## 1   Introduction

Semantics is the branch of linguistics devoted to the investigation of linguistic meaning, the interpretation of expressions in a language system. We do not attempt a comprehensive survey of the many different approaches to semantics in recent linguistics but choose instead to introduce a particular framework in some detail. Many of the concepts and analytical techniques we introduce have their origins in logic and the philosophy of language; we apply them to the study of actual human languages.

When we say that our focus is on semantics as a branch of linguistics, we are adopting a particular conception of the methods and goals of linguistic inquiry. That conception is rooted in the generative paradigm that began to reshape the field of linguistics in fundamental ways over forty years ago. Noam Chomsky's *Syntactic Structures*, published in 1957, introduced the three key ideas that we take to be definitive of that paradigm.

The first is the idea that a grammar of a language can be viewed as a set of abstract devices, rule systems, and principles that serve to characterize formally various properties of the well-formed sentences of that language. The grammar, in this sense, generates the language. This idea was already established in the study of various artificial languages within logic and the infant field of computer science; what was novel was Chomsky's claim that natural languages—the kind we all learn to speak and understand in early childhood—could also be generated by such formal systems. In a sense, when linguists adopted this view, they adopted the idea that theoretical linguistics is a branch of (applied) mathematics and in this respect like contemporary theoretical physics and chemistry.

Few generative linguists, however, would be completely comfortable with such a characterization of their discipline. A major reason for their finding it inadequate lies in the second key idea Chomsky introduced, namely, that generative grammars are psychologically real in the sense that they constitute accurate models of the (implicit) knowledge that underlies the actual production

and interpretation of utterances by native speakers. Chomsky himself has never spoken of linguistics as part of mathematics but has frequently described it as a branch of cognitive psychology. It is the application of mathematical models to the study of the cognitive phenomenon of linguistic knowledge that most generative linguists recognize as their aim. Again, the parallel with a science like physics is clear. To the extent that their interest is in mathematical systems as models of physical phenomena rather than in the formal properties of the systems for their own sake, physicists are not mathematicians. A single individual may, of course, be both a mathematician and a linguist (or a physicist). But as linguists, our focus is on modeling the cognitive systems whose operation in some sense "explains" linguistic phenomena. Linguistics is an empirical science, and in that respect it is like physics and unlike (pure) mathematics.

The third idea we want to draw from the generative paradigm is intimately connected to the first two: linguistics cannot be limited to the documentation of what is said and how it is interpreted—our actual *performance* as speakers and hearers—any more than physics can limit its subject matter to the documentation of measurements and meter readings of directly observable physical phenomena. The linguistic knowledge we seek to model, speakers' *competence*, must be distinguished from their observable linguistic behavior. Both the linguist and the physicist posit abstract theoretical entities that help explain the observed phenomena and predict further observations under specified conditions.

The distinction between competence and performance has sometimes been abused and often misunderstood. We want to emphasize that we are not drawing it in order to claim that linguists should ignore performance, that observations of how people use language are irrelevant to linguistic theory. On the contrary, the distinction is important precisely because observations of naturally occurring linguistic behavior are critical kinds of data against which generative linguists test their theories. They are not, however, the only kinds of data available. For example, linguists often ask native speakers (sometimes themselves) for intuitive judgments as to whether certain strings of words in a given language constitute a well-formed or grammatical sentence of that language. Such judgments are also data, but they seldom come "naturally."

Our approach to semantics lies in the generative tradition in the sense that it adopts the three key ideas sketched above: (1) that generative grammars of formal (artificial) languages are models of the grammars of natural languages, (2) which are realized in human minds as cognitive systems (3) that are distinct from the directly observable human linguistic behavior they help to explain. This tradition started, as we have noted, with important advances in the study of syntax; fairly soon thereafter it bore fruit in phonology. There was important semantic work done by generative grammarians from the early sixties on, but it was not until the end of the sixties that systematic ways of linking the semantic methods developed by logicians to the generative enterprise were found. In our view, this development constitutes a breakthrough of enormous significance, one whose consequences linguists will be exploring for some time. One of our main aims in this book is to introduce the concepts and methods that made the breakthrough possible and to indicate some of the ways logical semantics so conceived contributes to the generative enterprise in linguistics.

We begin by considering some of the linguistic phenomena that one might ask a semantic theory to account for, the range of data that seem at first glance centrally to involve meaning. Our first observation may discourage some readers: there is not total agreement on exactly which facts comprise that range. But this is hardly surprising. Recent discussions of epistemology and the philosophy of science repeatedly claim that there are no "raw" or "pure" data, that abstract principles come into play even in preliminary individuation of a given constellation of facts. Thus, identifying phenomena is itself inescapably theory-laden. We will try, however, to introduce data here that are bound to our particular theoretical hypotheses only weakly. That is, accounting for (most of) these data seems a goal shared by many different approaches to semantics.

