1 The Nature of Semantic Theory

Semantics is the study of linguistic meaning. Ultimately, its goal is to provide theoretical descriptions and explanations of all of the phenomena of linguistic meaning. In this chapter we introduce the general subject matter of semantic theory and attempt to chart its place and responsibilities in the study of language. Our aim will be to answer certain broad but fundamental questions regarding the data of semantic theory, the object of semantic investigation, the nature of semantic principles and rules, and the task of semantics in relation to other, sister sciences.

1.1 The Pretheoretical Domain of Semantics

Like any scientist, the semantic theorist must begin with whatever data are pretheoretically available and seek a satisfying explanation of them; the semanticist must start by addressing the pretheoretical domain. We begin our explorations by examining the kinds of phenomena that initially present themselves as facts about meaning, and hence the kinds of things our semantic theory might plausibly be responsible for. Considering the data in a rough, intuitive way, we may distinguish three general subdomains. First, there are facts about linguistic expressions themselves, including various properties that they have and various relations holding among them.[1] Second, there are facts about the relationships between linguistic expressions and the world we live in, discuss, and sometimes argue about. And finally, there are facts about the relationships between linguistic expressions and the speakers who use them to formulate thoughts, communicate ideas, persuade, and act. Let us examine these domains briefly.

1.1.1 Semantic Properties and Relations

Among the semantic properties that we would ascribe to natural-language expressions, the most obvious are surely the **actual meanings** that those expressions have. The following, for example, are simple and immediate facts about the actual meanings of three sentences of English and French:

(1) a. The English sentence *Camels have humps* means that camels have humps.
b. The French sentence *Les chameaux ont des bosses* means that camels have humps.
c. The English sentence *Camels have humps* does not mean that reptiles have wings.

Clearly, facts of this kind represent primary data that we would want any semantic theory to account for. We want to understand what it means for a sentence or phrase to have the particular meaning or range of meanings that it does, and how it comes to have that meaning rather than some other.

A second semantic property that we recognize immediately is **ambiguity**, the property of having more than one meaning. The sentences in (2) give sample facts of this kind:

(2) a. The English sentence *Pedro jumped from the top of the bank* has two meanings.
b. The English sentence *Mad dogs and Englishmen go out in the noonday sun* has two meanings.
c. The English sentence *John saw her duck* has two meanings.

The ambiguities arise from different sources in the three cases. Sentence (2a) is ambiguous according to whether we understand the word *bank* as referring to a financial institution or a fluvial embankment. Ambiguity here arises from one of the component words. By contrast, (2b) is ambiguous according to whether *mad* is taken to apply to both *dogs* and *Englishmen* or to *dogs* alone. Here ambiguity arises not in the words of the clause but rather in how we understand those words as combining. Finally, (2c) involves a combination of what occurs with (2a) and (2b). The pronoun *her* is ambiguous between a possessive form (as in *John saw her book*) and a simple object form (as in *John saw her*). The word *duck* is ambiguous between a noun and a verb. And the phrase *her duck* is ambiguous between a sentencelike phrase meaning she

ducked and a nominal phrase meaning the duck that she owns. Again, an account of ambiguity and also of how ambiguity can arise would appear to be the province of semantic theory.

A third semantic property that we might wish to recognize is **anomaly**: the property having an aberrant meaning. Anomaly is illustrated by the famous sentence in (3), from Chomsky 1957, and by the lines of the children's rhyme in (4), drawn from Leech 1974:

(3) Colorless green ideas sleep furiously.

(4) I went to the pictures tomorrow,
I took a front seat at the back;
I fell from the pit to the gallery,
And broke a front bone in my back.
A lady she gave me some chocolate,
I ate it and gave it her back;
I phoned for a taxi and walked it,
And that's why I never came back.
(Opie, *The Lore and Language of School Children*, p. 25)

In (3) and in each of the verses of (4), we can identify some form of oddness; the sentences are all nonsensical in some way. It seems reasonable to think that semantics should tell us about anomalies and their sources.

Along with semantic properties that hold of individual expressions, we also recognize various semantic relations that hold among them. These include, for example, **logicosemantic relations** such as contradiction, implication, and synonymy. Examples (5) through (7) illustrate:

(5) a. *John believes that the Earth is flat* contradicts *John doubts that the Earth is flat.*
b. *John claims that the Earth is flat* contradicts *John denies that the Earth is flat.*
c. *Some mice migrate* contradicts *No mice migrate.*

(6) a. *John is a human* implies *John is a mammal.*
b. *Mary was laughing and dancing* implies *Mary was dancing.*
c. *Mary usually takes the train* implies *Mary sometimes takes the train.*
d. i. *This is a blue gun* implies *This is blue* and *This is a gun.*
ii. *This is a small moon* implies *This is a moon* but not *This is small.*
iii. *This is a toy gun* implies *This is a toy* and *This is not a gun.*

(7) a. *John sold a car to Mary* is synonymous with *Mary bought a car from John.*
b. *Felicity is a female fox* is synonymous with *Felicity is a vixen.*
c. *John saw Mary* is synonymous with *Mary was seen by John.*
d. *Alice gave a present to Frank* is synonymous with *Alice gave Frank a present.*

The relations in (5) through (7) arise from different sources. Sometimes they issue from relations between particular pairs of words, such as *believe/doubt, claim/deny, usually/sometimes, buy/sell.* In other cases they come from from individual words, like *and.* In still other cases, the relations arise from pairs of sentence forms, such as active/passive, (7c), and an oblique dative versus a double object, (7d). The examples of (6d) make it clear that such data are in fact rather subtle: although each of the right-hand sentences involves an adjective-noun combination and all these combinations are superficially similar in form, the implicational relations are quite different in the three cases.

