):

(10) a. We enjoyed the movie with Cher.

b. The room was filled with noisy children and animals.

c. People with children who use drugs should be locked up.

d. I saw the astronomer with a telescope.

In the first sentence, it can be us or the movie that is 'with Cher'; in the second, it can be either just the children or both the children and the animals that are noisy; in the third, it can be the children or their parents who use drugs, and so forth. None of these ambiguities can be plausibly attributed to a lexical ambiguity. Rather, they seem to result from different ways of grouping the words into phrases.

In short, the fundamental defect of regular expressions as a theory of grammar is that they provide no means for grouping sequences of elements together to form a unit. The same holds true of several other formalisms that are provably equivalent to regular expressions (including what is known as 'finite state grammar').

The recurrent strings we have been seeing are usually called 'phrases' or '(syntactic) constituents'.[4] Phrases, like words, come in different types. All of the phrases in (8) – (9) above obligatorily include a noun, so they are called 'Noun Phrases'. The next natural enrichment of our theory of grammar is to permit our regular expressions to include not only

[4]There is a minor difference in the way these terms are used: linguists often use 'phrase' in contrast to 'word' to mean something longer, whereas words are always treated as a species of constituent.

words and parts of speech, but also phrase types. Then we also need to provide (similarly enriched) regular expressions to provide the patterns for each type of phrase. The technical name for this theory of grammar is 'Context-free Phrase Structure Grammar' or simply 'Context-free Grammar', sometimes abbreviated as CFG. CFGs, which will also let us begin to talk about structural ambiguity, form the starting point for most serious attempts to develop formal grammars for natural languages.

2.3 Context-free Phrase Structure Grammar

The term 'grammatical category' now covers not only the parts of speech, but also types of phrases, such as noun phrase and prepositional phrase. To distinguish the two types, we will sometimes use the terms 'lexical category' (for parts of speech) and 'non-lexical category' or 'phrasal category' to mean types of phrases. For convenience, we will abbreviate them, so that 'NOUN' becomes 'N', 'NOUN PHRASE' becomes 'NP', etc.

A context-free phrase structure grammar has two parts:

- A LEXICON, consisting of a list of words, with their associated grammatical categories.[5]

- A set of RULES of the form A $\rightarrow \varphi$ where A is a non-lexical category, and 'φ' stands for a regular expression formed from lexical and/or non-lexical categories; the arrow is to be interpreted as meaning, roughly 'may consist of'. These rules are sometimes called 'phrase structure rules'.

Intuitively, the left-hand side of each rule specifies a phrase type (including the sentence as a type of phrase), and the right-hand side gives a possible pattern for that type of phrase. Because phrasal categories can appear on the right-hand sides of rules, it is possible to have phrases embedded within other phrases. This permits CFGs to express regularities that seem like accidents when only simple regular expressions are permitted.

A CFG normally has one or more phrasal categories that are designated as 'initial symbols'. These are the types of phrases that can stand alone as sentences in the language. Most simple CFGs have just one initial symbol, namely 'S'. Any string of words that can be derived from one of the initial symbols by means of a sequence of applications of the rules of the grammar is licensed (or, as linguists like to say, 'generated') by the grammar. The language a grammar generates is simply the collection of all of the sentences it generates.

2.3.1 An Example

Consider the following CFG. (We use 'D' for 'Determiner', which includes what we have up to now been calling 'articles'.)

[5]This conception of a lexicon is rather impoverished. In particular, it leaves out information about the meanings and uses of words, except what might be generally associated with the grammatical categories. While this impoverished conception is what is standard in the formal theory of CFG, attempts to use CFG to describe actual natural languages have had lexicons that also included semantic information.

(11) a. **Rules**:
$S \rightarrow NP\ VP$
$NP \rightarrow (D)\ A^*\ N\ PP^*$
$VP \rightarrow V\ (NP)\ (PP)$
$PP \rightarrow P\ NP$

b. **Lexicon**:
D: the, some
A: big, brown, old
N: birds, fleas, dog, hunter
V: attack, ate, watched
P: for, beside, with

This grammar generates infinitely many English sentences. Let us look in detail at how it generates one sentence: *The big brown dog with fleas watched the birds beside the hunter.* We start with the symbol S, for 'Sentence'. This gets expanded to the sequence NP VP, by the first rule. The second rule allows a wide range of possibilities for the NP, one of which is D A A N PP. This PP must be expanded as P NP, by the fourth rule, and the NP so introduced can be expanded as D N. The third rule allows VP to be expanded as V NP PP, and this NP expands simply as N. Lastly, the final PP is expanded as P NP, and this NP, in turn, expands as D N. Putting these steps together yields the string D A A N P D N V N P D N, which can be converted into the desired sentence by inserting appropriate words in place of their lexical categories. This derivation can be summarized in the following figure (called a 'tree diagram'):

(12)

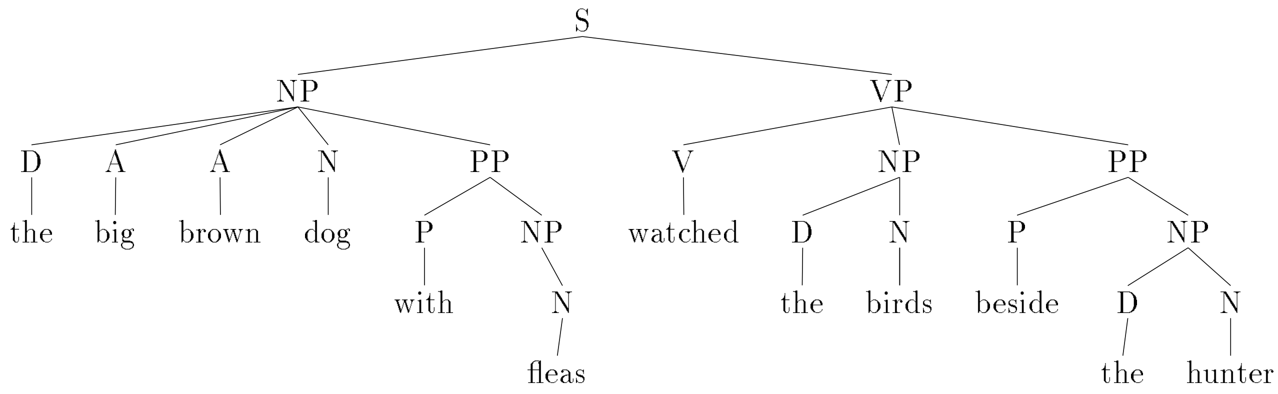

It is important to note that certain sentences generated by this grammar can be associated with more than one tree. (Indeed, the example just given is one such sentence, but finding the other tree will be left as an exercise). This illustrates how CFGs can overcome the second defect of regular expressions pointed out at the end of the previous section. Recall the ambiguity of (13):

(13) I saw the astronomer with a telescope.

The two distinct interpretations of this sentence ('I used the telescope to see the astronomer'; 'What I saw was the astronomer who had a telescope') correspond to two distinct tree structures that our grammar will assign to this string of words. The first interpretation corresponds to the tree where the PP *with a telescope* hangs from the VP; the latter is the meaning associated with the tree structure where that PP is part of the NP constituent: *the astronomer with a telescope*. CFG thus provides us with a straightforward mechanism for expressing such ambiguities, whereas grammars that use only regular expressions don't.

The normal way of talking about words and phrases is to say that certain word strings 'form a constituent'. What this means is that these strings function as units for some purpose (for example, the interpretation of modifiers) within the sentences in which they appear. So in (12), the sequence *with the cats* forms a PP constituent and *the big brown dog with the cats* forms an NP, the sequence *dog with the cats* forms no constituent. Structural ambiguity arises whenever a string of words can form constituents in more than one way.

Problem 1: Practice with CFG Assume the CFG grammar given in (11). Draw the tree structure for the other interpretation (i.e. not the one shown in (12)) of *The big brown dog with fleas watched the birds near the hunter.*

◇**Problem 2: More Practice with CFG** Assume the grammar rules given in (11), but with the following lexicon:

A:	big, unusual, young
D:	a, the
N:	cat, dog, hat, man, woman
P:	in, on, with
V:	admired, disappeared, put, relied

a. Give three grammatical sentences sanctioned by this grammar. Draw the tree structures that the grammar assigns to them.

b. Give a sentence that is structurally ambiguous according to this grammar. Draw the two distinct tree structures.

c. Give three grammatical English sentences (using only the words from this grammar) that are not covered by this grammar. They should be examples that differ in their trees, not just in the lexical entries they produce.

d. Explain what prevents each of the examples in c. from being covered.

e. Give three sentences sanctioned by this grammar that are incorrectly predicted to be grammatical. Again, make them interestingly different.

f. Discuss how the grammar might be revised to correctly exclude your examples (without simultaneously excluding good sentences).

g. How many sentences does this grammar admit?

h. How many would it admit if the NP rule were replaced by the following rule?
NP $\rightarrow$ (D) A N (PP) Explain your answer.

2.3.2 CFG as a Theory of Natural Language Grammar

As was the case with regular expressions, the formal properties of CFG are extremely well studied (see Hopcroft and Ullman (1979: chaps. 4–6) for a summary). In the early 1960s, several scholars published arguments purporting to show that natural languages exhibit properties beyond the desciptive capacity of CFGs. The pioneering work in the first two decades of generative grammar was based on the assumption that these arguments were sound. Most of that work can be viewed as the development of extensions to CFG designed to deal with the richness and complexity of natural languages.

The most celebrated proposed extension was a kind of rule called a 'transformation', invented by Zellig Harris in the early 1950s and further refined by (his student) Chomsky. Transformations are mappings from phrase structure representations to phrase structure representations (trees to trees, in our terms) that can copy, delete, and permute parts of trees, as well as insert specified new material into them. For example, in early work on transformations, it was claimed that declarative and interrogative sentence pairs (such as *The sun is shining* and *Is the sun shining?*) were to be derived from the same underlying phrase structure by a transformation that moved certain verbs to the front of the sentence. Likewise, passive sentences (such as *The cat was chased by the dog*) were derived from the same underlying structures as their active counterparts (*The dog chased the cat*) by means of a passivization transformation. The initial trees were to be generated by a CFG. The name 'transformational grammar' is sometimes used for theories positing rules of this sort. See Appendix B for more discussion of varieties of transformational grammar.

Around 1980, the earlier arguments against the adequacy of CFG as a theory of natural language structure were called into question by Geoffrey Pullum and Gerald Gazdar in an important (1982) paper. This led to a flurry of new work on the issue, culminating in new arguments that natural languages were not describable with CFGs. The mathematical and empirical work that resulted from this controversy substantially influenced the theory of grammar presented in this text. Many of the central papers in this debate were collected together by Savitch, Bach, Marsh, and Safran-Naveh (1987); of particular interest are Pullum and Gazdar's paper and Shieber's paper.

While the question of whether natural languages are IN PRINCIPLE beyond the generative capacity of CFGs is of some intellectual interest, working linguists tend to be more concerned with determining what sort of formalisms can provide elegant and enlightening accounts of linguistic phenomena IN PRACTICE. Hence, the arguments that tend to carry the most weight are ones about what formal devices are needed to capture linguistically significant generalizations. In the next chapter, we will consider some phenomena in English that suggest that the simple version of CFG introduced above should be extended.

2.4 Modern Phrase Structure Grammar

Accompanying the 1980s revival of interest in the mathematical properties of natural languages, considerable attention was given to the idea that, with an appropriately designed theory of syntactic features and general principles, context-free phrase structure grammar could serve as an empirically adequate theory of natural language syntax. This proposition was explored in great detail by Gazdar, et al. (1985), who developed the theory known as 'Generalized Phrase Structure Grammar' (or GPSG). Because CFGs are known to be computationally tractable and computer scientists have a good deal of experience working with them, GPSG attracted considerable interest from computational linguists. A number of natural language processing projects based on GPSG were started in laboratories around the world.

One of the GPSG implementation projects was begun in 1981 at Hewlett-Packard Laboratories in Palo Alto, California, with the authors of this text and a number of colleagues, including Geoff Pullum (UC Santa Cruz), Carl Pollard (now at Ohio State University), Dan Flickinger (Stanford), and Mark Gawron (now at SRI International). The opportunity to develop working systems based on the theoretical work of this group generated considerable excitement no only at HP, but also in the interdisciplinary community centered at Stanford's Center for the Study of Language and Information (CSLI). It was here that many of us learned for the first time how far the rhetoric of theoretical linguistics can be from the reality of working grammars.

By about 1986, the framework that grew out of this interdisciplinary and interinstitutional activity had evolved to such a degree that its name was changed to Head-driven Phrase Structure Grammar, reflecting the increased importance of information encoded in the lexical heads of syntactic phrases. The theoretical aspects of HPSG have been developed in considerable detail in two books (Pollard and Sag 1987, 1994) and a number of major articles. From its inception, HPSG was developed as a conscious effort to integrate ideas from diverse linguistic theories and related research traditions into a mathematically precise foundation for linguistic theory. See Appendix B for a brief survey of the most influential of those theories and some discussion of how their results and insights relate to the analyses presented in this book.

2.5 Further Reading

The standard reference work for the basic mathematical results on formal languages (including regular expressions and context-free languages) is Hopcroft and Ullman (1979). Partee et al. (1990) covers much of the same material from a more linguistic perspective. Classic works arguing against the use of context-free grammars for natural languages include Chomsky (1963) and Postal (1967). Papers questioning these arguments, and other papers presenting new arguments for the same conclusion are collected in Savitch, et al. (1987). For (somewhat dated) surveys of theories of grammar, see Sells (1985) and Wasow (1989). A more detailed presentation of GPSG is Gazdar et al. (1985). The history of generative grammar is presented from different perspectives by Matthews (1993), Newmeyer (1986), Harris (1993), and Huck and Goldsmith (1995).

Chapter 3

Applying Context-Free Phrase Structure Grammar

3.1 Introduction

In the previous chapter, we introduced the formalism of context-free grammar and showed how it permitted us to generate infinite collections of English sentences with simple rules. We also showed how it could provide a rather natural representation of certain ambiguities we find in natural languages. But the grammar we presented was just a teaching tool, designed to illustrate certain properties of the formalism; it was not intended to be taken seriously as an attempt to analyze the structure of English. In this chapter, we begin by motivating some phrase structure rules for English. In the course of doing this, we develop a new test for determining which strings of words are constituents. We also introduce a new abbreviatory convention that permits us to collapse many of our phrase structure rules into rule schemas.

3.2 Some Phrase Structure Rules for English

For the most part, we will use the traditional parts of speech, such as noun, verb, adjective, and preposition. In some cases, we will find it useful to introduce grammatical categories that might be new to readers, and we may apply the traditional labels somewhat differently than in traditional grammar books. But this classification of words into types has proved to be an extremely useful categorization over the past two millenia, and we see no reason to abandon it wholesale.

We turn now to phrases, beginning with noun phrases.

3.2.1 Noun Phrases

Nouns can appear in a number of positions, such as the positions of the three nouns in *Dogs give people fleas.* These same positions also allow sequences of an article followed by a noun, as in *The boy gave the dog a bath.* Since the position of the article can also be filled by demonstratives (e.g. *this, these*), possessives (e.g. *my, their*), or quantifiers (e.g. *each*, *some*, *many*), we use the more general term 'determiner' (abbreviated D) for this category. We can capture these facts by positing a type of phrase we'll call NP (for 'noun phrase'), and the rule NP $\rightarrow$ (D) N. As we saw in the previous chapter, this rule will need to be elaborated later to include adjectives and other modifiers. First, however, we wish to consider a type of construction we have not yet discussed.

3.2.2 Coordination

To account for examples like *A dog, a cat, and a wombat fought*, we want a rule that allows sequences of NPs, with *and* before the last one, to appear where simple NPs can occur. A rule that does this is NP $\rightarrow$ NP$^+$ CONJ NP.

Whole sentences can also be conjoined, as in *The dog barked, the donkey brayed, and the pig squealed.* Again, we could posit a rule like S $\rightarrow$ S$^+$ CONJ S. But now we have two rules that look an awful lot alike. We can collapse them into one rule schema as follows, where 'X' can be replaced by any grammatical category name:

(1) X $\rightarrow$ X$^+$ CONJ X.

Now we have made a claim that goes well beyond the data that motivated the rule, namely, that elements of ANY category can be conjoined in the same way. If this is correct, then we can use it as a test to see whether a particular string of words should be treated as a phrase. In fact, coordinate conjunction is widely used as a test for constituency. Though it is not an infallible diagnostic, we will use it as one of our sources of evidence for constituent structure.

We now have three types of arguments for saying a given string is a constituent:

- It exemplifies a pattern that shows up in multiple environments.
- Calling it a constituent helps us account for structural ambiguity.
- It can be a coordinate conjunct.

3.2.3 Verb Phrases

Consider (2):

(2) The man yelled, chased the cat, and gave the dog a bone.

(2) contains the coordination of strings consisting of V, V NP, and V NP NP. According to (1), this means that all three strings are constituents of the same type. Hence, we posit a constituent which we'll call VP, described by the rule VP $\rightarrow$ V (NP) (NP). VP is introduced by the rule S $\rightarrow$ NP VP.[1]

[1]There are other kinds of coordinate sentences that we are leaving aside here – in particular elliptical sentences that involve coordination of non-constituent sequences:

(i) Chris likes blue and Pat green.

(ii) Leslie wants to go home tomorrow, and Terry, too.

3.2.4 Prepositional Phrases

Expressions like *in Rome* or *at noon* that denote places or times ('locative' and 'temporal' expressions, as linguists would say) can be added to almost any sentence, and to NPs, too. For example:

(3) a. The fool yelled at noon.

b. This disease gave Leslie a fever in Rome.

c. A man in Rome laughed.

These are constituents, as indicated by examples like *The fool yelled at noon and at midnight, in Rome and in Paris.* We can get lots of them in one sentence, e.g. *A man laughed on the street in Rome at noon on Tuesday.* These facts can be incorporated into the grammar in terms of the phrasal category PP (for 'prepositional phrase', and the rules:

(4) a. PP $\rightarrow$ P NP

b. VP $\rightarrow$ VP PP

Since the second rule has VP on both the right and left sides of the arrow, it can apply to its own output. (Such a rule is known as a RECURSIVE rule[2]). Each time it applies, it adds a PP to the derivation. Thus, this recursive rule permits arbitrary numbers of PPs within a VP.

As mentioned earlier, locative and temporal PPs can also occur in NPs, e.g. *A painter on the street in Rome on Tuesday at noon laughed.* The most obvious analysis to try for this would be a rule that said: NP $\rightarrow$ NP PP. However, we're going to adopt a slightly more complex analysis. We posit a new nonlexical category, which we'll call NOM, and we replace our old rule: NP $\rightarrow$ (D) N with the following:

(5) a. NP $\rightarrow$ (D) NOM

b. NOM $\rightarrow$ N

c. NOM $\rightarrow$ NOM PP

The category NOM will be very useful later in the course. For now, we will justify it with the following sentences:

(6) a. The love of my life and mother of my children would never do such a thing.

b. The museum displayed no painting by Miro or drawing by Klee.

Notice that this kind of sentence, which will not be treated by the coordination rule discussed in the text, has a characteristic intonation pattern – the elements after the conjunction form separate intonational units separated by pauses.

[2]More generally, we use the term RECURSION whenever rules permit a constituent to occur within a larger constituent of the same type.

(6b) means that the museum displayed neither paintings by Miro nor drawings by Klee. That is, the determiner *no* must be understood as having scope over both *painting by Miro* and *drawing by Klee*. The most natural noun phrase structure to associate with this interpretation is:

(7) no [painting by Miro or drawing by Klee]

This, in turn, is possible with our current rules if the bracketed string is a conjoined NOM. It would not be possible without NOM.

Similarly, for (6a), *the* has scope over both *love of my live* and *mother of my children*.

3.2.5 Summary of Grammar Rules

Our grammar now has the following rules:

(8) S $\rightarrow$ NP VP
NP $\rightarrow$ (D) NOM
VP $\rightarrow$ V (NP) (NP)
NOM $\rightarrow$ N
NOM $\rightarrow$ NOM PP
VP $\rightarrow$ VP PP
PP $\rightarrow$ P NP
X $\rightarrow$ X^+ CONJ X

3.3 Trees Revisited

In grouping words into phrases and smaller phrases into larger ones, we are assigning internal structure to sentences. As noted in the previous chapter, this structure can be represented in a tree diagram. For example, our grammar so far generates the following tree:

(9)

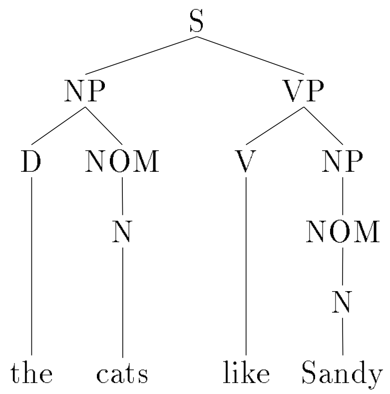

A tree is said to consist of NODES, connected by BRANCHES. A node above another on a branch is said to DOMINATE it. The nodes at the bottom of the tree – that is, those that do not dominate anything else – are referred to as TERMINAL nodes (or occasionally, as "leaves"). A node right above another node on a tree is said to be its MOTHER and to IMMEDIATELY DOMINATE it. A node right below another on a branch is said to be its DAUGHTER. Two daughters of the same mother node are, naturally, referred to as SISTERS.

One way to think of the way in which a grammar of this kind defines (or generates) trees is as follows. First, we appeal to the lexicon to tell us which LEXICAL trees are well-formed. So if *cats* is listed in the lexicon as belonging to the category N, and *like* is listed as a V, and so forth, then lexical structures like the following are well-formed.

(10) 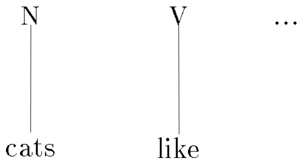

And the grammar rules are equally straightforward. They simply tell us how well-formed trees (some of which may be lexical) can be combined into bigger ones:

(11) 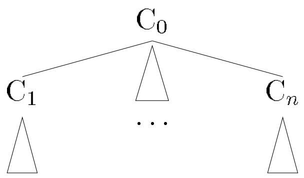

is a well-formed nonlexical tree just in case

are well-formed trees, and

$C_0 \rightarrow C_1 \ \ldots \ C_n$ is a grammar rule.

So we can think of this our grammar as generating sentences in a 'bottom-up' fashion – starting with lexical trees, and then using them to build bigger and bigger phrasal trees, until we build one whose top node is S. The set of all sentences that can be built that have S as their top node is the set of sentences the grammar generates. But note that our grammar could just as well have been used to generate sentences in a 'top-down' manner, starting with S. The set of sentences generated in this way is exactly the same. In fact, a CFG is completely neutral with respect to top-down and bottom-up perspectives on analyzing sentence structure. There is also no particular bias toward thinking of the grammar in terms of generating sentences or in terms of parsing.

These design properties – direction neutrality and process neutrality – stem from the fact that the rules and lexical entries simply provide constraints on well-formed structure. As we will suggest in chapter 9, these are in fact important design features of this theory (and of those we will develop that are based on it), as they facilitate the direct embedding of the abstract grammar within a model of language processing.

The lexicon and grammar rules together thus constitute a system for defining not only well-formed word strings (i.e. sentences), but also well-formed tree structures. Although our statement of the relationship between the grammar rules and the well-formedness of trees is at present rather trivial, as we modify our theory of grammar, our attention will increasingly turn to a more refined characterization of which trees are well-formed.

3.4 Worksection on Phrase Structure Grammar

Two of our arguments against overly simple theories of grammar in the previous chapter were that we wanted to be able to account for the infinity of language, and that we wanted to be able to account for structural ambiguity. The purpose of this section is to explore how our grammar so far handles these.

Problem 1: Structural Ambiguity Show that our grammar can account for the ambiguity of each of the following sentences by providing at least two trees licensed for each one, and explain briefly which interpretation goes with which tree:

a. Bo saw the group with the telescope.

b. Most dogs and cats with fleas live in this neighborhood.

c. The pictures show Superman and Lois Lane and Clark Kent.

Note: We haven't provided a lexicon, so technically, our grammar doesn't generate any of these. You can assume, however, that all the words in them are in the lexicon, with the obvious category assignments.

Problem 2: Infinity Our grammar has two mechanisms each of which permits us to have infinitely many sentences: the Kleene operators (plus and star), and recursion (categories that can dominate themselves). Construct arguments for why we need both of them. That is, why not use recursion to account for the unboundedness of coordination or use Kleene-star to account for the possibility of arbitrary numbers of PPs? [Hint: Consider the different groupings into phrases – that is, the different tree structures – provided by the two mechanisms. Then look for English data supporting one choice of structure over another.]

3.5 Heads

As we have seen, CFGs can provide successful analyses of quite a bit of natural language. But if our theory of natural language syntax was nothing more than CFG, then our theory would fail to predict the fact that certain kinds of CF rules are much more natural than others. For example, consider the rules in (12):

(12) VP $\rightarrow$ P NP
NP $\rightarrow$ PP S

As far as we are aware, no linguist has ever wanted to write rules like these for any human language. However, nothing in the formalism of CFG indicates that there is anything wrong – or even unusual – about such rules.

What is it that we don't like about the rules in (12)? An intuitive answer is that the categories on the left of the rules don't seem appropriate for the sequences on the right. For example, a VP should have a verb in it. This then leads us to consider why we named NP, VP, and PP after the lexical categories N, V, and P. In each case, the phrasal category was named after a lexical category that is an obligatory part of that kind of phrase. At least in the case of NP and VP, all other parts of the phrase may sometimes be missing (e.g. *Dogs bark*).

The lexical category that a phrasal category derives its name from is called the HEAD of the phrase. This notion of head plays a crucial role in all human languages and this fact points out a way in which natural language grammars differ from arbitrary CFGs. A significant number of the phrase structures of natural languages are headed structures, and we will want to incorporate such a notion into our theory in the process of refining our theory of grammar. In the process of doing this, we will also be able to simplify our grammar significantly.

3.6 Subcategorization

The few grammar rules we have so far cover only a small fragment of English. What might not be so obvious, however, is that they also overgenerate – that is, they generate strings that are not well-formed English sentences. Both *denied* and *disappeared* would be listed in the lexicon as members of the category V. This classification is necessary to account for sentences like (13):

(13) a. The defendant denied the accusation.

b. The problem disappeared.

But this would also permit the generation of the ungrammmatical sentences in (14).

(14) a. *The defendant denied.

b. *The teacher disappeared the problem.

Similarly, the verb *handed* must be followed by two NPs, but our rules allow a VP to be expanded in such a way that V can be followed by only one NP, or no NPs at all. That is, our current grammar fails to distinguish among the following:

(15) a. The teacher handed the student a book.

b. *The teacher handed the student.

c. *The teacher handed a book.

d. *The teacher handed.

In order to rule out the ungrammatical examples in (15), we need to distinguish among verbs that cannot be followed by an NP, those that must be followed by one NP, and those that must be followed by two NPs. These are often referred to as INTRANSITIVE, TRANSITIVE, and DITRANSITIVE verbs, respectively. In short, we need to distinguish subcategories of the category V.

One possible approach to this problem is simply to conclude that the traditional category of 'verb' is too coarse-grained for generative grammar, and that it must be replaced by at least three distinct categories, which we can call IV, TV, and DTV. We can then replace our earlier phrase structure rule

VP $\rightarrow$ V (NP) (NP)

with the following three rules:

(16) a. VP $\rightarrow$ IV

b. VP $\rightarrow$ TV NP

c. VP $\rightarrow$ DTV NP NP

◊**Problem 3** This grammar embodies the claim that IVs, TVs, and DTVs are entirely different categories. Hence, it provides no reason to expect that they would have more in common than any other collection of three lexical categories, say, N, P, and D. But these three types of verb do behave alike in a number of ways. For example, they all exhibit agreement with the subject of the sentence; this is discussed in the next section.

a. List as many other properties as you can think of that are shared by intransitive, transitive, and ditransitive verbs.

b. Construct arguments against the analysis in (16) using the properties you found in part a.

3.7 Transitivity and Agreement

Most nouns and verbs in English have both singular and plural forms. In the case of nouns, the distinction between, say *bird* and *birds* indicates whether the word is being used to refer to just one fowl or a multiplicity of them. In the case of verbs, distinctions like the one between *sing* and *sings* indicate whether the verb's subject denotes one or many individuals. In present tense English sentences, the plurality marking on the head noun of the subject NP and that on the verb must be consistent with each other. This is referred to as SUBJECT-VERB AGREEMENT (or sometimes just 'agreement', for short). It is illustrated in (17).

(17) a. The bird sings.

b. Birds sing.

c. *The bird sing.[3]

d. *Birds sings.

[3]There are dialects of English in which this is grammatical, but we will be analyzing the more standard dialect in which this kind of agreement marking is obligatory.

Perhaps the most obvious strategy for dealing with agreement is the one considered in the previous section. That is, we could divide our grammatical categories into smaller categories, distinguishing singular and plural forms. We could then replace the relevant phrase structure rules with more specific ones. In examples like (17), we could distinguish lexical categories of N-SING and N-PLU, as well as IV-SING and IV-PLU. Then we could replace the rule

S $\rightarrow$ NP VP

with two rules:

S $\rightarrow$ NP-SING VP-SING and
S $\rightarrow$ NP-PLU VP-PLU.

Since the marking for number appears on the head noun and head verb, other rules would also have to be changed. Specifically, the rules expanding NP and VP all would have to be divided into pairs of rules expanding NP-SING, NP-PLU, VP-SING, and VP-PLU. Hence, we would need all of the following:

(18) a. NP-SING $\rightarrow$ (D) NOM-SING

b. NP-PLU $\rightarrow$ (D) NOM-PLU

c. NOM-SING $\rightarrow$ NOM-SING PP

d. NOM-PLU $\rightarrow$ NOM-PLU PP

e. NOM-SING $\rightarrow$ N-SING

f. NOM-PLU $\rightarrow$ N-PLU

g. VP-SING $\rightarrow$ IV-SING

h. VP-PLU $\rightarrow$ IV-PLU

This set of rules is cumbersome, and clearly misses linguistically significant generalizations. The rules in this set come in pairs, differing only in whether the category names end in '-SING' or '-PLU'. Nothing in the formalism or in the theory predicts this pairing. The rules would look no less natural if, for example, the rules expanding -PLU categories had their right-hand sides in the reverse order from those expanding -SING categories. But languages exhibiting this sort of variation in word order do not seem to exist.

Things get even messier when we consider transitive and ditransitive verbs. Agreement is required irrespective of whether the verb is intransitive, transitive, or ditransitive. Thus, along with (17), we have (19) and (20).

(19) a. The bird devours the worm.

b. The birds devour the worm.

c. *The bird devour the worm.

d. *The birds devours the worm.

(20) a. The bird gives the worm a tug.

b. The birds give the worm a tug.

c. *The bird give the worm a tug.

d. *The birds gives the worm a tug.

If agreement is to be handled by the rules in (21),

(21) a. S $\rightarrow$ NP-SING VP-SING

b. S $\rightarrow$ NP-PLU VP-PLU

then we will now need to introduce lexical categories TV-SING, TV-PLU, DTV-SING, and DTV-PLU, along with the necessary VP-SING and VP-PLU expansion rules (as well as the two rules in (21)). What are the rules for VP-SING and VP-PLU when the verb is transitive or ditransitive? For simplicity, we will look only at the case of VP-SING with a transitive verb. Since the object of the verb may be either singular or plural, we need two rules:

(22) a. VP-SING $\rightarrow$ TV-SING NP-SING

b. VP-SING $\rightarrow$ TV-SING NP-PLU

Similarly, we need two rules for expanding VP-PLU when the verb is transitive, and four rules each for expanding VP-SING and VP-PLU when the verb is ditransitive (since each object can be either singular or plural). Alternatively, we could make all objects of category NP and introduce the two following two rules:

(23) a. NP $\rightarrow$ NP-SING

b. NP $\rightarrow$ NP-PLU

This would keep the number of VP-SING and VP-PLU rules down to three each (rather than seven each), but it introduces extra noun phrase categories. Either way, the rules are full of undesirable redundancy.

It should be clear, furthermore, that as additional coverage is incorporated – such as adjectives modifying nouns – the redundancies will proliferate. The problem is that we want to be able to talk about nouns and verbs as general classes, but we have now divided nouns into two categories (N-SING and N-PLU) and verbs into six categories (IV-SING, IV-PLU, TV-SING, TV-PLU, DTV-SING, and DTV-PLU). To make agreement work, this multiplication of categories has to be propagated up through at least some of the phrasal categories. The result is a very long and repetitive list of phrase structure rules.

What we need is a way to talk about subclasses of categories, without giving up the commonality of the original categories. That is, we need a formalism that permits us to refer straightforwardly to, e.g. all verbs, all singular verbs, all ditransitive verbs, or all singular ditransitive verbs. In the next chapter, we introduce a device that will permit us to do this.

3.8 Conclusion

In this chapter, we have isolated two ways in which Context-Free Grammars are inadequate as a theory of natural language:

1. CFGs are arbitrary. They fail to predict the 'headedness' that is characteristic of many types of phrase in natural language.
2. CFGs are redundant. Without some way to refer to kinds of categories rather than just individual categories, there is no way to eliminate the massive redundancy that will be required in order to analyze the agreement patterns of natural languages.

For these reasons, we cannot accept CFG alone as a theory of grammar. As we will see in the next few chapters, however, it is possible to retain much of the character of CFG as we seek to remedy its defects.

3.9 Further Reading

Perhaps the best discussions of the basic phrase structures of English are to be found in good descriptive grammars, such as Quirk, et al (1972) or Greenbaum (1996). Important discussions of the notion of 'head' and its role in phrase structure can be found in Chomsky (1970) and Gazdar and Pullum (1981). A detailed taxonomy of the subcategories of English verbs is provided by Levin (1993).

Chapter 4

Analyzing Features of Grammatical Categories

4.1 Introduction

In the last chapter, we saw that there are constraints on what words can go together (what linguists call 'cooccurrence restrictions') that are not efficiently described using the standard formalism of context-free grammar. Some verbs must take an object; others can never take an object; still others (e.g. *put*, *inform*) require both an object and another phrase of a particular kind. These cooccurrence restrictions, as we have seen, cause CFGs to become quite redundant. In addition, different forms of a given verb impose different conditions on what kind of NP can precede them (i.e. on what kind of subject they cooccur with). For example, *walks* requires a third person singular NP as its subject; *walk* requires a plural subject, or else one that is first or second person singular. As we saw in the last chapter, if we try to deal with this complex array of data by dividing the category V into more specific categories, each with its unique cooccurrence restrictions, we end up with a massively redundant grammar that fails to capture linguistically significant generalizations.

In addition to this redundancy, we isolated a second defect of CFGs, namely that they allow rules that are arbitrary. Nothing in the theory of CFG reflects the fact that the phrases of human language usually share certain key properties (nounhood, verbhood, preposition-hood, etc.) with a particular daughter within them – their head. We must somehow modify the theory of CFG to allow us to express the property of headedness.

Our solution to the problem of redundancy is to make grammatical categories decomposable into component parts. CFG as presented so far treats each grammatical category symbol as atomic – that is, without internal structure. Two categories are either identical or different; there is no mechanism for saying that two categories are alike in some ways, but different in others. However, words and phrases in natural languages typically behave alike in certain resepcts, but not others. For example, the two words *deny* and *denies* are alike in requiring an NP object (both being forms of a transitive verb). But they differ in terms of the kind of subject NP they take: *denies* requires a third-person-singular subject like *Kim* or *she*, while *deny* accepts most any NP subject EXCEPT the third-person-singular kind. On the other hand, *denies* and *disappears* both take a singular subject NP, but only the former can cooccur with a following object NP. An adequate formalism needs to be able to characterize the fact that words are organized into classes defined in terms of cross-cutting properties.

To accommodate this intuition, we will develop the view that grammatical categories are not atomic, but rather are complexes with internal structure. This innovation, much like the decomposition of molecules into atoms, or of atoms into subatomic particles, will allow us to talk precisely about how categories are the same in certain respects, while remaining different in others.

4.2 Feature Structures

Informally, we speak of verbs differing in their transitivity. More generally, linguists talk about elements that have different combinatoric potential in terms of differing 'valence'.[1] Likewise, we talk of the number (singular or plural) of a noun, the part of speech of a word (whether it's a noun, verb, etc.), or a verb's form (e.g. whether it is a present participle, an infinitive, etc.). Instead of associating words in the lexicon with a single atomic category, we can treat a lexical category as a complex of grammatical properties. To model such complexes, we use the notion standardly referred to as FEATURE STRUCTURE.

Intuitively, a feature structure is just a specification of grammatical information. Formally, a feature structure is nothing more than a specification of a set of features (which we will write in upper case), each of which is paired with a particular VALUE. Feature structures can be thought of in at least two more or less equivalent ways. For example, they may be conceived of as functions (in the mathematicians' sense of the word), specifying a value for each of a set of features, or else as directed graphs where feature names label arcs that point to appropriately labelled nodes. For grammatical purposes, however, it will be most useful for us to focus on DESCRIPTIONS of feature structures, which we will write in a square bracket notation, as shown in (1):

(1) $$\begin{bmatrix} \text{FEATURE}_1 & \text{VALUE}_1 \\ \text{FEATURE}_2 & \text{VALUE}_2 \\ \quad\ldots & \\ \text{FEATURE}_n & \text{VALUE}_n \end{bmatrix}$$

For example, we might treat the category of the word *bird* in terms of a feature structure that specifies both part of speech and number. We may assume such a category includes appropriate specifications for two appropriately named features: its part of speech (POS) is noun, and its number (NUM) is singular (sg). The lexical entry for *bird*, then, would be a pair consisting of a form and a feature structure description, roughly as shown in (2):[2]

(2) $$\left\langle \text{bird}, \begin{bmatrix} \text{POS} & \text{noun} \\ \text{NUM} & \text{sg} \end{bmatrix} \right\rangle$$

[1] This term, borrowed from chemistry, refers to the capacity to combine with other elements.

[2] Throughout this book, we will describe linguistic forms in terms of standard English orthography. In fact, a lexical entry such as this should contain a phonological description that will play a role in the word's phonological realization, a topic we will not consider in detail here.

Implicit in our use of feature structures is a commitment to developing a theory of what kinds of features go together, what values are appropriate for each particular feature, etc. – that is, a commitment to specifying which feature structure categories are well-formed and which are not. Note that this enterprise is also naturally viewed as providing a theory of what kind of linguistic entities exist in a given domain, and what properties those entities exhibit. Much of our grammar development will be concerned with formulating a natural theory of linguistic generalizations in terms of the constraints that govern the feature structure categories we are led to posit.

One of the first things we will want to do in developing a theory of grammar is to classify linguistic entities in various ways. To this end, it is particularly useful to introduce the notion of TYPE. This concept is really quite simple: if we think of a language as a system of linguistic entities (words, phrases, categories, sounds, and other more abstract entities that we will motivate as we go along), then types are just classes of those entities. We assign entities to these classes on the basis of certain properties that they share. Naturally, the properties we employ in our type classification will be those that we wish to refer to in our descriptions of the entities. Thus, each grammatical type will be associated with particular features and sometimes with particular values for those features.

Let us make this very abstract discussion more concrete by considering the use of feature structures to describe a simple non-linguistic domain. Imagine that we used feature structures to model universities and the people who are associated with them. We'll start from the assumption that the people and the other entities are really 'out there' in the real world. Our first step then in constructing a theory of this part of the world is to develop a model. A simple model will be a set of mathematical entities that we assume to be in correspondence with the real ones. Our theory will be successful to the extent that we can show that the properties that our theory ascribes to our modelling entitites (through stipulation or deduction from the stipulations) also hold of the real world entities they are assumed to be in correspondence with.

The most general kind of entity in the domain at hand should include universities, departments, and individuals (people). We might want to talk about certain properties of these entities, for example their name or telephone number. In this case, we would declare the existence of a general type called *entity* and say that the features NAME and TELEPHONE are appropriate features for all entities (tokens) of this type. So for each university, department or person in this university world, we would hypothesize a distinct feature structure model that we could describe as follows:

(3) a. $\begin{bmatrix} \textit{entity} & \\ \text{NAME} & \text{Stanford University} \\ \text{TELEPHONE} & \text{415-723-2300} \end{bmatrix}$

b. $\begin{bmatrix} \textit{entity} & \\ \text{NAME} & \text{Gerhard Casper} \\ \text{TELEPHONE} & \text{415-723-2481} \end{bmatrix}$

c. $\begin{bmatrix} \textit{entity} & \\ \text{NAME} & \text{Stanford Linguistics} \\ \text{TELEPHONE} & \text{415-723-4284} \end{bmatrix}$

Note that we use type names, written in italics, as labels on the top line within feature structures.