A second point to remember is that phenomena that pretheoretically involve meaning may prove not to be homogeneous. This too is unsurprising. Linguists have long recognized the heterogeneity of linguistic phenomena and so have divided the study of linguistic forms minimally into phonology and syntax and have further articulated each of these fields. And, of course, it is recognized that syntax and phonology themselves interact with other cognitive systems and processes in explaining, for example, how people arrange and pronounce words in producing utterances. Similarly, the study of meaning is bound to be parcelled out to a variety of disciplines and perhaps also to different branches of linguistics. A major aim of this book is to explore the question of how linguistic

investigations of meaning interact with the study of other cognitive systems and processes in our coming better to understand what is involved in the production and interpretation of utterances by native speakers of a language.

It seems very likely that certain aspects of utterance meaning fall outside the realm of semantic theorizing. It has been argued, for example, that some aspects of meaning are primarily to be explained in terms of theories of action. Several different sorts of pragmatic theory adopt this approach. Speech act theories, for example, focus on what people are doing in producing utterances: asserting, questioning, entreating, and so on. Such theories can help explain how people manage to mean more than they actually say by looking at the socially directed intentional actions of speakers.

Here is an example where what is meant might go beyond the meaning of what is said. Suppose Molly is at a restaurant and says to her waiter, "I'd like a glass of water." In a clear sense Molly has not directly asked the waiter to bring her a glass of water, yet she means much the same thing by her utterance as if she had said, "Bring me a glass of water." But if Molly utters "I'd like a glass of water" to her hiking companions as they ascend the final hundred feet of a long trail from the bottom to the top of the Grand Canyon, the interpretation is different. In the latter case she probably means simply to report on her desires and not to make a request of her fellow hiker. How do we know this? Presumably in part because we know that Molly cannot be expecting her words to move her walking companion to produce a glass of water for her, whereas she might well intend those same words so to move the waiter in the restaurant. This knowledge has to do with our experience of restaurants and hiking trails and with general expectations about people's motives in speaking to one another.

Understanding what Molly means by her utterance to a particular addressee seems, then, to involve at least two different kinds of knowledge. On the one hand, we must know the meaning of what she has explicitly said—in this case, what the English sentence "I'd like a glass of water" means. Roughly, semantics can be thought of as explicating aspects of interpretation that depend only on the language system and not on how people put it to use. In slightly different terms we might say that semantics deals with the interpretation of *linguistic expressions*, of what remains constant whenever a given expression is uttered. On the other hand, we will

not understand what Molly means in uttering that sentence unless we also know why she has bothered to utter it in the particular surroundings in which she and her addressee are placed—in this case, whether she is trying to do more than update her addressee on her internal state. Pragmatics is the study of *situated uses* of language, and it addresses such questions as the status of utterances as actions with certain kinds of intended effects. Since direct experience with interpretation of language is experience with interpreting uses, however, we cannot always be sure in advance which phenomena will fall exclusively in the domain of semantics and which will turn out to require attention to pragmatic factors as well.

As our adoption of the generative paradigm implies, we take linguistics to include not only the study of languages and their interpretations as abstract systems but also the study of how such systems are represented in human minds and used by human agents to express their thoughts and communicate with others. Thus we develop our semantic theory with a view to its interaction with a pragmatic theory. We will consider not only what linguistic expressions themselves mean (semantics in the strict sense) but also what speakers mean in using them (pragmatics). In this chapter, unless a distinction is explicitly drawn, semantic(s) should be thought of as shorthand for semantic(s)/pragmatic(s).

For most of our initial discussion we can safely ignore the important theoretical distinction between interpreted linguistic forms on the one hand (what, say, the English sentence "I'd like a glass of water" means) and interpreted utterances on the other (what Molly's utterance of "I'd like a glass of water" means). The issue of just how semantics should be related to more pragmatically oriented theories of information processing is wide open, however, and we will return to it at various points.

What should semantics, broadly construed, take as its subject matter? The rest of this chapter addresses this question. Our discussion is intended not to be exhaustive but only indicative of the range of language-related phenomena relevant to inquiry about meaning.

The third section of this chapter considers implication relations between sentences that speakers seem to recognize on the basis of their knowledge of meaning. The fourth and final section considers a number of other semantic properties and relations that speakers' intuitive judgments reveal, some of which are in some

sense parasitic on implication relations. Such judgments are often very subtle, and learning how to tap semantic intuitions reliably and discriminate among the distinct phenomena that give rise to them is an important part of learning to do semantics. In a real sense, such intuitive judgments constitute the core of the empirical data against which semantic theories must be judged.

## al Constraints on Semantic Theory

Before we can fruitfully consider particular varieties of intuitive judgments of semantic properties and relations, we need to consider some general properties of semantic competence.

### 2.1 The productivity of linguistic meaning

It is a familiar but no less remarkable fact that indefinitely many syntactically complex linguistic expressions in a language can have linguistic meanings associated with them. This is simply the semantic analogue of the fact that indefinitely many complex linguistic expressions can be classed as syntactically well-formed by the grammar.

We have no trouble whatsoever in grasping the meaning of sentences even if we have never encountered them before. Consider

(1) I saw a pink whale in the parking lot.