A second and slightly less familiar family of semantic relations is the group of **thematic relations** illustrated by the examples in (8) to (11), from Jackendoff 1983. Each of these triples displays a common pattern in meaning, a semantic parallelism. Thus the sentences in (8) all express the idea of an object undergoing a change of some kind. In (8a) it is a change of location, in (8b) a change of possession, and in (8c) a change of properties. The sentences in (9) express the common idea of an object traversing some path. The sentences in (10) all express the idea of an object extending over some path. And the sentences in (11) express the idea of an object being oriented along a path:[2]

(8) a. The train traveled from London to Paris.
b. The inheritance passed from John to Mary.
c. The substance changed from liquid to gas.

(9) a. John ran into the house.
b. The mouse skittered toward the clock.
c. The train rambled along the river.

(10) a. The highway extends from Denver to Indianapolis.
b. The flagpole reaches (up) toward the sky.
c. The sidewalk goes around the tree.

(11) a. The sign points to Philadelphia.
b. The house faces away from the mountains.
c. The cannons aim through the tunnel.

It is a fact about English that the sentences in these groups express the same (or significantly similar) thematic relations. A semantic theory should attempt to account for them.

1.1.2 The External Significance of Language

The second main area within the pretheoretical domain of semantics concerns the relation between language and the world, the "external significance of language," to use a phrase from Barwise and Perry 1983.[3] For example, it appears to be a fact about certain words that they make **reference** to specific objects, and a central part of learning these words lies in learning what object or objects they refer to. Thus a crucial aspect of the competent use of the name *Marilyn Monroe* lies in using it to refer to the same Hollywood actress that other people use it to refer to, (12a). Similarly, mastery of the French term *Les Etats Unis* crucially involves learning that it refers to the United States of America, (12b):

(12) a. *Marilyn Monroe* refers to Marilyn Monroe.
b. *Les Etats Unis* refers to the United States of America.

Reference is not a relation like contradiction or implication, which holds between linguistic expressions; rather, it is a relation between expressions and extralinguistic objects, such as people and countries.

A similar class of semantic facts involving language and the world are facts about **truth** and **falsity**. Many of the sentences in any natural language have the property of being either true or false, and for any of these sentences, the property of being true or false depends crucially on meaning. More exactly, the property of being true or false depends on two things: what the sentence means, and how things are in the extralinguistic world. Thus *Camels have humps* is true because it means what it does and because camels do, in fact, have humps. If camels did not have humps, the sentence would be false. And if *Camels have humps* meant that fish are reptiles, again the sentence would be false.

The relationship between meaning and truth is a complex and intimate one, and we will examine it in detail in the next chapter. However, even at this preliminary stage we can see important connections. Observe that this relationship supports the two inference schemata shown in (13):

(13) a. S means that *p*
S is true
---
p

b. S means that *p*
p

S is true

Thus if *Camels have humps* means that camels have humps and *Camels have humps* is true, then camels have humps. And if *Camels have humps* means that camels have humps, and camels do indeed have humps, then the sentence *Camels have humps* is true.

That meaning and truth are related in the way shown in (13) underlies the obvious but enormously significant fact that a natural language is **informative**, both about the world and about its speakers. Suppose that Barry telephones us from southern California and utters the English sentence *The air in L.A. is absolutely foul.* Knowing what this sentence means and also knowing that Barry is a truthful person (i.e., a person who speaks true sentences), we learn something about atmospheric conditions in a place thousands of miles from our homes. We do this by employing scheme (13a) as in (14):

(14) *The air in L.A. is absolutely foul* means
that the air in L.A. is absolutely foul.
The air in L.A. is absolutely foul is true.

The air in L.A. is absolutely foul.

Similarly, suppose that we are trying to determine the truth of a statement whose meaning we know, but whose truth we have cause to doubt. We do this by investigating the world. If Florence tells us that choosy mothers chose Jiff peanut butter and we know what this sentence means, then we proceed by employing scheme (13b) as in (15):

(15) *Choosy mothers chose Jiff* means
that choosy mothers chose Jiff.
Choosy mothers chose Jiff.

Choosy mothers chose Jiff is true.

Upon discovering whether choosy mothers do indeed choose Jiff peanut butter, we are in a position to assess the truth of Florence's statement.

The observation that language has external significance—that we can learn things about the world through speech and that we can learn things about sentences (their truth or falsity) through the world—is, of course, a commonplace, one that we rely on constantly in our daily lives. What is important not

to miss, however, is the central role that meaning plays in this: It is meaning that allows us to reason from from facts about language to facts about the world, and vice versa. It is meaning that is responsible for the essential informativeness of language. This is clearly a fact about meaning that we would want any semantic theory to capture and explain.

1.1.3 The Internal Significance of Language

Finally, the pretheoretical semantic data about natural language also include facts about the relations between language and its speakers. Meaning relates language not only outward to the world of smog and peanut butter but also inward to the mental lives of speakers. Meaning is connected not only with world-oriented notions like truth and reference but also with speaker-oriented notions like saying, believing, asking, ordering, doubting, etc. Again, the relationship between what sentences mean and the assertions, beliefs, questions, and commands that speakers express by using sentences is complex and subtle. However, in paradigm cases some connections are clear. For example, if a woman sincerely utters (16) under typical circumstances, then it is very likely that she is asserting that it is snowing, that she believes that it is snowing, and that she wishes to inform the audience that it is snowing:

(16) It is snowing.

Similarly, if a man utters the interrogative form (17), then under typical circumstances he would be asking whether it is snowing, indicating that he does not know whether it is snowing and that he wishes to know whether it is snowing:

(17) Is it snowing?

Of course, matters are often less straightforward. When someone is being sarcastic, that person may utter a sentence that means that *p* without themselves wishing to assert that *p*, without believing that *p*, or without wishing to inform the audience that *p*. Equally, someone might utter (17) knowing full well that it is snowing: the addressee might, for example, have promised to fix the central heating next time it snowed, and the utterance of (17) might serve as a timely, if somewhat oblique, reminder.