Of course 'entity' is a very general way to classify things like this – our theory would not have progressed very far if it recognized no more specific kinds of things. So in fact, we would want our theory to include the fact that there are different SUBTYPES of the type *entity*. Let's call these new types *university*, *department*, and *individual*. Entities belonging to each of these types have their own special properties. For example individual people have birthdays, but universities and departments don't (or not in the same sense). Similarly, departments have chairs (or 'heads of department'), but neither universities nor individuals do. Finally, only universities have presidents. We can accommodate these facts by declaring each of the relevant features (BIRTHDAY, CHAIR, PRESIDENT) to be appropriate for the right one of our new subtypes. This formal declaration is just a precise way of saying that the members of the relevant subclasses have certain properties that distinguish them from other entitites in the system. The resulting descriptions that we write will be appropriately more specific, e.g. as in (4):

(4) a. $\begin{bmatrix} \textit{university} & \\ \text{NAME} & \text{Stanford University} \\ \text{PRESIDENT} & \text{Gerhard Casper} \\ \text{TELEPHONE} & \text{415-723-2300} \end{bmatrix}$

b. $\begin{bmatrix} \textit{individual} & \\ \text{NAME} & \text{Gerhard Casper} \\ \text{BIRTHDAY} & \text{12-25-1937} \\ \text{TELEPHONE} & \text{415-723-2481} \end{bmatrix}$

c. $\begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Stanford Linguistics} \\ \text{CHAIR} & \text{Stanley Peters} \\ \text{TELEPHONE} & \text{415-723-4284} \end{bmatrix}$

Note that each of these descriptions reflects the hierarchical organization of types. Each type of entity has its own constellation of features – some of them were declared appropriate for the indicated subtype; others were sanctioned by the supertype *entity*. This is a simple illustration of how a hierarchical classification system works. A given feature structure contains only those features that are declared appropriate by one of its types, i.e. by its LEAF type[3] or one of its supertypes. As we will see, a feature structure also inherits any type constraints, i.e. constraints on feature values, that may be associated with its supertypes.

[3]The leaf types are the most 'basic' types in a hierarchy, i.e. the types that have no subtypes.

Articulating a type hierarchy and the feature structures associated with each type is an important component of a theory that uses typed feature structures as models.

In order to talk about merging information from more than one source, we're going to need to introduce a precise method for combining feature structure descriptions. A standard method for doing this is called UNIFICATION.[4] Unification corresponds intuitively to the notion of identity that we have already used, indicated by two occurrences of a single variable (e.g. the category variable 'X' used in our coordination rule). Up to now, we have used these identities only for atomic values. However, once we have introduced feature structures, which may be only partially described in some description, we need a general method for allowing two compatible partial descriptions to amalgamate the information they contain. Two feature structures descriptions can unify only if they are consistent – that is, unless they specify conflicting types or different atomic values for the same feature. The unification of two feature structure descriptions is simply the description obtained by combining all of the information from both of them.

For example, the feature structures in (5) cannot unify

(5) a. $\begin{bmatrix} \textit{university} & \\ \text{NAME} & \text{Stanford University} \end{bmatrix}$

b. $\begin{bmatrix} \textit{university} & \\ \text{NAME} & \text{Harvard University} \end{bmatrix}$

because they differ in the value they assign to the feature NAME. Intuitively, these two descriptions cannot describe the same entity. Similarly, the feature structures in (6) cannot unify

(6) a. $\begin{bmatrix} \textit{individual} \\ \text{TELEPHONE 415-555-4284} \end{bmatrix}$

b. $\begin{bmatrix} \textit{department} \\ \text{TELEPHONE 415-555-4284} \end{bmatrix}$

because they specify incompatible types – *individual* and *department*, and hence cannot describe the same entity. But the feature structure in (7) unifies with any of those in (8a)–(8c).

(7) $\begin{bmatrix} \text{TELEPHONE 555-234-5789} \end{bmatrix}$

(8) a. $\begin{bmatrix} \textit{university} \end{bmatrix}$

b. $\begin{bmatrix} \textit{individual} & \\ \text{NAME} & \text{Xena: Warrior Princess} \end{bmatrix}$

[4] For a detailed discussion of the formal properties of unification and its use in grammatical description, see Shieber (1986).

c. $$\begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Metaphysics} \\ \text{CHAIR} & \text{Alexius Meinong Jr.} \end{bmatrix}$$

In each case, the result of unification is the feature structure containing the information from each of the two structures unified, and nothing more. Thus the unification of (7) and (8b) is (9):

(9) $$\begin{bmatrix} \textit{individual} & \\ \text{NAME} & \text{Xena: Warrior Princess} \\ \text{TELEPHONE} & \text{555-234-5789} \end{bmatrix}$$

We will often use unification to indicate that two different feature structures in fact specify the same value for a given feature. To take a simple example, we might want to indicate that two different departments have the same phone number in the world we are describing. We will indicate this by two identical occurences of a given boxed integer – or 'tag', as illustrated in (10).

(10) $$\begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Metaphysics} \\ \text{TELEPHONE} & \boxed{2} \end{bmatrix} \qquad \begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Philosophy} \\ \text{TELEPHONE} & \boxed{2} \end{bmatrix}$$

And of course we might want to simultaneously indicate a feature identity and the value of the feature, e.g. as in (11):

(11) $$\begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Metaphysics} \\ \text{TELEPHONE} & \boxed{2}\text{800-374-6786} \end{bmatrix} \qquad \begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Philosophy} \\ \text{TELEPHONE} & \boxed{2} \end{bmatrix}$$

Note that it would make no difference if we wrote the phone number after the other occurrence of $\boxed{2}$ in (11). The intended interpretation would be exactly the same. It also makes no difference what order we write the features in. So (11) is equivalent, for example, to the following:

(12) $$\begin{bmatrix} \textit{department} & \\ \text{TELEPHONE} & \boxed{2}\text{800-374-6786} \\ \text{NAME} & \text{Philosophy} \end{bmatrix} \qquad \begin{bmatrix} \textit{department} & \\ \text{NAME} & \text{Metaphysics} \\ \text{TELEPHONE} & \boxed{2} \end{bmatrix}$$

Since the value of the feature TELEPHONE is atomic (i.e. it can bear no feature specifications of its own), the unification of values of this feature is much like the simple identity we expressed in the previous chapter by means of multiple occurrences of a variable (e.g. the variable X in the coordination schema). But in the chapters that follow, we will have occasion to unify feature structure descriptions that are embedded as the value of some feature.

4.3 The Linguistic Application of Feature Structures

4.3.1 Feature Structure Categories

So how does this help us with our linguistic deliberations? Instead of saying that there is just one kind of linguistic entity, which must bear a value for every feature we recognize in our feature structures, we will often want to say that a given entity is of a certain type and that certain features are appropriate only for things of that type. We will use typing in many ways, to make sure that [NUM sg] (or [NUM pl]) can only be specified for certain kinds of words (nouns and pronouns, for example), but not, for instance, for prepositions or adjectives.[5] Likewise, we will eventually want to distinguish auxiliaries (helping verbs like *will* and *have*) from other kinds of verbs in terms of a feature AUX. However, we will not want to say that nouns are all redundantly specified as [AUX –], but rather that the feature AUX just isn't appropriate for nouns. We can use types as a basis for classifying the feature structures we assume and the constraints we use to describe them. In so doing, we can then simply state that particular features only go with certain types of feature structures. This amounts to the beginnings of a linguistic ONTOLOGY – that is, the types tell us what kind of linguistic entities our theory posits, and the features associated with those types tell us what general properties those entitites exhibit. In this way, we also make feature structures compact, i.e. we make them reflect only information that is appropriate for the kind of linguistic entity they are associated with.

The hierarchical organization of linguistic types in our theory is significant. As illustrated above, this enables us to classify feature structures in more subtle ways that will allow intermediate level categories of various sorts. For example, verbs may be classified as intransitive or transitive; and transitive verbs may then be further subclassified as strict transitive (those taking a direct object and nothing else) or ditransitive. A hierarchical type system will let us talk about the properties shared by two distinct types (e.g. strict transitive and ditransitive) by associating a feature or a constraint with their common supertype (transitive).

To start with, let us draw a very intuitive type distinction – between WORD and PHRASE. Our grammar rules (i.e. our phrase structure rules) all specify the properties of phrases; the lexicon provides a theory of words. Continuing our practice of indicating types as labels on feature structures, we will replace (2) above by (13):

(13)
$$\left\langle \text{bird}\ ,\ \begin{bmatrix} \textit{word} & \\ \text{POS} & \text{noun} \\ \text{NUM} & \text{sg} \end{bmatrix} \right\rangle$$

We will now reformulate our grammar in terms of typed feature structures. We do so in two stages: first, we introduce the feature VAL (for 'valence'), without worrying about agreement. Then we present our first feature-based treatment of subject-verb agreement. Both of these analyses will be refined in subsequent chapters.

[5] Many such restrictions are language-particular. For example, adjectives are distinguished according to number (agreeing with the noun they modify) in many languages. Even prepositions exhibit agreement inflection in some languages (e.g. Modern Irish) and need to be classified in similar terms.

4.3.2 Handling Valence with a Feature

We can identify our earlier categories IV, TV, and DTV with the following three feature structures:

(14) $\text{IV} = \begin{bmatrix} \textit{word} \\ \text{POS verb} \\ \text{VAL itr} \end{bmatrix} \quad \text{TV} = \begin{bmatrix} \textit{word} \\ \text{POS verb} \\ \text{VAL tr} \end{bmatrix}$

$\text{DTV} = \begin{bmatrix} \textit{word} \\ \text{POS verb} \\ \text{VAL dtr} \end{bmatrix}$

These three categories all share the type *word* and the feature specification [POS verb]. This is just the combination of types and features that we would naturally identify with the category V. That is, by analyzing categories in terms of types and features, we can distinguish the different valence possibilities among verbs, while still having a general category to refer to that includes all verbs. We do this simply by leaving the value of the VAL feature unspecified. That is,

$\text{V} = \begin{bmatrix} \textit{word} \\ \text{POS verb} \end{bmatrix}$

This method is commonly referred to in linguistics as UNDERSPECIFICATION.

The category VP differs from the category V only with respect to its type assignment. So the feature structure for VP is the following:

(15) $\text{VP} = \begin{bmatrix} \textit{phrase} & \\ \text{POS} & \text{verb} \end{bmatrix}$

And so if we wanted to refer to the class that includes just verbs and verb phrases, we would refer to it as the underspecification in (16):

(16) $\begin{bmatrix} \text{POS verb} \end{bmatrix}$

Similarly, we can analyze the categories N and NP as follows:

(17) $\text{N} = \begin{bmatrix} \textit{word} \\ \text{POS noun} \end{bmatrix}$

$\text{NP} = \begin{bmatrix} \textit{phrase} & \\ \text{POS} & \text{noun} \end{bmatrix}$

Our method of analysis in fact allows underspecification of various kinds. Our goal is to provide compact descriptions for those categories that our grammar will actually need to refer to, what linguists usually call NATURAL CLASSES.

We will continue to use labels like V, N, VP, and NP, but they should now be regarded as abbreviations for the typed feature structures just described. Notice that the feature analysis we have just sketched does not accommodate the category NOM: since there are only two syntactic types, this system does not permit the sort of three-level phrases we posited for NPs in the previous chapter. In the next chapter, we will revise our treatment of the internal structure of categories in such a way that NOM can be distinguished both from N and from NP.

Turning now to the phrase structure rules, we can reformulate our VP rules in terms of our new feature structure categories. Consider the following way of stating the rules:

(18)
$$\begin{bmatrix} \textit{phrase} \\ \text{POS} \ \boxed{1} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} \\ \text{POS} \ \boxed{1} \\ \text{VAL itr} \end{bmatrix}$$

$$\begin{bmatrix} \textit{phrase} \\ \text{POS} \ \boxed{1} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} \\ \text{POS} \ \boxed{1} \\ \text{VAL tr} \end{bmatrix} \ \text{NP}$$

$$\begin{bmatrix} \textit{phrase} \\ \text{POS} \ \boxed{1} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} \\ \text{POS} \ \boxed{1} \\ \text{VAL dtr} \end{bmatrix} \ \text{NP} \ \ \text{NP}$$

The two occurences of $\boxed{1}$ in each of these rules are telling us that the POS value of the mother and that of the first daughter must be unified (which amounts to simple identity of atomic values in this case). Since the rules in (18) were introduced as VP rules, the obvious value to assign to $\boxed{1}$ is 'verb'. But by stating the rules in this underspecified way, we can use them to cover some other structures as well. The first rule, for intransitives, can be used to introduce nouns, which can never be sisters to NPs. This is done simply by instantiating $\boxed{1}$ as 'noun', which will in turn cause the mother to be an NP. To make this work right, we will have to require that lexical nouns, like intransitive verbs, be specified as [VAL itr]. Similarly, the second rule can subsume our PP expansion rule, if $\boxed{1}$ is instantiated as 'prep' and prepositions are lexically marked [VAL tr], thus allowing PPs to be built from a preposition and an NP object.

Another place where we would like to use underspecification in order to collapse two phrase structure rules into one is in the recursive rules introducing PPs, that is, the rules NOM $\rightarrow$ NOM PP and VP $\rightarrow$ VP PP. At this point, we can't do this, because we haven't yet integrated NOM into our system of types and features. But we would like to be able to replace these two rules with something like the following[6]:

(19)
$$\begin{bmatrix} \textit{phrase} \\ \text{POS} \ \boxed{2} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{phrase} \\ \text{POS} \ \boxed{2} \end{bmatrix} \ \text{PP[...]}$$

[6]Note that the choice of a particular tag here is entirely arbitrary – all that is crucial is that there be two occurrences of the same tag.

◊**Problem 1: Other Uses for Rule (19)** The rule for PP modifiers given in (19) interacts with the coordination rule presented in the previous chapter. The coordination rule, as you may recall, is the following:
$X \rightarrow X^{+}$ CONJ X

In particular, the interaction of these rules predicts that a PP can modify a coordinate structure, as in examples like: *Kim walks and reads books without difficulty* or *the poetry and the music on the program.*

a. Draw the tree structure defined by our grammar for each of these examples. Make sure the PP modifies the entire coordinate structure.
b. Draw a distinct tree structure defined by our grammar for each of these same examples. Make sure that the PP does not modify the entire coordinate structure.

The new version of the PP modifier rule also predicts that other categories besides VP and NP can occur with PP modifiers.

c. Using coordination as part of your argument, construct crucial examples demonstrating that S can also cooccur with PP modifiers.
d. Find at least one more category that can be modified by a PP, providing examples to support your claim. [Note: You may consider types of phrases we have not yet discussed, as well as ones that have appeared in our rules.]

The collapsing of phrase structure rules across POS values will be carried still further in the next chapter, when we revise our analysis of the valence feature. First, however, let us examine how features might be employed in the analysis of subject-verb agreement.

4.3.3 Handling Agreement with Features

One device introduced in the last section lends itself well to dealing with the problem of agreement. This is the use of a tag linking two feature values, the effect of which is to force two distinct nodes in a tree admitted by a rule to have identical values for a given feature. We can also employ this device to handle agreement by tagging values for the feature NUM. In the rule expanding S, we could require that the NP and the VP both have the same value of NUM, and we could use the same technique to ensure that the NUM value of a given phrase is identical to the NUM value of the lexical head of that phrase. With these revisions, the rules in question look like (20).

(20) a.
$$S \rightarrow \begin{bmatrix} \textit{phrase} & \\ \text{POS} & \text{noun} \\ \text{NUM} & \boxed{1} \end{bmatrix} \begin{bmatrix} \textit{phrase} & \\ \text{POS} & \text{verb} \\ \text{NUM} & \boxed{1} \end{bmatrix}$$

b. $$\begin{bmatrix} \textit{phrase} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \\ \text{VAL} & \text{itr} \end{bmatrix}$$

c. $$\begin{bmatrix} \textit{phrase} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \\ \text{VAL} & \text{tr} \end{bmatrix} \quad \text{NP}$$

d. $$\begin{bmatrix} \textit{phrase} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \end{bmatrix} \rightarrow \begin{bmatrix} \textit{word} & \\ \text{POS} & \boxed{1} \\ \text{NUM} & \boxed{2} \\ \text{VAL} & \text{dtr} \end{bmatrix} \quad \text{NP} \quad \text{NP}$$

4.4 The Head Feature Principle

The last three rules in (20) require that the mother and one of the daughters bear identical (unified) values for both POS and NUM. In fact, the constituent on the right-hand side that carries the matching feature values is always the head daughter. As we will see in later chapters, there are a number of properties of phrases that are also properties of their lexical heads.[7] In this section, we will develop general mechanisms for capturing the generalization that certain properties are characteristically shared by phrases and their heads.

Rather than stipulating identity of features in an ad hoc manner on both sides of the rules (as in (20)), our analysis will recognize that in a certain type of phrase – a HEADED PHRASE – one daughter is assigned special status as the HEAD DAUGHTER. Once such a notion is incorporated into our theory (thus providing a remedy for the second defect of standard CFGs noted in the last chapter), we can factor out the identity constraint that we need for all the headed phrases, making it a general principle. We will call this generalization the Head Feature Principle (HFP).

Certain rules, such as those in (20), introduce an element that functions as the head of the phrase characterized by the rule. We will call such rules HEADED RULES. To indicate which element introduced in a headed rule is the head daughter, we will label one element on the right hand side of the rule with the letter 'H'. So a headed rule will have the following general form:

(21) $$[\textit{phrase}] \rightarrow \ \ldots \quad \text{H}[\ ] \quad \ldots$$

[7]Alternative formulations might be that some properties of lexical heads are inherited by the phrases they 'project' or that properties of phrases are marked on the lexical heads. While it is often helpful to think of information as propagating up or down through a tree, this is just a metaphor. Our formulation of the generalization in the text avoids attributing directionality of causation in the sharing of properties between phrases and their heads.

Our goal is to formulate a general theory of what features the head daughter shares with its mother in a headed phrase, i.e. what features will always be the same for the element labeled 'H' in a headed rule and the phrase licensed by the rule.

Before proceeding, we need to reflect for a moment about parts of speech. As we noted above, there are certain features that are appropriate for certain parts of speech, but not others. For example, CASE is appropriate only for nouns (in English), while the feature AUX is specifiable only for verbs (to distinguish helping verbs from all others). Likewise, here we will use the features PER(SON) and NUM(BER) only for nouns, verbs and determiners.[8] To guarantee that only the right features go with the right parts of speech, we will treat parts of speech not as atomic values of the POS feature, but rather in terms of a set of types.[9] Then we can declare which features are appropriate for each part of speech type.

We therefore introduce the types *noun, verb, adj, prep, det*, and *conj* for the six lexical categories we have so far considered. We then make all of these subtypes of a type called *part-of-speech* (*pos*). Our grammar must also specify the appropriate values for each feature it employs. We will make the traditional assumption throughout this book that NUM takes either 'sg' (singular) or 'pl' (plural) as its value and that the values of PER are '1st', '2nd', or '3rd'.

Having eliminated the old feature POS, we now introduce a new one called HEAD. HEAD will always take as its value a part of speech, i.e. a feature structure assigned to some subtype of the type *pos*. In this way, HEAD does the work formerly assigned to POS; but it also does more, i.e. it provides a way for us to begin to provide an account of which features are appropriate for which parts of speech.

In making this change, it should be noted, we have also introduced a significant innovation into our theory of feature structures. Previously, all of our features' values were atoms (e.g. 'itr' or 'sg') with no internal structure. By introducing complexes of type *noun*, *verb*, etc. as values of HEAD, we have introduced complex values for features, i.e. feature structures within feature structures. This is a technique that will serve us well in the chapters to come. Moreover, it will be of immediate use in providing us with a simple way to express the relation between a headed phrase and its head daughter. That is, the Head Feature Principle (given below) can be stated simply in terms of the unification of HEAD specifications.

In a similar vein, we can now improve our treatment of agreement by introducing a new feature called AGR, whose value will be a feature structure containing the features NUM and PER. That is, AGR contains just the information that matters for agreement.[10] Since we need to have agreement features specified on phrases as well as on their heads, AGR must be a head feature. Hence, it shows up in feature structures like (22).

---

[8] This analysis will in fact be revised in the next chapter, where these features are used only for nouns and determiners.

[9] We might instead introduce some mechanism for directly stipulating dependencies between values of different features – such as a statement that the existence of a value for AUX implies that the value for POS is 'verb'. (For a theory that incorporates just such a mechanism, see Gazdar et al. (1985).) But such a mechanism is unnecessary, given the availability of types in our theory.

[10] Formally, this will be specified by defining a type *agreement-category* and saying that the features NUM and PER are appropriate only for entities of this type.

(22) $$\begin{bmatrix} \text{HEAD} \begin{bmatrix} \textit{noun} \\ \text{AGR} \begin{bmatrix} \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Now we have a more compact way to say that two elements agree with respect to ALL agreement features: we say that their AGR specifications are unified.

The type hierarchy for the parts of speech introduced in this section is summarized in (23), which also indicates the features declared to be appropriate for each individual type.

(23)

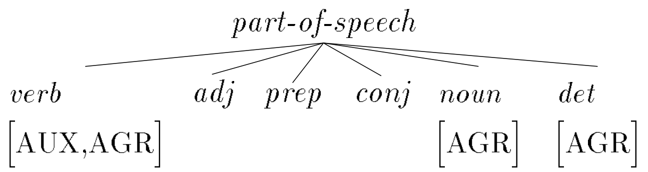

So far, we have done two things: (1) we have identified the head daughter in a headed rule and (2) we have bundled together (within the HEAD value) all the feature specifications that the head daughter must share with its mother. With these two adjustments in place, we are now in a position to simplify the grammar of headed phrases.

First we simplify all the headed rules: they no longer mention anything about number or the part of speech – the information specified within the HEAD value.

(24) a. $$\begin{bmatrix} \textit{phrase} \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{VAL itr} \end{bmatrix}$$

b. $$\begin{bmatrix} \textit{phrase} \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{VAL tr} \end{bmatrix} \ \text{NP}$$

c. $$\begin{bmatrix} \textit{phrase} \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{VAL dtr} \end{bmatrix} \ \text{NP} \quad \text{NP}$$

Recall that the element labeled 'H' in the above rules is the head daughter.

Second, we state the Head Feature Principle as a general constraint governing all trees built by headed rules.

(25) Head Feature Principle (HFP):

In any description of a headed phrase, the HEAD value of the mother and the HEAD value of the head daughter must be unified.

The HFP makes our rules simpler by factoring out those properties common to all headed phrases, and making them conditions that will quite generally be part of the trees defined by our grammar. By formulating the HFP in terms of unification of HEAD values, we allow information specified by the rule, information present on the daughter or the mother, or information required by some other constraint all to be amalgamated, as long as that information is compatible.

4.5 Trees and Phrase Structure

4.5.1 The Formal System: an informal account

At this point, we must address the general question of how rules, lexical entries and principles like the HFP interact to define linguistic structures. Our earlier discussion of the question in chapter 3 requires some revision, now that we have introduced feature structures and types. In the case of simple context-free grammars, there was no particular reason to draw a careful distinction between descriptions and the structures they describe: in CFG, each local subtree (that is, a mother node with its daughters) corresponds in a straightforward fashion to a rule of the grammar. All of the information in that local subtree come directly from the rule. There is no reason to draw a careful distinction between the linguistic objects and the grammar's descriptions of them. But now that rules, lexical entries and principles like the HFP all contribute partial information about linguistic tokens, we must take care to specify how these partial descriptions are amalgamated and how the grammar specifies which expressions are grammatical.

The distinction between descriptions and the structures they describe is fundamental. We use feature structures as our models of linguistic entities. Consider what this meant for the feature structures we used to model universities, departments and individuals. Each such model is assumed to have all the properties relevant to understanding the university system; this includes (for individuals) a name, a birthday, a telephone number (let's assume), and so forth. The objects we take as models are thus complete in relevant respects.[11] Contrast this with descriptions of university individuals. These come in various degrees of completeness. A description may be partial in not specifying values for every feature, in specifying only part of the (complex) value of a feature, in failing to specify a type, or in specifying nothing at all. A complete description of some entity will presumably be satisfied by only one thing – the entity in question. An empty description is satisfied by all the entities in the modelling domain. Any nonempty partial description is satisfied by some things in the modelling domain, and not by others.

Our theory of language works the same way. We use feature structures and trees to model expressions like words and phrases; we also use feature structures to model other things like grammatical categories. Since the structures we use are models, they too are complete with respect to all linguistically relevant properties. So a given word (token), say a pronoun, is a feature structure that has a determinate value for number – it is either singular or plural. Likewise, it is determinately 1st, 2nd or 3rd person. The only kind of thing that could be indeterminate – underspecified for this kind of information – is a feature structure description. Grammars are linguistic descriptions; feature structures and trees are the kinds of things that these descriptions describe. The grammar is successful to the extent that it can be shown that its models have properties that are in correspondence with our observations about how the language ('out there' in the world, in society, or in people's heads) really is.

[11] Of course what distinguishes a model from the thing it is a model of are discrepancies with respect to certain IRRELEVANT properties. Our models of university individuals should omit lots of irrelevant properties that all such individuals presumably have, ranging from hair color to grandmothers' middle names to disposition with respect to Indian food.

As we start to get more precise about our grammar, it will be important not to confuse linguistic structures and their descriptions. With this goal in mind, let us introduce a notational distinction. We will continue to use square bracket matrices (or ATTRIBUTE VALUE MATRICES) for feature structure descriptions, and whenever we want to make reference to a complete feature structure model, we will use box diagrams. Thus the lexical entry for the noun *fish* will include a category description like (26).

(26)
$$\begin{bmatrix} \textit{word} & \\ \text{HEAD} & \begin{bmatrix} \textit{noun} \\ \text{AGR} \begin{bmatrix} \text{PER 3rd} \end{bmatrix} \end{bmatrix} \\ \text{VAL} & \text{itr} \end{bmatrix}$$

And this word description will have two feature structure models, corresponding to the category of the singular word *fish* and that of the plural word *fish*, which we write as in (27):[12]

(27) a.
$$\boxed{\begin{array}{ll} \textit{word} & \\ \text{HEAD} & \boxed{\begin{array}{ll} \textit{noun} & \\ \text{AGR} & \boxed{\begin{array}{ll} \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \end{array}} \\ \text{VAL} & \text{itr} \end{array}} \end{array}}$$

b.
$$\boxed{\begin{array}{ll} \textit{word} & \\ \text{HEAD} & \boxed{\begin{array}{ll} \textit{noun} & \\ \text{AGR} & \boxed{\begin{array}{ll} \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{array}} \\ \text{VAL} & \text{itr} \end{array}} \end{array}}$$

It is the goal of a grammar to enumerate all these linguistic entities (structures assumed to be in correspondence to bits of reality) and to correctly describe their grammatical properties. We will provide an informal sketch of how this can be done here, postponing a more formal presentation until the end of chapter 6.

The basic idea is relatively simple. First we specify how lexical entries give rise to descriptions of word structures. Second, we characterize how descriptions of smaller structures (words or phrases) get amalgamated into descriptions of larger structures, including sentences. Then we describe the relation between a structural description and the set of phrase structure models that satisfy that description. We will continue to use trees throughout. Thus words and phrases will be uniformly characterized in terms of trees, as they are in CFG. We will refer to the entities described by our grammar as word structures and phrase structures.

Most of this book will be concerned with providing descriptions of word structures and phrase structures. We will refer to such descriptions as STRUCTURAL DESCRIPTIONS – SDs, for short. Descriptions of word structures will be called LEXICAL SDs, and descriptions of phrase structures will be called PHRASAL SDs.

[12] Actually, once our grammar is more fully developed, there will be more than two models for this description, each exhibiting a unique determination for such properties as mass vs. count, nominative vs. accusative case, and so forth.

A lexical SD is a tree diagram with two nodes: the lower node is a form that appears in the lexicon, and the upper node is the feature structure description associated with that form in its lexical entry.[13] Put more formally:

(28) Lexical Tree:

$$\begin{array}{c} \mathrm{C} \\ | \\ \omega \end{array}$$

is a well-formed lexical structural description just in case $\langle \omega, C \rangle$ is a lexical entry.

Note that the mother in such a description typically specifies only part of the information that any given word exhibits.

For phrases, the basic idea is again simple, but there are a number of conditions that have to be stated precisely. The simple idea is this: if we have already built up some SDs with our grammar, say,

(29) $\delta_1 \quad \ldots \quad \delta_n$

$\triangle \qquad\quad \triangle$

then we can use a grammar rule to construct a phrasal SD using these as daughters, just in case a number of conditions are met.

What are these conditions?

(30) a. First of all, each δ_i in (29) must be compatible with the corresponding daughter in the rule. This is just common sense, if something is built up as a VP description, you can't use a rule that says PP $\rightarrow$ P NP to construct anything from that description directly.

b. Second, the mother in the new SD needs to include all the information that the rule says is true of its mother node. Again, this is just common sense: the information in the SD should include what the rule says it should include.

c. Third, the rule might add information, 'unifying in' some feature information that was not already in δ_i. We will see examples of NP descriptions, for instance, that are unspecified for the feature CASE until they combine with some other element to build a larger tree.

d. And finally, the new tree description has to satisfy all the principles of our theory – for example, any SD built in accordance with a headed rule must satisfy the Head Feature Principle.

If these conditions are all satisfied, then the SDs in (29) can be fit into a phrasal SD like the following:

[13] We are ignoring compound words throughout.

(31)

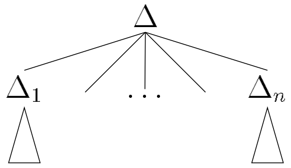

Notice that each Δ_i contains all the information in the corresponding δ_i. But Δ_i can also contain information that is not in δ_i. This can happen through the operation described in (30c or d) above. Indeed, in most cases, the interaction of the rules, the lexical entries, and the principles will result in nodes in the new description that have more information in them than what comes from any one of these sources alone.

Our lexicon, rules, and principles interact to generate an infinite set of phrasal SDs. We define the phrase structures of the language indirectly in terms of these descriptions:

(32) A phrase structure S satisfies a structural description D just in case:

(i) S has exactly the same configuration as D and

(ii) the feature structures labelling the nodes of S satisfy all the constraints specified in D.

4.5.2 An Example

Consider the sentence *Fish swim.* Let's suppose that the lexical entry for *fish* is underspecified for number, as shown above (repeated in (33)), and that the lexical entry for the plural form *swim* is underspecified for person as shown in (33b):

(33) a.

$$\left\langle \text{fish}\ ,\ \begin{bmatrix} \textit{word} \\ \text{HEAD} & \begin{bmatrix} \textit{noun} \\ \text{AGR} & \begin{bmatrix} \text{PER 3rd} \end{bmatrix} \end{bmatrix} \\ \text{VAL} & \text{itr} \end{bmatrix} \right\rangle$$

b.

$$\left\langle \text{swim}\ ,\ \begin{bmatrix} \textit{word} \\ \text{HEAD} & \begin{bmatrix} \textit{verb} \\ \text{AGR} & \begin{bmatrix} \text{NUM pl} \end{bmatrix} \end{bmatrix} \\ \text{VAL} & \text{itr} \end{bmatrix} \right\rangle$$

Given these two lexical entries, the following are both legal lexical SDs, according to our theory:

(34) a.
$$\begin{array}{c}\begin{bmatrix} \textit{word} & \\ \text{HEAD} & \begin{bmatrix}\textit{noun} \\ \text{AGR}\begin{bmatrix}\text{PER 3rd}\end{bmatrix}\end{bmatrix} \\ \text{VAL} & \text{itr}\end{bmatrix} \\ | \\ \text{fish}\end{array}$$

b.
$$\begin{array}{c}\begin{bmatrix} \textit{word} & \\ \text{HEAD} & \begin{bmatrix}\textit{verb} \\ \text{AGR}\begin{bmatrix}\text{NUM pl}\end{bmatrix}\end{bmatrix} \\ \text{VAL} & \text{itr}\end{bmatrix} \\ | \\ \text{swim}\end{array}$$

These SDs can now be embedded within bigger SDs sanctioned by the rule in (24ab), as illustrated in (35a,b).

(35) a.
$$\begin{array}{c}\begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \boxed{1}\begin{bmatrix}\textit{noun} \\ \text{AGR}\begin{bmatrix}\text{PER 3rd}\end{bmatrix}\end{bmatrix}\end{bmatrix} \\ | \\ \begin{bmatrix} \textit{word} & \\ \text{HEAD} & \boxed{1} \\ \text{VAL} & \text{itr}\end{bmatrix} \\ | \\ \text{fish}\end{array}$$

b.
$$\begin{array}{c}\begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \boxed{2}\begin{bmatrix}\textit{verb} \\ \text{AGR}\begin{bmatrix}\text{NUM pl}\end{bmatrix}\end{bmatrix}\end{bmatrix} \\ | \\ \begin{bmatrix} \textit{word} & \\ \text{HEAD} & \boxed{2} \\ \text{VAL} & \text{itr}\end{bmatrix} \\ | \\ \text{swim}\end{array}$$

Both of these obey the HFP, as the values of HEAD (and hence the values for all features within HEAD, e.g. AGR) on the mother match those on the head daughter. Note further that the AGR values of the two phrases are distinct, but compatible.

And finally, assuming we enforce agreement as before, i.e. by constraining our S rule as shown in (36), then we will build sentence descriptions like the one shown in (37).

(36)
$$\text{S} \rightarrow \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \begin{bmatrix}\textit{noun} \\ \text{AGR}\ \boxed{1}\end{bmatrix}\end{bmatrix} \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \begin{bmatrix}\textit{verb} \\ \text{AGR}\ \boxed{1}\end{bmatrix}\end{bmatrix}$$

(37)

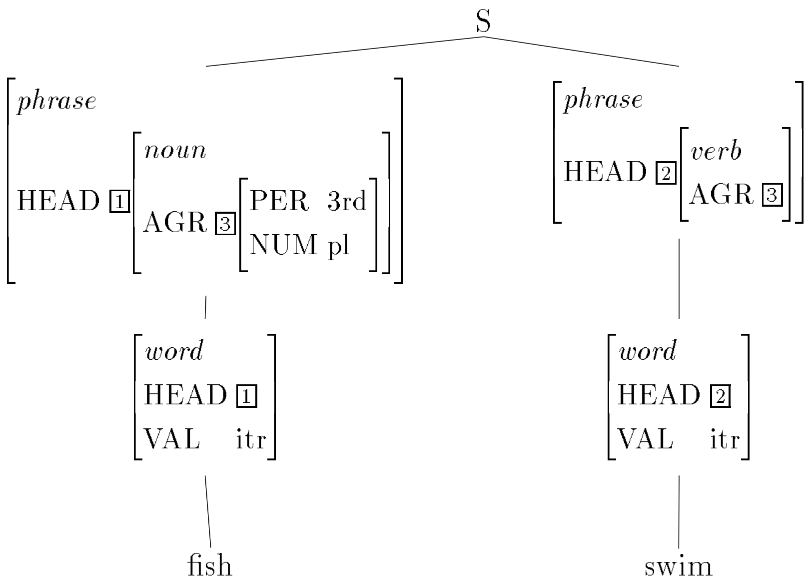

The symbol 'S' here is of course just an abbreviation for a feature structure description of a sort yet to be determined. It is at this level that the real effect of unification is seen. The AGR values of the two daughters of S in (37) are unified, as they must be because (i) the values for these features are unified in the rule (36) and (ii) our theory therefore requires that these values be unified in any tree constructed in accordance with that rule. The unified AGR value contains the two different pieces of compatible information, one coming from each of the phrase's daughters. Crucially, this is not possible in an SD like the following, where the NP and VP daughters contain conflicting AGR information that cannot be unified:

(38)

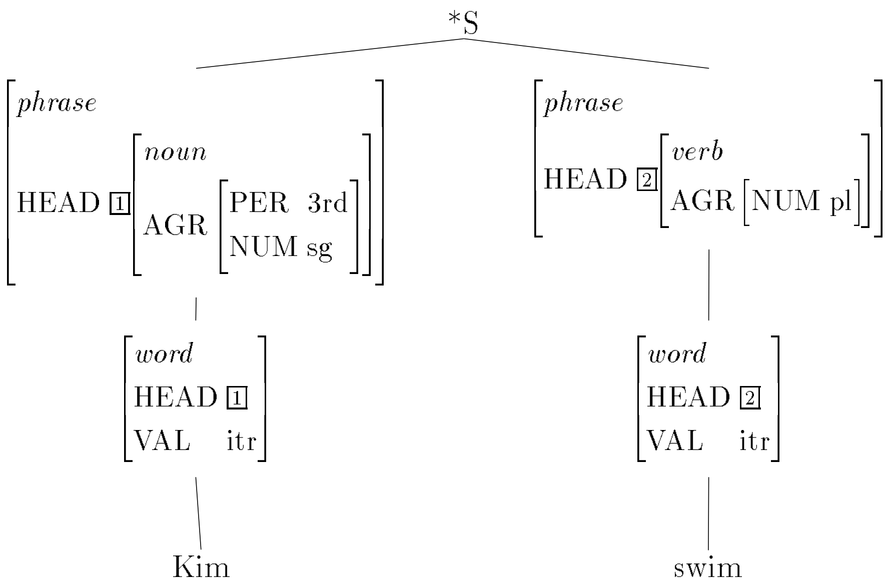

In this way, the various constructs of our theory work together to rule out certain sentences as inconsistent with the simultaneous satisfaction of all relevant constraints.

◊Problem 2: Applying the Rules

a. Explain how the grammar just illustrated rules out ungrammatical examples like **They sings.*

b. Formulate a feature structure version of the following rule (note that the head daughter is a word): NP $\rightarrow$ (D) N
Make sure that your rule is compatible with the grammar just illustrated and that it specifies all identities of feature values required to answer the remaining parts of this problem.

c. Give precise lexical entries for the determiners ([HEAD *det*]) *a* and *many*.

d. Demonstrate precisely how your answer to b. and c. work together to account for the following data:

(i) *a birds

(ii) *many bird

e. Explain precisely how your analysis of NPs interacts with the rule expanding S to explain the following examples.

(iii) *The bird sing.

(iv) *Many birds sings.

4.5.3 Completing the Picture

We can now ascertain precisely which word and phrase structures satisfy the descriptions just illustrated. The lexical SD in (34a), for example, is satisfied by both of the word structures in (39)

(39) a.

$$\boxed{\begin{array}{ll} \textit{word} & \\ \text{HEAD} & \boxed{\begin{array}{ll} \textit{noun} & \\ \text{AGR} & \boxed{\begin{array}{ll} \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \end{array}} \\ \text{VAL} & \text{itr} \end{array}} \end{array}}$$

fish

b.

$$\boxed{\begin{array}{ll} \textit{word} & \\ \text{HEAD} & \boxed{\begin{array}{ll} \textit{noun} & \\ \text{AGR} & \boxed{\begin{array}{ll} \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{array}} \\ \text{VAL} & \text{itr} \end{array}} \end{array}}$$

fish

But only (39b) fits into the larger phrase structure in (40), which satisfies the sentence description in (37).

(40)

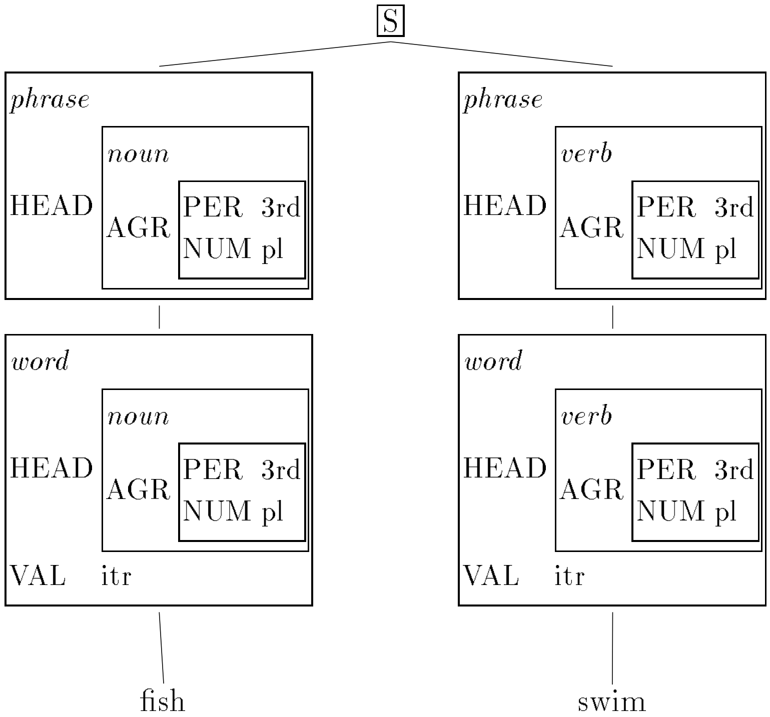

In short, the theory we have begun to sketch can provide a precise MODEL-THEORETIC account of which expressions of a language are grammatical, which potential expressions are ungrammatical, which expressions are ambiguous, and so forth. The account proceeds by expressing generalizations in terms of a CFG-like conception of grammar rule, underspecified descriptions, and general constraints like the Head Feature Principle.