Once again, the observation that language has internal significance—that we can discover what people believe, want, demand, and might do on the basis

of what they say—is a commonplace that we draw on implicitly every day of our lives. And once again, meaning plays a crucial role. In paradigm cases, the meaning of a sentence directly reflects the content of an internal, psychological state. A speaker utters a sentence meaning that *p*. By that act, the speaker asserts that *p*, indicates a belief that *p*, indicates a wish to inform the audience that *p*, and so on. In these cases the speaker says exactly what he or she means. And even in other cases there are evidently crucial connections between the meaning of the sentence uttered and the content that the speaker expresses. In the case of the speaker who uses (17) as a reminder of a promise, it is surely no accident that a sentence normally used to ask whether it is snowing is chosen to convey the meaning.

1.1.4 The Pretheoretical Domain as *Pre*theoretical

Our three subdomains present a fund of data and apparent data to which a nascent semantic theory should be addressed. But it is important to keep in mind that a mature semantic theory will not necessarily take these data as given and explain them under their pretheoretical guise. True, we must begin with a pretheoretical domain of facts. But like any scientist, we must bear in mind that what initially presents itself as relevant data may turn out not to be so further down the road. In the process of constructing a rigorous and explicit theory, we must be prepared for elements in the pretheoretical domain to be reanalyzed and redescribed in various ways.

For example, some data may end up being rejected altogether as merely apparent data, "facts" whose initial inclusion in the domain was based on false beliefs. The daily motion of the sun and planets is a case of this sort. At its inception, the pretheoretical domain of astronomy included the "fact" that heavenly bodies rise in the east, traverse the sky along the ecliptic, and set in the west. But, of course, this fact turned out to be merely an apparent one: the planets do not actually circle the Earth each day; rather, the Earth turns.

Other pretheoretical facts may end up being retained as genuine data but reassigned to some other discipline or collection of disciplines for explanation. The motion of the tides illustrates this case. The explanation for the regular ebbing and flooding of large water bodies was once thought to belong to geological theory: tidal motions were believed to be explained by the local properties of the earth.[4] Later, following the work of Kepler and Newton, it was recognized that tidal movement is not properly a geological phenomenon but

rather an astronomical one: tides are not caused by the motion or inner workings of the Earth but rather by its gravitational interaction with extraterrestrial bodies. Tidal motion was retained as a genuine fact, but the responsibility for its explanation was transferred to planetary astronomy and gravitational physics.

Still other data may be kept as part of the domain in question but significantly redescribed in the process of being brought within a precise theory. Pretheoretical data about the weight and heat of bodies are examples of this kind. While pre-eighteenth-century physics routinely used these familiar, homey notions, modern physical theory significantly redescribes them. The phenomena of weight are reinterpreted in terms of abstruse concepts like gravitational fields and inertial mass. The phenomena of heat are likewise redescribed in terms of the motion of molecules and other elementary constituents of matter.

Thus although our three semantic domains provide a fund of initial data, we must allow that the developed theoretical domain of semantics may count and classify these data in very different terms than those we began with. The process of constructing a systematic and explicit semantic theory from the pretheoretical domain may well involve a considerable amount of rejection, reassignment, and redescription.

1.2 Semantic Theory as a Theory of Linguistic Knowledge

Semantic facts like those surveyed above are verified by the judgments of native speakers. As native speakers of English, we judge the presence of ambiguity and anomaly, of implication and contradiction, and of thematic parallelisms. We judge the reference of terms, and the truth and falsity of sentences. And we judge what people have asserted, queried, or denied on the basis of what they have said, asked, or written. In our view, these judgments do not merely confirm the data of semantics but actually constitute them in an important sense. Human languages are, after all, the products of human minds. Languages have, to a large extent, just those semantic properties that their speakers ascribe to them. It is because English speakers take the string of words *Camels have humps* to mean that camels have humps that those words have this meaning in English. If English speakers all took the sentence to mean that reptiles have wings, then this is what it would mean.

Our ability to make linguistic judgments clearly follows from our knowing the languages that we know. Monolingual speakers of French or Warlpiri are not able to assess the semantic relations between the English sentences in (8) through (11), they are not able to judge the truth or falsity of English sentences like *Camels have humps*, and they are not able to assess the beliefs and thoughts of English speakers on the basis of what they say. We are able to make the judgments that we make because we English speakers have come to know the language that we know. And those who lack this knowledge lack the corresponding abilities. Given this simple observation, one way to construe semantic theory—the theory that attempts to describe and explain semantic facts—is as a theory of knowledge of one special kind. We can see semantics as a theory of the knowledge that underlies our ability to make semantic judgments. Semantic theory addresses one part of our linguistic knowledge: *knowledge of meaning*.[5]

To view the subject matter of semantics as linguistic knowledge is to locate the place of semantic theory within the general enterprise initiated by Noam Chomsky (1965, 1975, 1986a), for whom linguistic theory is a theory of the real knowledge of speakers. This project contrasts with a variety of other commonly held views of the subject matter. For example, some have taken semantics to be a theory of the semantic relations holding between expressions (including inferential and thematic relations).[6] Many others have construed semantics as a theory of the relations holding between language and the world.[7] Still others have insisted that since languages are abstract objects, linguistics (including linguistic semantics) should be pursued as a branch of mathematics.[8] Our conception differs from all of these. On our view, semantics is part of a theory of speakers' knowledge. Facts about language-to-language and language-to-world relations may furnish important clues about the content of this knowledge—they may furnish data—but they are not the object of inquiry itself. The object of inquiry is knowledge of language.[9]

As Chomsky (1986a) has especially emphasized, once we focus on knowledge of language as the object of investigation, three natural questions present themselves:

1. What do we know?
2. How do we come to know it?
3. How is this knowledge used?

In our view, these questions should be approached from a **cognitivist perspective**, according to which knowledge of language is knowledge of a body of (largely unconscious) rules and principles that assign representations and meanings to the physical forms of signs (be they phonetic, visual, or tactile).[10] On this conception, an answer to question 1 would specify these rules and principles and show how they affect the required mapping from sign to structure and meaning. An answer to question 2 would specify how such rules and principles are acquired, including what knowledge the language learner has at the outset and how this interacts with experience to yield adult knowledge. An answer to question 3 would specify how the rules and principles known by speakers are deployed in speech and understanding, and how they interact with other systems of thought and with action. Let us examine these points more carefully.