Somewhere along the line, of course, an adequate grammar will need to specify the well-formed 'stand alone' utterances of the language, i.e. the phrases that can be used in isolation to express a complete message are those whose mother is of the category S. But we are not quite ready yet to define the notion 'S'. This will emerge clearly only after we consider the feature structures of phrases a bit more carefully, as we do in the next chapter.

4.6 Conclusion

The introduction of features has given us a formal mechanism for talking about ways in which sets of words (and phrases) behave alike, while still recognizing that they behave differently in other respects. By allowing embedded feature structures, underspecifying categories and formulating general constraints stating unifications that must hold in well-formed trees, we have been able to make our phrase structure rules more general and reduce their number. This in turn has led us to revise slightly the way we conceive of the relationship between the grammar rules and structural descriptions.

The theory we are developing is still closely related to standard CFG, yet it is somewhat more abstract. We no longer think of our phrase structure rules as specifying all the information labelling the nodes in trees. Rather, the rules, the lexicon, and some general principles – of which the HFP is the first example – all place certain constraints on trees, and any imaginable tree is well-formed so long as it conforms to these constraints. In this way, our grammar continues to be constraint-based, with the rules, lexical entries and general principles all providing partial constraints on well-formed structural descriptions, and hence on the structures of the language.

But the changes introduced in this chapter are not yet sufficient. They still leave us with three different rules that have too much in common. Moreover, as we will see in the next chapter, we have simplified the facts of agreement too much. And, as we mentioned above, our new feature-based system of rules doesn't allow us to reconstruct the category NOM. These problems will be dealt with in the next chapter, where we enrich our conception of features and of the lexicon, in order to allow still further simplification of the phrase structure rules.

4.7 Further Reading

Chomsky (1965) provides one of the earliest explicit discussions of syntactic features in generative grammar. [Harman? Yngve?] Early uses of complex feature structures (that is, features with feature structures as their values) can be found in Kay (1979), Bresnan (1982), and Gazdar et al. (1985). A detailed formal treatment of typed feature structures was developed by Carpenter (1992). The mathematics of tree structures is discussed by Partee et al. (1990).

Chapter 5

Complex Feature Values

5.1 Introduction

By analyzing grammatical categories into features, we were able to generalize some of our phrase structure rules and eliminate others. This is not only a more compact way of representing syntactic information, it is also a way to systematically encode the fact that phrases of different types exhibit structural parallelisms. In addition, it predicts that different phrase types will generally be similarly structured. In particular, our rules in the previous chapter suggest that lexical head daughters in English uniformly occur at the left edge of their phrases[1].

Of course, while VPs and PPs are consistently head-initial, it's not completely clear that NPs exhibit the same head-first pattern. For example, NPs may take determiners and adjectives before the lexical head, as in *the noisy dogs*. Recall that in our earlier treatment, using NOM, we could say that NOM always started with its lexical head. But this is not true of NP. Hence, that analysis would allow us to maintain the generalization that the lexical head daughter is always the leftmost element under its mother node. It seems desirable, then, to find a way of reconstructing the NOM analysis within our new, feature-based theory.

Another motivation for revising our current analysis is that our rules are still not maximally general. We have three distinct rules introducing lexical heads, one for each of the three valences. We would like to find a way of consolidating these. Moreover, these three valences are far from the only possible environments lexical heads may require. Consider the examples in (1).

(1) a. Pat relies on Kim.

b. *Pat relies.

c. The child put the toy on the table.

d. *The child put the toy.

e. The teacher became angry with the students.

[1]This is not true in some other languages, e.g. Japanese.

f. *The teacher became.

g. The jury believed the witness lied.

Examples (1a,b) show that some verbs require following PPs; (1c,d) show that some verbs must be followed by both an NP and a PP; (1e,f) show a verb that may be followed by a kind of phrase we have not yet discussed, called an adjective phrase (AP); and (1g) shows a verb that may be followed by an S. We say only that *became* MAY be followed by an AP and *believed* MAY be followed by an S because they can also appear in sentences like *Pat became an astronaut* and *Pat believed the story*, in which they are followed by NPs. In fact, it is extremely common for verbs to be able to appear in multiple environments. Similarly, (2) shows that *ate*, like many other English verbs, can be used either transitively or intransitively.

(2) The guests ate (the cheese).

Facts like these show that the number of values of VAL must be far greater than three. This in itself would not be problematic, except that our current formalism, as developed in the previous chapter, requires a separate phrase structure rule for each value of VAL. This is an unwanted redundancy – the lexical distinctions are being encoded twice – once in the phrase structure rules and once in the (many) new values of VAL that would be required.

5.2 Complements

Intuitively, we would like to have our rule simply say that a phrase (a VP, in the cases above) may consist of a lexical head (a V, in these cases) followed by whatever other phrases the lexical head permits. We could then relegate to the lexicon the specification of what the following environment of a word is. In this section, we develop a way to do just this. It involves enriching our conception of features in a way somewhat analogous to what we did with grammatical categories in the previous chapter.

Before we begin the discussion of this analysis, however, let us consider briefly the status of the kinds of cooccurrence restrictions we have been talking about. For this discussion (and much else that follows), it will be convenient to have a term for the elements that characteristically cooccur with a lexical head – that is, for things like the phrases that occur after the verbs in (1). The term COMPLEMENT (of the head) is more or less standard and will be used throughout this text.

It has sometimes been argued that the number and type of complements a verb takes is fully determined by its meaning. For example, the verb *disappear* is used to describe events involving a single entity (expressed by its subject); *deny*'s semantics involves events with two participants, one typically human and the other a proposition; and an event described by *hand* must include three participants: the person who does the handing, the thing handed, and the recipient of the transaction. Correspondingly, *disappear* takes no complements, only a subject; *deny* takes a subject and a complement, which may be either an NP (as in *The defendant denied the charges*) or an S (as in *The defendant denied he was guilty*); and *hand* takes a subject and two NP complements (or one NP and one PP complement).

It is undeniable that the semantics of a verb is intimately related to its valence. There is, however, a certain amount of syntactic arbitrariness to it, as well. For example, the words *eat, dine*, and *devour* all denote activities necessarily involving both a consumer of food and the food itself. Hence, one might expect that all three would be simple transitives, requiring a subject and an NP complement (that is, a direct object). But this expectation would be wrong: *dine* is intransitive, *devour* is obligatorily transitive, and (as noted above), *eat* can be used intransitively or transitively.

(3) a. The guests devoured the meal.

b. *The guests devoured.

c. *The guests dined the meal.

d. The guests dined.

e. The guests ate the meal.

f. The guests ate.

Thus, while recognizing the links between meaning and valence, we will continue to specify valence syntactically. We will say more about the connection between meaning and valence – and more generally about the syntax-semantics interface – in later chapters.

5.3 The COMPS Feature

In our current grammar, the lexical entry for a verb like *deny* would specify that it is [VAL tr]. This ensures that it can only appear in lexical structures specified as [VAL tr], and such nodes are licensed only by the rule of our grammar that introduces an immediately following NP. Hence, *deny* has to be followed by an NP.[2]

In effect, the function of the value of the VAL feature in this system is as a pointer to the relevant phrase structure rule for licensing the verb in question. In fact, the values of VAL are redundantly in one-to-one correspondence with the phrase structure rules that match up the lexical head daughters with the appropriate kinds of complements. It would be desirable to eliminate this redundancy from our grammar.

An alternative approach to complement selection is to use features directly in licensing – that is, to have a feature whose value specifies what the complements must be. This intuitive idea is what we will now make explicit. The first step is to recall that, in the last chapter, we allowed some features (e.g. HEAD, AGR) to take values that are feature structures themselves. If we replace VAL with such a feature, we can allow its value to state directly what the word's complement must be. So we propose to replace VAL with a feature-valued feature, which we call COMPLEMENTS (COMPS). The value of COMPS for *deny* can simply be NP, i.e. the structural description (SD) in (4):

[2]As noted earlier, we have not dealt with the other possible environment for *deny*, namely the one where it is followed by a clause. We ignore this problem for the moment, but the analysis developed in this chapter provides a way of dealing with it.

(4) $\begin{bmatrix} \textit{phrase} \\ \text{HEAD } \textit{noun} \end{bmatrix}$

Similarly, we can indicate that a verb takes another type of complement: *rely, become*, and *believe*, for example, can take COMPS values of PP, AP, and S, respectively.[3] Optional complements, such as the object of *eat* can be indicated using parentheses; that is, the lexical entry for *eat* can specify [COMPS (NP)]. Likewise, we can indicate alternative choices for complements using the vertical bar notation introduced in the discussion of regular expressions in chapter 2. So the entry for *deny* or *believe* include the specification: [COMPS NP | S].

Of course there is a problem with this proposal: it does not cover verbs like *hand* and *put* that require more than one complement. But it's not hard to invent a straightforward way of modifying the COMPS analysis to let it encompass multiple complements. Instead of treating the value of COMPS as a single feature structure description, we will let it be a LIST of such descriptions. Intuitively, the list specifies a sequence of categories corresponding to the complements that the word combines with. So, for example, the COMPS values for *deny, become*, and *eat* will be lists of length one. For *hand*, the COMPS value will be a list of length two, namely ⟨NP,NP⟩.[4] For verbs taking no complements, like *disappear*, the value of COMPS will be ⟨ ⟩ (a list of length zero). We interpret this as meaning that trees containing the verb in question will be well-formed only if the sisters of the V-node are compatible with the categories specified on the list. For example, *rely* will only be allowed in trees where the VP dominates a V and a PP.

Now we can collapse all of the different rules expanding a phrase into a lexical head (H) and other material. We can just say:

(5) Head-Complement Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{COMPS } \langle\ \rangle \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{COMPS } \boxed{1} \end{bmatrix} \quad \boxed{1}$$

Note that here we are using the tag $\boxed{1}$ to designate neither an atomic value nor a feature structure, but rather a list of feature structures. Thus if a word is lexically specified as [COMPS ⟨AP⟩], it must appear with exactly one AP complement; if it is [COMPS ⟨NP,NP⟩], it must cooccur with exactly two NP complements, and so forth. Also, if the head is lexically specified as [COMPS ⟨ ⟩], then it must appear without any complements at all. Finally, the mother of any tree structure sanctioned by (5), which we will term a HEAD-COMPLEMENT PHRASE, must be specified as [COMPS ⟨ ⟩], because that mother must include all information specified on the left-hand side of the rule.[5]

[3] We have not yet said anything about how S is analyzed in terms of feature structures. Later in this chapter, however, we will present a treatment of S as an abbreviation for a feature structure, just like other category names.

[4] We use angle brackets to designate lists.

[5] Note that by underspecifying the complements introduced by this rule – not even requiring them to be phrases, for example, we are implicitly leaving open the possibility that some complements will be nonphrasal. This will become important in the analysis of negation presented in chapter 13.

In short, the COMPS list of a lexical entry specifies a word's cooccurrence requirements; and the COMPS list of a phrasal node registers those cooccurrence requirements imposed by the head that are not discharged in the subtree dominated by that node. So, in particular, if a V has sisters that match all the feature structures in its COMPS value, then the VP that it heads will have the empty list as its COMPS value and hence cannot combine with further complements. In fact, because the head introduced in rule (5) is of type *word*, all of the complements will have to be realized as sisters to the lexical head.[6]

If you think in terms of building the tree bottom-up, starting with the verb as head, then the verb has certain demands that have to be satisfied before a complete, or 'saturated', constituent is formed. Under this conception, the complements can be thought of as being 'cancelled off' of the head daughter's COMPS list in the process of building a headed phrase. We illustrate with the VP *put flowers in a vase*: the verb *put* requires both a direct object NP and a PP complement, so its COMPS value is ⟨NP,PP⟩. The requisite NP and PP will both be sisters to the V, as in (6), as they combine to form a VP, i.e. a verbal phrase whose complement requirements have been fulfilled.

(6)

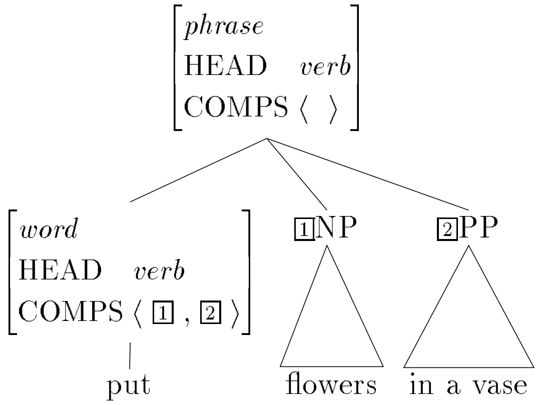

As is evident from this example, we assume that the elements in the value of COMPS occur in the same order as they appear in the sentence. We will continue to make this assumption, though a more sophisticated treatment of linear ordering of phrases in sentences may ultimately be necessary.

Notice, by the way, that although we have motivated our treatment of complements entirely in terms of verbs and verb phrases, we have formulated our analysis to be more general. In particular, our grammar of head-complement structures allows adjectives, nouns, and prepositions to take complements of various types. The following examples suggest that, like verbs, these kinds of words exhibit a range of valence possibilities.

(7) Adjectives

a. The children are happy (with the ice cream).

[6]This appears well-motivated for English, but our general theory allows us to write a head-complement rule for some other language that does not require a lexical head daughter. Such grammatical variations might be regarded as 'parameters' that are set differently in particular languages. That is, it may be that all languages manifest the head-complement rule, but there are minor differences in the way languages incorporate that rule into their grammar. The order of the head and the complements is another such parameter of variation.

b. The children are fond of ice cream.

c. *The children are fond.

d. The children are happy (that they have ice cream).

e. *The children are fond that they have ice cream.

f. *The children are fond.

(8) Nouns

a. A magazine (about crime) appeared on the newsstands.

b. **Newsweek* about crime appeared on the newsstands.

c. *Newsweek* appeared on the newsstands.

d. The report (that crime was declining) surprised many people.

e. *The book that crime was declining surprised many people.

d. The book surprised many people.

(9) Prepositions

a. The storm arrived after the picnic.

b. The storm arrived after we ate lunch.

c. The storm arrived during the picnic.

d. *The storm arrived during we ate lunch.

e. *The storm arrived while the picnic.

f. The storm arrived while we ate lunch.

♢**Problem 1: Valence Variations** In this problem, you may use NP, VP, etc. as abbreviations for the feature structures on COMPS lists.

a. What does the grammaticality of sentences like *Kim put the book here/there* suggest about the COMPS and HEAD values of the words *here* and *there*?

b. What is the COMPS value for the adjective *fond*?

c. Assume that motion verbs like *jump*, *move*, etc. take an optional PP complement. Given that, what do the following examples tell you about the COMPS values of the lexical entries of the prepositions *out*, *from* and *of*:

(i) Kim jumped out of the bushes.

(ii) Bo jumped out from the bushes.

(iii) Lee moved from under the bushes.

(iv) Leslie jumped out from under the bushes.

(v) Dana jumped from the bushes.

(vi) Chris ran out the door.

(vii) *Kim jumped out of from the bushes.

(viii) Kim jumped out.

(ix) *Kim jumped from.

d. Based on the following data, sketch the lexical entries for the words *grow* (in the 'become' sense, not the 'cultivate' sense), *seem*, *happy*, and *close*.

(i) They seemed happy (to me).

(ii) Lee seemed an excellent choice (to me).

(iii) They grew happy.

(iv) *They grew a monster (to me).

(v) *They grew happy to me.

(vi) They grew close to me.

(vii) They seemed close to me to Sandy.

NB: APs have an internal structure that is analogous to that of VPs. Though no adjectives select NP complements (in English), there are some adjectives that select PP complements (e.g. *to me***), and some that do not.**

5.4 Specifiers

Cooccurrence restrictions are not limited to complements. As we have noted in earlier chapters, certain verb forms appear with only certain types of subjects. That is, in the present tense, English subjects and verbs must agree in number. Likewise, certain determiners cooccur only with nouns of a particular number.

(10) a. This dog barked.

b. *This dogs barked.

c. *These dog barked.

d. These dogs barked.

Moreover, some determiners are restricted to occur only with 'mass' nouns (e.g. *furniture*), and others only with 'count' nouns (e.g. *chair*):

(11) a. Much furniture was broken.

b. *A furniture was broken.

c. *Much chair was broken.

d. A chair was broken.

We can handle such cooccurrence restrictions in much the same way that we dealt with the requirements that heads impose on their complements. First, we introduce the term SPECIFIER to refer to both subjects and determiners. We then introduce the feature SPECIFIER (SPR), where we can state the cooccurrence restrictions between heads and the specifiers they select. On analogy with COMPS, we make the value of SPR a list. This decision may strike some readers as odd, since sentences only have a single subject and NPs never have more than one determiner. But making SPR list-valued provides a uniform way of formulating of the idea that a particular valence requirement is unfulfilled (the valence feature – that is, COMPS or SPR – has a nonempty value) or else is satisfied (the value of the valence feature is the empty list).

The category NOM now becomes definable in terms of the following feature structure descriptions:

(12)
$$\text{NOM} = \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \textit{noun} \\ \text{COMPS} & \langle\ \rangle \\ \text{SPR} & \langle \begin{bmatrix}\text{HEAD } \textit{det}\end{bmatrix} \rangle \end{bmatrix}$$

Notice the similarity between (12) and (what is now) the feature specification for VP:

(13)
$$\text{VP} = \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \textit{verb} \\ \text{COMPS} & \langle\ \rangle \\ \text{SPR} & \langle \text{ NP } \rangle \end{bmatrix}$$

Both (12) and (13) are of type *phrase* with empty COMPS lists and a single element in their SPR lists. Both are intermediate between lexical categories (type *word*, with possibly nonempty COMPS lists) and 'saturated' phrases – that is, phrases with both their COMPS and SPR lists empty.

Similarly, we can now introduce a verbal category that is analogous in all relevant respects to the saturated category NP. It is just the feature structure analogue of the familiar category S.

(14)
$$\text{NP} \;=\; \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \textit{noun} \\ \text{COMPS} & \langle\ \rangle \\ \text{SPR} & \langle\ \rangle \end{bmatrix} \qquad \text{S} \;=\; \begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \textit{verb} \\ \text{COMPS} & \langle\ \rangle \\ \text{SPR} & \langle\ \rangle \end{bmatrix}$$

Because NP and S now have a parallel formulation in terms of feature structures and parallel constituent structures, we may collapse our old rules for expanding these categories (given in (15)) into a single rule, shown in (16), that sanctions all HEAD-SPECIFIER PHRASEs.

(15) a. S $\rightarrow$ NP VP

b. NP $\rightarrow$ (D) NOM.

(16) Head-Specifier Rule:
$$\begin{bmatrix} \textit{phrase} & \\ \text{COMPS} & \langle\,\rangle \\ \text{SPR} & \langle\,\rangle \end{bmatrix} \;\rightarrow\; \boxed{2} \quad \text{H}\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \langle\,\boxed{2}\,\rangle \end{bmatrix}$$

Having consolidated the rules in this way, we will need to elaborate our treatment of the SPR feature so as to account for various cooccurrence restrictions between heads and specifiers. These include the fact that the specifiers of sentences (i.e. subjects) are NPs, whereas the specifiers of NPs are Ds. They also include facts like those in (10)–(11), and the agreement between subjects and verbs, which will be dealt with in sections 6 and 7 below.

The first of these is relatively simple: the value for SPR in the lexical entries for nouns is the list $\langle$ D $\rangle$, and in the lexical entries for verbs, it is the list $\langle$ NP $\rangle$. Notice that this analysis entails that the lexical head's value for the SPR feature must be available at the non-lexical level where the specifier phrase is attached. Thus in order for the new rule in (16) to work right, we will have to modify our head-complement rule so that it 'passes' the SPR value of the lexical head 'up' to its mother.[7] We might thus add a stipulation to this effect, as shown in (17):

(17) Head-Complement Rule:
$$\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \boxed{2} \\ \text{COMPS} & \langle\,\rangle \end{bmatrix} \;\rightarrow\; \text{H}\begin{bmatrix} \textit{word} & \\ \text{SPR} & \boxed{2} \\ \text{COMPS} & \boxed{1} \end{bmatrix} \quad \boxed{1}$$

This solves the problem of getting the SPR selection information from the lexical head up to VP or NOM – the phrase that will combine directly with the specifier (via our new rule (16)), but it does so at the cost of adding a stipulation to our rules. Moreover, more such stipulations are needed if we consider additional rules. In particular, recall the rule

[7] At first glance, one might be tempted to accomplish this by making SPR a head feature, but in that case the statement of the HFP would have to be complicated in order to allow rule (16) to introduce a discrepancy between the HEAD value of a mother and its head daughter.

for introducing PP modifiers, which was discussed in the previous chapter. Because no complements or specifiers are added by this rule, we do not want any cancellation from either of the head daughter's valence features. Hence, we would end up complicating the rule so as to transmit values for both valence features up from the head daughter to the mother, as shown in (18):

(18) $\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \boxed{1} \\ \text{COMPS} & \boxed{2} \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \boxed{1} \\ \text{COMPS} & \boxed{2} \end{bmatrix} \quad \text{PP}$

Without some such constraint, we can't have our modifiers combine with a VP head daughter to build another VP. It is time to contemplate a more general theory of how the valence features behave in headed phrases.

5.5 The Valence Principle

The intuitive idea in the previous section is quite straightforward, namely, that certain lexical entries specify what they can cooccur with by listing (as the value of the features COMPS and SPR) the particular kinds of dependents they select. And we formulated rules simply stating that all the head's COMPS members are 'discharged' in a head-complement phrase, while the item in the SPR value is discharged in a head-specifier phrase. But in order to make these rules work right, we had to add constraints preserving valence specifications in all other instances: the mother in the head-specifier rule preserves the head's COMPS value (the empty list); the mother in the head-complement rule preserves the head's SPR value, and the mother in the head-modifier rule must preserve both of the head's valence specifications. The operant generalization that can be factored out of our rules can be expressed as the following principle which, like the HFP, constrains the structural descriptions sanctioned by our grammar rules:

(19) The Valence Principle:

Unless the rule says otherwise, the mother's SPR and COMPS values are unified with those of the head daughter.

Two lists are unified only if all their members are unified. Hence, the effect of the Valence Principle is that (1) the appropriate elements mentioned in particular rules are cancelled from the relevant valence specifications of the head daughter in head-complement or head-specifier phrases, and (2) all other valence specifications are simply passed up from head daughter to mother. Once we factor these constraints out of our headed rules and put them into a single principle, it again becomes possible to simplify our grammar rules. This is illustrated in (20):

(20) a. Head-Specifier Rule:

$\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \langle\,\rangle \end{bmatrix} \rightarrow \boxed{1} \quad \text{H}\begin{bmatrix} \textit{phrase} & \\ \text{SPR} & \langle \boxed{1} \rangle \end{bmatrix}$

b. Head-Complement Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{COMPS}\ \langle\ \rangle \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{COMPS}\ \boxed{1} \end{bmatrix}\ \ \boxed{1}$$

c. Head-Modifier Rule:

$$\begin{bmatrix} \textit{phrase} \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{phrase} \end{bmatrix}\ \text{PP[...]}$$

Our work is not yet done. We will modify the head-modifier rule in the next chapter; in addition there are further modifications of (20a,b) yet to be introduced. In spite of this, the simplicity of the rules as formulated in (20) is striking.

Note that since head-specifier structures require a phrasal head and the head-complement rule requires a lexical head, it follows that head-complement phrases must be embedded within head-specifier phrases (rather than vice versa), as shown in (21):

(21)

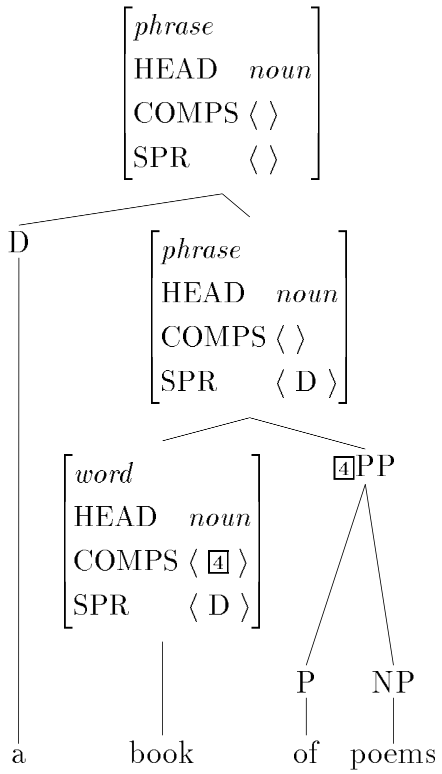

Thus COMPS cancellation happens at the NOM level and SPR cancellation happens at the NP level. The same is true of verbal structures: complements must be introduced lower (within VP) than specifiers (which appear as sister to a VP head within S).

5.6 Subject-Verb Agreement Revisited

Let us return now to the problem of subject-verb agreement. Our earlier analysis had assigned the feature NUM to both nouns and verbs, and one of our grammar rules had stipulated that the NUM values of VPs and their subjects had to match. But our new analysis has no need to specify NUM values for VPs. Since VPs all select for their subject via the SPR feature, what are traditionally referred to as 'plural verbs' can be treated simply in terms of the specification [SPR ⟨NP[NUM pl]⟩]; 'singular verbs' can similarly be specified just as [SPR ⟨NP[NUM sg]⟩]. Additional agreement features on verbs, VPs (and on S!) are simply unnecessary.

However, subject-verb agreement is a bit more complex, because it does not depend only on number. More specifically, English agreement also depends on PERSON. As noted earlier, we analyze person in terms of varying specifications for the feature PER. [PER 1st] is our notation for first person, that is, the pronouns *I* and *we*. [PER 2nd] denotes second person, which in English is always *you*. [PER 3rd] covers all non-pronominal NPs, as well as *he*, *she*, *it*, and *they*. Most present tense English verbs have one form when their subjects are third person singular (namely a form ending in *-s*) and another form covering all other persons and numbers. The only verb that makes finer distinctions than this is *be*, which has a special first person singular form, *am*.

We can couch our analysis of these verb forms in terms of a distinction between two kinds of values for the feature AGR. Suppose we call these two types *3sing* and *non-3sing*. The distinction among the third-person singular pronouns *he*, *she*, and *it* is attributed to a feature GEND(ER), with values *masc*, *fem*, and *neut*. GEND also differentiates among *him*, *her*, *it*, as well as among *himself*, *herself*, and *itself*. There is no motivation in English for assigning GEND to anything other than words that are third-person and singular.[8]

What this means is that we can set up our system so that the values for the feature AGR obey restrictions on the combinations of values for the features PER and NUM. That is, the only possible values for the feature AGR (in English) are those described in (22).

(22) Possible AGR Values:

$$\begin{bmatrix} \textit{3sing} & \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \\ \text{GEND} & \text{fem} \end{bmatrix} \begin{bmatrix} \textit{3sing} & \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \\ \text{GEND} & \text{masc} \end{bmatrix} \begin{bmatrix} \textit{3sing} & \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \\ \text{GEND} & \text{neut} \end{bmatrix}$$

$$\begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{bmatrix} \begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{1st} \\ \text{NUM} & \text{sg} \end{bmatrix} \begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{1st} \\ \text{NUM} & \text{pl} \end{bmatrix}$$

$$\begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{2nd} \\ \text{NUM} & \text{sg} \end{bmatrix} \begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{2nd} \\ \text{NUM} & \text{pl} \end{bmatrix}$$

A pronoun may specify one of these possibilities in its lexical entry; other elements, e.g. proper names, will have lexical entries whose AGR value is specified only as *3sing*, making these elements compatible in principle with any gender specification.

[8]This can be taken as independent evidence for the existence of *3sing* as a separate type – we can express this restriction by declaring the feature GEND as appropriate only for feature structures of the type *3sing*.

This treatment of the AGR values of nouns and NPs enables us to streamline our analysis of verbs. We can require that the lexical entries for words like *walks*, *runs*, or *is* (all derived in later chapters by lexical rule) include the following specification:

(23)
$$\left[\text{SPR}\left\langle\begin{bmatrix}\text{HEAD} & \begin{bmatrix}\textit{noun}\\ \text{AGR } \textit{3sing}\end{bmatrix}\\ \text{COMPS} & \langle\,\rangle \\ \text{SPR} & \langle\,\rangle\end{bmatrix}\right\rangle\right]$$

Because AGR is a head feature, any NP that serves as the subject of such a verb will pass the type specification [AGR *3sing*] down to its lexical head noun, courtesy of the HFP. Thus verbs whose lexical entries include (23) will occur only with NPs headed by third-person singular nouns.

It is often assumed that it is necessary to posit separate lexical entries for present tense verb forms that take plural subjects and those taking singular, non-third person subjects, as sketched in (24a,b):

(24) a.
$$\left[\text{SPR}\left\langle\begin{bmatrix}\text{HEAD} & \begin{bmatrix}\textit{noun}\\ \text{AGR}\begin{bmatrix}\text{NUM pl}\end{bmatrix}\end{bmatrix}\\ \text{COMPS} & \langle\,\rangle \\ \text{SPR} & \langle\,\rangle\end{bmatrix}\right\rangle\right]$$

b.
$$\left[\text{SPR}\left\langle\begin{bmatrix}\text{HEAD} & \begin{bmatrix}\textit{noun}\\ \text{AGR}\begin{bmatrix}\text{NUM sg}\\ \text{PER } \ 1\text{st} \mid 2\text{nd}\end{bmatrix}\end{bmatrix}\\ \text{COMPS} & \langle\,\rangle \\ \text{SPR} & \langle\,\rangle\end{bmatrix}\right\rangle\right]$$

But such an analysis would fail to explain the fact that the former type of verb would always be identical in form to the latter: a suspicious loss of generalization in the lexicon.

Once we bifurcate the types of AGR values, as described above, this problem disappears. We need only have a single kind of verb, one that includes the following lexical information:

(25)
$$\left[\text{SPR}\left\langle\begin{bmatrix}\text{HEAD} & \begin{bmatrix}\textit{noun}\\ \text{AGR } \textit{non-3sing}\end{bmatrix}\\ \text{COMPS} & \langle\,\rangle \\ \text{SPR} & \langle\,\rangle\end{bmatrix}\right\rangle\right]$$

Verbs so specified project VPs that take subjects whose head nouns must bear *non-3sing* AGR values, and these, as described above, must either be plural, or else first or second person.

Problem 2: The SPR Value of ***am*** What would the SPR value in the lexical entry for *am* be?

5.7 Determiner-Noun Agreement

We have just seen how our new analysis of specifiers, taken together with the Head Feature Principle, provides an account of the fact that third person singular verb forms (e.g. *walks*) take subject NPs headed by third person singular nouns. But, as we saw in section 4, the specifiers of the phrases projected from these nouns also participate in number agreement. In fact, the relevant data are slightly more complicated here too. English has determiners like *this* and *a*, which only appear with singular nouns, plural determiners like *these* and *few*, which only appear with plural nouns, and other determiners like *the*, which go either way. These facts are illustrated in (26) – (28):

(26) a. This dog barked.

b. *This dogs barked.

c. A dog barked.

d. *A dogs barked.

(27) a. *These dog barked.

b. These dogs barked.

c. *Few dog barked.

d. Few dogs barked.

(28) a. The dog barked.

b. The dogs barked.

In addition to the feature 'passing' that is guaranteed by the HFP, there is systematic agreement of person and number between heads and specifiers within the NP. To express this fact generally, we may add the following constraint on the (*part-of-speech*) type *noun*:

(29) Nominal SPR Agreement (NSA):

$$\textit{noun}: \begin{bmatrix} \text{AGR } \boxed{1} \\ \text{SPR } \langle\ (\begin{bmatrix}\text{AGR } \boxed{1}\end{bmatrix})\ \rangle \end{bmatrix}$$

The NSA requires that the AGR value of a noun be unified with that of its determiner, if the latter is realized. It thus makes determiner-noun agreement a lexical fact about nouns. This account presupposes that determiners and nouns both bear AGR specifications, as illustrated in (30)

(30) person, boat, a, this: $\left[\text{AGR}\begin{bmatrix}\text{PER} & \text{3rd}\\ \text{NUM} & \text{sg}\end{bmatrix}\right]$

people, boats, few, these: $\left[\text{AGR}\begin{bmatrix}\text{PER} & \text{3rd}\\ \text{NUM} & \text{pl}\end{bmatrix}\right]$

fish, the: $\left[\text{AGR}\begin{bmatrix}\text{PER 3rd}\end{bmatrix}\right]$

These lexical specifications, taken together with the NSA and the HFP are sufficent to account for all the agreement data in (26)– (28) above.

In section 4 above, we also observed that some determiners are restricted to occur only with 'mass' nouns (e.g. *furniture*), and others only with 'count' nouns (e.g. *chair*):

(31) a. Much furniture was broken.

b. *A furniture was broken.

c. *Much chair was broken.

d. A chair was broken.

The cooccurrence restriction illustrated in (31) – that is, the count/mass noun distinction – might of course be a totally semantic matter. The contrasting judgements in (31) are particularly striking, however, and may well warrant our treating these as ungrammatical. Assuming the syntactic account to be desirable (and such matters are often difficult to decide), we could analyze the data in (31) by introducing a feature COUNT. Certain determiners (e.g. *a* and *few*) will be lexically specified as [COUNT +] and others (e.g. *much*) will be lexically treated as [COUNT –]. Still others, such as *the*, will be lexically unmarked for this feature.

Once these lexical entries are in place, the facts are accounted for by treating nouns as follows. The SPR value of a count noun like *chair* would be ⟨D[COUNT +]⟩, forcing such nouns to cooccur with a count determiner. And the SPR value of a mass noun like *furniture* would be ⟨(D[COUNT –])⟩. As usual, the parentheses here designate optionality, and they appear only with mass nouns, because singular count nouns require determiners:

(32) a. (The) furniture is expensive.

b. The chair is expensive.

c. *Chair is expensive.

Notice that, under this analysis, the feature COUNT is marked only on the determiners, not on the nouns themselves. Count and mass nouns are distinguished only by whether their SPR lists contain a [COUNT +] or [COUNT –] element. In this way, the matching of count nouns and count determiners is handled entirely by the Valence Principle, which will identify the SPR value of NOM (the head of the NP) and the category of its specifier. No new machinery need be introduced. Notice that this analysis also predicts that the COUNT specification is never passed up to the NP, and hence can never be selected by the VP. Thus there should be no verbs in English that require a count or mass subject or object.

Problem 3: COUNT and NUM Is this last prediction right? That is, are there any verbs in English that require a count subject or object NP or a mass subject or object NP? Provide examples in support of your answer.

◇**Problem 4: COUNT and NUM** An alternative to the analyses just presented would eliminate the feature COUNT and assign three values to the feature NUM: sg, pl, and mass. That is, mass nouns like *furniture* would be given the value [NUM mass]. Use the following data to provide an argument favoring the analysis given in the text over this alternative:

(i) We don't have much rice/oats.

(ii) *We don't have many rice/oats.

(iii) The rice is in the bowl.

(iv) *The rice are in the bowl.

(v) The oats are in the bowl.

(vi) *The oats is in the bowl.

[NOTE: You may speak a variety of English that accepts *many oats* as a well-formed NP. There are some other nouns that are like *oats* in the relevant respects in at least some dialects, including *grits* (as a kind of cereal), *mashed potatoes*, and (somewhat distastefully, but grammatically more clearly) *feces*. If you can find a noun that patterns as we claim *oats* does in examples (i)-(vi), work the problem using that noun. If your dialect has no such nouns, then work the problem for the dialect described here, ignoring your own judgements.]

5.8 Worksection on Case Marking

We close this chapter with a series of problems dealing with grammatical case. The tools developed in this chapter should suffice to provide an analysis of the phenomenon of case, in a manner analogous to the treatment of agreement.

5.8.1 Case Marking in English

Consider the following data:

(33) a. Dogs like him.

b. *Dogs like he.

c. *Him likes dogs.

d. He likes dogs.

These sentences exemplify what linguists call CASE MARKING: *him* is what is called an ACCUSATIVE (or OBJECTIVE) case pronoun, and *he* is in what is called the NOMINATIVE (or SUBJECTIVE) case.
For the following problems, assume that there is a feature CASE with the values 'acc' and 'nom', and assume that English pronouns have CASE values specified in their lexical entries.

◊**Problem 5: Assessing the Facts** As noted in Chapter 2, NPs appear in a variety of positions in English, including subject of a sentence, object (direct or indirect) of a verb, and object of a preposition. For each of these NP positions, determine which case the pronouns in that position must have. Give examples to support your claims.

◊**Problem 6: A Rule-based Analysis** Assuming that nominative and accusative pronouns are entered in the lexicon with appropriate case values, indicate how we might change the head-complement and head-specifier rules in (20) to make sure that pronouns always appear in the correct case. Don't forget that an adequate analysis will also have to get the case values right in a sentence like *Kim gave me them.*

◊**Problem 7: A Lexical Analysis** How could we handle case entirely in the lexicon without changing our grammar rules? [Hint: This will involve modifying the lexical entries of elements that combine with NPs.]

Problem 8: Choosing Between the Analyses Provide arguments for choosing between the two analyses you have just presented. That is, say why you think case marking should be handled entirely in the lexicon rather than by adding constraints to the grammar rules (or vice-versa). [Hint: Redundancy (e.g. splitting rules into two, unnecessarily multiplying lexical entries) is always something to be avoided.]

Problem 9: Case and Coordination There is considerable variation among English speakers about case marking in coordinate NPs. Consult your own intuitions (or those of a friend, if you are not a native English speaker) to determine what rule you use to assign case to pronouns in coordinate structures. State the rule informally – that is, give a succinct statement, in English, of a generalization covering case in coordinate NPs in your dialect. Provide both grammatical and ungrammatical examples in support of your rule.

5.8.2 Case Marking in Icelandic

Icelandic is a language closely related to English, but it has much more elaborate and interesting case system. For one thing, it has four cases: nominative, accusative, genitive, and dative. Second, case is marked not just on pronouns, but also on nouns. A third difference is illustrated in the following examples[9]:

(34) a. *Drengurinn kyssti stúlkuna.*
the-boy-NOM kissed the-girl-ACC
'The boy kissed the girl.'

b. *Drengina vantar mat.*
the-boys-ACC lacks food-ACC
'The boys lack food.'

c. *Verkjanna gætir ekki.*
the-pains-GEN is-noticeable not
'The pains are not noticeable.'

d. *Barninu batnaði veikin*
the-child-DAT recovered-from the-disease-NOM
'The child recovered from the disease.'

♢**Problem 10: Choosing Analyses in Icelandic** Discuss how the examples in (34) bear on the analysis of case marking. In particular, how can they help us choose between handling case in the grammar rules and in the lexicon?

5.8.3 Agreement and Case Marking in Wambaya

In Wambaya, a language of Northern Australia, nouns are divided into four genders: masculine (masc), feminine (fem), vegetable (veg), and neuter (neut). They are also inflected for case, such as nominative (nom) and accusative (acc). Consider the following Wambaya sentences, paying attention only to the agreement between the determiners and the nouns (you do not have to worry about accounting for, or understanding, the internal structure of these words or anything else in the sentence).[10]

(35) (i) Ngankiyaga bungmanyani ngiya-ngajbi yaniyaga darranggu.
that.fem.ERG woman.fem.ERG she-saw that.neut.ACC tree.neut.ACC
'That woman saw that tree.'

(ii) Ngankiyaga bungmanyani ngiya-ngajbi mamiyaga jigama.
that.fem.ERG woman.fem.ERG she-saw that.veg.ACC yam.veg.ACC
'That woman saw that yam.'

[9] In the glosses, NOM stands for 'nominative', ACC for 'accusative', GEN for 'genitive', and DAT for 'dative'

[10] In fact, the Wambaya data presented here are simplified in various ways: only one of the numerous word order patterns is illustrated and the auxiliary plus verb sequences (e.g. *ngiya-ngajbi*) are here presented as a single word, when in fact the auxiliary is an independent verb in 'second' position. We are grateful to Rachel Nordlinger, who constructed this problem, in addition to conducting the field work upon which it is based.