1.2.1 Properties of Semantic Rules

If knowledge of meaning is knowledge of a body of rules and principles, what sorts of rules and principles would these be? A natural idea is that some of them would tell you about the meanings of individual words and morphemes and others would tell you how these meanings interact when their corresponding expressions are put together in a sentence. Consider the simple example in (18) from this point of view:

(18) Boris kissed Natasha.

Roughly speaking, understanding (18) would involve rules specifying what *Boris* and *Natasha* refer to and what relation is expressed by the verb *kissed.* It would also involve rules that identify the subject of the transitive verb as the agent of the relation (so that it's Boris who does the kissing) and the object of the verb as the patient (so that it's Natasha who is kissed). Semantic rules of this kind are said to be **compositional**. They give the semantic content of a sentence by specifying the semantic contributions of its parts and the semantic significance of putting those parts together according to a definite mode of syntactic combination.

The hypothesis that we know a set of compositional semantic rules and principles is a highly attractive one having a great deal of explanatory power. In particular, it accounts for three notable and closely related features of linguistic competence. First, it explains why *our understanding of sentences is sys-*

tematic—why there are definite, predictable patterns among the sentences we understand. For example, we would confidently predict that anyone who understands (18) will also understand (19), and vice versa:

(19) Natasha kissed Boris.

This is explained by compositionality. Once you know the rules yielding the meaning of (18), you already know enough to understand (19). The same rules allow for the interpretation of both sentences.

Second, the hypothesis accounts for the obvious but important fact that *we can understand new sentences*, sentences that we have never come across before. This too is easily explained if we have a body of rules that allow us to infer the meanings of new sentences from prior knowledge of the meanings of their parts and from knowledge of the semantic significance of their combination.

Third, the hypothesis accounts for the slightly less obvious but equally important fact that *we have the capacity to understand each of an indefinitely large number of sentences*. Consider, for example, the set of sentences in (20), from Platts 1979:

(20) a. The horse behind Pegasus is bald.
b. The horse behind the horse behind Pegasus is bald.
c. The horse behind the horse behind the horse behind Pegasus is bald.
d. The horse behind the horse behind the horse behind the horse behind Pegasus is bald.
⋮

Clearly this list could be extended indefinitely. Yet in some obvious sense, we seem to be able to understand them all. Of course, our actual capacity to use or react to these sentences is limited in certain ways. When the sentences get too long, we cannot get our minds around them: we forget how they began, or we get distracted, or we simply lose track. Consequently, we cannot show our understanding in the usual ways. For example, we cannot explain what the sentences mean or use them in inferences. But it seems that these limitations reflect constraints on such things as memory and attention span and have little to do with specifically linguistic abilities. If we had unlimited attention spans, memories, and so on, we would presumably be able to understand all the sentences in the set.

Such collections of examples show that our linguistic competence is rich enough to yield an infinite number of distinct judgments. But since our brains are finite objects with finite storage capacity, it is clear that this competence must have a finite form. The compositionality hypothesis accounts for this in a straightforward way. If we know a finite vocabulary and a finite number of the right kinds of rules for deriving what sentences mean from what words mean, then we could generate infinitely many results. For example, suppose that the words and phrases in (20a) are grouped together in the way shown by the traditional tree diagram in (21):[11]

(21)

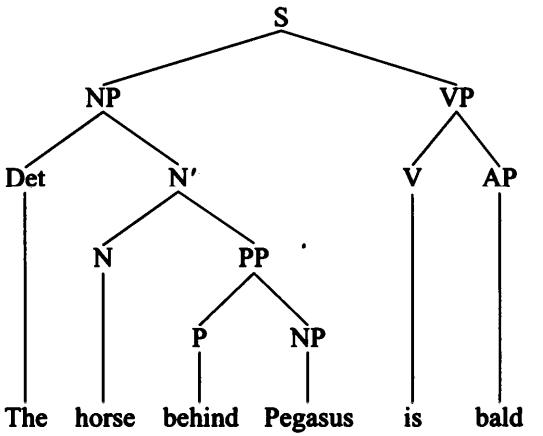

Suppose further that we could assign some kind of semantic contribution or value to each of the leaves of this tree (i.e., to each of the individual words), and suppose that we could give a way of combining these semantic values for each of the branches.[12] Thus, we would have a semantic value for *bald*, for *Pegasus*, for *behind*, etc. And we would also have a general way of combining the values of nominals (N′) and prepositional phrases (PP) in the configuration $[_{N'}$ N PP], a way of combining the values of verbs (V) and adjectives (AP) in the configuration $[_{VP}$ V AP], a way of combining the values of noun phrases (NP) and verb phrases (VP) to yield the meanings of sentences, and so on.

If we could give such semantic values and combination schemes, then we could directly account for our ability to assign meanings to the unbounded sequence of sentences in (20). Consider the tree underlying (20b), depicted in (22):

(22)

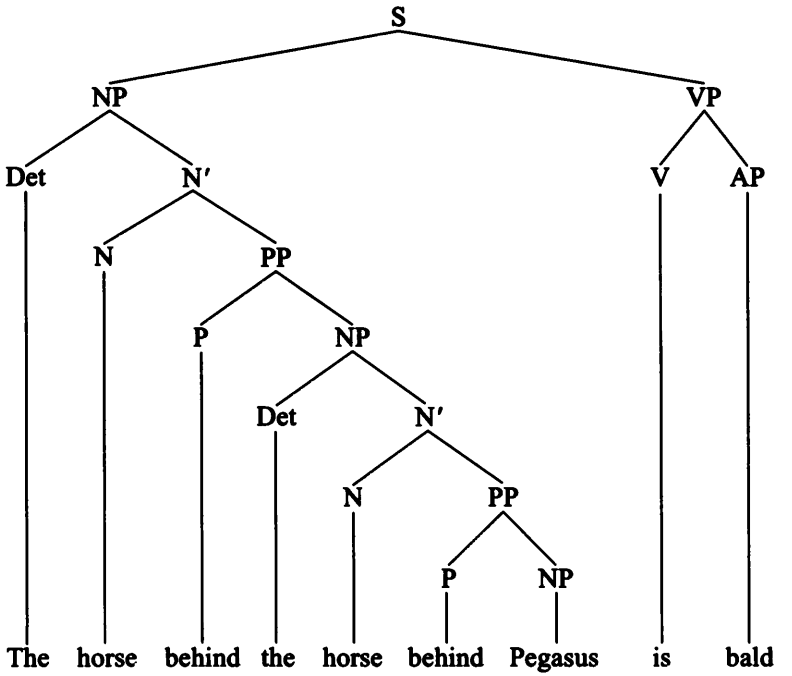

This tree differs from (21) in having extra [$_{NP}$ Det N′] and [$_{N'}$ N PP] branches involving the lexical items *the*, *behind*, and *horse*. But all of these elements already occur in (21). That is, (22) involves recycling the configurations [$_{NP}$ Det N′] and [$_{N'}$ N PP]. It is clear how this recycling occurs. NP introduces an N′ node, which in turn can introduce a PP node, which in turn can introduce an NP again. Configurations of this kind are called **recursive**, and the rules that generate such configurations, rules that have the property that they can apply to their own output, are called **recursive rules**. Given the recursive nature of the configuration in (21), it follows that if we have the semantic resources for computing the meaning of (21), we will "automatically" have the resources for computing the meaning of (22). A compositional semantic theory specifying values and rules for combining them will deliver meanings for (20a), (20b), and indeed *all* the sentences in the list.

It is not difficult to figure out in rough intuitive terms the rules involved in the understanding (18) and (19). But the situation is much less clear with many other examples. Consider the pairs in (23) and (24), adapted from Chomsky 1986a:

(23) a. Boris ate an apple.
b. Boris ate.

(24) a. Boris is too clever to catch Rocky.
b. Boris is too clever to catch.

Sentence (23b) means that Boris ate something or other edible. Comparing (23a) and (23b), then, one might think that whenever a transitive verb appears without a direct object, the patient, or undergoer of the action, is taken to be arbitrary: some typical example of an appropriate patient for the verb. But the pair in (24) shows that this rule is much too simple. Sentence (24b) does not mean that Boris is too clever to catch something or other catchable. Rather, it means that Boris is too clever for anyone to catch him. Interestingly, the implicit subject of the verb *catch* is now taken to be arbitrary, and the direct object is understood as *Boris*.

These examples reveal two important further points about knowledge of language. The first is that some of the knowledge that goes into language understanding is unconscious or **tacit**. We are simply not aware of why we assign grammatical relations in (24) the way we do. Introspection does not reveal the rules that we are following, and indeed there is no obvious explanation of the features noted in (24). We can easily imagine a language in which (24b) meant that Boris is to clever to catch anyone. This would be a perfectly respectable meaning for the sentence to have. But (24b) does not have this meaning in English. Why not? The answer to this question turns out to be exceedingly intricate. What accounts for the meaning of (24b) lies deeply buried in our minds, does not emerge after a few moments of reflection, and indeed very probably cannot be accessed through any form of introspection.

A second point illustrated by (23) and (24) is that to understand a sentence, we are compelled to view it as more than a superficial string of words. In particular, we are led to attribute *structure* to sentences, or better, to take sentences as structured things. To understand (24a), you need to know that *Boris* is the subject of the verb *catch*. To understand (24b), you need to know that the subject of *catch* is not *Boris* but an implicit element meaning approximately the same as *one* or *anyone*, and that *Boris* is the object. Understanding a spoken or written string of words thus involves at least two components: there must be an assignment of structure, and there must be an assignment of meaning based on the meanings of the individual elements and the effect of combining them in that particular structure. The rules that account for the

first component are syntactic rules, and those that account for the second are semantic rules.[13] We will discuss the relationship between syntax and semantics in more detail in chapter 3.

1.2.2 Acquisition of Semantic Rules

Knowledge of language is something acquired over a period of time. When born into a linguistic environment, a normal human infant will, after a few years, gain knowledge of the rules of language. This distinguishes human infants from rocks, chimpanzees, and indeed everything else. Human infants thus have some property *P* lacking in rocks, chimpanzees, etc., that enables them to learn language when placed in a suitable environment. *P*, together with the infant's experiences, account for its acquiring linguistic competence.

An answer to our second question about knowledge of meaning—how we come to know it—will explain how adult knowledge of semantic rules and principles is achieved and the nature of the property underlying this achievement. More precisely, it will specify what semantic knowledge is available to the child at the outset of the learning period, and how this knowledge interacts with experience to yield adult competence. This task is a formidable one, since, as in so many other domains of human learning, there is a considerable gap between what adults come to know and the evidence by which they come to know it.