(iii) Ngankiyaga bungmanyani ngiya-ngajbi iniyaga bungmaji.
that.fem.ERG woman.fem.ERG she-saw that.masc.ACC man.masc.ACC
'That woman saw that man.'

(iv) Ninkiyaga bungmanyini gina-ngajbi naniyaga bungmanya.
that.masc.ERG man.masc.ERG he-saw that.fem.ACC woman.fem.ACC
'That man saw that woman.'

(v) Ninkiyaga bungmanyini gina-ngajbi yaniyaga darranggu.
that.masc.ERG man.masc.ERG he-saw that.neut.ACC tree.neut.ACC
'That man saw that tree.'

(vi) Ninkiyaga bungmanyini gina-ngajbi mamiyaga jigama.
that.masc.ERG man.masc.ERG he-saw that.veg.ACC yam.veg.ACC
'That man saw that yam.'

Ergative (ERG) is the standard name for the case of the subject of a transitive verb in languages like Wambaya, where intransitive and and transitive subjects show different inflection patterns. Nothing crucial in this problem hinges on the distinction between nominative and ergative case. Note that the agreement patterns in (35) are the only ones possible, i.e. changing *mamiyaga* to *yaniyaga* in (vi) would be ungrammatical, etc. Note also that the verbs are selecting for the case of the subject and object NPs, so, e.g. *gina-ngajbi* must take a nominative subject and and accusative object, as in English.

Problem 11: **(a)** Since verbs select subjects and objects of a particular case and this case shows up in terms of the inflection of the head noun, what does this illustrate (minimally) about the feature CASE (i.e. where must it go in our feature geometry)?

(b) Explain how our analysis of English determiner-noun agreement would have to be modified in order to account for Wambaya determiner-noun agreement.

(c) Illustrate your analysis with lexical entries for *bungmanyani*, *ngankiyaga*, *bungmaji*, and *iniyaga*.

5.9 Summary

In the previous chapter, we already saw that cross-categorial generalizations about phrase structure can be expressed in terms of schematic phrase structure rules and categories specified in terms of feature structures. In this chapter, the real power of feature structure grammars has started to emerge. We have begun the process of providing a unified account of the generalizations about complementation and specifier selection in terms of the list-valued features COMPS and SPR. These features have enabled us to eliminate further redundancy from grammatical rules. Moreover, as we have already seen in this chapter, key generalizations about agreement and case marking can also be expressed in terms of a highly compact rule system, once we rely on categories modelled as feature structures.

5.10 Further Reading

The idea of schematizing phrase structure rules across parts of speech was introduced into generative grammar by Chomsky (1970). For a variety of perspectives on grammatical agreement, see Barlow and Ferguson (1988). A good discussion of Icelandic case is provided by Andrews (1982).

Chapter 6

Semantics and Summary

6.1 Introduction

Before we can return to the distribution of reflexive and non-reflexive pronouns, which will be the topic of the next chapter, we need to enter into some discussion of the nature of reference and coreference – topics that are fundamentally semantic in nature (i.e. that have to do in large part with the meaning of the sentences we are considering). And before we can do that, we need to say something about meaning in general and to sketch how meaning will be represented in our grammar.

The study of meaning is even older than the study of grammar, and there is little hope of doing justice to problems of semantics in a textbook whose primary concern is grammatical structure. Minimally, however, if the study of grammar is going to play any role in modelling real language use, then grammar has to include an analysis of the meaning of individual words and a treatment of how these combine – that is an account of how meanings of phrases and sentences are built up from the meanings of their parts. So let us begin by contemplating the nature of sentence meaning.

6.2 Semantics and Pragmatics

Meaning is inextricably bound up with actions – people intentionally using language for all kinds of communicative purposes. Some sentences are used to convey questions; others simple assertions; still others conventionally convey commands (or 'directives', as they are sometimes called). Even a piece of a sentence, say an NP like *the student sitting behind Leslie*, can be used in isolation to perform a simple act of referring to an individual.

The kind of meaning that (a particular use of) a sentence conventionally conveys depends crucially on its syntactic form. For example, a simple 'inverted' sentence like (1), where there is an auxiliary verb before the subject NP, conventionally conveys a question.

(1) Is Sandy tall?

And the question posed by (1) is closely related to the proposition that is asserted by an utterance of the noninverted sentence in (2).

(2) Sandy is tall.

In fact, uttering (2) is a perfectly good way of answering (1).

In order to even get started in dealing with semantic problems such as these, we are going to have to clarify what the units of meaning are and how particular kinds of sentences or smaller phrases are tied to particular types of meaning by linguistic conventions. We will make the standard assumption that communication has two components: linguistic meaning and reasoning about communicative goals). On this view, when a linguistic expression is uttered, its linguistic meaning makes a significant contribution to, but does not fully determine the communicative function of the utterance. Consider, for example, an utterance of (3).

(3) Do you have a quarter?

The linguistic meaning of this sentence is a familiar kind of semantic object: a question. And a question of this form has a determinate answer: yes or no. However, an utterance of (3) might serve to communicate much more than such a simple factual inquiry. In particular, in addition to posing a financial question to a given hearer, an utterance of (3) is very likely to convey a further message – that the speaker was making the following request of the addressee.

(4) Please give me a quarter!

The question conveyed by an utterance of (3) is generally referred to as its LITERAL or CONVENTIONAL meaning. A request like (4) is communicated as a further message above and beyond the message coming directly from the literal meaning of the question in (3). We will leave the account of such embellished communication (even the routine 'reading between the lines' that occurs more or less effortlessly in cases like this) to a more fully developed theory of language use – i.e. to a theory of linguistic PRAGMATICS. The inference from question to request is pragmatic in nature.

By contrast, the fact that a sentence like (3) must have a question as its literal meaning is semantic in nature. SEMANTICS is the study of linguistic meaning, i.e. the contribution to communication that derives directly from the conventions of the language. Pragmatics is the more general study of how linguistic meaning interacts with situational factors and the plans and goals of conversational participants to achieve more subtle, often elaborate communicative effects.

The semantic analysis that a grammar provides serves as input for a theory of language use. Such a theory sets as its goal to explain what actually gets communicated via pragmatic inference derived from the linguistic meaning of an utterance. For example, pragmatic theory might include a principle like (5):[1]

(5) Quantity Principle (simplified):

If X is weaker than Y, then asserting X implies the denial of Y.

[1](5) relies on the undefined term 'weaker'. In some cases (such as the example that follows), it is intuitively obvious what 'weaker' means. But a full-fledged pragmatic theory that included (5) would have to provide a precise definition of this term.'

This principle leads to pragmatic inference via 'proofs' of the following kind (justifications for steps of the proof are given in parentheses):

(6)
- A says to B: *Two things bother Pat.*
- A uttered something whose linguistic meaning is:
  'Two things bother Pat.' (semantic analysis)
- 'Two things bother Pat.' is weaker than 'Three things bother Pat.' (a fact in the context; possibly a more general fact)
- B assumes that A also meant to communicate: 'It's not the case that three things bother Pat.' (Quantity Principle)

Note that exactly the same pragmatic inference would arise from an utterance by A of any semantically equivalent sentence, e.g. *There are two things that bother Pat.* or *Pat is bothered by two things.*. This is because pragmatic theory works from the linguistic meaning of an utterance (as characterized by our semantic analysis) and hence is indifferent to the form by which such meanings are expressed.[2]

There is much more that could be said about the fascinating topic of pragmatic inference. Here, the only purpose has been to show that the semantic analysis that must be included in any adequate grammar in fact plays an essential role, albeit an indirect one, in explaining the communicative function of language in context.

6.3 Linguistic Meaning

6.3.1 Kinds of Meaning

When we ask a question, make an assertion, or even issue a command, we are also making reference to something that that is often called a SITUATION or EVENT.[3] If you utter a sentence like *Kim is running*, for example, you assert that there is some running situation in the world that involves someone named Kim. The proposition that you assert is either

[2]This is not quite true. Sometimes the MANNER in which something is said (the form of an utterance) can make some pragmatic contribution to an utterance, but a discussion of such cases would take us too far afield.

[3]Although the term 'event' is often used in a general sense in semantic discussions, this terminology can be misleading, especially in connection with circumstances like the following, where nothing very event-like is happening:

(i) Bo knows baseball.

(ii) Dana is aggressive.

(iii) Sydney resembles Terry.

(iv) Chris is tall.

(v) 37 is a prime number.

We find it much more intuitive to discuss such sentences in terms of 'situations' and hence have adopted this as our official terminology for the semantics of sentences.

true or false depending on a number of things, e.g. whether this situation is a running event (maybe Kim is moving too slowly for it to really qualify as running), whether the runner is someone named Kim (maybe the person you have in mind is really Leslie), whether the running situation is really happening now (maybe Kim is supposed to be running a race on another continent and my watch stopped several hours ago). If any of these 'maybes' turns out to be the case, then what you have said is false – the situation you are referring to fails to satisfy the restrictions specified by the linguistic meaning of the sentence.

An important part of the business of semantics is specifying truth conditions such as these, i.e. the restrictions that must be satisfied by particular situations involved in the propositions that people assert. Our grammar will thus be incomplete, unless we introduce (i) some way of representing the linguistic meanings of words and (ii) a set of constraints that allows us to correctly predict the linguistic meanings of phrase structures in terms of the meanings of their parts (their subconstituents).

Consider what this means in the case of *Kim is running*. What we need to guarantee is that this sentence gets a semantics that is a proposition (not a question or a directive, for example) specified in terms of the following conditions:[4]

(7) a. there is a situation s

b. s is a running situation

c. the runner is some individual i

d. i is named Kim

e. s is temporally located around the time of utterance

If there is some situation and some individual i such that all the conditions in (7) are satisfied, then the proposition expressed by *Kim is running* is true. If not, then that proposition is false.

So in order to take care of semantic business, we will need a way of ensuring that the various pieces of this sentence, e.g. the noun *Kim*, the verb *is*, and the verb *running* each make their appropriate contribution to the set of constraints summarized in (7) and that the grammar specifies how such propositions are built up from the substructures of the sentence. Our account must also be sufficiently general so as to assign the correct semantic description to all sentences of the language. So for example, a sentence like *Is Kim running?* should be assigned a semantics of a different type – a question – but a question about whether or not there is a situation s and an individual i such that all the conditions in (7) are satisfied.

In this book, we will build up meaning descriptions by providing constraints on grammar rules that specify how the phrase's (the mother's) semantics is built up from the semantics of its immediate constituents (the daughters). Our method for doing this will thus be the same as the way we constrain the syntactic feature specifications of the mother node in a tree description. The semantic objects of our grammar will be classified in terms of four

[4] The exact meaning of the progressive (*be...-ing*) construction is a fascinating semantic topic with a considerable literature that we cannot do justice to here. We have adopted clause (7e) as a convenient first approximation of the truth conditional contribution of the present progressive in English.

SEMANTIC MODES – that is, the four basic kinds of meanings that are enumerated and illustrated in (8).

(8)

SEMANTIC MODE	KIND OF PHRASE	EXAMPLE
proposition	noninverted sentence	Kim is happy.
question	inverted sentence	Is Kim happy?
directive	imperative phrase	Be happy!
reference	NP	the weather in San Francisco

To achieve this classification, we will represent meanings for all kinds of linguistic expressions in terms of feature structures that specify three things: a semantic mode, an index corresponding to the situation or individual being referred to, and a restriction that specifies a list of conditions that the situation or individual has to satisfy. Semantic structures then will quite generally look like (9):

(9) $$\begin{bmatrix}\text{MODE}\\ \text{INDEX}\\ \text{RESTR(ICTION)}\end{bmatrix}$$

Propositions are analyzed in terms of feature structures like the following one.

(10) $$\begin{bmatrix}\text{MODE} & \text{prop(osition)}\\ \text{INDEX} & \text{s}\\ \text{RESTR} \ \langle \ ... \ \rangle & \end{bmatrix}$$

A proposition like (10) will be true just in case there is some real situation s (and there exist appropriate other individuals corresponding to whatever variables – like 'i' or 's' – are present in (10)) such that the constraints specified in the RESTR value of (10) are all satisfied. These restrictions, the nature of which will be explained in the next section, must include all those that are relevant to the meaning of the sentence, e.g. all the constraints just mentioned in conjunction with the truth or falsity of *Kim is running*. Our grammatical analysis must make sure that we end up with exactly the right constraints in the RESTR list of a sentence's semantics, so that we associate exactly the right meaning with any sentence sanctioned by our grammar.

Questions and directives have a similar analysis, though the intuition behind their meaning is somewhat different. A question like *Is Kim running?* must be assigned a semantics just like the one assigned to *Kim is running*, except that the MODE value must be 'question', rather than 'prop':

(11) $$\begin{bmatrix}\text{MODE} & \text{ques(tion)}\\ \text{INDEX} & \text{s}\\ \text{RESTR} \ \langle \ ... \ \rangle & \end{bmatrix}$$

The question of whether a situation s satisfies a set of restrictions is the kind of semantic object that can be resolved positively or negatively in a given context. We can thus talk about true or false answers to questions, but questions themselves are neither true nor false.

Neither are directives, which are represented as in (12):

(12) $\begin{bmatrix} \text{MODE} & \text{dir(ective)} \\ \text{INDEX} & \text{s} \\ \text{RESTR} & \langle \ ... \ \rangle \end{bmatrix}$

A directive is rather something that can be fulfilled or not; and a directive is fulfilled only if the hearer of an utterance acts so as to bring about a situation that satisfies all the specified restrictions.

A *reference* is similar to the kinds of meanings just illustrated, except that it can be used to pick out (or DENOTE) all kinds of entities – not just situations. We use what are usually called INDICES (notated with the letters *i*, *j*, *k*, etc.) as the INDEX value for the semantics of nominal expressions. These function pretty much in the same way as variables in algebra or in mathematical logic. Formally, the INDEX values we write as 's' are also indices, but a particular type of index that refers only to situations. More generally, indices are free to be associated with any kind of entity found in a discourse context.[5] So the semantics we assign to a referring NP has the following form:

(13) $\begin{bmatrix} \text{MODE} & \text{ref(erence)} \\ \text{INDEX} & \text{i} \\ \text{RESTR} & \langle \ ... \ \rangle \end{bmatrix}$

As we have just seen, there are a number of differences among the various semantic modes we have assumed. Despite these differences, there is a fundamental commonality of reference. Every kind of linguistic expression we have considered, irrespective of its semantic mode, makes reference to something that must satisfy an indicated list of restrictions. Our approach to semantics expresses this general fact by treating all expressions in terms of a single type of semantic object that introduces a referential index of one kind or another. The semantic work of distinguishing the different ways that the individuals and situations referred to contribute to linguistic meaning is left to the differing values of the feature MODE.

6.3.2 Predications

The business of semantics is to account for the role of language in communication. As we have just seen, we can begin to take care of this business in terms of a semantic analysis that recognizes diverse semantic modes and reference to individuals and situations. Much of the interesting work in linguistic semantics is done by the conditions that particular linguistic expressions impose on such situations and individuals. In terms of our analysis, this means

[5]There are any number of intriguing puzzles that are the subject of ongoing semantic inquiry. For example, what does an NP like *a page* refer to in the sentence *A page is missing from this book.*? How about *the unicorn that Chris is looking for* in the sentence: *The unicorn that Chris is looking for doesn't exist.*?

that there will be a particular importance attached to the contributions made by individual words to the values of the feature RESTR and to the way that the RESTR values of phrases are built up from those of their parts.

The restrictions that individuals or situations must satisfy come in many varieties, which concern what properties some individual has, who did what to whom in some situation, when, where, or why some situation occurred, and so forth. That is, semantically relevant restrictions involve specifications of the properties that must hold of individuals and situations and of certain relations that must hold among them.

To represent this sort of information, we must introduce into our semantics some way of specifying relations among entities quite generally. We do this by introducing a type of feature structure called *predication* (*pred*). The features of a predication specify (i) what kind of relation is involved and (ii) who or what is participating in the relation. Examples of feature structures of type *pred* are given in (14):[6]

(14)

$$\text{a. } \begin{bmatrix} \text{RELN} & \text{love} \\ \text{SIT(UATION)} & \text{s} \\ \text{LOVER} & \text{i} \\ \text{LOVED} & \text{j} \end{bmatrix} \quad \text{b. } \begin{bmatrix} \text{RELN} & \text{walk} \\ \text{SIT} & \text{s} \\ \text{WALKER} & \text{i} \end{bmatrix} \quad \text{c. } \begin{bmatrix} \text{RELN} & \text{give} \\ \text{SIT} & \text{s} \\ \text{GIVER} & \text{i} \\ \text{RECIPIENT} & \text{j} \\ \text{GIFT} & \text{k} \end{bmatrix}$$

$$\text{d. } \begin{bmatrix} \text{RELN} & \text{book} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{k} \end{bmatrix} \quad \text{e. } \begin{bmatrix} \text{RELN} & \text{happy} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \quad \text{f. } \begin{bmatrix} \text{RELN} & \text{under} \\ \text{SIT} & \text{s} \\ \text{LOWER} & \text{i} \\ \text{HIGHER} & \text{j} \end{bmatrix}$$

The predications in (14) are meant to correspond to conditions such as: 's is a situation wherein i loves j', 's is a situation wherein i walks', 's is a situation where i gives k to j', 's is a situation wherein k is a book', 's is a situation wherein i is happy', and 's is a situation wherein i is under j', respectively. We will henceforth make frequent use of predications like these, without taking the time to present a proper theory of relations, predications, and the features that go with them. Note that the restriction associated with many nouns and adjectives (*book*, *happy*, etc.) includes a predication of only one (non-situation) argument. In such cases – for example, (14d,e) – we use the feature INSTANCE.

Almost all words specify restrictions that involve predications of one kind or another, including verbs, adjectives, prepositions, and nouns. In order for the restrictions specified in the lexical entries of such words to get inherited by the phrases that contain them, we will need to posit constraints that (minimally) guarantee that the RESTR values of a phrase's daughters are part of that phrase's RESTR value. Only in this way will we end up with

[6] The kind of event-based semantic analysis we employ was pioneered by the philospher Donald Davidson in a number of papers. Our simplified representations differ from certain popular formulations where all the talk of existence is represented via explicit existential quantification, i.e. in terms of representations like (i):

(i) there is an event s and an individual i such that: s is a running event, the runner of s is i, i is named Kim, and s is temporally located around the time of utterance

We will leave all such existential quantification as implicit in our semantic descriptions.

a sentence whose RESTR value includes all the necessary restrictions on the relevant event participants.

For example, for a simple sentence like (15), we will want a semantic description like the one sketched in (16).

(15) A woman rescued a man.

(16)
$$\left[\begin{array}{ll} \text{MODE} & \text{prop} \\ \text{INDEX} & s_1 \\ \text{RESTR} & \left\langle \left[\begin{array}{ll} \text{RELN} & \text{rescue} \\ \text{SIT} & s_1 \\ \text{RESCUER} & \text{i} \\ \text{VICTIM} & \text{j} \end{array}\right], \left[\begin{array}{ll} \text{RELN} & \text{woman} \\ \text{SIT} & s_2 \\ \text{INSTANCE} & \text{i} \end{array}\right], \left[\begin{array}{ll} \text{RELN} & \text{man} \\ \text{SIT} & s_3 \\ \text{INSTANCE} & \text{j} \end{array}\right] \right\rangle \end{array}\right]$$

The conditions on s_1 come from the lexical entry for the verb *rescue*, the constraint that i – the rescuer – must be a woman comes from the noun *woman*, and the constraint that j – the victim in the rescuing situation – must be a man comes from the lexical entry for the noun *man*. By associating (15) with the feature structure in (16), our semantic analysis says that the linguistic meaning of (15) is the proposition that will be true just in case there is a real situation satisfying the complex condition that it involves a rescuing of a man by a woman. But in order to produce the right set of conditions in the sentence's semantic description, the conditions of the parts of the sentence have to be amalgamated into a single list of conditions. Note in addition that the main situation of the sentence is derived from that introduced by the verb. It is in general true that the semantics of a phrase will crucially involve the semantics of its head daughter.

6.4 How Semantics Fits In

In earlier chapters, we considered only the syntactic properties of linguistic expressions. To accommodate the basic analysis of linguistic meaning just sketched, we need some way of introducing semantic structures into the feature structures we use to analyze words and phrases. We do this by adding two new features – SYN(TAX) and SEM(ANTICS) – and adding a level of embedding within our feature structures, as illustrated in (17):

(17)
$$\left[\begin{array}{ll} \text{SYN} & \left[\begin{array}{ll} \text{HEAD} & [\ldots] \\ \text{VAL} & \left[\begin{array}{ll} \text{SPR} & \ldots \\ \text{COMPS} & \ldots \end{array}\right] \end{array}\right] \\ \text{SEM} & \left[\begin{array}{l} \text{MODE} \\ \text{INDEX} \\ \text{RESTR}\ \langle\ \ldots\ \rangle \end{array}\right] \end{array}\right]$$

There is now a syntactic side and a semantic side to all feature structures like (17), which we will assign to the type called *synsem-struc(ture)*. Although we will add a few more features as we progress, this is in essence the feature geometry that we will adopt in the remainder of the book. We will frequently have occasion in what follows to refer to the synsem-struc of a phrase. What we mean by this is the full feature structure that is the top node of the SD of that phrase, including both its SYN and SEM values.

This changes the way lexical entries look, of course; and this is illustrated in (18):

(18) a.

$$\left\langle \text{woman}\ ,\ \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} \\ \text{AGR}\ \boxed{1}\textit{3sing} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} & \langle\ \text{Det}[\text{AGR}\ \boxed{1}]\ \rangle \\ \text{COMPS} & \langle\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{ref} \\ \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{woman} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

b.

$$\left\langle \text{Kim}\ ,\ \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} \\ \text{AGR}\ \textit{3sing} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} & \langle\ \rangle \\ \text{COMPS} & \langle\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{ref} \\ \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{name} \\ \text{SIT} & \text{s} \\ \text{NAME} & \text{Kim} \\ \text{NAMED} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

c.
$$\left\langle \text{love}, \left[\begin{array}{ll} \text{SYN} & \left[\begin{array}{ll} \text{HEAD} & [\textit{verb}] \\ \text{VAL} & \left[\begin{array}{ll} \text{SPR} & \langle \text{NP} \rangle \\ \text{COMPS} & \langle \text{ NP[acc] } \rangle \end{array}\right] \end{array}\right] \\ \text{SEM} & \left[\begin{array}{ll} \text{MODE} & \text{prop} \\ \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \left[\begin{array}{ll} \text{RELN} & \text{love} \\ \text{SIT} & \text{s} \\ \text{LOVER} & \text{i} \\ \text{LOVED} & \text{j} \end{array}\right] \right\rangle \end{array}\right] \end{array}\right] \right\rangle$$

Under our assumption that all NPs[7] have a referential index as their INDEX value, it is straightforward to associate the semantics of NPs with the particular roles in a verb's predication. If, as assumed above, the role arguments within predications are indices, then we can achieve the desired result by letting a verb (and other elements with predicational semantics) specify links between the indices of its dependents (specifiers and complements) and the role arguments in the predication on the RESTR list of its own semantics. This is illustrated in (19) for the verb *love*.

NB: Here and throughout, we use $\mathbf{NP_i}$ as a shorthand for an NP whose SEM value's INDEX is i.

(19)
$$\left\langle \text{love}, \left[\begin{array}{ll} \text{SYN} & \left[\begin{array}{ll} \text{HEAD} & [\textit{verb}] \\ \text{VAL} & \left[\begin{array}{ll} \text{SPR} & \langle \text{NP}_i \rangle \\ \text{COMPS} & \langle \text{ NP[acc]}_j \rangle \end{array}\right] \end{array}\right] \\ \text{SEM} & \left[\begin{array}{ll} \text{MODE} & \text{prop} \\ \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \left[\begin{array}{ll} \text{RELN} & \text{love} \\ \text{SIT} & \text{s} \\ \text{LOVER} & \text{i} \\ \text{LOVED} & \text{j} \end{array}\right] \right\rangle \end{array}\right] \end{array}\right] \right\rangle$$

In this way, as the verb combines with a particular NP object, the index of that NP is identified with the value of the feature LOVED in the verb's semantics. The verb's INDEX value is identified with that of the VP it heads and the verb's restrictions are included in the VP's RESTR value, so that when that VP combines with the subject, the identities specified in (19) will also guarantee that the index of the actual subject of the sentence will be the

[7]Other than 'dummy' NPs like *there*, which we will turn to in chapter 11.

LOVER in the loving predication on the VP's RESTR list. This will in turn be included in the RESTR value of the S (see section 6 below), and hence the unifications specified in (19) will cause the conditions specified by the sentence's subject and object to be conditions about the lover and loved participants in the predication that restricts the situation that the sentence as a whole is about.

We will illustrate how the SEM value of complex expressions relate to their parts in more detail in section 6. Before concluding this section, however, we must point out that we are glossing over one important issue (*inter alia*) in this cursory presentation of semantic analysis. This is the matter of quantification. Sentences like those in (20) require a semantic treatment that goes beyond simple reference to individuals.

(20) a. Everyone liked *Sense and Sensibility.*

b. Most people are reading a book by Austen.

c. Few people who have met a friend of yours from Texas say there is nothing unusual about you.

The approach to NP semantics sketched here can easily be extended to deal with sentences containing QUANTIFIED NPs (*everyone*, *most people*, etc.), basically by augmenting our feature structures to allow more complex propositions (as well as questions and directives) that represent quantification over individuals explicitly (in terms of 'binding' of indices). We will leave this entire topic unexplored here, however, noting only that it is possible to introduce explicit quantification over situations as well.

6.5 Modification

Suppose that we introduce a HEAD feature called MOD and that the MOD value of a word specifies the kind of thing the word must modify. Then we could make it a lexical property of adjectives that they were [MOD NOM] (or [MOD NP]) and a lexical property of adverbs that they were [MOD VP] (or [MOD S]). Since these modifiers would then 'pass up' their MOD value, courtesy of the HFP, to any phrase that they projected (i.e. that they were the head daughter of), it would then become possible to use a single head-modifier rule, say the one in (21), to account for both nominal and verbal modification.

(21) The Head-Modifier Rule:

$$[\mathit{phrase}] \;\rightarrow\; \mathrm{H}\boxed{1}[\mathit{phrase}] \; \begin{bmatrix} \mathit{phrase} \\ \mathrm{HEAD}\begin{bmatrix}\mathrm{MOD}\ \boxed{1}\end{bmatrix} \end{bmatrix}$$

That is, rule (21) will license a NOM just in case the head-daughter is a phrase of category NOM and the modifier daughter's MOD value is also of category NOM:

(22)

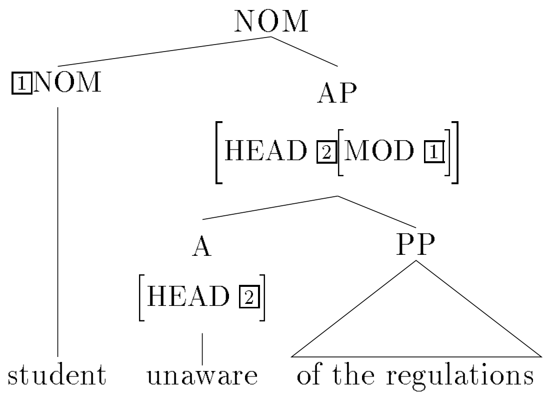

This NOM can combine with a D as specifier to build an NP like (23):

(23) a student unaware of the regulations

The head-modifier rule in (21) will also license the verb phrase in (24), under the assumption that adverbs are lexically specified as [MOD VP]:

(24)

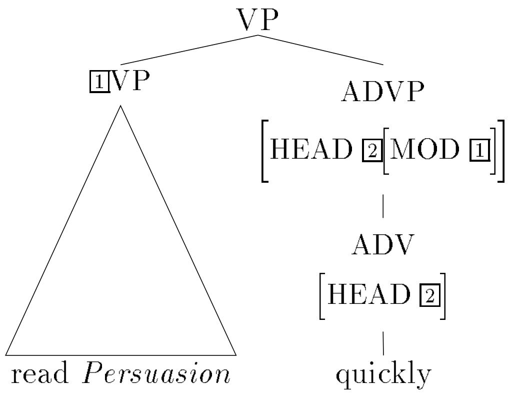

And a VP satisfying this description can combine with a subject like the one in (23) to build a sentence like (25).

(25) A student unaware of the regulations read *Persuasion* quickly.

Note further that many prepositional phrase modifiers can modify either nominal or verbal constituents:

(26) a. The [reporter *in Rome*]...

b. We [went shopping *in Rome*]

(27) a. The [weather *on Sunday*]...

b. We [went shopping *on Sunday*]

The bare bones treatment of modification just sketched can easily account for these examples simply by allowing underspecified or disjunctively specified MOD values for such prepositions.

6.6 The Semantics Principle

Not only is our grammar now able to analyze sentences of considerable complexity, but it is in addition now possible to treat the meanings of complex sentences by adding semantic constraints to the trees defined by our rules. The most general of these semantic constraints is given in (28):

(28) Semantics Principle:

In any well-formed structural description of a headed phrase,

(i) the mother's RESTR value is the sum of the RESTR values of the daughters, and

(ii) the mother's MODE and INDEX values are unified with those of the head daughter.

Clause (i) of (28) simply states that all restrictions from all the daughters in a headed phrase are collected into the RESTR value of the mother. The term 'sum' has a fairly obvious meaning here: the sum of the RESTR values of the daughters is the list whose members are those values, taken in order.[8] We will use the symbol '$\oplus$' to designate the sum operator. Clause (ii) further guarantees that the semantic MODE and INDEX of a phrase are identified with those of the head daughter, thus ensuring that semantics too has a 'head-driven' character to it.

The effect of the Semantics Principle is illustrated in the following simple example:

[8] That is, the sum of lists $\langle$ A $\rangle$, $\langle$ B , C $\rangle$, and $\langle$ D $\rangle$ is the list $\langle$ A , B , C , D $\rangle$.

(29)

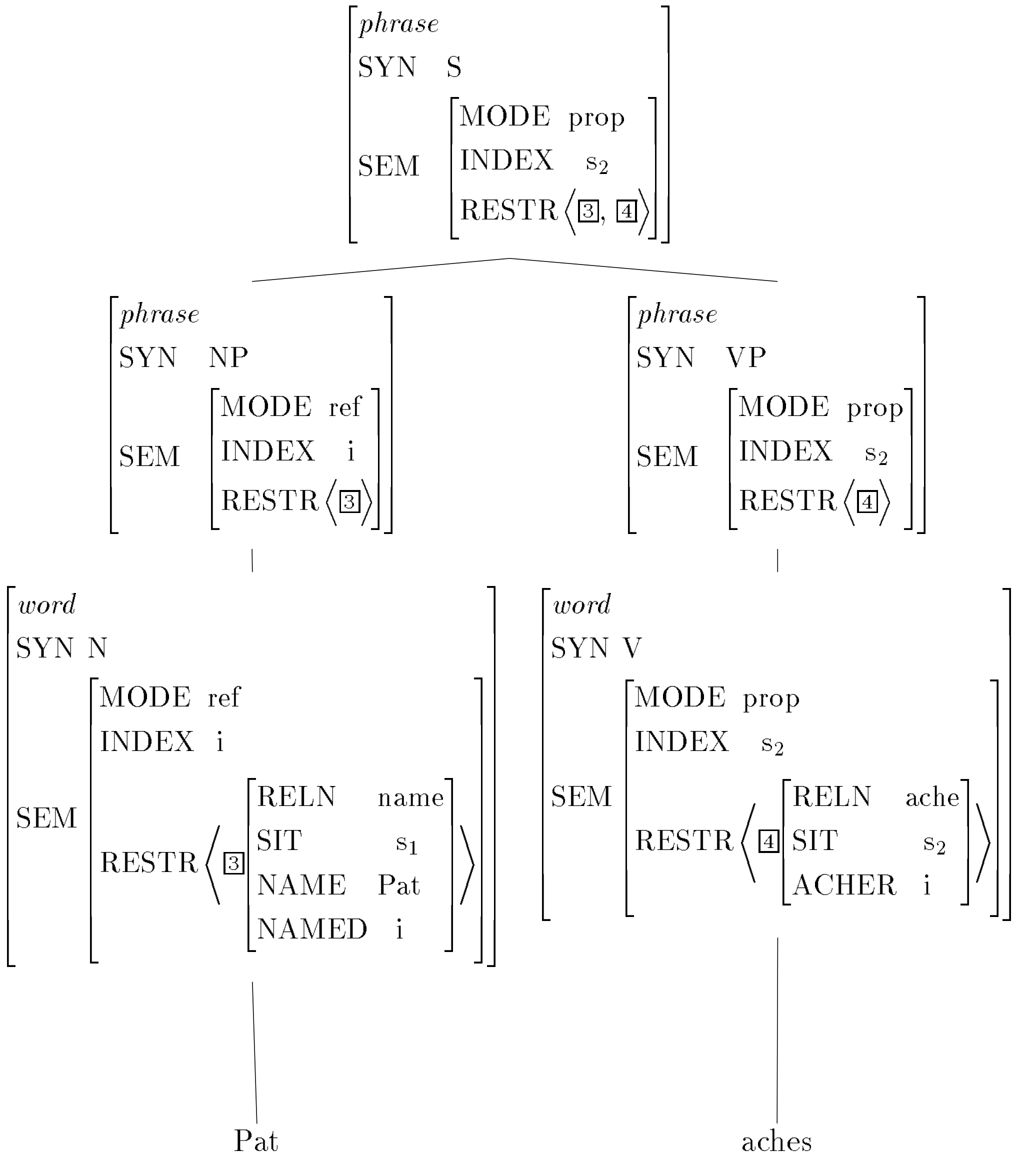

The two clauses of the Semantics Principle can be clearly observed in the S node on top. The MODE is 'prop', inherited from its head daughter, the VP node, (and ultimately from the head verb, *aches*) by clause (ii). Similarly, the INDEX value 's_2' comes from the verb, through the VP, by clause (ii). The RESTR value of the S node, [RESTR $\langle \boxed{3}, \boxed{4} \rangle$], is the sum of the RESTR values of the NP and VP nodes, as specified by clause (i) of the Semantics Principle.

Note that here, as before, we use abbreviations like 'NP', 'S', and 'V' to abbreviate feature structure descriptions specifying purely syntactic information. Since nodes in tree descriptions will now be labelled by feature structures that also include semantic information, the notational abbreviations should henceforth be reinterpreted accordingly.

With the Semantics Principle in place, we can now complete our account of modification. Let us assume that an adverb like *today* has a lexical entry like the one in (30).

(30)

$$\left\langle \text{today}\ , \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{adv} \\ \text{MOD VP[SEM [INDEX } \boxed{1}\text{]]} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} & \langle\ \rangle \\ \text{COMPS} & \langle\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{prop} \\ \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{today} \\ \text{SIT} & \text{s} \\ \text{ARG} & \boxed{1} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

The key thing to see here is that the MOD value tags the index of the VP to be modified as $\boxed{1}$ and this is identified with the argument of the relation 'today' in the semantic restriction. This means that once the adverb combines with a VP, the (situational) index of that VP IS the argument of 'today'.

The Semantics Principle, the Head-Modifier rule, and the lexical entry in (30) thus interact to SDs like the following (only SEM values are indicated):

(31)

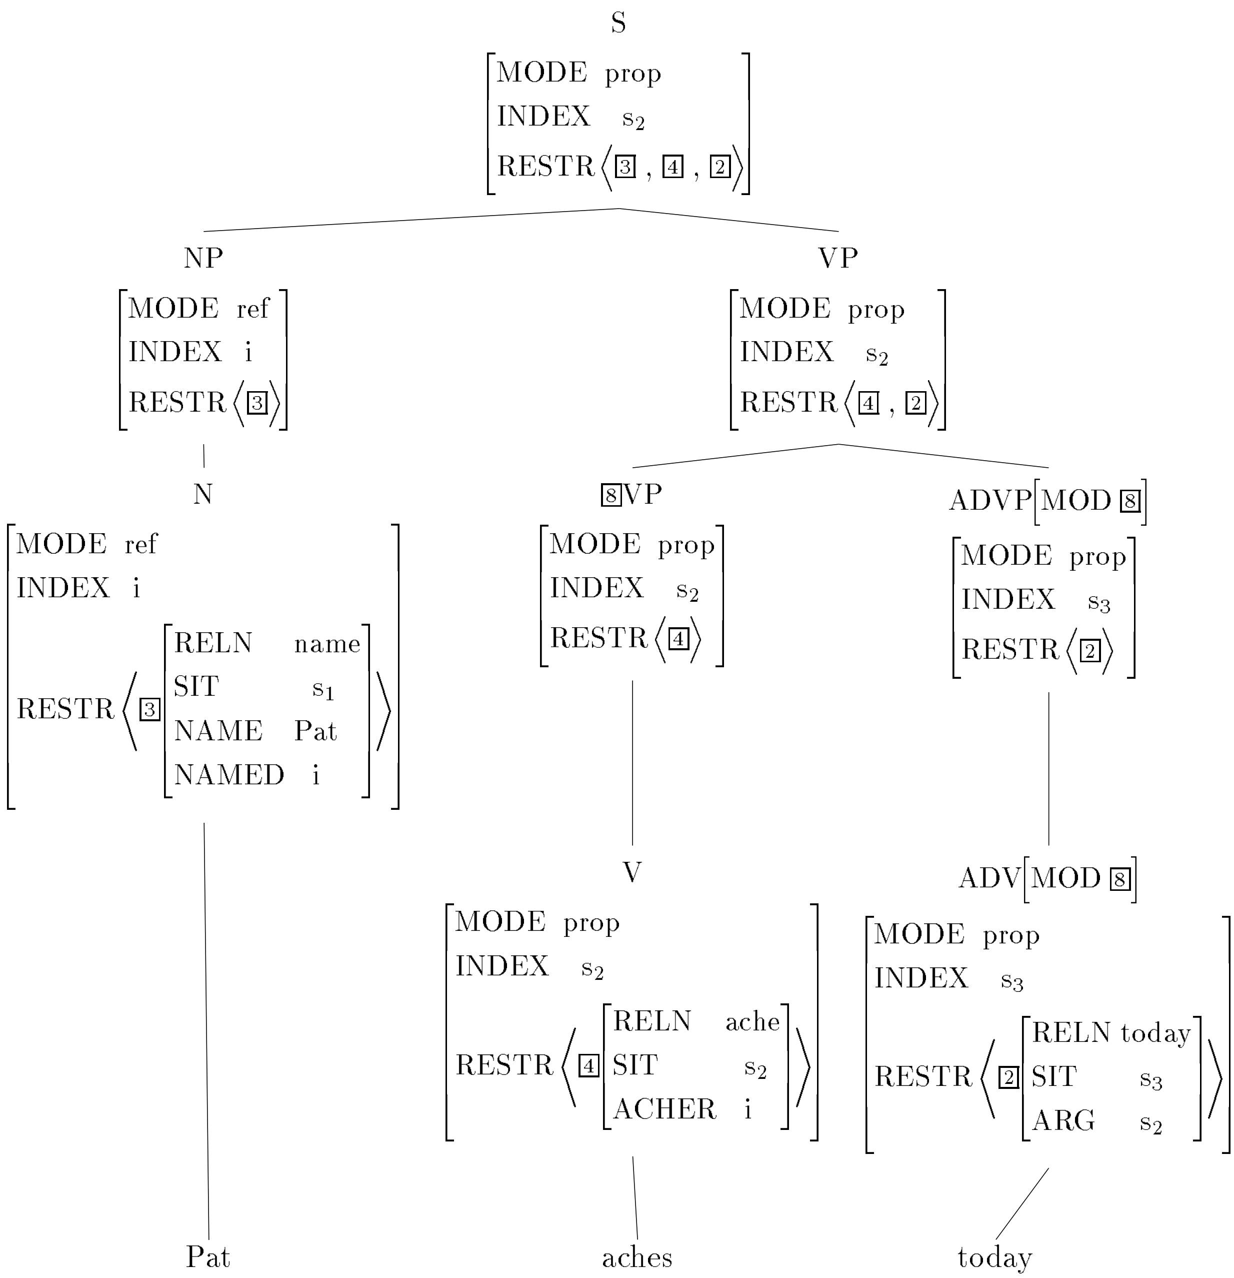

In this way, our analysis provides a general account of how meanings are constructed. The Semantics Principle embodies a simple and yet powerful theory of the relation between the structures of our grammar and the meanings they convey.

6.7 Summary

Three chapters ago, we began modifying the formalism of context-free grammar to make it better adapted to the sorts of generalizations we find in natural languages. We broke grammatical categories down into features, and then we broke the values of features down into features, as well. In the process, we moved more and more syntactic information out of the grammar rules and into the lexicon. In effect, we changed our theory of grammar so that the rules only give very general patterns that cut across grammatical categories. Details about what can go with what are specified in lexical entries in terms of valence features (i.e. in VAL).

With the expanded ability of our new feature structure complexes to express cross-categorial generalizations, our four remaining grammar rules cover a very wide range of cases. Two of them – the rules introducing complements and specifiers – were discussed extensively in the last chapter. The third one – a generalization of our old rules introducing PP modifiers to VP and NOM – was illustrated in the previous section. The fourth is the coordination schema. The formal statements of these rules are given in (32), along with informal translations, given in italics below the rules.[9]

(32) a. Head-Complement Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{COMPS } \langle\,\rangle \end{bmatrix} \rightarrow \text{H}\begin{bmatrix} \textit{word} \\ \text{COMPS } \boxed{1} \end{bmatrix} \quad \boxed{1}$$

A phrase may consist of a lexical head followed by all its complements.

b. Head-Specifier Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle\,\rangle \end{bmatrix} \rightarrow \boxed{1} \quad \text{H}\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle \boxed{1} \rangle \end{bmatrix}$$

A phrase may consist of a phrasal head preceded by its specifier.

c. The Head-Modifier Rule:

$$[\textit{phrase}] \rightarrow \text{H}\boxed{1}[\textit{phrase}] \quad \begin{bmatrix} \textit{phrase} \\ \text{HEAD} \begin{bmatrix} \text{MOD } \boxed{1} \end{bmatrix} \end{bmatrix}$$

A phrase may consist of a phrasal head followed by a compatible modifier phrase.

d. The Coordination Rule:

$$\begin{bmatrix} \text{SYN } \boxed{1} \end{bmatrix} \rightarrow \begin{bmatrix} \text{SYN } \boxed{1} \end{bmatrix}^{+} \quad \text{CONJ} \quad \begin{bmatrix} \text{SYN } \boxed{1} \end{bmatrix}$$

Any number of occurrences of elements of the same syntactic category may be conjoined to make a coordinate element of that category.

In addition to our grammar rules, we must provide (as we did in the case of CFGs) some characterization of the 'initial symbols', corresponding to the types of phrase that can stand alone as sentences of the language. We postpone a characterization of this until chapter 9, when we will have introduced a method for distinguishing finite (that is, tensed) clauses from others.

The richer feature structures we now employ, together with our highly schematized rules,

[9]Notice that the rule in (32c) does not cover all kinds of modifiers. In particular, some modifiers – such as adjectives inside NPs – precede the heads that they modify. To accommodate such modifiers, we would need an additional rule.

have required us to refine our notion of how a grammar is related to the structural descriptions it generates and how those descriptions are related to the fully determinate phrase structures of the language. Intuitively, here is how it works.

First, each lexical entry can license a lexical SD which is, as before, a nonbranching tree. The mother in a lexical SD is the feature structure specified in the lexical entry, and the daughter is the word itself. These lexical SDs form the bottom layer of well-formed phrasal SDs. They can be combined[10] into larger SDs in the ways permitted by the grammar rules, filling in features as required by our three principles: the Head Feature Principle, the Valence Principle, and the Semantics Principle. This process can apply to its own output, making ever larger phrasal SDs. So long as each local SD that we construct is licensed by a grammar rule and conforms to the three principles, it is well-formed. Typically, each node in a well-formed SD will contain some information that was stipulated by a rule and other information that percolated up (metaphorically speaking) from lower nodes (and ultimately from the lexical entries) via the principles.

For most of the material presented in the chapters that follow, an intuitive grasp of how the lexicon, grammar rules, and principles interact to license SDs may suffice. Nevertheless, for completeness we include an explicit definition of the SDs admitted by our grammar as an appendix to this chapter.

The intended effect of this definition is that an SD contains exactly as much information as is obtained by unifying the constraints that percolated up from the daughters (represented in the formal statement by δ_i's), the constraints that are specified in each rule, and those constraints that are part of our general theory (the HFP, etc.).

We have formulated our theory so that, as successively larger SDs are produced, the descriptions of the (nonterminal) nodes expand through unification. That is, information that is left underspecified in an SD often needs to be included if it is to be embedded into a larger SD. This important side effect of unifying SDs can be illustrated simply. Consider the SD in (33), where the notation '...' is used to indicate that some information has been omitted.

(33)
$$\begin{array}{c} \text{NP} \\ \begin{bmatrix}\text{HEAD } \boxed{2} \\ \ldots\end{bmatrix} \\ | \\ \text{N} \\ \begin{bmatrix} \text{HEAD } \boxed{2} \begin{bmatrix} \textit{noun} \\ \text{AGR } \textit{3sing}\end{bmatrix} \\ \ldots \end{bmatrix} \\ | \\ \text{Leslie}\end{array}$$

[10]Our informal discussion is worded in terms of a process of building trees up from the bottom. This is a conceptually natural way of thinking about it, but it should not be taken too literally. The formal definition of well-formed SD that we give below is deliberately non-procedural.

This contains information about the HEAD value of the phrase (unified with the HEAD value of the N dominating *Leslie*). But when this SD is embedded within a larger one like (34), licensed by the Head-Complement Rule, the result is as shown in (34).

(34)

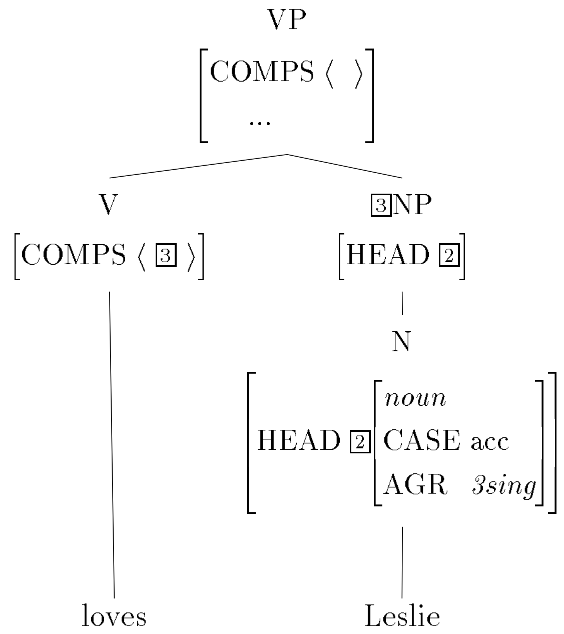

Because the Head-Complement Rule unifies the head daughter's COMPS list with the list of (the feature structures of the) complement daughters, the accusative case specification IS part of the SD of the object NP's HEAD value. And since that NP's HEAD specification is unified with the HEAD specification of its head daughter (namely, the N dominating *Leslie*), it follows that the accusative case specification is part of the SD of this embedded category as well.

The information specified by our rules and lexical entries is thus PARTIAL information. Each rule says, in effect, that subtrees of a certain type are sanctioned, but the rule only specifies SOME of the constraints that the trees that it licenses must obey. Likewise, a lexical entry says that certain trees dominating the phonological form in that entry are sanctioned, but the entry only specifies SOME of the information relevant at higher levels of structure. The general principles of our theory constrain the ways in which feature values can be distributed in well-formed trees. Unification of constraints is the basic method that allows the job of determining well-formedness to be distributed among the pieces of our grammatical system in a parsimonious way.

In short, we have arrived at a particular factorization of the information necessary for a precise account of grammatical descriptions. By far the richest source of information in this factorization is the lexicon. That is, our account embodies the claim that the problem of determining which strings of words constitute well-formed sentences and the problem of specifying the linguistic meaning of sentences both depend mostly on the nature of words. Of course, it must also be recognized that there are many regularities about what words go together (and how they go together). The theoretical constructs summarized here capture a number of such regularities; subsequent chapters will provide ways of capturing more.

◊**Problem 1: English Possessives** Consider noun phrases like *Kim's brother* and *the president's salary*. One traditional treatment of the possessive *'s* in English is to call it a case marker. The principal objection to a treatment of *'s* as a case marker is that, in conjunction with our analysis of case (in particular our assumption that CASE is a HEAD feature), it makes the wrong predictions about examples like the following:

(i) The Queen of England's crown disappeared.

(ii) *The Queen's of England crown disappeared.

(a) Explain what the '*'s*-as-case-marker' analysis predicts about the grammaticality of examples like (i) and (ii), and why.

An alternative analysis of possessive *'s* is to say that it is a determiner that builds a determiner phrase (abbreviated DP). On this analysis, *'s* selects for no complements, but it obligatorily takes an NP specifier. The word *'s* thus has a lexical category that is similar to that of an intransitive verb.

This analysis is somewhat unintuitive, perhaps, because (1) it requires that we have an independent lexical entry for *'s*, which seems more like a piece of a word, phonologically, and (2) it makes the nonword *'s* the head of a phrase! However, this analysis does a surprisingly good job of covering the facts of English possessives, so we would like to adopt it, at least for purposes of this text.

(b) Give the lexical entry for *'s* under the '*'s*-as-determiner' analysis, and draw a phrasal SD (that is, a tree description) for the NP, *Kim's brother* (Go ahead and use abbreviations when necessary, for example: NP[3rd, sg] or [SPR ⟨NP[nom]⟩]).

(c) Explain how your lexical entry gets the facts right in the *Queen of England* examples above.

(d) How does this analysis handle recursion in possessives, e.g. *Kim's brother's wife*, or *Kim's brother's wife's parents*? Provide at least one tree description to illustrate your explanation (again you can abbreviate complex feature structures if the omitted details are clearly irrelevant).

(e) Possessive pronouns like *my*, *your* etc. function as determiners in NPs like *my books*, *your mother*, etc. You might think we should treat possessive pronouns as determiners that have the same AGR value as the corresponding nonpossessive pronoun. That is, you might think that *my* should be specified as:

$$\begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{det} & \\ \text{AGR} & \begin{bmatrix} \text{PER} & \text{1st} \\ \text{NUM} & \text{sg} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Explain, providing appropriate examples, why this analysis will fail to provide an adequate account of *my books*, *your mother*, etc.

◊**Problem 2: A Sentence** Give an analysis, using the grammar as developed so far, of the following sentence:

> They sent us a letter.

Giving an analysis involves the following parts:

(a) Draw a tree description, with the head and valence features indicated on each node. (You may use the usual abbreviations, such as 'NP' and 'S' in the values for the valence features).

(b) Identify which grammar rule licenses each of the phrasal SDs.

(c) For each word description, explain which features came from the word's lexical entry and which were 'unified in' as a phrasal SD was constructed.

(d) For each phrasal SD, explain how the grammar rules, the HFP, the Valence Principle and the Semantics Principle interact to determine the values of SPR, COMPS, HEAD, and SEM.

6.8 Appendix: A Precise Definition of Well-Formed Tree Descriptions

Note: the symbol $\sqcup$ here designates unification. Formally, the unification of two tree descriptions is just the description that results from unifying the constraints specified by the two descriptions for each node of the tree.

(35) Lexical Structural Descriptions:

$$\begin{array}{c}\Delta\\ |\\ \omega\end{array}$$

is a well-formed lexical structural description just in case $\langle \omega, D \rangle$ is a lexical entry.

(36) Phrasal Structural Descriptions:

If $\begin{array}{ccc}\delta_1 & \ldots & \delta_n\\ \triangle & & \triangle\end{array}$ are well-formed SDs (that is, lexical or phrasal structural descriptions),

and R is an instantiation of a grammar rule[11] of the form: $D_0 \rightarrow D_1 \ldots D_n$.

[11]The term 'instantiation' of a grammar rule is used because grammar rules like the Head-Complement Rule in (32a) and the Coordination Rule in (32d) are formulated schematically so as to allow any number of daughters. Each phrasal SD, of course, is in correspondence with an instance of such rules that has a fixed number of daughters.

then R sanctions a phrasal SD Δ just in case:

1. If R is unheaded, then Δ =

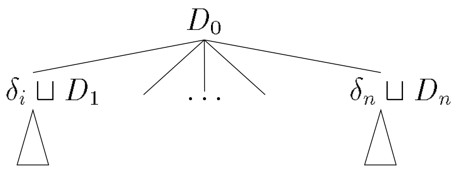

2. If D_h is the head of D_0 in R, then Δ is the result of unifying the SD in 1. with all of the following SDs:

   a. (Head Feature Principle:)

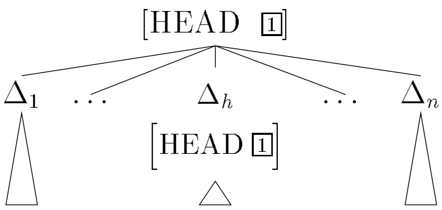

   b. (Valence Principle [unless R says otherwise]:)

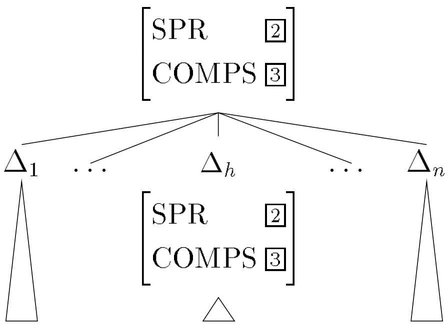

   c. (Semantics Principle (i):)

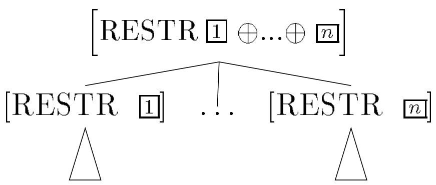

   d. (Semantics Principle (ii):)

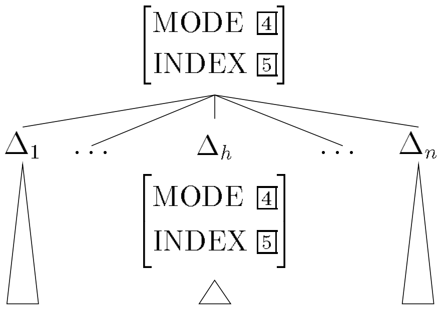

6.9 Further Reading

The seminal work in modern resarch on natural language semantics was Frege's (1892) essay, 'Über Sinn und Bedeutung' (usually translated as 'On Sense and Reference'), which has been translated and reprinted in many anthologies. More recently, the papers of Richard Montague (Montague, 1970) had a revolutionary influence, but they are extremely technical and difficult. An elementary presentation of his theory is given by Dowty, Wall, and Peters (1981). Good general introductory texts in semantics include Chierchia and McConnell-Ginet (1990) and Gamut (1991). Shorter overviews of semantics include Bach (1989), Barwise and Etchemendy (1989) and Partee (1995).

Chapter 7

Binding Theory

7.1 Introduction

In this chapter, we revisit a topic introduced very informally in chapter 1, namely, the distribution of reflexive and nonreflexive pronouns. In that discussion, we quickly noticed that the well-formedness of sentences containing reflexives usually depended crucially on whether there was another expression in the sentence which had the same referent as the reflexive; we called such an expression the 'antecedent' of the reflexive. With nonreflexive pronouns, the issue was typically whether a particular NP could have the same referent as (or, as linguists often put it, 'be COREFERENTIAL with') a given pronoun – that is, whether that NP could serve as the antecedent for that pronoun.

In discussions of these phenomena, we will use the notation of subscripted indices to mark which expressions are intended to have the same referent and which are intended to have distinct referents. Two expressions with the same index are to be taken as coreferential, whereas two expressions with different indices are to be understood as having distinct referents.

Thus, the markings in (1) indicate that *himself* must refer to the same person as *John*, and that the referent of *her* must be someone other than Susan.

(1) a. John$_i$ frightens himself$_i$.

b. *Susan$_i$ frightens her$_i$.

c. Susan$_i$ frightens her$_j$.

As discussed in the previous chapter, the subscript notation is shorthand for the value of the feature INDEX.

In examples like (1a), the reflexive *himself* is often said to be 'bound' by its antecedent. This terminology derives from the analogy between natural language pronouns and variables in mathematical logic. The principles governing the possible pairings of pronouns and antecedents are often called BINDING PRINCIPLES, and this area of study is commonly referred to as BINDING THEORY.[1] The term ANAPHORA is also used for the study of expressions (including pronouns) whose interpretation requires them to be associated with other elements in the discourse.

With this notation and terminology in place, we are now ready to develop a more precise and empirically accurate version of the binding theory we began to work on in chapter 1.

7.2 The Chapter 1 Binding Theory Revisited

In chapter 1, on the basis of examples like (2)–(9), we formulated the hypothesis in (10).

(2) a. Susan$_i$ likes herself$_i$.

b. *Susan$_i$ likes her$_i$.

(3) a. Susan$_i$ told herself$_i$ a story.

b. *Susan$_i$ told her$_i$ a story.

(4) a. Susan$_i$ told a story to herself$_i$.

b. *Susan$_i$ told a story to her$_i$.

(5) a. Susan$_i$ devoted herself$_i$ to linguistics.

b. *Susan$_i$ devoted her$_i$ to linguistics.

(6) a. Nobody told Susan$_i$ about herself$_i$.

b. *Nobody told Susan$_i$ about her$_i$.

(7) a. *Susan$_i$ thinks that nobody likes herself$_i$.

b. Susan$_i$ thinks that nobody likes her$_i$.

(8) a. *Susan$_i$'s friends like herself$_i$.

b. Susan$_i$'s friends like her$_i$.

[1] Much of the literature on binding theory actually restricts the term 'binding' to elements in certain syntactic configurations. Specifically, an element A is often said to bind an element B if and only if: (i) they have the same index; and (ii) A c-commands B. The technical term 'c-command' has been defined in several (non-equivalent) ways in the literature; the most commonly used definition is the following: node A in a tree c-commands node B if and only if every branching node dominating A dominates B. Intuitively, this means very roughly that A is at least as high in the tree as B. Our investigations into binding theory will not impose any such configurational limitation, deriving a similar, arguably superior characterization of constraints on binding in terms of ARG-ST lists (see below).

Note that we are interested in determining the conditions governing the pairing of pronouns and antecedents in a sentence. We will not, however, consider what possible things outside the sentence (be they linguistic expressions or entities in the world) may serve as antecedents for pronouns.

(9) a. *That picture of Susan_i offended herself_i.

b. That picture of Susan_i offended her_i.

(10) Reflexive pronouns must be coreferential with a preceding argument of the same verb; nonreflexive pronouns can't be.

Our task in this chapter is to reformulate something close to the generalization in (10) in terms of the theoretical machinery we have been developing in the last four chapters. We would also like to extend its empirical coverage to deal with examples our informal statement did not adequately handle. Toward this end, let us divide (10) into two principles, one for reflexive pronouns and the other for non-reflexive pronouns. Our first try at formulating them using the new binding terminology is then simply the following:

(11) **Principle A** (version I): A reflexive pronoun must be bound by a preceding argument of the same verb.

Principle B (version I): A nonreflexive pronoun may not be bound by a preceding argument of the same verb.

7.3 A Feature-Based Formulation of Binding Theory

Both of our binding principles make use of the phrase 'a preceding argument of the same verb.' We need to explicate this informal notion within the theory we have been developing. The place in our feature structures where we encode information about what the arguments to a verb are is the valence features (i.e. the values of VAL). Though we have not said much about the linear ordering of arguments, we have been listing our COMPS values in the order in which they appear in the sentence. Hence, to the extent that precedence information is represented in the feature structures, it is in the valence features. So the valence features seem like a natural place to start trying to formalize the binding principles.

7.3.1 The Argument Structure List

There is, however, a problem. For examples like (2)–(5), the binding in question involves the subject NP and one of the non-subject NPs; but our valence features separate the subject (specifier) and the non-subject (complements) into two different lists. To facilitate talking about all of the arguments of a verb together, we will posit a new list-valued feature ARGUMENT-STRUCTURE (ARG-ST), consisting of the sum (in the sense introduced in chapter 6) of the SPR value (the subject) and the COMPS value (the complements).

Words obey the following generalization, where '$\oplus$' again denotes the operation we have called 'sum', appending one list onto another:[2]

[2]Recall that $A \oplus B$ is the list consisting of the members of list A, taken in order, followed by the members of list B, also taken in order.

(12) **Argument Realization Principle (Version I)**:

A word's value for ARG-ST is $\boxed{1} \oplus \boxed{2}$, where $\boxed{1}$ is its value for SPR and $\boxed{2}$ is its value for COMPS.

So, if a verb is specified as [SPR ⟨ NP ⟩] and [COMPS ⟨ NP ⟩], then the verb's argument structure list is ⟨ NP, NP ⟩. And if another verb is specified as [SPR ⟨ NP ⟩] and [COMPS ⟨ PP, VP ⟩], then that verb's argument structure list is ⟨ NP, PP, VP ⟩, and so forth. Of course we mean real identity between the members of these lists, as shown by the specifications in (13):

(13) a.
$$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{VAL} \begin{bmatrix} \text{SPR} & \langle\, \boxed{1} \,\rangle \\ \text{COMPS} & \langle\, \boxed{2} \,\rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \langle\, \boxed{1}\text{NP} \,,\, \boxed{2}\text{NP} \,\rangle \end{bmatrix}$$

b.
$$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{VAL} \begin{bmatrix} \text{SPR} & \langle\, \boxed{1} \,\rangle \\ \text{COMPS} & \langle\, \boxed{2} \,,\, \boxed{3} \,\rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \langle\, \boxed{1}\text{NP} \,,\, \boxed{2}\text{PP} \,,\, \boxed{3}\text{VP} \,\rangle \end{bmatrix}$$

These identities are crucial, as they ensure that the binding properties of the complements are actually merged into the verb's argument structure, where they are governed by our binding principles. For example, the Head-Specifier Rule identifies a subject's synsem-struc with the synsem-struc of the sole member of the VP's SPR list. This means that the subject's synsem-struc is identical to the feature structure on the VP's SPR list. It follows (by the Valence Principle) that the subject's synsem-struc is also the sole member of the verb's SPR list. This, in turn, entails (by the ARP) that the subject's synsem-struc is the first member of the verb's ARG-ST list. Thus once the distinctions relevant to binding theory are encoded (by whatever means) in the feature structures of reflexive and nonreflexive NPs, this same information will be present in the ARG-ST of the lexical head of the sentence, where the binding principles can be enforced.

The generalization in (12) holds of words only; in fact, only lexical entries have the feature ARG-ST. Despite its close relationship to the valence features, ARG-ST is not itself part of VAL. Intuitively, this is because VAL's function (with the help of the Valence Principle) is to keep track of elements that a given expression has not yet combined with. As successively larger phrase structure descriptions are constructed, the list values of features specified within VAL get shorter. By contrast, the argument structure list exists only as part of a word description embedded within a larger description, and it always contains the same number of elements. But as larger phrases enter the picture, the ARG-ST list gradually acquires more and more information via unification.

ARG-ST, then, is a feature we only find on lexical heads, and the ordering of the elements in its value imposes a ranking on the phrases in the phrase structure that correspond to those elements. A bit more precisely, we can say:

(14) A phrase A outranks a phrase B just in case A's *synsem-struc* precedes B's *synsem-struc* on some argument structure (ARG-ST) list.

We can now reformulate our binding principles as follows:

(15) **Principle A** (version II): A reflexive pronoun of an ARG-ST list must be outranked by a coindexed element.

Principle B (version II): A nonreflexive pronoun of an ARG-ST list must not be outranked by a coindexed element.

7.3.2 Two Problems for Binding Theory

These formulations have certain problems, requiring some discussion and further refinement.

Pronominal Agreement

First, (15) says nothing about agreement between pronouns and antecedents, but we do not want Principle A to license examples like (16).

(16) a. *I enjoy yourself.

b. *He enjoys themselves.

c. *She enjoys himself.

We could rule these out by adding a stipulation to Principle A, requiring a reflexive and its antecedent to agree. But this ad hoc approach wouldn't explain much of anything. It is intuitively clear why coindexed elements should exhibit a form of agreement: coindexation indicates that the expressions denote the same entity, and the properties indicated by agreement features are characteristically properties of the entity referred to (the expression's DENOTATION). Thus, for example, singular NPs normally denote single entities, whereas plural NPs denote collections. Hence, a singular pronoun cannot be coindexed with a plural NP because they cannot have the same denotation.

We will consequently refrain from any mention of agreement in the binding principles. Instead, we adopt the following general principle:

(17) **Anaphoric Agreement Principle (AAP)**:

Coindexed elements agree.

By 'agree', we mean have the same values for AGR. Recall that AGR was introduced in chapter 5 as a feature whose value is a feature structure that specifies values for the features PER (person), NUM (number), and GEND (gender). Only PER and NUM matter for the purposes of subject-verb agreement, but pronouns must also agree with their antecedents in gender, as illustrated in (16c). Since GEND is part of AGR, it is covered by the AAP.

One advantage of leaving agreement out of the formulation of binding principles themselves is that the AAP also covers agreement between nonreflexive pronouns and their antecedents. Since Principle B only says which expressions must NOT be coindexed with nonreflexive pronouns, it says nothing about cases in which such pronouns ARE legally coindexed with something. The AAP rules out examples like (18), which were not covered by our formulation of Principle B.

(18) *I_i thought that nobody liked him_i.

It is important to realize that coindexing is not the same thing as coreference, though the former entails the latter. There are some tricky cases that might seem to be counterexamples to the AAP, and all of these turn on this distinction. One such example is the following:

(19) An interesting couple walked in. He was four foot nine; she was six foot two.

Here, we would say that the NP *an interesting couple* refers to one entity – the couple, which is a collection of two individuals. As the collection is introduced into the discourse, however, it also makes salient each individual that is in the collection, and it is these individuals that the pronouns in the next sentence refer to. Thus in this discourse, the NP *an interesting couple*, the pronoun *he* and the pronoun *she* actually all refer to different things. Therefore they have distinct indices and need not agree, according to the AAP.

Similar examples involve collective nouns like *family*, which can denote a single entity, as shown by the singular verb agreement in (20), but which can, as a 'side effect', introduce a collection of entities that can serve as the antecedent for a subsequent plural pronoun:

(20) My family hates cornflakes. But they love granola.

Again there are two distinct entities being referred to by distinct indices.[3]

And there are even some cases that we will probably want to treat as two nonagreeing NPs that nonetheless denote the same thing:

(21) The solution to this problem is rest and relaxation.

Here the singular NP *the solution to this problem* appears to refer to the same thing as the plural NP *rest and relaxation*. And indeed we would say that the two NPs are coreferential, but they are not coindexed. Thus while coindexing and coreference usually go hand in hand, they don't in this case. The whole point of identity sentences of this kind is to convey the information that two distinct (i.e. distinctly indexed) expressions refer to the same thing. If you are familiar with mathematical logic, this may remind you of situations in which two distinct variables are assigned the same value (making, for example '$x = y$' true). Indices are like variables; thus binding theory constrains variable identity, not the assignments of values to variables.

Binding in Prepositional Phrases

A second problem with our formulation of the binding principles is that reflexives and their antecedents can be objects of prepositions. A PP that consists of a prepositional head daughter like *to* or *about* and a reflexive NP object can then become a complement of the

[3]For some speakers, this is even possible in the context of reflexive pronouns, i.e. in examples like (i).

(i) Pat's family is enjoying themselves.

The theory we develop does not allow for examples of this sort.

verb; and when this happens, the reflexive NP inside the PP enters into binding relations with the other arguments of the verb. The same is true of a nonreflexive pronoun when it functions as a prepositional object. Thus we find the pattern of binding contrasts illustrated in (22) and (23).

(22) a. They$_i$ talk [to themselves$_i$].

b. *They$_i$ talk [to them$_i$].

(23) a. Nobody told Susan$_i$ [about herself$_i$].

b. *Nobody told Susan$_i$ [about her$_i$].

And in similar examples, the prepositional object can serve as the binder of a reflexive, but not of a nonreflexive:

(24) a. Nobody talked [to Susan$_i$] [about herself$_i$].

b. *Nobody talked [to Susan$_i$] [about her$_i$].

In examples like these, the binding principles as formulated above make the wrong predictions: the Argument Realization Principle (henceforth ARP) requires that the verb's ARG-ST contain the synsem-struc of the PP, not that of the prepositional object NP within the PP. Hence, if a reflexive pronoun is inside a PP that is a complement to a verb, the reflexive's synsem will not appear on the same ARG-ST list with the verb's subject and object NPs. The binding theory, as formulated, thus fails to take into account the fact that certain prepositions seem to be transparent for binding purposes. That is, if prepositions such as these were simply not there, then the prepositional object would be an object of the verb and binding theory would then make just the right predictions about (22)–(24) and related examples.

This problem raises both empirical and formal questions. The empirical question is the issue of precisely when objects of prepositions can enter into binding relations with elements of higher ARG-ST domains. As we noted in our initial discussion of binding theory in chapter 1, there is some variability about the binding possibilities of objects of prepositions. This is illustrated in (25).[4]

(25) a. The house$_i$ had a fence around it$_i$/*itself$_i$.

b. To make a noose, you wind the rope$_i$ around itself$_i$/*it$_i$.

c. Susan$_i$ wrapped the blanket around her$_i$/herself$_i$.

[4]It is likely that there is some cross-speaker variation as to whether examples like (25c) are acceptable or not.

These examples also show that it is not simply the choice of preposition that determines whether a prepositional object can be a reflexive.

One possible explanation of such differences is based on the intuitive idea underlying our binding theory, namely that reflexives and their antecedents are always arguments of the same predicate. It seems plausible to claim that English prepositions have two distinct semantic functions. In some uses, they function very much like verbs, introducing new predicates, with their own argument structures. In other uses, they are simply functioning as argument markers – that is, they indicate what role their object plays in the situation denoted by the verb of the clause they appear in. The clearest examples of this argument-marking use of prepositions are sentences like (4), in which *to* is used to mark what traditional grammarians called the indirect object. It seems plausible to claim that, in (25a), the preposition functions as a separate predicate (making the sentence mean roughly, 'The house had a fence, and the fence was around the house'), whereas in (25b), the preposition simply marks one of the arguments of the verb *wind*. Cases like (25c), then, would have to be treated as having prepositions that are ambiguous between being independent predicates and argument markers.

Let us now formalize this intuition. First, we introduce a new feature, which we will call ANAPHOR (ANA). ANA is used to distinguish reflexive pronouns (and reciprocals – that is, *each other*) from nonreflexive pronouns – and from other elements as well. Reflexives are [ANA +]; nonreflexive pronouns and nonpronouns are all [ANA –]. This will allow us to reformulate Principle A in terms of the feature ANA, keeping open the possibility that reflexives and reciprocals may not be the only elements subject to the principle. Since ANA specifications must be percolated up from the head noun to the NP it projects, we will treat ANA as a HEAD feature, as shown in the following lexical entries.

(26) a.
$$\left\langle \text{yourself}\ ,\ \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} & \\ \text{AGR} & \begin{bmatrix} \text{PER} & \text{2nd} \\ \text{NUM} & \text{sg} \end{bmatrix} \\ \text{ANA} & + \end{bmatrix} \\ \ldots & \end{bmatrix} \\ \text{SEM} & \ldots \end{bmatrix} \right\rangle$$

b.
$$\left\langle \text{them}\ ,\ \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} & \\ \text{CASE} & \text{acc} \\ \text{AGR} & \begin{bmatrix} \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{bmatrix} \\ \text{ANA} & - \end{bmatrix} \\ \ldots & \end{bmatrix} \\ \text{SEM} & \ldots \end{bmatrix} \right\rangle$$

c.
$$\left\langle \text{Terry}\ , \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} & \\ \text{AGR} & \begin{bmatrix} \text{PER} & \text{3rd} \\ \text{NUM} & \text{sg} \end{bmatrix} \\ \text{ANA} & - \end{bmatrix} \\ \ldots & \end{bmatrix} \\ \text{SEM} & \ldots \end{bmatrix} \right\rangle$$

Now we are in a position to distinguish the two types of prepositions discussed above. For the purposes of binding theory, nothing new needs to be said about the prepositions that function as independent predicates. If the object of such a preposition is a [ANA +], then Principle A will require it to be coindexed with something that outranks it on the preposition's ARG-ST list. This is not the case in (25a).[5] If the prepositional object is [ANA –], it must not be coindexed with anything that outranks it on the preposition's ARG-ST list. Since the subject of the sentence in (25a) does not appear on the ARG-ST list of *around*, Principle B permits a nonreflexive pronoun *it* coindexed with *the house* to appear as the object of *around*.

For prepositions that function as argument markers, however, we need to provide some way by which they can transmit information about their object NP up to the PP that they project. The higher verb that takes such a PP as its complement will then have the information about the PP's object NP in its ARG-ST, within the PP's synsem-struc. Note that without some method for transmitting this information up to the PP, the information about the preposition's object is invisible to the higher verb selecting the PP as its complement. The COMPS list of the PP, for example, is empty.

The method we use to transmit this information is a HEAD feature we will call P-OBJ. Only argument-marking prepositions like *to*, *about* and *of* will allow non-empty P-OBJ specifications. And the non-empty values of this feature are constrained to be identical to the feature structure of the preposition's object, as shown in the lexical entry in (27):

(27)
$$\left\langle \text{to}\ , \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{prep} & \\ \text{P-OBJ} & \boxed{1} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR}\ \langle\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \boxed{1}\text{NP[acc]} \right\rangle \\ \text{SEM} & \ldots \end{bmatrix} \right\rangle$$

Because P-OBJ is a HEAD feature, it is projected up as shown in (28):

[5] We leave open for now the question of how many arguments such predicational prepositions have. If *around* in (25a) has two arguments (as seems intuitive from its relational meaning), then the first argument should be identified with *a fence*; hence, *itself* could still not be coindexed with *the house*. In Chapter 12, we will investigate mechanisms by which different ARG-ST lists can have elements with the same index.

(28)

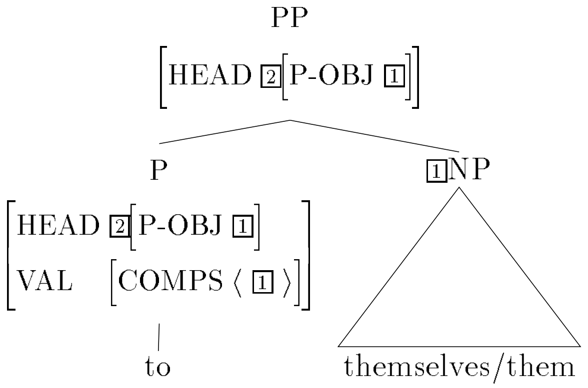

A PP like this will be selected by a verb like *tell* (in one of its VALENCE manifestations). Hence, the PP on its ARG-ST list will contain the P-OBJ value within it. Put another way, all the information about the object of the preposition is available in the ARG-ST list.

(29)
$$\left\langle \text{tell}, \begin{bmatrix} \text{SYN} & \begin{bmatrix}\text{HEAD} & \textit{verb}\\ \text{VAL} & \begin{bmatrix}\text{SPR} & \langle \boxed{1}\text{NP}\rangle\\ \text{COMPS} & \left\langle \boxed{2}\text{NP}, \boxed{3}\text{PP}[\text{P-OBJ}\ \boxed{4}]\right\rangle\end{bmatrix}\end{bmatrix}\\ \text{SEM} & \ldots \\ \text{ARG-ST} & \langle \boxed{1}, \boxed{2}, \boxed{3}\rangle\end{bmatrix}\right\rangle$$

In order to get the right binding results, we now need to make our binding principles sensitive to the presence of a P-OBJ specification within one (or more) elements of an ARG-ST list. We do this by making a slight modification to our definition of 'outranks'. In particular, we now say that a PP and its P-OBJ are 'of equal rank', and we treat the notion of 'ranking' as transitive in the obvious way. More precisely:

(30) (i) A PP and its P-OBJ value are of equal rank.
(ii) If there is an ARG-ST list on which A precedes B, then A has a higher rank than (i.e. outranks) B.

With these innovations in place, we may finally reformulate the binding principles in terms of the feature ANA. (We assume our modified definition of 'outrank').

(31) **Principle A** (version III): An [ANA +] synsem-struc must be outranked by a coindexed synsem-struc.

Principle B (version III): An [ANA −] synsem-struc must not be outranked by a coindexed synsem-struc.

7.3.3 Examples

So far, this section has motivated several technical innovations in our theory (ARG-ST, P-OBJ, and the distinction between the two types of prepositions). In this subsection, we go through two examples to illustrate the formal machinery we have been discussing.

Consider first (4a), repeated here for convenience.

(32) Susan$_i$ told a story to herself$_i$.

The SD that our grammar must license is the following (omitting irrelevant details):

(33)

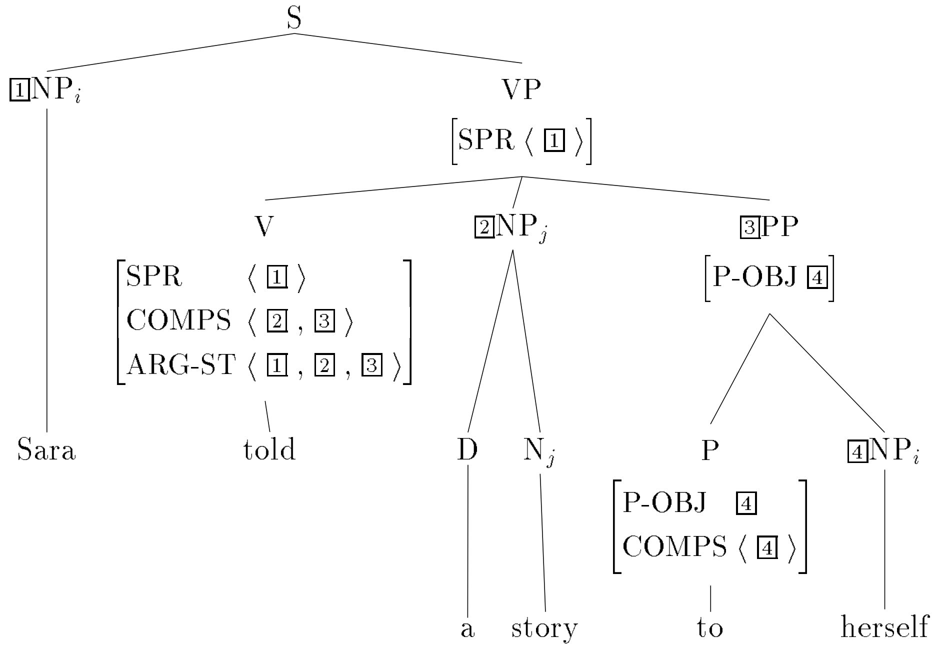

The geometry of the SD is licensed by our phrase structure rules in ways that are by now familiar. The aspect of the SD we are concerned with here is the coindexing of the nodes, indicated by the subscripted *i* and the resulting argument structure of the verb *told*, which is displayed in (34):

(34) $\left[\text{ARG-ST} \quad \left\langle \text{NP[ANA –]}_i \text{ , NP[ANA –]}_j \text{ , PP[P-OBJ NP[ANA +]}_i\text{]} \right\rangle\right]$

This ARG-ST conforms to the binding theory: the [ANA +] NP and the PP are of equal rank because the former is the P-OBJ value of the latter; both are therefore outranked by the first [ANA –] NP, which precedes the PP on the ARG-ST list; since this first NP and the [ANA +] NP are coindexed, Principle A is satisfied. Further, neither [ANA –] NP is outranked by a coindexed element. Notice that the effect of Principle A, though a local constraint on the verb's ARG-ST list, is to require coindexing between the subject NP one level higher in the SD and the prepositional object one level lower.

Principle A would also be satisfied if the anaphor were coindexed with the direct object NP:

(35) $\left[\text{ARG-ST} \quad \left\langle \text{NP[ANA –]}_j \text{ , NP[ANA –]}_i \text{ , PP[P-OBJ NP[ANA +]}_i\text{]} \right\rangle\right]$

Although this is implausible with *told* (because of the nonlinguistic fact that people are not the kind of thing that get told to others), it is much easier to contextualize grammatically analogous sentences with the verb *compared*

(36) a. We compared him_i [to himself_i] (at an earlier age)

b. We compared them_i [to each other_i].

Thus in the ARG-ST lists of both (34) and (35), the anaphor is outranked by some coindexed element. This prediction seems correct, as far as grammar is concerned, though there are orthogonal factors of plausibility that interact to diminish the acceptability of many grammatical examples.

Now consider (4b), repeated here for convenience.

(37) *Susan_i told a story to her_i.