Consider first an example from the domain of word meaning due to Susan Carey (1984). Imagine a child in the process of acquiring the meaning of the English noun *rabbit*.[14] The child's parent shows it a number of pink, velour, stuffed rabbits. Each time the parent picks up a rabbit, it says the word *rabbit*. What conjectures can the child make about the meaning of *rabbit*? One possibility, of course, is that *rabbit* means rabbits, the concept that adult speakers eventually acquire. But clearly there are other hypotheses fitting the data that the parent has provided; indeed, there are indefinitely many such hypotheses. For example, *rabbit* could mean toy rabbit, or pink velour, or pink, or rabbit or chicken, or rabbit or planet, or rabbit or prime number lying between 8642 and 8652, and so on. Some of these hypotheses might be tested and corrected against further data, but clearly for a great many possibilities (like the last), such checking would be difficult to do. Furthermore, there are some hypotheses that no such empirical tests will serve to distinguish. For example, suppose the child is attempting to decide whether the word *rabbit* means rabbit versus undetached rabbit part, that is, piece of a rabbit connected to a

whole rabbit.[15] Since picking up a rabbit always entails grasping some undisconnected piece of it (a paw, a tail, an ear, etc.), any evidence for the former hypothesis will also be evidence for the latter. No simple visual presentation will tell them apart.

In a similar vein, consider the child exposed to verbs like *chase*, *buy*, and *give* in a context where the parent is using hand puppets to act out various scenes while uttering sentences like *Big Bird is chasing Oscar*, or *Ernie is buying a block from Kermit*, or *Fozzy is giving a block to Bert*. In these circumstances the child might indeed conjecture that *chase* means to chase, *buy* means to buy, and *give* means to give. However, as Gleitman (1990) points out, for each of these predicates, there is another available alternative. Any situation in which *Big Bird is chasing Oscar* is true will also be one in which *Oscar is fleeing Big Bird* is true; any situation in which *Ernie is buying a block from Kermit* is true is also one in which *Kermit is selling a block to Ernie* is true; and any situation in which *Fozzy is giving a block to Bert* is true is one in which *Bert is getting a block from Fozzy* is true. *Chase*, *buy*, and *give* each have a symmetric counterpart describing a closely related situation involving similar roles for the participants. In view of this, it will also be plausible for the child to conjecture that *chase* means to flee, *buy* means to sell, and that *give* means to get. In and of themselves, the scenes presented by the parent will not tell these conjectures apart.[16]

Finally, consider a child in the process of fixing the meaning of nominal structures of the form [$_{NP}$ NP*'s* N'], as they occur in examples like *Rosa's book* or *Kyoko's picture of Max*. One possible conjecture is that [$_{NP}$ NP*'s* N'] refers to the object NP possesses having the property described by N' (thus, the object Rosa possesses that is a book, the object Kyoko possesses that is a picture of Max, etc.). This is roughly the interpretation rule that adults come to know. But an indefinite number of other hypotheses are also possible. Thus children might conjecture that [$_{NP}$ NP*'s* N'] means the object that is NP and is an N' (e.g., Rosa, who is a book), or they might take it to refer to the object that NP either possesses *or* likes, having the property described by N', etc. Again, some of these hypotheses will be eliminable through simple stimulus data, but not all them. For example, suppose the child takes [$_{NP}$ NP*'s* N'] to refer to the object NP possesses that has the property described by N' and is either a number or else not a number. Then there will be no simple situations in which the "correct" adult rule yields a true sentence and the alternative rule yields a false one. That is, *Rosa's book is red* will be true under the adult rule exactly when it is true under the alternative rule.

These examples illustrate the basic **induction problem** facing theories of semantic acquisition: the rules that children ultimately adopt regarding word and phrase meaning are substantially underdetermined by the evidence that the learning environment provides. Given this disparity between knowledge and evidence, the conclusion seems inevitable that there are a priori principles in play allowing children to make the right choices. That is, in acquiring the meanings of words and sentences, it seems that children must be guided by antecedently known constraints permitting them to select appropriate hypotheses in the face of minimal data.

At present, the principles children deploy in acquiring word and phrase meanings are only beginning to be understood clearly. And indeed, different kinds of principles and data appear to be involved for different parts of the grammar. For example, in resolving whether a noun like *rabbit* refers to rabbits or undetached rabbit parts, it appears that children deploy principles involving knowledge of objects and their properties, and their familiarity with the word that accompanies the object.[17] Thus if children are presented with an unfamiliar object (a rabbit) under an unfamiliar term (*rabbit*), they will generally conjecture that the term applies maximally to the object as a whole (so that *rabbit* refers to rabbits and not rabbit parts). On the other hand, if a familiar object (a rabbit) composed of an unfamiliar substance (e.g., velour) is presented under an unfamiliar term (*velour*), they will conjecture that the new term refers to the substance of which the object is composed (*velour* refers to velour material). And so on. Such reasoning evidently requires children to have a substantial theory allowing them to isolate objects and their properties, and a variety of experimenters have argued that such a theory is available from a very early age.

In fixing the meanings of vocabulary items like verbs, more specifically linguistic data seem to be involved. A number of researchers have argued that children are able to deduce substantial aspects of the meanings of predicates from the syntactic structures in which they appear. Consider the case of *chase*, *buy*, and *give* versus *flee*, *sell*, and *receive*. One prominent difference between these predicates involves who is the agent in each. Although *Ernie is buying a block from Kermit* and *Kermit is selling a block to Ernie* are true in the same circumstances, we understand Ernie to be the agent in the former but Kermit to be the agent in the latter. Suppose, then, that children deploy universal linking rules specifying that in a structure of the relevant kind (that is, roughly, NP_1 V NP_2 P NP_3), the subject noun phrase (NP_1) must always refer to the individual playing the agent role. Then they will be able to distinguish *buy* and

sell correctly in terms of the environmental data and the syntactic forms that accompany their presentation. Hearing *Ernie is buying a block from Kermit*, observing the situation described, and knowing the linking rule just given, the child can reason that V denotes events of commercial exchange in which the recipient of the object exchanged is also the agent. That is, V denotes buyings. On the other hand, hearing *Kermit is selling a block to Ernie*, observing the situation described, and knowing the linking rule, the child can reason that V denotes events of commercial exchange in which the source of the object exchanged is also the agent. That is, V denotes sellings.

The picture that emerges from these results is a familiar one under the cognitive perspective. The child appears as as active agent who must deduce the rules of its language on the basis of available evidence. Given the underdetermination of hypotheses by data, such deductions evidently require the presence of substantial constraints limiting the space of possible conjectures. Without them, the language learner would be forever in search of data to rule out extravagant hypotheses. Since the principles in question apply to all human language learning, they will clearly constrain the possibilities of all learnable languages. Only languages whose rules conform to the principles will be humanly learnable. So, to pursue the above example of *rabbit*, it seems plausible to conclude that whatever principles intervene to fix the meaning of the English noun *rabbit* as rabbit, and not rabbit or prime number between 8642 and 8652, will also be the operating in French to fix the meaning of *lapin* as rabbit, and not rabbit or prime number between 8642 and 8652, and similarly for many other nouns and many other possible hypotheses. Such principles specify universal constraints on human language, and so are appropriately called **universal grammar** (UG). On the cognitivist view, property *P* is knowledge of UG, and a cognitivist answer to question 3 would specify UG and show how knowledge of particular languages results from the interaction of UG with experience.

As we commented above, it seems clear that the ability to learn languages must be a species-specific property of humans (Chomsky 1975, 1986a). No nonhuman animal or other earthly object learns language when placed in a typical human habitat. Moreover, children of all races and nationalities appear equally able to learn all languages: a child born of English parents but brought up by Japanese speakers will typically learn Japanese. To be sure, not every human can learn language. But those that cannot are atypical of the species in precisely that respect. So it is reasonable to conclude that it is part of our genetic endowment.[18]

1.2.3 Use of Semantic Rules

Knowing a language is not sufficient by itself to enable one to understand and to speak. Such knowledge must be applied to the problems of interpreting the sentences one hears and finding the right words and structures to express what one wants to say. An answer to our third question—how knowledge of meaning is used—would describe the mechanisms that access this knowledge and explain how they work. We will confine ourselves to some very brief remarks on the topic.

When we express ourselves in language, we move various parts of our bodies (our vocal chords, our tongue, our lips, our lungs, our hands) to produce the words we wish to produce. These motions result from the intention to produce just these words and no others. Very little is known about how such intentions are arrived at. Introspection is of little help. When we converse, it often appears as though we think our thoughts even as we speak them. If this appearance is correct in at least some cases, then in these cases the problem of explaining how we find the words to express ourselves (what Chomsky calls the "production problem") is a subproblem within the much larger general problem of how we arrive at the thoughts we have. If the appearance is not correct, then the alternative hypothesis is presumably that the thought is present in our minds in some form prior to its expression and some cognitive processes are responsible for selecting the best words of natural language to convey the thought.

It is clear that at least sometimes there is a gap between thought and its expression. To borrow an example from Higginbotham 1987, someone might be on the verge of using a particular sentence, realize that it is ambiguous, and so choose another one instead. In such a case, a small aspect of the subject's deploying knowledge of meaning is introspectively revealed: he or she knows the candidate sentence is ambiguous and reasons that it is unsuitable for communicating. But this casts little light on the larger process: what leads the subject to the initial idea of uttering the first sentence, and what leads to the choice of the next sentence? About such matters almost nothing is known.

The process of using knowledge of language in understanding perceived sentences is slightly less mysterious. We can distinguish at least three kinds of processes involved. First there is what is called **parsing**. When you hear a sentence of your own language, you need to identify the speech sounds for what they are, that is, you must identify its phonological form. You must identify the syntactic arrangement of its constituents. You must also identify the

meanings of its words and how they compose. This much seems clear and is widely accepted. What is not clear and is much more controversial is how all this is carried out. It seems possible a priori that the process is serial: first you identify the phonology, then the syntax, then the semantics. But it is also quite possible that the task is staggered or occurs fully in parallel, so that, for example, you compute the semantics at the very same time you are computing the phonology. Research in this area is very active at present, and both the serial and parallel views are under intense scrutiny and debate.[19]

Parsing concerns the application of strictly linguistic knowledge only. Linguistic knowledge will provide only what we might call the **context-independent meaning of an utterance**. To fully understand an utterance of a sentence, more than this is often required. For example, if someone utters (25), your knowledge of the language will not tell you who *she* refers to on this occasion or which place is identified by *here*:

(25) She is here.

These are context-dependent features of utterance meaning. Your knowledge of language will tell you, very roughly, that the utterance means that some independently identified female is near the speaker. The second kind of process involved in language understanding, then, concerns identifying the relevant features of context and combining them with knowledge of language to arrive at a full interpretation of the utterance.[20]

The third kind of process is the application of knowledge of general conversational principles—**pragmatics**. Knowledge of language provides only the literal meanings of sentences. But there is often a gap between what is said with a sentence (fully interpreted in a context) and what a speaker using the sentence intends to convey. Pragmatics is required to bridge the gap. Consider how you would understand the following testimonial by Professor Williams about Dr. Stevens, who is a candidate for a philosophy position:[21]

(26) Mr. Stevens command of English is excellent, and his attendance at tutorials has been regular.

By writing (26), Williams conveys the message that Stevens is not very good as a philosopher. How so? Grice (1975) suggests that there is a general principle of "quantity" governing linguistic communication: make your communication as informative as required. Using this principle and one's general knowledge of the circumstances, one can infer the conveyed message: Williams is surely in a position to comment on Stevens's philosophical ability and knows that this

is what she is supposed to be doing. Williams must therefore have some reason for not wishing to provide the required information explicitly. But since she has bothered to write, she must wish to get the information across in some other way. The natural conclusion is that she thinks that Stevens is no good at philosophy.