The tree structure that our grammar must rule out is the following:

(38)

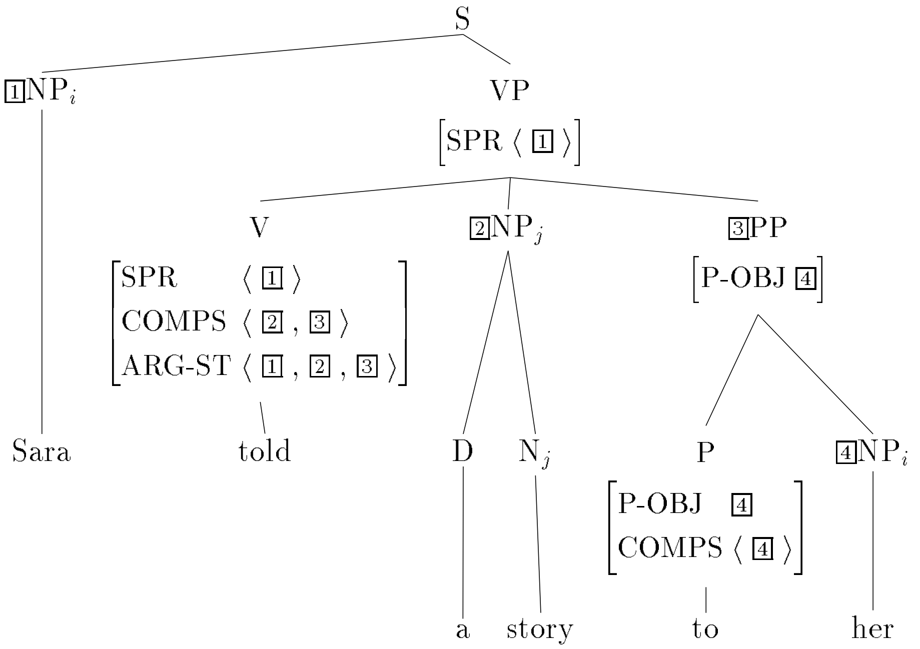

The lexical entry for *her* specifies that it is [ANA –], that is, that it is not a reflexive (or reciprocal) pronoun. As in the case of the previous example, the P-OBJ feature passes this information up to the PP, which carries it along into the verb's ARG-ST list, which looks like (39):

(39) $^*\Big[\text{ARG-ST} \quad \Big\langle \text{NP[ANA –]}_i\ ,\ \text{NP[ANA –]}_j\ ,\ \text{PP[P-OBJ NP[ANA –]}_i]\Big\rangle\Big]$

As before, the PP and the P-OBJ value (in this case, NP[ANA $-]_i$) are of equal rank, and the first NP outranks them. But in this case, since we have an [ANA –] NP that is outranked by a coindexed NP, Principle B is violated, and the tree is not licensed.

◊**Problem 1: Classifying Prepositions** We have divided prepositions into two sorts: those functioning as predicates and those functioning as argument-markers. Classify the italicized prepositions in each of the following sentences into one of these two sorts (or as being ambiguously in both), and justify your classification with data:

(i) The dealer dealt an ace *to* Bo.

(ii) The chemist held the sample away *from* the flame.

(iii) Alex kept a loaded gun *beside* the bed.

(iv) We bought flowers *for* you.

(v) The car has a scratch *on* the fender.

7.4 Imperatives

In chapter 1 we noted that the behavior of reflexive and non-reflexive pronouns in sentences like (40) is what one would expect if they had second person subjects.

(40) a. Protect yourself!

b. *Protect myself/himself!

c. *Protect you!

d. Protect me/him!

Sentences like these are known as IMPERATIVE sentences. Their characteristic properties are that they lack an overt subject, employ an uninflected form of the verb, and are used to express directives. Such sentences are sometimes said to have 'understood' second person subjects. The distribution of reflexives illustrated in (40) shows that imperatives do indeed behave in at least one way as if they had second person subjects.

The machinery introduced in this chapter provides a straightforward way of capturing the intuition that imperatives have understood subjects. First we are going to need to allow for verbal forms that convey imperative meaning. These forms, produced by a lexical rule discussed in the next chapter, have no inflectional endings and specify their semantic MODE as directive ('dir', rather than 'prop'). As also discussed in the next chapter, verb forms are distinguished from one another in terms of differing values for the HEAD feature FORM. Thus imperative forms will be distinguished from others by the specification [FORM imp]; in addition, they require their subject to be 2nd person (i.e. they are also specified as [SPR ⟨ [PER 2nd] ⟩].)

We also introduce a new grammar rule to analyze imperative sentences. This rule allows a sentence to consist of a single daughter: a VP specified as [FORM imp]. In requiring that the daughter be so specified, we thereby ensure that the lexical head of that phrase will be an imperative verbal form, i.e. an uninflected form like *be*, *get*, *run*, or *look*. The new rule we need for imperative sentences thus looks like (41).

(41) Imperative Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle\ \rangle \end{bmatrix} \rightarrow \quad \text{H}\begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \begin{bmatrix} \textit{verb} & \\ \text{FORM} & \text{imp} \end{bmatrix} \\ \text{SPR} & \big\langle [\ ] \big\rangle \end{bmatrix}$$

And this rule will sanction SDs like the one in (42):

(42)

$$\begin{array}{c} \text{S} \\ \begin{bmatrix} \text{SPR} & \langle\ \rangle \\ \text{COMPS} & \langle\ \rangle \\ \text{HEAD} & \boxed{1} \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{dir} \\ \text{RESTR} & \boxed{3} \end{bmatrix} \end{bmatrix} \\ | \\ \text{VP} \\ \begin{bmatrix} \text{SPR} & \langle\ \text{NP[PER} \ \ \text{2nd]}\ \rangle \\ \text{COMPS} & \langle\ \rangle \\ \text{HEAD} & \boxed{1}\text{[FORM} \ \ \text{imp]} \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{dir} \\ \text{RESTR} & \boxed{3} \end{bmatrix} \end{bmatrix} \\ \triangle \\ \text{get a job} \end{array}$$

Note that (42) also obeys the Head Feature Principle, The Semantics Principle and the Valence Principle. This is obvious in the case of the HFP; The Semantics Principle ensure the indicated inheritance of MODE and RESTR values; and the Valence Principle is satisfied because of its 'unless' clause, which allows a grammar rule to specify a non-canonical relationship between a node's VAL feature value (SPR in this case) and that of its head daughter.

The last thing to understand about the rule in (41) is that it explains the observations we have made about anaphor binding in imperative sentences. By requiring the first argument of imperative forms to be second person, we have ensured that in a structure like the following, Principle A will require a reflexive object to be coindexed with (and hence, by the AAP, to agree with) the second person subject:

(43)

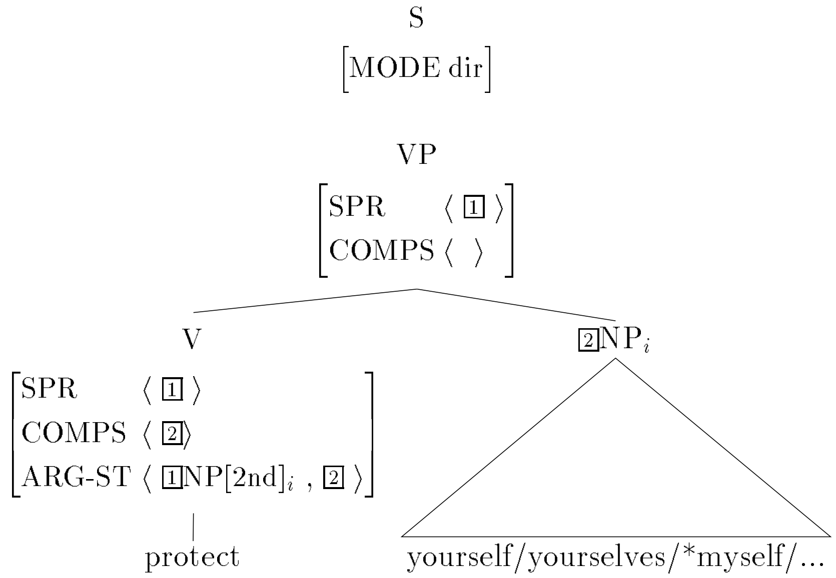

In this way, our treatment of imperatives interacts with our treatment of ARG-ST so as to provide an account of what are sometimes called ‘understood’ arguments. The values of VAL determine what arguments are expressed by overt NPs in the sentence, but ARG-ST may also include synsem-strucs that are not overtly expressed, i.e. which correspond to no overt phrase.

Note that we can use binding theory to confirm whether or not a given subjectless clause should involve an understood subject or not. For example, it would be a mistake to analyze sentences of the form *Damn NP!* along the lines just employed for imperatives. If we posited an understood NP in the ARG-ST of *damn*, it would license a reflexive pronoun (of the appropriate person, number, and gender) in the position after *damn*. But this is not possible:

(44) *Damn myself/yourself/herself/himself/itself/themselves!

Hence, *damn* in this use will have to be analyzed as being truly subjectless, both in terms of valence and in terms of argument structure. Examples like (44) are then ruled out because the reflexive element in the ARG-ST is not outranked by any coindexed element.

We have given a preview here of the analysis of verb forms that will be developed in the next chapter. There we will address the question of how the different forms are differentiated formally, and how to manage the proliferation of entries for different forms of the same word.

◊**Problem 2: Binding in Japanese** Japanese is a classic example of a ‘Subject-Object-Verb’ (SOV) language:[6]

(i) Taroo-ga hon-o yonda
Taroo-nom book-acc read-past
‘Taroo read the book’

[6]This problem is not concerned with the word-order difference between English and Japanese. It is mentioned only to help the reader understand the glosses of the Japanese sentences.

But in Japanese, both subjects and objects can in general be omitted. When such an element is omitted, the sentence usually has a meaning that we would express in English by using a pronoun to refer to the element that is left implicit in Japanese.

(ii) Taroo-ga yonda
Taroo-nom read-past
'Taroo read it/them'

(iii) hon-o yonda
book-acc read-past
'He/she/they read the book'

(iv) yonda
read-past
'He/she/they read it/them

Task 1: This phenomenon presents some apparent problems for the ARP, as we have formulated it. For purposes of this problem, assume that Japanese will require a somewhat different version of the ARP – possibly the weakened version introduced in Problem 2. With that assumption, show how the phenomenon illustrated in (ii)-(iv) could be handled. You will probably want to have separate lexical entries for those verbs whose arguments are missing. (Don't worry now about the lexical redundancy this appears to cause). Be explicit about the verbs' values for SPR, COMPS, and ARG-ST. How will you account for the various possible pronominal interpretations for unexpressed elements?

Now consider the following data:

(v) Hanako ga zibun o tatai-ta.
Hanako NOM pro-self ACC hit-PST
'Hanako$_i$ hit herself$_i$/*him$_j$/*her$_j$/*them$_j$.'

(vi) Taroo ga tatai-ta.
Taroo NOM hit-PST
'Taroo$_i$ hit *himself$_i$/him$_j$/her$_j$/them$_j$.'

Task 2: Does your analysis predict these data? If so, explain how. If not, modify your analysis to take these data into account.

Now consider some more data[7]

(vii) Taroo ga Hanako ga zibun o tatai-ta koto o hookokusi-ta.
Taroo NOM Hanako NOM pro-self ACC hit-PST that ACC tell-PST
'Taroo$_i$ said that Hanako$_j$ hit herself$_j$/him$_i$/*them$_k$/*him$_k$/*her$_k$.'

[7]As the word-by-word glosses indicate, *zibun* is a reflexive. However, because the Japanese binding principles are not quite the same as those for English, we sometimes have to translate it with a non-reflexive English pronoun.

(viii) Taroo ga Hanako ga tatai-ta koto o hookokusi-ta.
Taroo NOM Hanako NOM hit-PST that ACC tell-PST
'Taroo$_i$ said that Hanako$_j$ hit *herself$_j$/him$_{i,k}$/her$_k$/them$_k$.'

Here you may assume that the verb *hookokusi-ta* takes an S complement, much as the verb *say* does in English (though the verb in Japanese always follows the complements, no matter what they are).

Task 3: On the basis of these facts, how would you characterize the difference between the binding theory we need for Japanese and the one we have developed for English? Be sure to be explicit in your discussion of both principle A and principle B. (You may want to construct analogous examples involving reflexive and nonreflexive pronouns in the S complements of *say* in English, to see the difference more clearly).

7.5 The Argument Realization Principle Revisited

ARG-ST lists in general, and the ARP in particular, will play an increasingly important role in the chapters to come. We will place various constraints on the ARG-ST values of particular kinds of words, yet these would be vacuous without the ARP, which relates ARG-ST values to the values of the VAL features SPR and COMPS. This connection is central, if the constraints we place on lexical heads are really to interact with the elements that heads combine with syntactically. The binding theory presented in this chapter illustrates the importance of both ARG-ST and the ARP in our theory.

A further motivation for separating out ARG-ST in our lexical entries is that fact that the syntactic environments in which a verb may appear are determined largely (though not entirely) by its meaning. Attempts to formulate what is sometimes called a LINKING THEORY – that is, a set of principles to derive the syntactic properties of verbs from their meanings – have typically mapped some sort of semantic structure onto a single syntactic argument structure. Our approach is similar, except that our notion of ARG-ST is not purely syntactic, as it a list of feature structures, each of which contains both syntactic and semantic information.

In Chapter 6, the Head Feature Principle, Valence Principle, and Semantics Principle were all conceptualized (and formalized) as part of the definition of a well-formed structural description: SDs have to conform not only to the constraints defined by our rules and lexical entries, but also to the constraints embodied in our principles. The ARP can be conceived of in the same way – that is, as a constraint on the relation between lexical entries and lexical SDs. We would like our lexical entries not to have to specify values for VAL. Rather, we want the ARG-ST values in the lexical entries, together with the ARP to determine the VAL values in lexical SDs. That is, a lexical SD is well-formed only if its VAL values add up to its ARG-ST value.

This idea is easy to incorporate into the formal definition of well-formed SD given in the appendix to Chapter 6.

(45) Lexical SDs (revised):

$$\begin{array}{c}\Delta\\ |\\ \omega\end{array}$$

is a well-formed lexical SD just in case $\langle \omega, D \rangle$ is a lexical entry and Δ is the unification of D with the following feature structure description:

$$\left[\begin{array}{ll}\text{SYN} & \left[\text{VAL}\left[\begin{array}{ll}\text{SPR} & \boxed{1}\\ \text{COMPS} & \boxed{2}\end{array}\right]\right]\\ \text{ARG-ST} & \boxed{1} \oplus \boxed{2}\end{array}\right]$$

Now suppose a lexical entry specified nothing about VAL, e.g. as in (46).

(46)
$$\left\langle \text{loves}\ ,\ \left[\begin{array}{ll}\text{SYN} & \left[\text{HEAD}\left[\textit{verb}\right]\right]\\ \text{ARG-ST} & \langle\ \text{NP}_i\left[\text{AGR}\ \textit{3sing}\right],\ \text{NP}_j\ \rangle\\ \text{SEM} & \left[\begin{array}{ll}\text{MODE} & \text{prop}\\ \text{INDEX} & \text{s}\\ \text{RESTR} & \left\langle\left[\begin{array}{ll}\text{RELN} & \text{love}\\ \text{LOVER} & \text{i}\\ \text{LOVED} & \text{j}\end{array}\right]\right\rangle\end{array}\right]\end{array}\right]\right\rangle$$

In conjunction with the ARP, this will give rise to the well-formed lexical SD in (47), which includes a value for VAL.

(47)
$$\left[\begin{array}{ll}
\text{SYN} & \left[\begin{array}{ll}\text{HEAD} & \left[\textit{verb}\right] \\ \text{VAL} & \left[\begin{array}{ll}\text{SPR} & \langle \boxed{3} \rangle \\ \text{COMPS} & \langle \boxed{4} \rangle\end{array}\right]\end{array}\right] \\
\text{ARG-ST} & \langle\ \boxed{3}\text{NP}_i\left[\text{AGR } \textit{3sing}\right],\ \boxed{4}\text{NP}_j\ \rangle \\
\text{SEM} & \left[\begin{array}{ll}\text{MODE} & \text{prop} \\ \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \left[\begin{array}{ll}\text{RELN} & \text{love} \\ \text{LOVER} & \text{i} \\ \text{LOVED} & \text{j}\end{array}\right] \right\rangle\end{array}\right]
\end{array}\right]$$

$$\text{loves}$$

Given what we have said so far, (47) is not the only lexical SD that would be licensed by (46). The ARP would also allow a lexical SD in which both of the elements of the argument structure list in (46) were identified with complements – that, both $\boxed{3}$ and $\boxed{4}$ occur on the COMPS list. Similarly, they could both be on the SPR list. Such lexical SDs will need to be ruled out. In the next chapter, we introduce a constraint requiring verbs to have exactly one element on their SPR lists. This will ensure that (47) is the only lexical SD licensed by (46).

7.6 Further Reading

The binding of anaphors has been the topic of an extensive literature since the late 1960s. A seminal and very readable paper was Lasnik (1976). To our knowledge, the first proposal to treat reflexive binding in terms of a hierarchy of the verb's arguments was made by Johnson (1978). The Binding Theory of Chomsky (1981) distilled many of the insights of the research of the preceding decade into three principles; this theory was developed further in a number of works within the Government and Binding theory of grammar. A detailed account of binding within Lexical Functional Grammar is presented by Dalrymple (1993). The theory of binding presented in this chapter is based on Pollard and Sag (1992, 1994). One of the most detailed attempts to date at formulating a linking theory is Davis (1996), whose theory of the alignment of semantics and argument structure would allow a further streamlining of all our lexical descriptions.

Chapter 8

The Structure of the Lexicon

8.1 Introduction

Several chapters ago, we began with context-free grammar and motivated numerous modifications to develop a formalism better tailored to the description of natural language phenomena. In the course of doing this, we have put more and more of the descriptive burden of our theory into the lexicon. Lexical entries have evolved from simple pairings of phonological forms with grammatical categories into elaborate information structures, in which phonological forms are now paired with more articulated feature structure (descriptions). This has permitted us to reduce our inventory of grammatical rules to a few very general schemata and to account for a range of syntactic phenomena.

Since our theory relies heavily on rich lexical representations, we need to consider what kind of internal organization the lexicon might have. In particular, we do not want to claim that all information contained in lexical entries is simply stipulated. A great deal of what we are now putting into lexical entries is predictable, so stipulating all of it would miss linguistically significant generalizations. For example, we handle subject-verb agreement by having the lexical entries for certain verbs (basically, those with the suffix *-s*) specify that their SPR values have to be third person and singular. Aside from that specification, these entries are essentially identical to those for other forms of what is intuitively the same verb: their part of speech is *verb*; they have the same COMPS values; and their semantics is virtually the same. This is no accident, nor is the fact that the same suffix is used to mark almost all third-person singular present tense verb forms.

Notice, by the way, that capturing such generalizations is motivated not only by general considerations of parsimony, but also by psychological considerations. On encountering a novel English verb (say, a recent coinage such as *email* or an obscure word like *cark*), any competent speaker will add the suffix *-s* when using it in the present tense with a third-person singular subject. In short, speakers know that there are systematic relationships among different forms of the same word, and our grammar should reflect this fact. The focus of the present chapter is to develop mechanisms for expressing regularities within the lexicon.

8.2 Lexical Types

One mechanism for reducing the redundancy of the lexicon has already been introduced: a hierarchy of types. Our initial motivation for introducing types was to define feature appropriateness, i.e. to avoid having to specify values for features that are irrelevant to particular classes (such as COUNT for prepositions). Later, we also used types for stating constraints on feature structures. In particular, in Chapter 5, we motivated the 'Nominal SPR Agreement' (NSA) constraint, which required unification of any noun's AGR specification with the AGR specification of an element on its SPR list – that is, nouns have to agree with their specifiers.

We also introduced the idea that some types are subtypes of others, in the following sense:

(1) If T_2 is a subtype of T_1, then

a. every feature specified as appropriate for T_1 is also appropriate for T_2, and

b. every constraint associated with T_1 affects all instances of T_2.

Formulated in this way, the inheritance of constraints in our type hierarchy is MONOTONIC, i.e. constraints on supertypes affect all instances of subtypes, without exception. An intuitive alternative to this conception is to allow for DEFAULT INHERITANCE of constraints, according to which contradictory information associated with a subtype takes precedence over (or OVERRIDES) constraints that would otherwise be inherited from a supertype. Default inheritance allows a type system to express the idea that language embodies general constraints with exceptions – subclasses with idiosyncratic constraints that violate the general ones.

This intuitive idea is quite simple to express: we need only allow constraints associated with lexical types to be marked as defaults. Suppose a constraint associated with a lexical type T_i is marked as a default. Then this constraint holds of any lexical entry of type T_i for which it is not explicitly overridden. It can be overridden in either of two ways. First, a subtype of T_i might have a constraint associated with it that contradicts the default. That is, there could be a type T_j that is a subtype of T_i and has a constraint associated with it that is incompatible with the default. Then the constraint on T_j takes precedence. Second, the information stipulated in a particular lexical entry might contradict the default. That is, a particular instance of a leaf type T_k (T_k a subtype of T_i) could contradict the default. In this case, too, the information associated with the lexical entry takes precedence over the default. But the default constraint is true of all instances of subtypes of T_i in which it is not overridden (as are all non-default constraints associated with those types).

Natural languages exhibit a great many regularities that have exceptions that are naturally modelled in type hierarchies of this sort. Returning to the example of agreement between nouns and specifiers, NSA was originally motivated by considering what are standardly referred to as COMMON NOUNS, specifically, the noun *dog*. Names, which are often called PROPER NOUNS, are, for the most part, irrelevant to NSA because they do not take specifiers. This is illustrated in (2)

(2) a. Cameron skates.

b. *A/*The Cameron skates.

Moreover, proper nouns are normally third-person singular.

(3) *Cameron skate.

These generalizations will be captured in our type system by introducing a type for proper nouns with the constraint (stated more formally below) specifying that the value of AGR must be of type *3sing* and that the ARG-ST (and hence both SPR and COMPS lists) must be the empty list. But there are exceptions to this constraint. In particular, there are several proper nouns in English naming mountain ranges that appear only in the plural, and only with a determiner.

(4) a. The Andes are magnificent.

b. *The Ande is magnificent.

c. Hannibal crossed the Alps.

d. *Hannibal crossed Alps.

This is a typical situation: many broad and productive generalizations in languages have exceptions, either idiosyncratic lexical items or exceptional subclasses thereof. For this reason, we shall adopt the method of default constraint inheritance in our type hierarchy. This will allow us both to restrict the number of types that are required and to keep our constraints simple, without precluding the possibility that some instances or subtypes might be exceptions to the constraints.

Problem 1: Mountains and NSA Does NSA apply to the lexical entries for proper nouns like *Alps, Andes*, and *Himalayas*? Explain your answer and provide data to justify it.

The lexicon itself can be treated in terms of a type hierarchy with *lexical-item* (*lex-item*) as the most general type of feature structure. We posit two immediate subtypes of *lex-item*, which we call *word* and *lexeme*. These two types correspond to what, in everyday English, are two different uses of the term 'word'. In some contexts, people informally distinguish, e.g. *runs* and *ran* as two different words: they are pronounced differently, have (subtly) different meanings, and have slightly different cooccurrence restrictions. But in other contexts, the same people would have no hesitation in referring to *runs* and *ran* as two forms of the word *run*. Clearly, these are two very different conceptions of a word: the first refers to a certain pairing of sounds and meanings, whereas the latter refers to a family of such pairings. In a formal theory of grammar, these two concepts must not be conflated. Our type *word* corresponds to the first usage (in which *runs* and *ran* are distinct words). The lexical entries that give rise to lexical SDs must all be of type *word*.

But we also want to capture what people have in mind when they use 'word' in the second sense. That is, we want to be able to express the relationship between *runs* and *ran* (and *run* and *running*). We do this by means of the type *lexeme*. A *lexeme* can be thought of as an abstract proto-word which, by means to be discussed in this chapter, gives rise to the actual words (that is, elements of type *word*).

Each word or lexeme of the language belongs to some leaf type. The various intermediate types then correspond to intermediate levels of classification where type constraints can express linguistic generalizations. Since this is a default inheritance hierarchy, we can provide a natural account of the fact that lexemes have many properties in common, but may differ from one another in terms of particular constraints that override the general constraints governing their supertypes.

The most basic lexical entries then are pairings of a form and a lexeme, which give rise to a set of word entries.[1] Lexemes thus serve as the atoms from which all linguistic descriptions are built. In the next section (and much of the remainder of this book) we will discuss lexical rules, the mechanism we employ to derive words from lexemes.

We will assume that *lexeme* and *word* are the two immediate subtypes of the type *lex-item*. Among lexemes, we draw a further distinction between those that give rise to a set of inflected forms and those that do not inflect. That is, we posit *inflecting-lexeme* (*infl-lxm*) and *constant-lexeme* (*const-lxm*) as two subtypes of *lexeme*.[2] This gives rise to the type hierarchy sketched in (5).

(5)
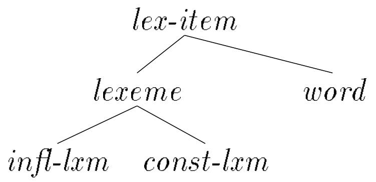

Many of the constraints we present here are stated solely in terms of ARG-ST specifications. These constraints affect lexemes directly, but they also indirectly affect words derived from those lexemes, as the lexical rules we will propose have the effect of preserving ARG-ST specifications. Because of the ARP, developed in the last chapter, the SPR and COMPS values in SDs formed from these word descriptions will be appropriately constrained as well.

To see how these constraints work, consider first the following constraint on the type *lexeme*:

(6)
$$lexeme: \begin{bmatrix} \text{SYN} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{ANA} \;\; / - \end{bmatrix} \end{bmatrix} \\ \text{SEM} \begin{bmatrix} \text{MODE} \;\; / \text{ prop} \end{bmatrix} \end{bmatrix}$$

Here we use the symbol '/' to indicate that a certain constraint is a default and hence can be overridden by some conflicting specification.[3] This type constraint thus makes it a general,

[1]More precisely, basic lexical entries are a pair consisting of a form and a lexeme description; these are used to derive lexical entries that pair a form with a word description. For ease of exposition, we will sometimes blur the distinction between 'lexeme', 'lexeme description', and 'lexical entry pairing a form with a lexeme description'. A similar blur will sometimes occur between 'word', 'word description', and 'lexical entry pairing a form and a word description'.

[2]We assume that even non-inflected words are derived from lexemes. Not much hangs on the choice to formalize our lexicon in this way, rather than to enter such words directly into the lexicon, with no corresponding lexemes.

[3]This notation for defaults is taken from Lascarides et al. (1996).

but DEFEASIBLE (that is, overrideable) property of lexemes that they have a propositional semantics and are not anaphors. These properties will then be inherited by all instances of the subtypes of *lexeme* except when they are overridden, as described above. These properties will also be inherited by words derived by some lexical rule (unless the lexical rule in question explicitly changes the default specification).

We will further classify inflecting lexemes in terms of the subtypes *noun-lexeme* (*noun-lxm*) and *verb-lexeme* (*verb-lxm*), with the former being further subclassified in terms of the types *proper-noun-lexeme* (*pn-lxm*) and *common-noun-lexeme* (*cn-lxm*), as shown in (7).

(7)

- *infl-lxm*
  - *noun-lxm*
    - *pn-lxm*
    - *cn-lxm*
  - *verb-lxm*

These types are associated with various constraints that need to be explicated.

First, noun lexemes are governed by the constraint in (8):

(8)
$$\textit{noun-lxm}: \begin{bmatrix} \text{SYN} & \begin{bmatrix}\text{HEAD } \textit{noun}\end{bmatrix} \\ \text{ARG-ST} & / \langle\ \rangle \\ \text{SEM} & / \begin{bmatrix}\text{MODE ref}\end{bmatrix} \end{bmatrix}$$

(8) ensures that all nouns, whether they are proper nouns or common nouns, have a HEAD value of the type *noun*, and hence bear specifications for CASE (a feature appropriate for *noun* objects, but not, say, for *verb* objects). The '–' default specification for ANA is inherited by instances of *noun-lxm*, but this value will be overridden in the entries for reflexive pronouns and reciprocals. (8) also states that nouns by default have an empty ARG-ST list, i.e. that they take no specifiers or complements.[4] Common nouns will override this constraint, a matter we turn to directly. Finally, (8) also specifies the semantic mode of nouns as 'ref' (by default) and hence overrides the MODE specification 'prop' that is associated with the superordinate type *lexeme*. This SEM specification is itself a default for reasons that will become clear in chapter 11.

The subtypes of *noun-lxm* are associated with the following constraints:[5]

(9) a. $\textit{pron-lxm}: \begin{bmatrix} ... \end{bmatrix}$

b. $\textit{pn-lxm}: \begin{bmatrix}\text{SYN} \begin{bmatrix}\text{HEAD} \begin{bmatrix}\text{AGR } \textit{3sing}\end{bmatrix}\end{bmatrix}\end{bmatrix}$

[4]This particular constraint is overridden by subordinate types not discussed here (e.g. types we might use for nouns like *picture*, that select for PPs headed by *of*) or by nouns with idiosyncratic complement selection properties of one sort or another (e.g. *difference*, which selects an optional PP argument headed by *between*).

[5]The notation '[ ]' is used to indicate the obligatory presence of a synsem-struc, whose features are left unspecified.

c.
$$\textit{cn-lxm}: \begin{bmatrix} \text{SYN} & \begin{bmatrix}\text{HEAD} & \begin{bmatrix}\text{AGR}\ \boxed{1}\end{bmatrix}\\ \text{VAL} & \begin{bmatrix}\text{SPR}\ /\ \langle [\ ]\ \rangle\end{bmatrix}\end{bmatrix}\\ \\ \text{ARG-ST} & /\left\langle \begin{array}{c}\text{DetP}\\ \begin{bmatrix}\text{AGR}\ \boxed{1}\end{bmatrix}\end{array}\right\rangle \end{bmatrix}$$

The type *pron-lxm* (used for personal pronouns and reflexives) has no constraints of relevance associated with it. The constraint in (9b) is a formal statement of the generalization stated earlier: that proper nouns take third-singular AGR values and that they do not select for a specifier. The latter constraint follows from the empty ARG-ST specification inherited from the supertype *noun-lxm*. (9c) requires that common nouns by default select agreeing DetPs as their specifier, thus incorporating NSA and overriding the default empty ARG-ST specification. Further subtypes of *cn-lxm* might be introduced to distinguish count and mass nouns,[6] constraining them to select for an appropriate ([COUNT +] or [COUNT −]) specifier. Note that (9) need not identify the member of the SPR list with the sole member of the ARG-ST list, as this unification (in any lexical SDs that lexemes of this type will give rise to) is already a consequence of the ARP.[7]

In consequence of the type hierarchy and constraints just outlined, the rather complex set of specifications that we want to associate with a particular lexeme can in the main be determined simply by associating it with the appropriate type in the lexical hierarchy. For example, from the simple fact that the lexical specification for *book* says it is assigned to the type *cn-lxm*, it follows that that the lexeme *book* is associated with all the constraints shown in (10).

(10)
$$\begin{bmatrix} \textit{cn-lxm} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD} & \begin{bmatrix}\textit{noun}\\ \text{ANA}\ -\\ \text{AGR}\ \boxed{1}\end{bmatrix}\\ \text{VAL} & \begin{bmatrix}\text{SPR}\ \langle [\ ]\ \rangle\end{bmatrix}\end{bmatrix}\\ \text{ARG-ST} & \left\langle \begin{array}{c}\text{DetP}\\ \begin{bmatrix}\text{AGR}\ \boxed{1}\end{bmatrix}\end{array}\right\rangle \\ \text{SEM} & \begin{bmatrix}\text{MODE ref}\end{bmatrix} \end{bmatrix}$$

All that remains to be stipulated about this lexeme is its phonological form and the particular predication in its semantic restriction – i.e. its form and its idiosyncratic meaning. The rest of its grammatical properties follow from the fact that it is assigned to a particular type to

[6] The distinction between count and mass nouns was discussed in chapter 5.

[7] The alert reader may have noticed that (9b)'s claim that specifiers are obligatory with common nouns appears to be inconsistent with the existence of plural and mass NPs that lack determiners. The analysis of such NPs is the topic of Problem 3 below.

which it is assigned in our system of grammatical description. This is precisely what lexical stipulation should be reduced to, whenever possible.

We require that all instances of the type *verb-lxm* must have both a HEAD value of type *verb* and a SPR list that contains one element. In addition, the argument structure of a lexeme of this type by default begins with an NP.[8] This information is specified in the following constraint on the type *verb-lxm*:

(11)
$$\textit{verb-lxm}: \left[\begin{array}{ll} \text{SYN} & \left[\begin{array}{ll} \text{HEAD} & \textit{verb} \\ \text{VAL} & \left[\text{SPR} \ \langle\ [\ ]\ \rangle\right] \end{array}\right] \\ \text{ARG-ST} & /\ \langle\ \text{NP}, \ldots\ \rangle \end{array}\right]$$

The various subtypes of *verb-lxm*, distinguished by their ARG-ST specifications, are organized as shown in (12):

(12)
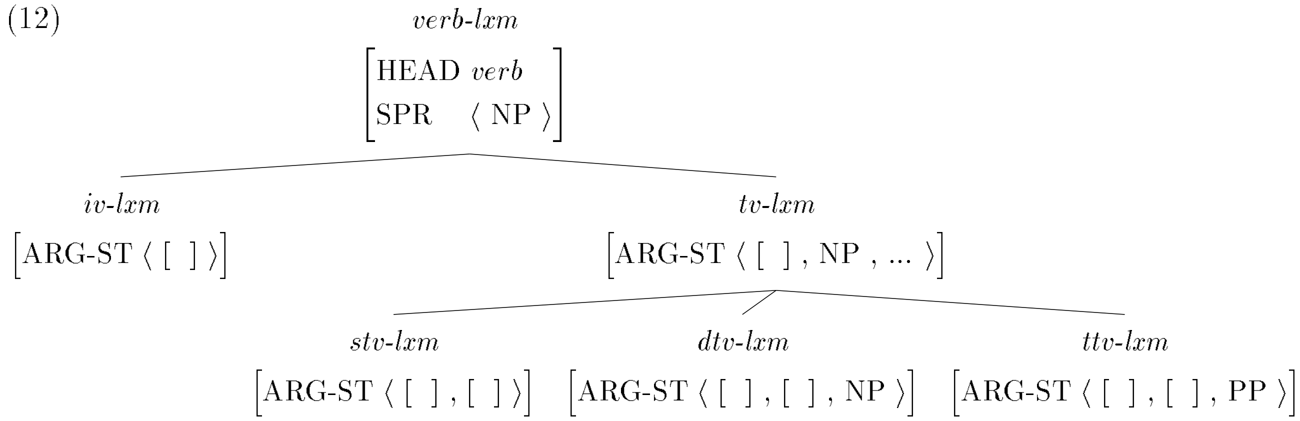

Here we have introduced the types *intransitive-verb-lexeme* (*iv-lxm*), *transitive-verb-lexeme* (*tv-lxm*) at the first level of subhierarchy. The transitive verb lexemes are then further classified in terms of the subtypes *strict-transitive-verb-lexeme* (*stv-lxm*), *ditransitive-verb-lexeme* (*dtv-lxm*), and *to-transitive-verb-lexeme* (*ttv-lxm*). As before, these types and their associated constraints allow us to eliminate lexical stipulation in favor of type-based inference. Thus from the simple statement that *give* belongs to the type *dtv-lxm*, we may infer that it has all the properties shown in (13):

(13)
$$\left[\begin{array}{ll} \textit{dtv-lxm} & \\ \text{SYN} & \left[\begin{array}{ll} \text{HEAD} & \left[\begin{array}{l} \textit{verb} \\ \text{ANA} - \end{array}\right] \\ \text{VAL} & \left[\text{SPR} \ \langle\ [\ ]\ \rangle\right] \end{array}\right] \\ \text{ARG-ST} & \langle\ \text{NP}\ ,\ \text{NP}\ ,\ \text{NP}\ \rangle \\ \text{SEM} & \left[\text{MODE prop}\right] \end{array}\right]$$

[8]In chapter 11, we discuss verbs that take non-NP subjects, and hence override this default.

Because all instances of *verb-lxm* are specified as [SPR ⟨[ ]⟩], words formed from these lexemes will also be so specified. This fact in turn interacts with the ARP to ensure that in verbal word structures, the first member of the ARG-ST list must also be on the SPR list. Without this [SPR ⟨ ⟩] constraint, there would be another way to satisfy the ARP, namely with an empty SPR list, and all the arguments on the COMPS list. As noted in the final section of the last chapter, this consequence must be avoided, as it is in the present analysis.

Let us turn now to non-inflecting lexemes, i.e. to the various subtypes of the type *const-lxm* illustrated in (14):

(14)

- *const-lxm*
  - *prep-lxm*
    - *pred-p-lxm*
    - *mark-p-lxm*
  - *adj-lxm*
  - *conj-lxm*
  - ...

These correspond to the various kinds of lexeme that do not undergo inflectional rules in English.[9] Here the types *preposition-lexeme*, *adjective-lexeme*, and *conjunction-lexeme* have their obvious abbreviations. Once again, we specify general constraints on elements of each subtype of *const-lxm*. For example, the general properties of prepositions are expressed in terms of the following type constraint:

(15) *prep-lxm*: $\left[\text{SYN}\left[\text{HEAD prep}\right]\right]$

In chapter 7 we distinguished between two kinds of prepositions – those that function as predicates and those that serve as argument markers. This distinction corresponds to the two types *predicational-preposition-lexeme* (*pdp-lxm*) and *marking-preposition-lexeme* (*mkp-lxm*) in (14). Recall that in our earlier discussion we distinguished these prepositions in terms of both their HEAD value and their VALENCE value. We may now incorporate this distinction into our lexicon by attaching the following type constraints to the two subtypes of *prep-lxm*:

(16) a.

$$\textit{pdp-lxm}: \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{MOD} & \{\text{NOM, VP, none}\} \\ \text{P-OBJ} & \text{none} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} \langle\ [\ ]\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} \langle\ \text{NP}\ ,\ \text{NP}\ \rangle & \end{bmatrix}$$

b.

$$\textit{mkp-lxm}: \begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{MOD} & \text{none} \\ \text{P-OBJ} & \boxed{1} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} \langle\ \ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} \langle\ \boxed{1}\ \rangle & \end{bmatrix}$$

[9] The type *adj-lxm* arguably should be classified as a subtype of *infl-lxm*, rather than as a subtype of *const-lxm*, in light of the fact that many adjectival lexemes give rise to comparative and superlative forms, e.g. *tall*, *taller*, *tallest*. We will not pursue this matter here.

The constraint in (16a) allows predicational prepositions to function as modifiers of nominal phrases or VPs, but it also allows for phrases that modify nothing, as is the case when such a PP occurs as the complement of a verb like the ones discussed in chapter 7:[10]

(17) I wrapped the blanket [around me].

Note that the first argument of a predicational preposition must be identified with the SPR element, in accordance with the ARP. What is functioning in these roles in (17) is the NP *the blanket*, which is also an argument of the verb *wrapped*. This is the first time we have seen one constituent serving as an argument of more than one predicate at the same time. This is a common phenomenon, however, as we will see in subsequent chapters. It is discussed in some detail in Chapter 12.

The argument marking prepositions, because of the constraint in (16b), project a non-modifying, [SPR $\langle$ NP $\rangle$] PP whose P-OBJ value is identified with the category of the preposition's NP object:

(18) He talks [to himself].

As described in chapter 7, the P-OBJ value of such a VP is incorporated into the ARG-ST list of the higher verb (e.g. *talks*), where the binding constraints must be satisfied. Finally, recall that some prepositions, e.g. *around* behave either as predicational or as argument marking. Hence the following example is also well-formed.

(19) I wrapped the blanket [around myself].

This pattern of optional reflexivization is now neatly accounted for by allowing *around* to live a double life – as either a predicational or argument marking preposition.

Some readers may have noticed that our type hierarchy posits two distinct types corresponding roughly to each of the traditional parts of speech. In addition to *noun, verb, adjective, preposition*, etc. – the subtypes of *pos* introduced in Chapter 4 – we now have *noun-lxm, verb-lxm, adj-lxm, prep-lxm*, etc., which are subtypes of *lexeme*. A careful examination of the way we use these two sets of types reveals that they serve rather different functions in our grammar. The subtypes of *pos* exist in order to specify what features are appropriate for what categories of words and phrases. These distinctions manifest themselves in the HEAD features, so that the *pos* subtypes show up in our feature structures as specifications of what type of HEAD value we have. The subtypes of *lexeme*, on the other hand, are used to introduce constraints on what combinations of feature values are possible, e.g. the NSA or the constraint that verbs take NP subjects (SPRs). These typically involve argument structure (and/or valence features), as well as HEAD features. Consequently, the *pos* subtypes (*noun, verb*, etc.) frequently show up inside of the constraints associated with the *lexeme* subtypes (*noun-lxm, verb-lxm*, etc.).