Semantic knowledge is thus used along with a variety of other kinds of knowledge in understanding language. It is used along with knowledge of phonology and syntax in the parsing processes that deliver the literal, context-independent meaning of a perceived sentence form. This knowledge is combined with knowledge of relevant features of context in further inferences to yield an interpretation of the sentence in the context. And it is combined also with knowledge of pragmatics in yet more inferential processes that provide conclusions about what people are saying.[22]

1.3 The Place and Responsibilities of Semantic Theory

Semantic theory as we have sketched it here is a component of the larger enterprise of cognitive linguistics. As such, it rests on two major empirical assumptions. Like cognitive linguistics as a whole, it assumes that linguistic competence consists of an unconscious body of knowledge, what Chomsky has termed the "language faculty." Furthermore, it assumes that the language faculty contains a specifically semantic module: a particular, isolable domain of linguistic knowledge beyond phonology, morphology, syntax, etc., that is concerned with meaning.

In our picture, the semantic module has contents of a very specific sort. The module contains specifications of meaning for the simplest expressions of the language and rules for deducing the meanings of complex expressions on the basis of the meanings of their parts and the structural configuration in which they occur. Moreover, the semantic module occupies a definite place both within the language faculty and within the larger cognitive system. In the language faculty, semantics is connected to syntax, yielding meanings for the structures that syntax provides. It is also connected to those modules mentioned under the heading "Use of Semantic Rules": the pragmatics module and the parser (figure 1.1).

In the larger cognitive domain, we will later see grounds for thinking that semantics is connected to a module containing knowledge of inferential principles (tacit logic), and to the modules containing our implicit theories of objects

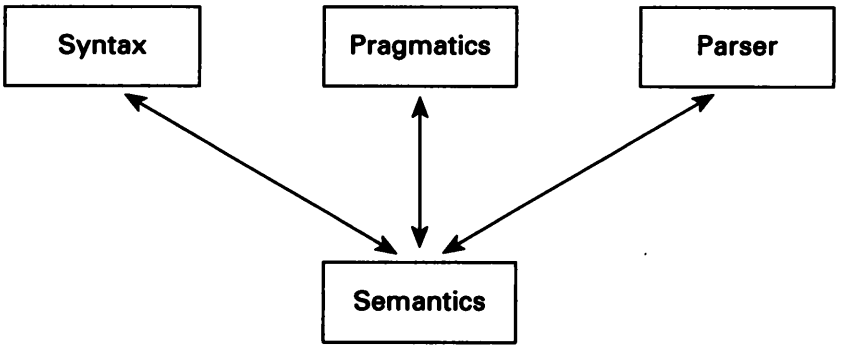

Figure 1.1 The place of semantics in the language faculty.

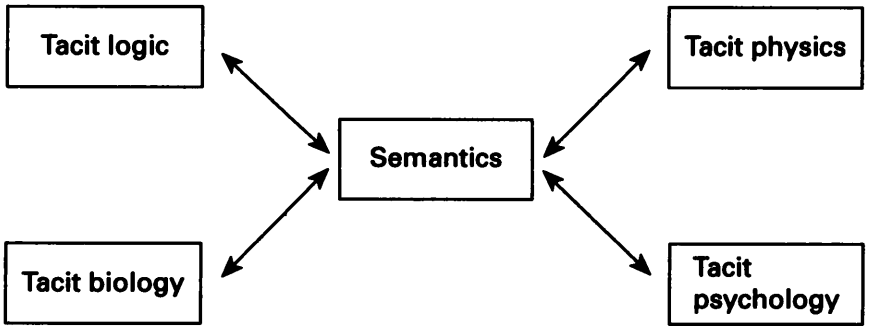

Figure 1.2 The place of semantics in the larger cognitive domain.

and forces (tacit physics), of creatures and their goals (tacit biology), and of people's cognitive states and their actions (tacit psychology) (figure 1.2).[23]

The task of semantic theory, then, is to specify precisely the contents of the semantics module. We must say what knowledge is contained in it, and we must show that this knowledge is sufficient to fulfill its purpose and explain semantic facts. If the enterprise fails, this will reflect back on one of the two assumptions on which the theory is built: Our efforts will fail because the whole enterprise of cognitive linguistics is mistaken. Or they will fail because there is no semantics module. It is worth reflecting on this latter point, since, however plausible it may seem, the assumption that there is a semantics module within the language faculty cannot be demonstrated a priori. It is conceivable that the genuinely modular areas of this faculty—the areas that contain a

definite set of rules and principles with a well-defined role in cognition—do not include a semantics.[24] In such a case, syntactic knowledge would feed directly into a variety of different processors containing different kinds of knowledge and serving different tasks involved in speech, understanding, and (no doubt) other matters as well.[25]

There is no quick way of establishing the truth of either of our guiding assumptions. In both cases we must wait on the development of the best overall theory. To establish the correctness of cognitivist linguistics as a whole, it is necessary to develop a relatively detailed and complete theory and to demonstrate, first, that it works—that if it were true, it would account for a very significant body of data—and, second, that it outperforms competitors (if there are any). To establish the existence of the semantics module, one must work within linguistics and show that the most successful linguistic theory overall is one that includes semantics as a part. It is to the latter task that this book is addressed. In the next chapter we articulate the general form of semantic theory, explaining exactly what kinds of rules and principles are involved in semantic knowledge and showing how a theory of this form could play the role assigned to it. In chapter 3 we specify in detail the relations between semantics and syntax. In the following chapters we deal with a variety of constructions in natural language and develop detailed and well-supported, specific theories of their semantic functioning. We conclude, in the final chapter, with a deeper and broader discussion of the conceptual issues and methodology underlying our approach.