[10]Note in addition that nothing in our analysis blocks the projection of subject-saturated PPs like *[My blanket [around me]]*. In English, these occur only in restricted circumstances, e.g. as 'absolutive' clauses:

(i) [My blanket around me], I was ready to curl up for the night.

The type hierarchy serves to simplify our descriptions in two ways: it saves us from having to assign values to features where they do no work, e.g. PER (person) in prepositions; and it allows us to stipulate common combinations of feature values only once, using (default) inheritance to account for their distribution. Our hierarchy has two sets of types corresponding roughly to the traditional parts of speech because the hierarchy is serving these two distinct functions.

Problem 2: Parts of Speech and Types What would happen if we tried to eliminate the *pos* subtypes like *noun* and *verb*? To answer this, you will need to consider where the features currently associated with the *pos* subtypes would have to be declared, and what consequences this would have for our feature structures. Be explicit.

Up to this point, we have made no mention of CASE specifications. In fact, NPs in English are accusative except when they are the subject of a finite verb form. One might think this is a constraint on lexemes, but in fact the various lexical rules we will formulate, e.g. passive, never have the effect of transferring accusative case to their first argument (the subject), as we will see in chapter 10. For this reason, we will treat the assignment of accusative case as a fact about words, not about lexemes. The easiest way to do this is to add the following constraint to our definition of well-formed lexical SDs:

(20) Case Constraint:

An NP in a noninitial position of a word's ARG-ST list is [CASE acc].

Note that this principle allows us to keep our constraints on verbal lexemes just as we formulated them above: they need make no mention of case.[11]

8.3 Lexical Rules

The lexical rule is a mechanism for further reducing redundancy and stipulation in the lexicon by using information in one lexical entry as the basis for generating another lexical entry. The idea is that the lexicon contains two types of entries: basic ones (lexemes) and those that are 'based on' or 'derived from' the basic ones. The lexical rules are used for deriving predictably related lexical entries, such as inflected forms of verbs and nouns.

It is traditional to think of words (or at least certain kinds of words) as being built up from smaller units through the addition of affixes. We have followed this tradition by using our notion of type to distinguish *lexeme* from *word*. For most nouns and verbs, we will assume that there is only one basic lexical entry, which is of type *lexeme*. We then derive all the nominal and verbal *word*s by applying lexical rules to *lexemes*.

Lexical rules have the general form $X \Rightarrow Y$. The intended interpretation of this is that for any lexical entry satisfying the description in X, there is another lexical entry satisfying the description in Y. The input and output entries are identical, except in those ways that the rule specifies. If a feature F_1 is given conflicting values in X and Y, it means that input

[11]However, argument marking prepositions, since they have only one NP argument, which hence falls outside the scope of the Case Constraint, must specify lexemically that that argument is [CASE acc].

and output entries must have different values for F_1. If X specifies a value v_2 for a feature F_2, but Y says nothing about F_2, then the rule applies only to entries whose value for F_2 is v_2; in this case, both input and output entries have the same value for F_2 (namely v_2). If Y specifies that feature F_3 has value v_3 but X says nothing about F_3, the situation is slightly more complex. In this case, the output's value for F_3 is the unification of v_3 with the input's value for F_3. Thus, if the input's value for F_3 is compatible with v_3 (and this includes the case where the input does not even mention F_3), the output's value for F_3 consists of both the input's value and v_3 (more precisely, their unification); but if the input's value for F_3 is incompatible with v_3, then there is no output (that is, the rule does not apply). A lexical rule applies to an input entry and produces as output a new entity whose (morphological) form, syntactic category and meaning are systematically related to the input. Lexical rules often add morphological endings in English, but not always.

Our first example of a lexical rule in fact does nothing to the phonological form of the base; but it adds a restriction requiring that the specifier's AGR value is of type *3sing*, and hence [PER 3rd] and [NUM sg], as described in chapter 5.

(21) Singular Noun Lexical Rule:

$$\left\langle \boxed{1}, \begin{bmatrix} \textit{cn-lxm} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \boxed{1}, \begin{bmatrix} \textit{word} \\ \text{SYN} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{AGR } \textit{3sing} \end{bmatrix} \end{bmatrix} \end{bmatrix} \right\rangle$$

This rule produces lexical entries containing singular forms of both count and mass nouns without mentioning anything about their form or meaning. Notice that the rule output only specifies the AGR value of the noun itself. But by NSA (which we incorporated into the constraints on the type *cn-lxm*, as discussed in the previous section), this is identified with the AGR value of the noun's specifier. An output of the rule in (21) such as the one illustrated in (22), must thus combine with a singular determiner.

(22)
$$\left\langle \text{book}\ , \begin{bmatrix} \textit{word} \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} \\ \text{ANA}\ - \\ \text{AGR}\ \boxed{1}\textit{3sing} \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \overset{\text{DetP}}{\begin{bmatrix} \text{COUNT} & + \\ \text{AGR} & \boxed{1} \end{bmatrix}} \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{ref} \\ \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{book} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

Now consider the closely related rule that maps nominal lexemes into lexical entries for their corresponding plural forms:

(23) Plural Noun Lexical Rule:

$$\left\langle \boxed{1}\ , \begin{bmatrix} \textit{cn-lxm} \\ \text{ARG-ST}\ \langle\ \boxed{3}[\text{COUNT}\ +]\ \rangle \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{NPL}(\boxed{1})\ , \begin{bmatrix} \textit{word} \\ \text{SYN} & \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{AGR} \begin{bmatrix} \text{NUM pl} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST}\ \langle\ \boxed{3}\ \rangle \end{bmatrix} \right\rangle$$

Here, F_{NPL} is a morphological function that applies to nominal bases in English, giving their plural forms. In most cases, it suffixes an *-s*, though its full specification would also include a listing of irregular plurals. Thus, F_{NPL}(*cat*) = *cats*, but F_{NPL}(*child*) = *children*. The rule says that, for every nominal lexeme that takes a [COUNT +] specifier, there is a corresponding lexical entry for a plural noun whose form is dictated by the function F_{NPL} and which takes a plural specifier. The requirement that the input take a [COUNT +] specifier keeps the rule from applying to mass nouns like *furniture*, so that there is no word **furnitures*.

A complete formulation of both these lexical rules would require the introduction of a fundamental difference between the semantics of singular and plural nouns. Since the Singular Noun Lexical Rule must affect mass and count nouns alike, it is likely that the singular semantics should be built into the lexical entry for the noun lexeme, with only the Plural Noun Lexical Rule introducing a change in the value of SEM. The details of such a treatment, however, are beyond the scope of this book.

Problem 3: Plural and Mass NPs Without Specifiers There is a problem with our treatment of common nouns. The type *cn-lxm* requires common nouns to have non-empty SPR lists, and this requirement is preserved in the Plural Noun Lexical Rule above. But specifiers are optional for plural nouns and mass nouns.

A. Give examples showing the optionality of specifiers for plural and mass nouns.

Two obvious approaches to dealing this problem are the following:

(i) Allow empty SPR lists in the lexical entries for plural and mass nouns

(ii) Introduce a new grammar rule to account for NPs with plural or mass heads and no specifiers.

Alternative (i) would involve modifying the Plural Noun Lexical Rule, as well as introducing a new subtype of *cn-lxm* for mass nouns. The rule in alternative (ii) would be analogous to the Imperative Rule given in Chapter 7, in that it would have only one constituent on the right hand side, and its function would be to license a constituent without a specifier, although its head has a non-empty SPR list.

B. Show how alternative (i) runs into trouble with the Argument Realization Principle. [Hint: Find examples of common nouns with complements but no specifiers.]

C. Formulate the rule required for alternative (ii). Be as precise as you can. [Hint: The trickiest part is stating the disjunction of plural or mass nouns. The place to do this is in the SPR list of the SD on the right hand side of the rule.]

We posit additional lexical rules to generate entries for the various inflected forms of verbs. Since the different forms have different distributions, it is useful to have a feature whose value distinguishes among them. We have already seen this feature, which we call FORM, in the analysis of imperatives presented in the previous chapter. Later, we will use this same feature (with different values) to distinguish among prepositions and to mark certain fixed expressions. For now, we will identify just the values of FORM that correspond to verbal inflectional categories for which we need to introduce lexical rules. They are[12]

[12]Other works that have made use of a feature FORM to distinguish verbal inflectional categories have not always posited precisely these values. In particular, what we call 'inf' has often been called 'base' (with [FORM inf] being reserved for the *to* of infinitives). In addition, a value 'ger' (for 'gerund') has sometimes been proposed for a kind of word not covered here. Like present participles, gerunds are suffixed with *-ing*, but unlike present participles, gerunds head phrases that have the distribution of NPs. The occurrences of *singing* in (i)-(iii) are present participles; those in (iv)-(vi) are gerunds.

(i) The birds are singing.

(ii) Anyone singing in class will be punished.

(iii) Ashley began singing Christmas carols in October.

(iv) Ashley's singing Christmas carols in October annoyed Jordan.

(v) We denied singing during class.

(vi) Don't even think about singing!

(24) a. inf – the bare uninflected form, as in *Andy would* **eat** *rice* or *Andy tried to* **eat** *rice*;

b. imp – the bare uninflected form, as in *Be reasonable!* or *Eat something!*;

c. fin – 'finite', i.e. present or past tense, as in *Andy* **eats** *rice* or *Andy* **ate** *rice*;

d. prp – 'present participle', suffixed with *-ing*, usually following some form of *be*, as in *Andy is* **eating** *rice*;

e. psp – 'past participle' (or 'perfect participle'), the form that follows *have*, as in *Andy has* **eaten** *rice*;

f. pass – 'passive', as in *Rice was* **eaten** *by Andy* (to be discussed in Chapter 10.).

Since FORM is a HEAD feature, VP nodes carry the FORM value of their lexical heads. This permits us to specify in the ARG-ST value of certain elements that they take VP complements whose lexical heads have a particular inflection. For example, the *to* that occurs before verbs (which we will analyze in more detail in Chapter 12) takes [FORM inf] VP complements, and the auxiliary verb *have* takes [FORM psp] VP complements. Such cooccurrence restrictions will be developed in detail in subsequent chapters. We mention them here only to indicate why FORM is needed, and why it is classified as a HEAD feature.

We are now in a position to state several lexical rules for verbal inflections.

(25) 3rd Singular Verb Lexical Rule:

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix}\text{RESTR } \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{3SG}(\boxed{3}), \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD } \begin{bmatrix}\text{FORM fin}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{CASE nom}\\ \text{AGR } \textit{3sing}\end{bmatrix}, \ldots \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX } \boxed{2} \\ \text{RESTR } \boxed{1} \oplus \left\langle \begin{bmatrix} \text{RELN t-overlap} \\ \text{ARG1 } \boxed{2} \\ \text{ARG2 now}\end{bmatrix}\right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

As with the Plural Noun LR, we have glossed over the morphological component of the 3sg Verb LR by simply giving it a name, F_{3SG}. The semantic effect of this rule is clear, however – it preserves the semantics of the input, except for the value of the feature RESTRICTION, which is augmented to include a further predication requiring that the INDEX value – the

The analysis of gerunds is beyond the scope of this text. Hence, we will not consider the question of whether there should be a FORM value for gerunds.

situation described by the verb – be in the temporal-overlap (t-overlap) relation with 'now', an atomic value that will always be designate the time of an utterance. Thus, according to this rule, using a 3rd singular form of a verb lexeme imposes the requirement that the situation introduced by the verb be located in some temporal interval that overlaps the time of the utterance. This is perhaps only a first-cut at the semantics of present tense verb forms, but our goal in presenting it is to make clear how a precise semantic analysis, however refined, can be integrated into lexical rules of the sort outlined here. What the rule in (25) says, then, is that for any verbal lexeme, there is a corresponding finite verb that takes a nominative, third person singular subject. Further, the morphology and semantics of the latter verb are systematically related to those of the base verb.

Problem 4: *-s* In most cases, F_{3SG} has the same effect as F_{NPL}, namely suffixing *-s*. In fact, both suffixes have multiple pronunciations, and the conditions under which they are pronounced like *s*, like *z*, or like *uz* are identical. (They depend on phonological properties of the preceding sound). Nevertheless, these two morphological functions are not identical. Why? [Hints: (1) Remember that a function is single-valued, i.e. it specifies only one output for each input. (2) Consider elements that can be used as both nouns and verbs.]

We turn next to the rule that generates the entries for finite verbs with subjects other than third person singular NPs. Because the type distinction we have drawn between the AGR values *3sing* and *non-3sing* already distinguishes third singular NPs from all others, this rule is almost identical to the last one, as shown in (26):

(26) Non-3rd-Singular Verb Lexical Rule:

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix}\text{RESTR } \boxed{1}\end{bmatrix}\end{bmatrix}\right\rangle \Rightarrow$$

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD} \begin{bmatrix}\text{FORM fin}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{CASE} & \text{nom}\\ \text{AGR} & \textit{non-3sing}\end{bmatrix}, \ldots \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \boxed{2} \\ \text{RESTR} & \boxed{1} \oplus \left\langle \begin{bmatrix}\text{RELN} & \text{t-overlap}\\ \text{ARG1} & \boxed{2} \\ \text{ARG2} & \text{now}\end{bmatrix}\right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

No change in form is introduced, the nominative subject of the output must have a *non-3sing* AGR value (see chapter 5 for further discussion) and the situation picked out by the verb must overlap the utterance time. A typical output, illustrated in (27) can never combine with a third person singular subject and will project a sentence whose semantics includes the restrictions illustrated in addition to those that arise from the particular dependents of the verb or from modifiers.

(27)

$$
\left\langle \text{give}\ ,\ \left[\begin{array}{ll}
\textit{word} & \\
\text{SYN} & \left[\begin{array}{ll}\text{HEAD} & \left[\begin{array}{ll}\textit{verb} & \\ \text{FORM} & \text{fin} \\ \text{ANA} & -\end{array}\right]\end{array}\right] \\
\text{ARG-ST} & \left\langle \overset{\text{NP}_i}{\left[\begin{array}{ll}\text{AGR} & \textit{non-3sing} \\ \text{CASE} & \text{nom}\end{array}\right]},\ \text{NP}_j\ ,\ \text{NP}_k \right\rangle \\
\text{SEM} & \left[\begin{array}{ll}
\text{MODE} & \text{prop} \\
\text{INDEX} & \text{s}_1 \\
\text{RESTR} & \left\langle \left[\begin{array}{ll}\text{RELN} & \text{give} \\ \text{SIT} & \text{s}_1 \\ \text{GIVER} & \text{i} \\ \text{GIVEN} & \text{j} \\ \text{GIFT} & \text{k}\end{array}\right],\ \left[\begin{array}{ll}\text{RELN} & \text{t-overlap} \\ \text{SIT} & \text{s}_2 \\ \text{ARG1} & \text{s}_1 \\ \text{ARG2} & \text{now}\end{array}\right] \right\rangle
\end{array}\right]
\end{array}\right] \right\rangle
$$

The two rules just discussed create present tense forms of verbs. The next rule creates the past tense forms. English makes no distinction between singular and plural in past tense forms (aside from *was* vs. *were*);[13] hence, only one rule is needed.

---

[13] Of course, something must be said about this exception, and about the first person singular form *am*. The fact that *be* makes finer distinctions among its verb forms than other verbs does not justify making these distinctions throughout the rest of the verbal system in English. Rather, it is more parsimonious to seek a way to make *be* an exception to these lexical rules, and to stipulate the individual forms in the lexicon. The fact that there is no morphological regularity to the forms of *be* lends credibility to this approach. We will not go into the question of what kind of formal machinery to use to specify that certain lexical entries are exceptions to certain lexical rules. It would not be hard to invent a way of doing this, but we will limit ourselves to noting that some mechanism to account for lexical exceptions is necessary.

(28) Past Tense Verb Lexical Rule:

$$\left\langle \boxed{2}\ , \begin{bmatrix} \textit{lexeme} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD } \textit{verb}\end{bmatrix} \\ \text{SEM} & \begin{bmatrix}\text{RESTR } \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{PAST}(\boxed{2})\ , \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD}\begin{bmatrix}\text{FORM fin}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{CASE nom}\end{bmatrix}, \ldots. \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX } \boxed{3} \\ \text{RESTR } \boxed{1} \oplus \left\langle \begin{bmatrix} \text{RELN t-precede} \\ \text{ARG1 } \boxed{3} \\ \text{ARG2 now} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(28) posits a function F_{PAST} to account for the morphological relation between verbal lexemes and their past tense forms; in most cases, this consists of suffixing *-ed*, though there are many exceptions (such as *sleep, eat*, and *put*). Like the lexical rules for present tense verbs, (28) requires its subject to be nominative (to rule out examples like **Him died*); but unlike the present tense rules, it puts no number or person restrictions on the subject, since English past tense verbs exhibit no agreement with their subjects. The semantic effect of the rule is parallel to that of the two present tense rules, except that the relation on the restriction added is 't-precede', rather than 't-overlap'. That is, the situation denoted by the index of the verb temporally precedes the time of utterance if the verb is in the past tense. Again, this is only an approximation of the semantics of the English past tense, but it is good enough for our purposes.

Let us now turn to the imperative forms posited in the previous chapter. These are the outputs of the following lexical rule:

(29) Imperative Lexical Rule:

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix}\text{MODE} & \text{prop} \\ \text{RESTR} & \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD}\begin{bmatrix}\text{FORM imp}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{AGR}\begin{bmatrix}\text{PER 2nd}\end{bmatrix}\end{bmatrix}, ... \right\rangle \\ \text{SEM} & \begin{bmatrix}\text{MODE} & \text{dir} \\ \text{RESTR} & \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle$$

The outputs of this rule, as desired, all have a 2nd person subject (which is unspecified for number), are specified as [FORM imp], and have a directive mode in their semantics.

Finally, we will need a trivial lexical rule for non-inflecting lexemes. This rule does nothing except promote such lexemes to wordhood so that the resulting entries can give rise to word SDs.

♢**Problem 5: More LRs** Write the lexical rules that generate the infinitival (base), present participle and past participle entries for verbs in English. You should be able to specify the syntactic effects of these rules quite precisely and (for the most part) straightforwardly. For the semantics, you will need to invent some relation names (that is, some new values for the feature RELN).

♢**Problem 6: Coordination and Verbal Form** FORM values must in general be preserved in coordinate structures. That is, if a VP[FORM $\boxed{1}$] is coordinated, each conjunct must also be VP[FORM $\boxed{1}$]. Demonstrate that this is true by constructing examples with coordinate VPs of various kinds, e.g. the VP head of a finite sentence or the VP complement of verbs like *will*, *have*, or (progressive) *be*.

Problem 7: Coordination and Tense Notice that we have not introduced distinct FORM values to distinguish between past and present tense verbs – we treat both as [FORM fin]. Using the result of the preceding problem, explain why this decision is right or wrong. Be sure to consider examples where finite VPs that differ in tense are coordinated.

8.4 Derivational Rules

The lexical rules we looked at in the previous section all map lexemes into lexical entries for inflected words. Such rules are usually called inflectional rules. In contrast to these, it is also traditional to assume that languages have another kind of lexical rule that maps lexemes into lexemes. This is pretty clearly what we want for certain elements built up by prefixes and suffixes from basic lexemes. So noun lexemes like *driver* or *eater* might be derived by the following lexical rule:

(30) Agent Nominalization Lexical Rule:

$$\left\langle \boxed{2}\ , \begin{bmatrix} \textit{verb-lxm} \\ \text{ARG-ST} \ \langle \ \text{NP}_i \ , \ ... \ \rangle \\ \text{SEM} \quad \begin{bmatrix} \text{RESTR} \left\langle \begin{bmatrix} \text{R} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{-er}(\boxed{2})\ , \begin{bmatrix} \textit{cn-lxm} \\ \text{SEM} \quad \begin{bmatrix} \text{INDEX} & \text{i} \end{bmatrix} \end{bmatrix} \right\rangle$$

Here the function F_{-er} adds the appropriate suffix to the form of the rule output. The input is a verb whose subject's index is identified with the index of the nominal output. Note that the change in type from *verb-lxm* to *cn-lxm* has many side effects in terms of HEAD and ARG-ST values and in terms of the MODE value within the semantics. The RESTR value remains unchanged, however, as the information present in the input is compatible with the type constraints associated with the output type.

So the verb lexeme *drive*, whose semantics is a proposition whose RESTR value contains a driving predication with the (referent) of the verb's subject assigned the driver role, gives rise to a noun lexeme *driver*, whose index is restricted to be the driver in a driving. This noun lexeme can now undergo both our nominal lexical rules, and so we derive the singular form *driver* and its plural analog *drivers*.

There are further semantic constraints that have to be placed on our derivational rule, however. For example, the subject in the input verb has to be sufficiently agentive – that is, it must play an active (usually volitional) role (given by the variable 'R' in (30)) in the situation. That's why nominalizations like *knower* or *resembler* sound funny. But the formulation in (30) is a reasonable first pass at the problem that gives you the idea of how phenomena like this can be analyzed within our framework.

There are many other cross-categorial relations that work this way in English. Noun lexemes, both common and proper, can be converted into verbal lexemes:

(31) a. Sandy *porched* the newspaper without difficulty.

b. The senator *houdinied* his way out of the accusations.

c. They have been *computering* me to death all morning.

This kind of affixless derivation, an instance of what is often called ZERO DERIVATION, could be handled by one or more derivational rules.

Finally, we should note that lexical rules are one traditional way of approaching the problem of VALENCE ALTERNATIONS, i.e. the fact that many verbs allow systematically related valence patterns. Among the most famous of these is the dative alternation illustrated in (32) – (33).

(32) a. Birch gave Brooke a book.

b. Birch gave a book to Brooke.

(33) a. Birch handed Brooke a book.

b. Birch handed a book to Brooke.

Rather than list two distinct verbal lexemes for *give*, *hand*, and a family of related elements, it makes much more sense to list only one (with one of the two valence patterns fixed) and to derive the other by a derivational rule. Note however, that there are certain other verbs or particular idiomatic uses that appear in only one of the two valence patterns[14].

(34) a. Kris donated a book to the library.

b. *Kris donated the library a book.

(35) a. Dale gave Brooke a hard time.

b. ??Dale gave a hard time to Brooke.

These underline once again the need for a theory of exceptions to lexical rules and lexical irregularity.

Other famous examples of valence alternation are illustrated in (36) – (40):

(36) a. The police sprayed the protesters with water.

b. The police sprayed water on the protesters. ('*spray/load*' alternations)

(37) a. The students drove cars.

b. These cars drive easily. ('middle' uses)

(38) a. Pat sneezed.

b. Pat sneezed the napkin off the table. ('caused motion' uses)

(39) a. The horse kicked me.

b. The horse kicked me black and blue. ('resultative' uses)

(40) a. Pat yelled.

b. Pat yelled her way into the meeting. ('the way' construction)

All of these patterns of valence alternation are governed by both semantic and syntactic constraints of the kind that could be described by finely tuned lexical rules.

◊**Problem 8: Arguments in Japanese** As noted earlier, Japanese word order differs from English in a number of ways, including the fact that it is a 'Subject-Object-Verb' (SOV) language. Here are a few relevant examples ('nom', 'acc', and 'dat' stand for nominative, accusative and dative case, respectively).

[14]Following standard practice among generative grammarians, we use designations '?', '??', and '?*' to indicate different levels of naturalness between full acceptability and complete unacceptability.

(i) hitorino otoko-ga sono hon-o yonda
one man-nom that book-acc read-past
'One man read that book'
[cf. *Yonda hitorino otoko-ga sono hon-o
*Hitorino otoko-ga yonda sono hon-o
*Otoko-ga hitorino sono hon-o yonda
*Hitorino otoko-ga hon-o sono yonda
*Hitorino otoko-ni/-o sono hon-o yonda
*Hitorino otoko-ga sono hon-ga/-ni yonda.]

(ii) Hanako-ga hon-o yonda
Hanako-nom book-acc read-past
'Hanako read the book(s)'
[cf. *Yonda Hanako-ga hon-o
*Hanako-ga yonda hon-o
*Hanako-ni/-o hon-o yonda
*Hanako-ga hon-ni/-ga yonda.]

(iii) sensei-ga Taroo-ni sono hon-o ageta
sensei-nom Taroo-dat that book-acc gave-past
'The teacher(s) gave that book to Taroo
[cf. *Ageta sensei-ga Taroo-ni sono hon-o
*Sensei-ga ageta Taroo-ni sono hon-o
*Sensei-ga Taroo-ni ageta sono hon-o
*Sensei-o/-ni Taroo-ni sono hon-o ageta
*Sensei-ga Taroo-ga/-o sono hon-o ageta
*Sensei-ga Taroo-ni sono hon-ga/-ni ageta.]

(iv) Sono hon-ga akai-desu
that book-nom red-copula
'That book is red'

As the contrasting ungrammatical examples show, the verb must appear in final position in Japanese. In addition, we see that verbs select for NPs of a particular case, much as in English. Assume that the nouns and verbs of Japanese are inflected words, derived by lexical rule from the appropriate lexemes.

Task 1. Given the data illustrated here, how could the head-complement and head-specifier rules be revised to deal with Japanese? Explain the effect of the difference(s) you have posited.

Task 2. Give the lexical entry for each of the verbs illustrated in (i)–(iv) [treat *akai-desu* as a simple intransitive verb]. The data given only permit you to specify some features; leave others unspecified. Make sure your entries interact with the rules you presented in Task 1 to account for all the above data.

Task 3. Assume that nouns like *Taroo*, *hon* etc. are entered in the Japanese lexicon as nominal lexemes. Give lexical entries for these lexemes (again limiting the features specified to those for which you have data).

Task 4. Formulate three lexical rules for deriving the inflected forms (the words ending in *-ga*, *-o*, and *-ni*) from the nominal lexemes.

Task 5. Explain the role of the Head Feature Principle in accounting for the case-selecting properties of verbs in (i)–(iv).

As you know from problem 3 of chapter 7, both subjects and objects in Japanese can in general be omitted. When such an element is omitted, the sentence usually has a meaning that we would express in English by using a pronoun corresponding to the unexpressed element:

(v) Taroo-ga yonda
Taroo-nom read-past
'Taroo read it/them'

(vi) hon-o yonda
book-acc read-past
'He/she/they read the book(s)'

(vii) yonda
read-past
'He/she/it/they read it/them'

(viii) Taroo-ga tatai-ta.
Taroo-NOM hit-PST
'Taroo$_i$ hit him$_j$/her$_j$/them$_j$/*himself$_i$.'

Task 6. Sketch an account of this phenomenon that uses lexical rules. Be explicit about the verbs' values for SPR, COMPS, and ARG-ST. You should assume that all complements can be omitted, not just those illustrated in (v) – (viii); i.e. your lexical rules should apply with maximum generality. Following the suggestion in Chapter 7, Problem 3, you should assume that the ARP is a default principle in Japanese (unlike English).

8.5 Conclusion

An important insight, going back at least to de Saussure, is that all languages involve arbitrary (that is, unpredictable) information. Most clearly, the association between the forms (sounds) and meanings of words is purely conventional, in the vast majority of cases. A grammar of a language must list those associations somewhere. The original conception of the lexicon in modern linguistics was simply as the repository of such arbitrary information.

This conception did not last long, however. Beginning in the early years of transformational grammar, linguists began enriching their conception of the lexicon to include information that was not idiosyncratic to individual words. This trend continued in a great deal of research carried out within a variety of grammatical frameworks.

In this text, we have to some extent recapitulated this history. We began with context-free grammar, in which the lexicon contained only idiosyncratic information, and we gradually enriched our lexical representations, including more and more information – much of it systematic and predictable – about the grammatical and semantic properties of words. Indeed, most of the information needed to determine the well-formedness of sentences is now encoded into our lexical entries.

With the increased expressiveness and concomitant complexity of lexical entries came a need to express succinctly certain generalizations about words. In this chapter, we have examined two formal mechanisms for capturing such generalizations. Structuring the lexicon as an inheritance hierarchy of types has made it possible to factor out information common to many lexical entries, thereby greatly reducing lexical redundancy. By allowing inherited information to be overridden, we can encode default values for features, while still allowing for lexical idiosyncracy. The second mechanism, the lexical rule, is an older idea, going back to work in transformational grammar of the 1970s. We will make extensive use of lexical rules in subsequent chapters. In fact, many of the phenomena that provided the motivation for transformations in the 1950s and 1960s can be reanalyzed in our theory using lexical rules. These include the passive construction – the topic of chapter 10 – and many of the properties of the English auxiliary verb system.

8.6 Further Reading

An important early paper on lexical rules was Jackendoff (1975). The idea of combining lexical rules with an inheritance hierarchy was first put forward by Flickinger et al. (1985). See also Pollard and Sag (1987, chap. 8). Briscoe et al. (1993) is a collection of more recent papers about lexical hierarchies, default inheritance, and related issues.

Chapter 9

Realistic Grammar

9.1 Introduction

In the preceding eight chapters, we have laid out the theory that we will apply to more complex data in the remainder of this book. Surprisingly perhaps, only one more grammatical principle is required; the other extensions we will present in subsequent chapters (aside from a minor reformulation of the ARP in chapter 15) all concern the lexicon. Before we proceed, however, it is perhaps useful to consolidate the components of our treatment of English grammar and to reflect on the strategy we have adopted for solving syntactic problems – to reflect on the motivation for the design of grammar.

The next section of this chapter presents a formal summary of everything we have covered so far, including types, lexical entries, grammar rules, the definition of well-formed structural descriptions (incorporating various principles), and lexical rules. The later sections of the chapter discuss some basic issues that arise when one asks what the kind of grammatical theory we have been developing here has to do with how people actually use language.

Section 2.1 presents the type hierarchy in two different formats. The first is in the shape of a tree diagram, showing which types inherit from which other types. The second is a table, providing details of the features and constraints associated with each type. Constraints on the possible values of features are introduced by enclosing the list of possible values in curly brackets, following the feature name. Almost all of the types and constraints listed in Section 2.1 were introduced in earlier chapters. We have added almost nothing that is new. Three new types are *gram-cat* (for 'grammatical category'), *sem-struc* (for 'semantic structure), and *val-cat* (for 'valence category'). The first of these is the type of the value of SYN – that is, feature structures of type *gram-cat* have the features HEAD and VAL. The second is the type of the value of SEM, which consists of the features MODE, INDEX, and RESTR. And the third of the three new types (*val-cat*) takes the features SPR and COMPS. All three of these have been implicit in the discussion of the preceding chapters. The type *verb* lists no features or constraints associated with it, but we have put in a placeholder '[...]' because this will change in subsequent chapters.

Section 2.2 gives some sample lexical entries. It is worth noting that most of what we have to stipulate in our entries is semantic. By virtue of having a richly structured lexicon, we are able to limit greatly the amount of syntactic information that has to be listed in individual entries.

Section 2.3 lists our familiar grammar rules from Chapter 6, together with the Imperative Rule introduced in Chapter 7.

Section 2.4 gives the formal version of our definition of well-formed structural descriptions, incorporating all of the general principles of grammar we have adopted so far. This version is slightly modified from the one given in Chapter 6, because the Argument Realization Principle has been incoporated into the first part (that is, the definition of lexical SDs). In addition, our binding theory and the Anaphoric Agreement Principle have been built in. This version of the definition also provides something that was only promised in Chapter 6: an initial symbol for our grammar. It says, in effect, that a well-formed stand-alone expression in English must be a finite or imperative sentence.

Section 2.5 simply states the relationship between structural descriptions and the phrase structures they describe.

Section 2.6 lists the lexical rules that were presented in Chapter 8.

9.2 The Grammar To Date

9.2.1 The Type Hierarchy

(1)

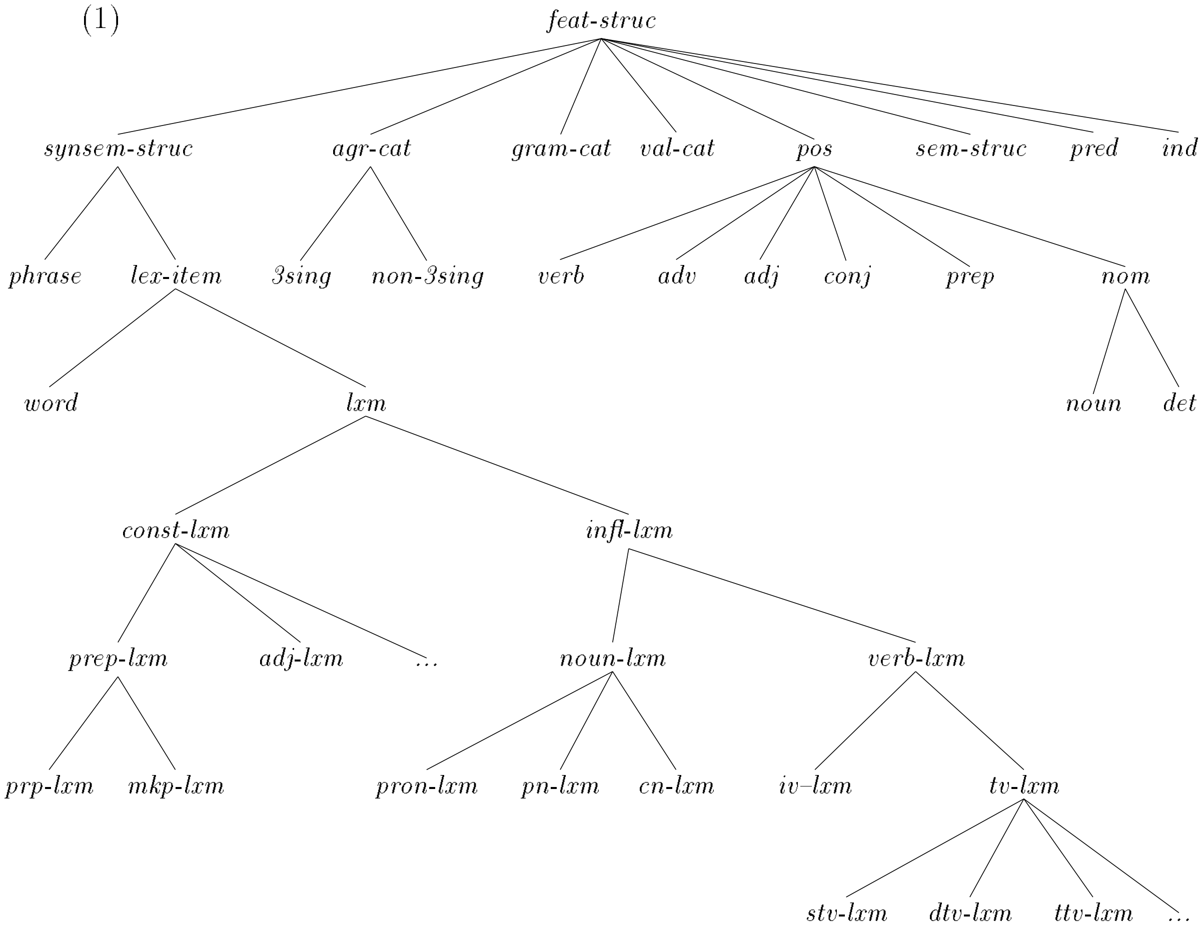

(2)

SOME GENERAL TYPES		
TYPE	FEATURES/CONSTRAINTS	ISA
feat-struc		
synsem-struc	$\begin{bmatrix}\text{SYN } \textit{gram-cat}\\ \text{SEM } \textit{sem-struc}\end{bmatrix}$	*feat-struc*
phrase		*synsem-struc*
lex-item	$\begin{bmatrix}\text{ARG-ST list}(\textit{synsem-struc})\end{bmatrix}$	*synsem-struc*
word		*lex-item*
lexeme	$\begin{bmatrix}\text{SYN }\begin{bmatrix}\text{HEAD }\begin{bmatrix}\text{ANA } / -\end{bmatrix}\end{bmatrix}\\ \text{SEM }\begin{bmatrix}\text{MODE } / \text{ prop}\end{bmatrix}\end{bmatrix}$	*lex-item*

(3)

LEXEME TYPES		
TYPE	CONSTRAINTS	ISA
infl-lxm, *const-lxm*		*lexeme*
noun-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{noun} \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & / \langle\ \rangle \\ \text{SEM} & / \begin{bmatrix} \text{MODE ref} \end{bmatrix} \end{bmatrix}$	*infl-lxm*
pron-lxm	$\begin{bmatrix} ... \end{bmatrix}$	*noun-lxm*
pn-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{AGR} & / \ \textit{3sing} \end{bmatrix} \end{bmatrix} \end{bmatrix}$	*noun-lxm*
cn-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{AGR} \ \boxed{1} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} & / \ \langle\ [\ ]\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & / \left\langle \begin{matrix} \text{DetP} \\ \begin{bmatrix} \text{AGR} \ \boxed{1} \end{bmatrix} \end{matrix} \right\rangle \end{bmatrix}$	*noun-lxm*
verb-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \textit{verb} \\ \text{VAL} & \begin{bmatrix} \text{SPR} \ \langle\ [\ ]\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & / \ \langle\ \text{NP}, \ldots\ \rangle \end{bmatrix}$	*infl-lxm*
iv-lxm	$\begin{bmatrix} \text{ARG-ST} \ \langle\ [\ ]\ \rangle \end{bmatrix}$	*verb-lxm*
tv-lxm	$\begin{bmatrix} \text{ARG-ST} \ \langle\ [\ ]\ ,\ \text{NP}\ ,\ \ldots\ \rangle \end{bmatrix}$	*verb-lxm*
stv-lxm	$\begin{bmatrix} \text{ARG-ST} \ \langle\ [\ ]\ ,\ [\ ]\ \rangle \end{bmatrix}$	*tv-lxm*
dtv-lxm	$\begin{bmatrix} \text{ARG-ST} \ \langle\ [\ ]\ ,\ [\ ]\ ,\ \text{NP}\ \rangle \end{bmatrix}$	*tv-lxm*
ttv-lxm	$\begin{bmatrix} \text{ARG-ST} \ \langle\ [\ ]\ ,\ [\ ]\ ,\ \text{PP}\ \rangle \end{bmatrix}$	*tv-lxm*
prep-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \textit{prep} \end{bmatrix} \end{bmatrix}$	*const-lxm*
pdp-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{MOD} & \{\text{NOM, VP, none}\} \\ \text{P-OBJ} & \text{none} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} \ \langle\ [\ ]\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \langle\ \text{NP}\ ,\ \text{NP}\ \rangle \end{bmatrix}$	*prep-lxm*
mkp-lxm	$\begin{bmatrix} \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{MOD} & \text{none} \\ \text{P-OBJ} & \boxed{1} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} \ \langle\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \langle\ \boxed{1}\ \rangle \end{bmatrix}$	*prep-lxm*
adj-lxm, *conj-lxm*, *det-lxm*, ...	...	*const-lxm*

OTHER GRAMMATICAL TYPES		
TYPE	FEATURES/CONSTRAINTS	ISA
gram-cat	$\begin{bmatrix}\text{HEAD} & \textit{pos}\\ \text{VAL} & \textit{val-cat}\end{bmatrix}$	*feat-struc*
val-cat	$\begin{bmatrix}\text{COMPS} & \text{list}(\textit{synsem-struc})\\ \text{SPR} & \text{list}(\textit{synsem-struc})\end{bmatrix}$	*feat-struc*
pos	$\begin{bmatrix}\text{FORM} & \{\text{fin, to,...}\}\\ \text{ANA} & \{+, -\}\end{bmatrix}$	*feat-struc*
verb	$[\ldots]$	*pos*
prep	$\begin{bmatrix}\text{P-OBJ} & \{\textit{synsem-struc}, \text{none}\}\\ \text{MOD} & \{\text{NOM, VP, none,...}\}\end{bmatrix}$	*pos*
nominal	$\begin{bmatrix}\text{AGR} & \textit{agr-cat}\end{bmatrix}$	*pos*
noun	$\begin{bmatrix}\text{CASE} & \{\text{nom, acc}\}\end{bmatrix}$	*nominal*
det	$\begin{bmatrix}\text{COUNT} & \{+, -\}\end{bmatrix}$	*nominal*
adv	$\begin{bmatrix}\text{MOD} & \{\text{VP, none,...}\}\end{bmatrix}$	*pos*
adj	$\begin{bmatrix}\text{MOD} & \{\text{NOM, none}\}\end{bmatrix}$	*pos*
conj		*pos*
agr-cat	$\begin{bmatrix}\text{PER} & \{\text{1st, 2nd, 3rd}\}\\ \text{NUM} & \{\text{sg, pl}\}\end{bmatrix}$	*feat-struc*
3sing	$\begin{bmatrix}\text{PER} & \text{3rd}\\ \text{NUM} & \text{sg}\\ \text{GEND} & \{\text{masc, fem, neut}\}\end{bmatrix}$	*agr-cat*
non-3sing		*agr-cat*
sem-struc	$\begin{bmatrix}\text{MODE} & \{\text{prop, ques, dir, ref}\}\\ \text{INDEX} & \textit{index}\\ \text{RESTR} & \text{list}(\textit{predication})\end{bmatrix}$	*feat-struc*
predication	$\begin{bmatrix}\text{RELN} & \{\text{love,walk,...}\}\\ \ldots & \end{bmatrix}$	*feat-struc*
index		*feat-struc*

9.2.2 The Basic Lexicon

Here are some sample lexical entries that are part of the basic lexicon. '...' means 'this is something we haven't dealt with but a complete grammar would have to'.

Nouns

(4)

$$\left\langle \text{she}\ ,\ \begin{bmatrix} \textit{pron-lxm} & \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{CASE} & \text{nom} \\ \text{AGR} & \begin{bmatrix} \textit{3sing} \\ \text{GEND fem} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{female} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(5)

$$\left\langle \text{him}\ ,\ \begin{bmatrix} \textit{pron-lxm} & \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{CASE} & \text{acc} \\ \text{AGR} & \begin{bmatrix} \textit{3sing} \\ \text{GEND masc} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{male} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(6)

$$\left\langle \text{themselves}\ ,\ \begin{bmatrix} \textit{pron-lxm} & \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{CASE} & \text{acc} \\ \text{ANA} & + \\ \text{AGR} & \begin{bmatrix} \textit{non-3sing} & \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{pl} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{group} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(7)
$$\left\langle \text{Kim}\ ,\ \begin{bmatrix} \textit{pn-lxm} & \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{name} \\ \text{SIT} & \text{s} \\ \text{NAME} & \text{Kim} \\ \text{NAMED} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(8)
$$\left\langle \text{book}\ ,\ \begin{bmatrix} \textit{cn-lxm} & \\ \text{SYN} & \begin{bmatrix} \text{VAL} \begin{bmatrix} \text{SPR}\ \langle\ [\text{COUNT}\ +]\ \rangle \end{bmatrix} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{i} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{book} \\ \text{SIT} & \text{s} \\ \text{INSTANCE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

Verbs

(9)
$$\left\langle \text{die}\ ,\ \begin{bmatrix} \textit{iv-lxm} & \\ \text{ARG-ST}\ \langle\ [\ \ ]_i\ \rangle & \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{die} \\ \text{SIT} & \text{s} \\ \text{CORPSE} & \text{i} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(10)
$$\left\langle \text{love}\ ,\ \begin{bmatrix} \textit{stv-lxm} & \\ \text{ARG-ST}\ \langle\ [\ \ ]_i\ ,\ [\ \ ]_j\ \rangle & \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{love} \\ \text{SIT} & \text{s} \\ \text{LOVER} & \text{i} \\ \text{LOVED} & \text{j} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(11)
$$\left\langle \text{give}_1 , \begin{bmatrix} \textit{dtv-lxm} & \\ \text{ARG-ST} \ \langle \ [\ ]_i \ , [\ ]_j \ , [\ ]_k \ \rangle & \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{give} \\ \text{SIT} & \text{s} \\ \text{GIVER} & \text{i} \\ \text{GIVEN} & \text{j} \\ \text{GIFT} & \text{k} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(12)
$$\left\langle \text{give}_2 , \begin{bmatrix} \textit{ttv-lxm} & \\ \text{ARG-ST} \ \langle \ [\ ]_i \ , [\ ]_k \ , [\text{P-OBJ} \ \ \text{NP}_j] \ \rangle & \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \text{s} \\ \text{RESTR} & \left\langle \begin{bmatrix} \text{RELN} & \text{give} \\ \text{SIT} & \text{s} \\ \text{GIVER} & \text{i} \\ \text{GIVEN} & \text{j} \\ \text{GIFT} & \text{k} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

Odds and Ends

(13)
$$\left\langle \text{the} , \begin{bmatrix} \textit{det-lxm} & \\ \text{SEM} & [...] \end{bmatrix} \right\rangle$$

(14)
$$\left\langle \text{few} , \begin{bmatrix} \textit{det-lxm} & \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{COUNT} & + \\ \text{AGR} & \begin{bmatrix} \textit{non-3sing} & \\ \text{NUM} & \text{pl} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{SEM} & [...] \end{bmatrix} \right\rangle$$

(15)
$$\left\langle \text{'s} , \begin{bmatrix} \textit{det-lxm} & \\ \text{SEM} & [...] \end{bmatrix} \right\rangle$$

(16)
$$\left\langle \text{to} , \begin{bmatrix} \textit{mkp-lxm} \\ ... \end{bmatrix} \right\rangle$$

(17)
$$\left\langle \text{and} , \begin{bmatrix} \textit{conj-lxm} \\ ... \end{bmatrix} \right\rangle$$

9.2.3 The Grammar Rules

(18) Head-Complement Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{COMPS}\ \langle\,\rangle \end{bmatrix} \quad \rightarrow \quad \text{H}\begin{bmatrix} \textit{word} \\ \text{COMPS}\ \ \boxed{1} \end{bmatrix} \quad \boxed{1}$$

(19) Head-Specifier Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle\,\rangle \end{bmatrix} \quad \rightarrow \quad \boxed{1} \quad \text{H}\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle\, \boxed{1}\,\rangle \end{bmatrix}$$

(20) Head-Modifier Rule:

$$[\textit{phrase}] \quad \rightarrow \quad \text{H}\boxed{1}[\textit{phrase}] \quad \begin{bmatrix} \textit{phrase} \\ \text{HEAD}\ \begin{bmatrix}\text{MOD}\ \boxed{1}\end{bmatrix} \end{bmatrix}$$

(21) Coordination Rule:

$$\begin{bmatrix}\text{SYN}\ \boxed{1}\end{bmatrix} \ \rightarrow \ \begin{bmatrix}\text{SYN}\ \boxed{1}\end{bmatrix}^{+} \quad \text{CONJ} \quad \begin{bmatrix}\text{SYN}\ \boxed{1}\end{bmatrix}$$

(22) Imperative Rule:

$$\begin{bmatrix} \textit{phrase} \\ \text{SPR} \quad \langle\ \rangle \end{bmatrix} \quad \rightarrow \quad \text{H}\begin{bmatrix} \textit{phrase} & \\ \text{HEAD} & \begin{bmatrix} \textit{verb} & \\ \text{FORM} & \text{imp} \end{bmatrix} \\ \text{SPR} & \left\langle [\ ] \right\rangle \end{bmatrix}$$

9.2.4 Well-Formedness of Structural Descriptions

(23) Lexical SDs:

$$\begin{array}{c}\Delta\\ |\\ \omega\end{array}$$

is a well-formed lexical SD just in case

1. $\langle \omega, D \rangle$ is a lexical entry and Δ is the unification of D with the following feature structure description (Argument Realization Principle):

$$\begin{bmatrix}\text{SYN} & \begin{bmatrix}\text{VAL} & \begin{bmatrix}\text{SPR} & \boxed{1}\\ \text{COMPS} & \boxed{2}\end{bmatrix}\end{bmatrix}\\ \text{ARG-ST} & \boxed{1} \oplus \boxed{2}\end{bmatrix}$$

   and

2. An NP in a noninitial position of a word's ARG-ST list is [CASE acc] (Case Constraint).

(24) Phrasal SDs:

If

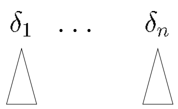

are well-formed SDs (that is, lexical or phrasal structural descriptions), and R is an instantiation of a grammar rule of the form: $D_0 \rightarrow D_1 \ldots D_n$, then R sanctions a phrasal SD Δ just in case:

1. If R is unheaded, then $\Delta =$

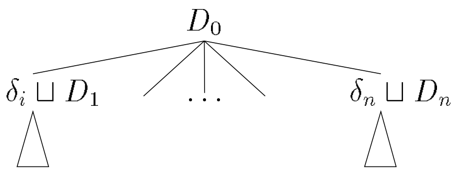

2. If D_h is the head of D_0 in R, then Δ is the result of unifying the SD in 1. with all of the following SDs:

   a. (Head Feature Principle:)

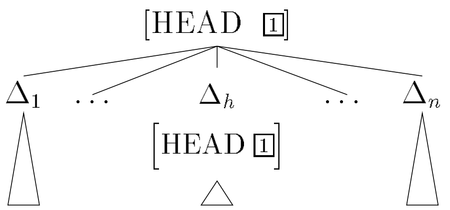

   b. (Valence Principle [unless R says otherwise]):

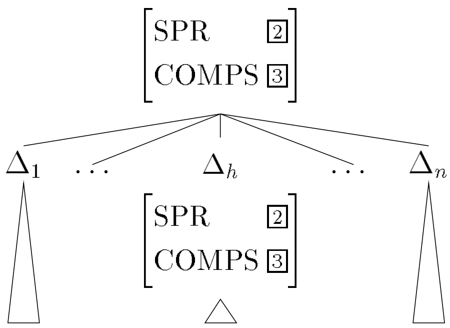

   c. (Semantics Principle (part i)):

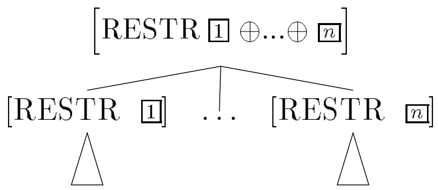

   d. (Semantics Principle (part ii)):

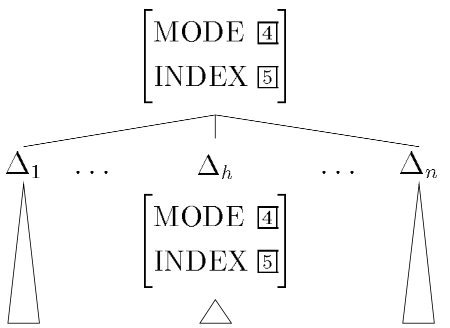

(25) 'Initial Symbol' Definition:

In order for a phrase to stand alone, its synsem-struc must unify with the following specification:

$$\begin{bmatrix} \textit{phrase} \\ \text{SYN} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{verb} \\ \text{FORM fin} \mid \text{imp} \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SPR} & \langle\ \rangle \\ \text{COMPS} & \langle\ \rangle \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

(26) The Binding Theory:

Principle A (version III): An [ANA +] synsem-struc must be outranked by a coindexed synsem-struc.

Principle B (version III): An [ANA –] synsem-struc must not be outranked by a coindexed synsem-struc.

(27) Anaphoric Agreement Principle:

Coindexed elements agree (i.e. share a single AGR value).

9.2.5 Satisfaction of Structural Descriptions

(28) A phrase structure S satisfies a structural description Δ just in case:

(i) S has exactly the same configuration as Δ and

(ii) the feature structures labelling the nodes of S satisfy all the constraints specified in Δ.

9.2.6 Lexical Rules

(29) Singular Noun Lexical Rule:

$$\left\langle \boxed{1}, \begin{bmatrix}\textit{cn-lxm}\end{bmatrix} \right\rangle \Rightarrow \left\langle \boxed{1}, \begin{bmatrix} \textit{word} \\ \text{SYN} \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{AGR } \textit{3sing} \end{bmatrix} \end{bmatrix} \end{bmatrix} \right\rangle$$

(30) Plural Noun Lexical Rule:

$$\left\langle \boxed{1}, \begin{bmatrix} \textit{cn-lxm} \\ \text{ARG-ST} \ \langle\ \boxed{3}[\text{COUNT} +]\ \rangle \end{bmatrix} \right\rangle \Rightarrow \left\langle \text{F}_{NPL}(\boxed{1}), \begin{bmatrix} \textit{word} \\ \text{SYN} & \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{AGR} \begin{bmatrix} \text{NUM pl} \end{bmatrix} \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \langle\ (\boxed{3})\ \rangle \end{bmatrix} \right\rangle$$

(31) 3rd Singular Verb Lexical Rule:

$$\left\langle \boxed{3}\ ,\ \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix}\text{RESTR}\ \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{3SG}(\boxed{3})\ ,\ \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD} \begin{bmatrix}\text{FORM fin}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{CASE nom} \\ \text{AGR}\ \ \textit{3sing}\end{bmatrix}, \ldots \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX}\ \boxed{2} \\ \text{RESTR}\ \boxed{1} \oplus \left\langle \begin{bmatrix}\text{RELN t-overlap} \\ \text{ARG1}\ \boxed{2} \\ \text{ARG2 now}\end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(32) Non-3rd-Singular Verb Lexical Rule:

$$\left\langle \boxed{3},\ \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix}\text{RESTR}\ \boxed{1}\end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \boxed{3},\ \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix}\text{HEAD} \begin{bmatrix}\text{FORM fin}\end{bmatrix}\end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix}\text{CASE nom} \\ \text{AGR}\ \ \textit{non-3sing}\end{bmatrix}, \ldots. \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX}\ \boxed{2} \\ \text{RESTR}\ \boxed{1} \oplus \left\langle \begin{bmatrix}\text{RELN t-overlap} \\ \text{ARG1}\ \boxed{2} \\ \text{ARG2 now}\end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(33) Past Tense Verb Lexical Rule:

$$\left\langle \boxed{2}, \begin{bmatrix} \textit{lexeme} & \\ \text{SYN} & \left[\text{HEAD } \textit{verb}\right] \\ \text{SEM} & \left[\text{RESTR } \boxed{1}\right] \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \text{F}_{PAST}(\boxed{2}), \begin{bmatrix} \textit{word} & \\ \text{SYN} & \left[\text{HEAD } \left[\text{FORM fin}\right]\right] \\ \text{ARG-ST} & \left\langle \left[\text{CASE nom}\right], ... \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{INDEX } \boxed{3} & \\ \text{RESTR } \boxed{1} \oplus \left\langle \begin{bmatrix} \text{RELN t-precede} \\ \text{ARG1 } \boxed{3} \\ \text{ARG2 now} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \right\rangle$$

(34) Present Participle Lexical Rule:

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \left[\text{RESTR } \boxed{1}\right] \end{bmatrix} \right\rangle \Rightarrow \left\langle \text{F}_{PRP}(\boxed{3}), \begin{bmatrix} \textit{word} & \\ \text{SYN} & \left[\text{HEAD } \left[\text{FORM prp}\right]\right] \\ \text{SEM} & \left[\text{RESTR } \boxed{1} \oplus \ \ldots \right] \end{bmatrix} \right\rangle$$

(35) Past Participle Lexical Rule:

$$\left\langle \boxed{2}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \left[\text{RESTR } \boxed{1}\right] \end{bmatrix} \right\rangle \Rightarrow \left\langle \text{F}_{PSP}(\boxed{2}), \begin{bmatrix} \textit{word} & \\ \text{SYN} & \left[\text{HEAD } \left[\text{FORM psp}\right]\right] \\ \text{SEM} & \left[\text{RESTR } \boxed{1} \oplus \ \ldots \right] \end{bmatrix} \right\rangle$$

(36) Infinitival Lexical Rule:

$$\left\langle \boxed{2}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \left[\text{RESTR } \boxed{1}\right] \end{bmatrix} \right\rangle \Rightarrow \left\langle \boxed{2}, \begin{bmatrix} \textit{word} & \\ \text{SYN} & \left[\text{HEAD } \left[\text{FORM inf}\right]\right] \\ \text{SEM} & \left[\text{RESTR } \boxed{1} \oplus \ \ldots \right] \end{bmatrix} \right\rangle$$

(37) Imperative Lexical Rule:

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{verb-lxm} & \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{prop} \\ \text{RESTR} & \boxed{1} \end{bmatrix} \end{bmatrix} \right\rangle \Rightarrow$$

$$\left\langle \boxed{3}, \begin{bmatrix} \textit{word} & \\ \text{SYN} & \begin{bmatrix} \text{HEAD} \begin{bmatrix} \text{FORM imp} \end{bmatrix} \end{bmatrix} \\ \text{ARG-ST} & \left\langle \begin{bmatrix} \text{AGR} \begin{bmatrix} \text{PER 2nd} \end{bmatrix} \end{bmatrix}, \ldots. \right\rangle \\ \text{SEM} & \begin{bmatrix} \text{MODE} & \text{dir} \\ \text{RESTR} & \boxed{1} \end{bmatrix} \end{bmatrix} \right\rangle$$

9.3 Constraint-Based Lexicalism

We turn now to some reflections on the relationship between grammatical descriptions like the ones in this text and what is known about the mental processes underlying the comprehension and production of language by people. We believe it is possible to draw at least preliminary conclusions about how the flexibility and incrementality of human language processing – language comprehension in particular – should inform the design of grammars that aspire to be models of human linguistic knowledge. Adopting the familiar terminology of Chomsky (1965), we will suggest that competence grammars, if they are ever to be DIRECTLY embedded within realistic models of how language is used (what Chomsky would call linguistic *performance*), should embody certain design properties that make them consistent with these basic facts about processing.

Let us start with three basic observations about the grammar we have been developing:

1. It is *surface-oriented.* Our grammar (like standard context-free grammars) provides a reasonably simple structure that is directly associated with the string of words that constitute each sentence. The ancillary structure that has to be computed in order to ascertain that a given sentence is grammatical expresses information that is straightforwardly derivable from properties of the words in the string. No additional abstract structures are posited.

2. It is *constraint-based.* There are no operations that destructively modify any representations. The principles of the theory, the grammar rules, and the lexical entries that result once the grammar is 'closed' under lexical rules are all just constraints that interact so as to define a set of phrasal SDs – those structural descriptions which simultaneously satisfy the relevant constraints of our grammar.

3. It is *strongly lexicalist.* We have localized most grammatical and semantic information within lexical entries. These lexical entries furthermore correspond directly to the words present in the sentence, which can be viewed as the key elements that drive the construction of the syntactic and semantic structure of the sentence.

Each of these three design properties, we believe, plays an important role when we turn to embedding our theory of linguistic competence within a theory of performance – specifically a theory of how language is processed. Any theory that has these design properties exemplifies a viewpoint that we will refer to as *Constraint-Based Lexicalism* (CBL).

9.4 Incremental Processing

We don't have to venture into a psycholinguistic laboratory to convince ourselves that language processing is highly incremental. We already saw this in Chapter 1, when we considered examples like (38).

(38) After finding the book on the atom, Sandy went into class, confident that there would be no further obstacles to getting that term paper done.

When we hear such a sentence, we process it as it comes – more or less word by word – building structure and partial interpretation incrementally, using nonlinguistic information as we can to make the right decisions at choice points, e.g. the point when we encounter the PP *on the atom* and have to decide whether it modifies VP or NOM. We make this decision 'on-line' it seems, using a plausibility assessment of the meaning that would result from each structure.

Psycholinguists have shown us that sentence processing sometimes goes astray. GARDEN PATH examples like (39a,b) are as remarkable today as they were they were first brought to the attention of language researchers.[1]

(39) a. The horse raced past the barn fell.

b. The boat floated down the river sank.

On first encountering such sentences, almost all English speakers judge them to be totally ungrammatical. However, after seeing them juxtaposed to fully well-formed examples like (40), speakers recognize that sentences like (39) are grammatical, but very hard to process.

(40) a. The horse that was raced past the barn fell.

b. The horse taken to the hospital died.

c. The boat that was floated down the river sank.

d. The boat seen down the river sank.

Experimental researchers thought at first that these garden paths showed that certain purely linguistic processing strategies (like trying to build an S out of the NP *the horse* and a VP beginning with *raced past*) were automatic - virtually impossible to turn off. But modern psycholinguistics has a very different story to tell.

[1]By Bever (1970).

First, note that in the right context, one can eliminate the garden path effect even with the sentences in (39). The right context can even make the NOM-modifying interpretation of *raced past the barn* the most natural one:[2]

(41) The horse that they raced around the track held up fine. The horse that was raced down the road faltered a bit. And the horse raced past the barn fell.

The presuppositions of the context here highlight the need to identify one horse among many, which in turn favors the meaning of the NOM-modifying structure of (39a).

Moreover, if we keep the same potential for ambiguity, but change the words, we can eliminate the garden path effect even without an elaborate preceding context. Consider examples like (42a,b)

(42) a. The evidence examined by the judge turned out to be unreliable.

b. The thief arrested by the police turned out to be our nephew.

As shown in a number of studies,[3] examples like these present no more processing difficulty than their unambiguous counterparts in (43).

(43) a. The evidence that was examined by the judge turned out to be unreliable.

b. The thief who was arrested by the police turned out to be our nephew.

That is, the examples in (42), even in the absence of biasing prior context, do not cause garden path effects.

The explanation for this difference lies in the relevant nonlinguistic information. Evidence (or, say, a particular piece of evidence) is a pretty implausible examiner and the sentence built out of a subject NP *the evidence* and the VP *examined...* would require precisely that interpretation. (Similar conclusions hold with respect to (42b), given that a thief is unlikely to be the one doing the arresting). That is, it is a fact about the world that examiners are animate, and since evidence is inanimate, that hypothesis about the interpretation of the sentence is implausible. The fact that the decision to reject that interpretation (and hence the associated sentential structure) is made so quickly as to be imperceptible (i.e. so as to produce no noticeable garden path effect) is evidence that language comprehension is working in a highly integrative and incremental fashion. Linguistic and nonlinguistic constraints on the interpretation are being interleaved in real time.

Language understanding is a process of constraint satisfaction. Competing interpretations exist in parallel, but are active to varying degrees. A particular alternative interpretation is active to the extent that evidence is available to support it as the correct interpretation of the utterance being processed. Note by the way, that frequency has a very important role to play in all this. The reason why the *horse raced past the barn...* example is such a strong garden path is that *raced* is much more frequent as a finite verb form than as the passive

[2]This kind of effect is discussed by Crain and Steedman (1985).

[3]See, for example, Trueswell, Tanenhaus and Garnsey (1992), Pearlmutter and MacDonald (l992) and Tabossi et al (1994).

participle of the transitive use of *race*, which is precisely what the NOM-modifying reading requires. Ambiguity resolution is a continuous process, where inherent degrees of activation, e.g. those correlating with gross frequency, fluctuate as further evidence for particular interpretations become available. Such evidence may in principle stem from any aspect of the sentence input or the local or discourse context. A garden-path sentence is one that has an interpretation strongly supported by initial evidence that later turns out to be incorrect.

9.5 A Performance-Plausible Competence Grammar

9.5.1 Surface-Orientation

We know that sentence processing proceeds in a more or less word-by-word (or even syllable-by-syllable) fashion. In comprehending utterances, hearers use their knowledge of language to build partial hypotheses about the intended meaning. These hypotheses become more or less active, depending on how plausible they are, that is, depending on how well their meaning squares with the hearers' understanding of what's going on in the discourse.

Sometimes the process even gets short-circuited. We have all had the experience of completing someone else's utterance (a phenomenon that is, incidentally, far more common than one might imagine, as shown, for example, by Wilkes-Gibbs (1986)) or of having to wait for someone to finish an utterance whose completion had already been made obvious by context. One striking example of this is English 'echo questions', as illustrated in the following kind of dialogue:

(44) a. [Speaker A:] Senora Maria Consuelo Bustamante y Bacigalupo is coming to dinner tomorrow night.

b. [Speaker B:]

ght?
o did you say is coming to dinner tomorrow ni
h
w
* * * * * * * * * * * * * * * * * * *

In a dialogue like this, it is quite likely that Speaker A may comprehend the intent of Speaker B's utterance well before it is complete, e.g. somewhere in the region indicated by *. Presumably, this is possible precisely because Speaker A can recognize that the remainder of B's utterance is a repetition of A's own utterance and can graft that bit of content onto the partial analysis A has performed through word-by-word processing of B's utterance. What examples like this show is that a *partial* linguistic analysis (e.g. the partial linguistic analysis of *who did you*, *who did you say* or *who did you say is*) is constructed incrementally, assigned a (partial) interpretation, and integrated with information from context to produce a complete utterance interpretation even before the utterance is complete. Amazing, if you think about it!

So if a grammar is to be realistic, i.e. to be directly embedded in a model of this kind of incremental and integrative language processing, then it needs to characterize linguistic knowledge in a way that allows for the efficient incremental computation of partial analyses. Moreover, the partial grammatical analyses have to be keyed in to partial linguistic meanings, because that is what touches base with other factors in processing.

The kind of grammar we are developing seems quite compatible with these performance-driven design criteria. The lexical information that comes with each word provides information about the structure of the sentence directly, i.e. about the phrases that the words are part of and about the neighboring phrases that they combine with syntactically. In addition, our words provide partial information about the meaning of those phrases and hence, since all phrases are built up directly from the component words and phrases in a context-free manner, there is useful partial semantic information that can be constructed incrementally, using our surface-oriented grammar.

Of course this is no rigorous proof that a precise performance model exists based on grammars like ours, but the context-free like architecture of the theory and the hybrid syntactic-semantic nature of the lexical data structures are very suggestive. Incremental computation of partial semantic structures, the key to modelling integrative sentence processing, seems very compatible with our grammars.

9.5.2 Constraint-Based Grammar

In addition to the incremental and integrative nature of human language processing, we may also observe that there is no fixed order in which particular kinds of information are considered. For example, it is not the case that syntactic information (e.g. agreement information that might rule out a particular parse) is always consulted before semantic information (e.g. semantic incompatibility that would favor or disfavor some potential interpretation of an utterance). In fact, it is possible to make an even stronger claim. In examples like (45), early accessing of morphological information allows the number of sheep under discussion to be determined incrementally, and well before the world knowledge necessary to select the 'fenced enclosure' sense of *pen*, rather than its 'writing implement' sense.

(45) The sheep that was sleeping in the pen stood up.

In (46), on the other hand, the relevant world information – that sheep might fit inside a fenced enclosure, but not inside a writing implement – seems to be accessed well before the relevant morphological information constraining the number of sheep.[4]

(46) The sheep in the pen had been sleeping and were about to wake up.

So the information accessed in on-line language processing is typically made available in an order determined by the input stream, not by the constructs of some grammatical theory. In comprehending these examples, for example, a hearer accesses morphological information earlier in (45) and later in (46) precisely because the order of access is tied fairly directly to the order of the words being processed.

[4]This pair of examples is due to Martin Kay.

Finally, we know that linguistic information, in the main, functions pretty uniformly in many diverse kinds of processing activity, including comprehension, production, translation, playing language games, and the like. By 'pretty uniformly' we mean that the set of sentences reliably produceable[5] by a given speaker-hearer is quite similar to, in fact bears a natural relation (presumably proper inclusion) to, the set of sentences that that speaker-hearer can comprehend. This might well have been otherwise. The fact that there is so close and predictable a relation between the production activity and the comprehension activity of any given speaker of a natural language argues strongly against any theory where production grammars are independent from comprehension grammars, for instance. Rather, this simple observation suggests that the differences between, say, comprehension and production should be explained by a theory that posits different kinds of processing regimes based on a single linguistic description. And that description should therefore be a process-neutral grammar of the language that can serve each kind of process that plays a role in on-line linguistic activity.[6]

Observations of this sort, i.e. that linguistic descriptions get used in an order-independent fashion, lead naturally to the conclusion that the constructs of linguistic theory should have an order-independent character. That is, the architecture of grammar should be such that linguistic constraints can be independently deployed in different situations. Finally, given that linguistic knowledge is process-independent, there should be no bias within a linguistic theory – whether overt or hidden; intentional or inadvertent - toward one kind of processing, rather than another.

So grammars whose constructs are truly process-neutral hold the best hope for the development of processing models. And the best known way to ensure process-neutrality is to formulate a grammar as a declarative system of constraints. Such systems of constraints fit well into models of processing precisely because all the information they provide is in principle on an equal footing. So what the observations we have been making add up to is a view of grammar as a set of constraints, each expressing partial information about linguistic structures, rather than a system employing destructive operations of any kind. Moreover, these constraints, as we have also seen, should exhibit certain further properties, e.g. order-independence, if performance-compatibility is to be achieved. But the grammar we are developing has just these design properties – all the constructs of the grammar (lexical entries, grammar rules, even our general principles) are nothing more than constraints that produce equivalent results no matter what order they are applied in.

9.5.3 Strong Lexicalism

While our theory partitions grammatical information into a number of different components, whose interaction determines the well-formedness of particular examples, by far the richest

[5]That is, sentences short enough to utter in a real language use situation. We also intend to rule out production errors.

[6]The fact that comprehension extends beyond systematic production can be explained in terms of differences of process - not differences of grammar. Speakers that stray very far from the grammar of their language run serious risk of not being understood; yet hearers that allow grammatical principles to relax when necessary, will understand more than those that don't. There is thus a deep functional motivation for why the two kinds of processing might differ as they appear to.

locus of such information is the lexicon. Our grammar rules are very simple in their formulation and general in their application, as are such universal aspects of our formal theory as the Head Feature Principle and the Valence Principle. Most of the details we need in order to analyze individual sentences are codified in lexical entries (though much of it need not be stipulated, thanks to lexical rules and inheritance through the type hierarchy).

Other divisions of grammatical labor are conceivable. Indeed, a number of theories with highly articulated rule systems and relatively impoverished lexicons have been developed in considerable detail (e.g. early transformational grammar and GPSG). We have argued for strong lexicalism on the basis of linguistic adequacy (along with general considerations of elegance and parsimony). It turns out, however, that the psycholinguistic evidence on language processing also points in the same direction.

Investigations of syntactic ambiguity resolution and garden path effects have shown that both are sensitive to a variety of different types of information. That is, the difficulty listeners exhibit in resolving such ambiguities and in overcoming garden paths is influenced by various factors other than the syntactic structure. These include semantic compatibility and pragmatic plausibility, type and valence of the words involved, and the frequencies with which individual words occur in particular constructions (see Tanenhaus & Trueswell (1995) for a survey of relevant results).

For example, a sentence beginning with the sequence NP_1 - V - NP_2 may be continued in any number of ways. In particular, NP_2 may be the object of the verb, or it may be the subject of a complement sentence. This is illustrated in (47a), which may be continued as in (47b) or (47c).

(47) a. Lou forgot the umbrella ...

b. Lou forgot the umbrella was broken.

c. Lou forgot the umbrella in the closet.

Hence, a listener or reader encountering (47a) must either postpone the decision about whether to attach the NP *the umbrella* to the VP, or may be forced to go back and reanalyze it later. Either way, this creates a burden on the parser in at least some cases. Various experimental paradigms have been used to verify the existence of this parsing difficulty, including measuring reading times and tracking the eye movements of readers.

However, not all verbs that could appear in place of *forgot* in (47a) can appear in both of the contexts in (47b) and (47c). This is illustrated in (48).

(48) a. Lou hoped the umbrella was broken.

b. *Lou hoped the umbrella in the closet.

c. *Lou put the umbrella was broken.

d. Lou put the umbrella in the closet.

The increased parsing load due to the ambiguity in (47a) is greatly reduced when the verb is changed to one whose valence makes it unambiguous, as in (48). This has been demonstrated using the same sorts of methods as were used to establish the complexity of the ambiguity in the first place (see Trueswell et al. (1993)). This provides clear evidence that people use lexical valence information incrementally as they process sentences.

Similarly, people use semantic and pragmatic information about the verb and the following NP to choose between possible attachment sites for the NP. For example, though *learn* may take either an NP object or a sentential complement, illustrated in (49), when the immediately following NP is not the sort of thing one can learn, people do not exhibit the same level of complexity effects in parsing that show up with (47).

(49) a. Dana learned the umbrella was broken.

b. Dana learned a new theorem in class.

The same sort of effect of lexical meaning on parsing shows up with PP attachment ambiguities, like those in (50).

(50) a. The sheriff shot the bandit with a rifle.

b. Lynn likes the hat on the shelf.

The structural ambiguity of such sentences causes parsing complexity, but this is substantially mitigated when the semantics or pragmatics of the verb and/or noun strongly favors one parse, as in (51).

(51) a. The sheriff shot the bandit with a knife.

b. Lynn bought the hat on the shelf.

After reviewing a number of studies dealing with the factors that influence syntactic ambiguity resolution, MacDonald, et al. (1994) discuss what information they believe needs to be lexically specified in order to account for the psycholinguistic results. Their list includes:

- valence;
- 'course-grained semantic information' (that is, the sort of information about who did what to whom that is given in our SEM feature); and
- 'grammatically relevant features' such as 'tense..., finiteness..., voice (active or passive), number..., person..., and gender...'.

They also mention grammatical category, which we represent in our lexical entries by means of types. The elements MacDonald, et al. list correspond remarkably well to the information that we list in our lexical entries. In short, the information that we have be led to posit in our lexical entries has been independently found to play a role in language processing.

9.6 Universal Grammar: A Mental Organ?

The preceding sections have argued that the design features of our grammatical theory comport well with existing evidence about how people process language. There is another sort of psycholinguistic consideration that has played a central role in much work in generative grammar, namely learnability. In this section, we address briefly the question of evaluating our theory by this criterion.

As noted in Chapter 1, Chomsky has argued that the most remarkable fact about human language – and the one he thinks linguists should be primarily concerned with explaining – is that virtually all children become fluent speakers of a language, with little apparent effort or instruction. The puzzle, as Chomsky sees it, is how people can come to know so much about language so quickly and easily. His solution, in a nutshell, is to argue that people's knowledge of language is for the most part innate, not learned. This entails that much linguistic structure – namely those aspects that are innate – must be common to all languages. Consequently, a central goal of much work in modern syntactic theory has been to develop a rich conception of universal grammar that permits the descriptions of particular languages to be as simple as possible.

Chomsky's strong claims about the role of innate knowledge in language acquisition are by no means uncontroversial among developmental psycholinguists. In particular, many scholars disagree with his position that the human language faculty is highly task-specific – that is, that people are born with a 'mental organ' for language which is distinct in its organization and functioning from other cognitive abilities (see, for example, Bates and MacWhinney 1989 and Tomasello 1992 for arguments against Chomsky's position).

There can be little doubt that biology is crucial to the human capacity for language; if it were not, family pets would acquire the same linguistic competence as the children they are raised with. It is far less clear, however, whether the human capacity for language is as specialized as Chomsky says. A range of views on this issue are possible. At one end of the spectrum is the idea that the language faculty is a fully autonomous module, unrelated to general cognitive capacity. At the other end is the idea that there are no specifically linguistic abilities – that our capacity to learn language arises essentially as a side-effect of our general intelligence or of other abililities. Chomsky's view is close to the former; Tomasello (1992) argues for something close to the latter. Other scholars have defended views somewhere in between.

The participants in this debate often seem to be talking past one another. Opponents of task-specificity tend to take a simplistic view of linguistic structure, emphasizing basic communicative functions while ignoring the intricacies of syntax that are the bread and butter of generative grammar. On the other hand, proponents of task-specificity have a tendency to leap from the complexity of their analyses to the conclusion that the knowledge involved must be innate, and unique to language.

We find much of the argumentation on both sides of this controversy unconvincing, and we hence take no position in this book. Nevertheless, the theory presented here can contribute to its resolution. By making analyses explicit, we make it possible to formulate more precisely what is at issue in debates about task-specificity. Moreover, providing formal representations of our data structures and their interactions permits us to see more clearly where there are analogues in other cognitive domains. Our position is that the grammatical formalism we have been developing in this text is well suited to the formulation of a theory of universal grammar, whether or not that theory turns out to be highly task-specific, and that its explicitness can be helpful in resolving the task-specificity question.

To justify this claim, we will consider in turn various components of our theory, namely: the phrase structure rules, the features and their values, the type hierarchy with its feature declarations and constraints, the definition of a well-formed SD (incorporating the Head Feature Principle, the Valence Principle, and the Semantics Principle), the binding theory, and our lexical rules. We will find that most of these have elements that are very likely universal, and that our formulations do not prejudge the issue of task-specificity.

- Our grammar rules (with the exception of the Imperative Rule) are sufficiently general that, aside from the linear ordering of the constituents, they are natural candidates for universality. It would not be hard to factor out the ordering, so that versions of these rules could be posited as part of universal grammar. The sort of hierarchical structure induced by the rules, which we represent with trees, is arguably not unique to language: it also seems appropriate, for example, to aspects of mathematical reasoning. On the other hand, the concepts of 'head', 'complement', 'specifier', and 'modifier', which are crucial to our formulation of the rules, appear to be rather specialized to language. If it should turn out, however, that they can be shown to be instances of some more generally applicable cognitive relations, this would in no way undermine our analysis.

- Most of the features we have posited have obvious cross-linguistic application. It seems at least plausible that a more fully worked out version of the theory being presented here could include an inventory of features from which the feature structures of all languages must be constructed. In later chapters, we will identify the values of some features with particular English words, a practice that is not consistent with saying that the set of possible feature values is part of universal grammar. It might be possible, however, to restrict feature values to come either from the set of morphological forms of the language or from a universally specifiable set. Some features (e.g. PER, GEND, COUNT) clearly reflect properties of the world or of human thought, whereas others (e.g. ANA, P-OBJ, FORM) seem rather specifically linguistic. Our formalism is neutral on the question of whether grammatical features will ultimately be reducible to more general aspects of cognition.

- Like the features, the types we have proposed could arguably be drawn from a fixed universal inventory. The feature declarations associated with the types are likewise probably quite similar across languages. The constraints introduced by some types (such as NSA), on the other hand, appear to be more language particular. Our types are arranged in a default inheritance hierarchy, a kind of structure that very likely plays an important role in how people organize many kinds of information. Indeed, the use of such hierarchies in linguistics was inspired by earlier work in artificial intelligence, which suggested this sort of structure for taxonomies of concepts. The particular types we have posited look quite task-specifically linguistic, though, again, this is not built into our formalism.

- Our definition of a well-formed Structural Description involves both universal and English-specific elements. For example, as noted earlier, the Argument Realization Principle may well differ across languages. Even more clearly, the Case Constraint as we have formulated it applies only to English. On the other hand, the Head Feature and Semantics Principles are intended to apply to all languages Some parts of the SD definition make reference to specifically linguistic constructs (such as grammar rules, heads, and particular features)y, but the idea of unifying information from diverse sources into a single phrase structure description is one that might have non-linguistic applications as well.

- All languages evidently have some binding principles, and they are quite similar. Characteristically, there is one type of element that must be bound within a local domain and another type that cannot be locally bound. But there is cross-language variation in just what counts as 'local' and in what can serve as the antecedents for particular elements. Our particular binding theory is thus not part of universal grammar. Ideally, a grammatical theory should delineate the range of possible binding principles, of which the ones presented in Chapter 7 would be instances. While these principles appear to be quite language-specific, it is conceivable that they could be explained in terms of more general cognitive principles governing identity of reference.

- The lexical rules presented in the previous chapter are clearly parochial to English. A general theory of what kinds of lexical rules are possible would be desirable, but is beyond the scope of this book. The contents of these rules are quite specific to language, but their general form is one that one might expect to find in many domains: if a database contains an feature structure of form X, then it also contains one of form Y.

Summing up this very superficial survey of the components of our theory, we see that it contains many elements (the grammar rules, the tree definition, the features and types) that are plausible candidates for being part of universal grammar. Moreover, some other elements (the binding principles, some lexical rules) probably have close analogues in many other languages. Although our central purpose in this book is to present a precise framework for the development of descriptively adequate grammars for human languages, rather than to account for the puzzle of language learnability through the development of a highly restricted theory of universal grammar, the framework we have presented here is nevertheless quite compatible with the latter goal.

Further, our grammatical theory suggests a number of parallels between the kinds of information structures needed to account for linguistic competence and those employed in other cognitive domains. However, we need not commit ourselves on the question of task-specificity; rather, we offer the hope that increasingly precise linguistic descriptions like those that are possible within the framework developed here will help to clarify the nature of this controversy and its resolution.

9.7 Conclusion

Chomsky's famous distinction between knowledge of language ('competence') and use of language ('performance') has allowed syntacticians to concentrate on relatively tractable problems, abstracting away from many features of the way people actually speak. But most generative grammarians have agreed that an optimal theory of competence will play a role in explaining many features of linguistic performance. To the extent that a theory of grammar attains this ideal, we call it 'realistic'.

We have argued in this chapter that the theory we are developing in this book does rather well by this criterion. Our theory, by virtue of being surface-oriented, constraint-based, and strongly lexicalist, has properties that fit well with what we know about how people process utterances and extract meaning from them. Our understanding of the mechanisms underlying linguistic performance is incomplete at present, and many of the points discussed in this chapter remain controversial. Nevertheless, a preliminary examination of what is known about processing provides grounds for optimism about our approach to syntactic theory. Considerations of learnability also support such a favorable assessment.

9.8 Further Reading

Many of the issues raised in this chapter are discussed at a relatively elementary level in the essays in Gleitman and Liberman (1995).