Few if any of our readers will have heard or seen this particular sentence before. Yet you can quite easily understand it. How is this feat possible? The experience of understanding a newly encountered sentence like (1) seems much like the experience of adding two numbers we have never summed before, say

(2) $1437.952 + 21.84$

We can do the sum in (2) and come up with 1459.792 because we know something about numbers and have an algorithm or rule for adding them together. For instance, we may break each of the two numbers to be summed into smaller pieces, adding first the digits in the thousandths place (having added a 0 in that place to the second number), moving on to the hundredths place, and so on. All we really have to know are the numbers (on this approach, the significance of the decimal representation of each number in a base

ten system) and how to sum single digits, and we are then in business. By the same token, we presumably understand a sentence like (1) because we know what the single words in it mean (what *pink* and *whale* mean, for example) and we have an algorithm of some kind for combining them. Thus part of the task of semantics must be to say something about what word meaning might be and something about the algorithms for combining those word meanings to arrive at phrasal and sentential meanings.

Whatever linguistic meaning is like, there must be some sort of compositional account of the interpretation of complex expressions as composed or constructed from the interpretations of their parts and thus ultimately from the interpretations of the (finitely many) simple expressions contained in them and of the syntactic structures in which they occur. We will speak of the simplest expressions as words, except when we want to recognize semantically relevant morphological structure internal to words. Sentences are complex expressions of special importance, but smaller phrases are also semantically relevant. We also briefly look at interpretive phenomena that go beyond single sentences and involve discourse.

In theory the semantically relevant structure of a complex expression like a sentence might bear little or no relation to the syntactic structure assigned to it on other linguistic grounds (on the basis, for example, of grammaticality judgments and intuitions about syntactic constituency). In practice, many linguists assume that semantics is fed fairly directly by syntax and that surface syntactic constituents will generally be units for purposes of semantic composition. And even more linguists would expect the units of semantic composition to be units at some level of syntactic structure, though perhaps at a more abstract level than the surface.

Logicians used to be notorious among linguists for their pronouncements on the "illogicality" of natural language surface syntax. More recently, however, logical approaches to semantics have proposed that the surface syntactic structure of natural language is a much better guide to semantic constituency than it might at first seem to be. Both syntax and the relevant areas of logic have developed rapidly in recent years, but it is still an open question just how close the correspondence is between the structure needed for constructing sentential meanings (what we might think of as semantic structure) and that needed for constructing sentences as syntactic objects. There is also a vigorous debate about

whether more sophisticated approaches to semantics and syntax make it possible to dispense with multiple levels of syntactic structure.[1]

Certainly, however, interpretations of both words and syntactic constructions will play a role in any systematic account of how sentences (and larger discourse texts) are assigned interpretations. An important test of a semantic theory is set by compositionality. Can the theory generate the required interpretations for complex expressions from a specification of interpretations for the basic items? As we will see, explicit specification of how word meanings are combined to produce sentential meanings is not a trivial task.

## 2.2  Semantic universals

A fundamental concern of generative linguistics is to specify what characteristics seem to be constitutive of the human language capacity. In what ways are languages fundamentally alike? We may also be able to say some very interesting things about the ways in which that linguistic capacity constrains possible differences among languages, about the parameters of variation.

There is rather little that might count as semantic typology or as a direct analogue to the parametric approach in syntax.[2] There has, however, been some attention to semantic universals. In the late sixties and early seventies, quite interesting attempts to get at universal semantic principles came from the so-called generative semanticists. Working in the generative tradition, these linguists claimed that semantics was fundamentally just a very abstract level of syntax where a universally available stock of basic words or concepts were combined. The syntax of this universal semantic base was simple, involving a very few categories and rules for combining them. Getting from these abstract structures to the surface sentences of a natural language involved, among other things, replacing complex structures with single words. It was hypothesized, for example, that something like the structure in (3) is the source of English *kill*; a lexical substitution rule replaces the tree with the single word *kill*. Small capital letters indicate that the words represented are from the universal semantic lexicon. (Generative semanticists used V for simple verbs and for other predicate expressions, including predicate adjectives and the negative particle *not*.)

(3)
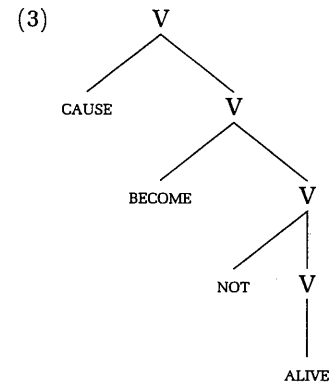


From this standpoint, it is natural to look to syntactic structures for constraints on what might possibly get lexicalized. McCawley (1971), for example, claimed that there could not be a word, say *flimp*, meaning to kiss a girl who is allergic to ..., that is, that no sentence of form (4*a*) could mean what is meant by (4*b*).

(4)  *a.* Lee flimped garlic.
     *b.* Lee kissed a girl who is allergic to garlic.

The explanation he offered was that lexical substitution rules have to replace single constituents and *kiss a girl who is allergic to* is not a single constituent. Of course, since the replaced elements come from a universal language that is not spoken by anyone, it is not easy to be sure that something with the meaning in question might not be expressible as a single constituent. The verb *flimp* might be introduced in a group that thinks that kissing a girl allergic to a certain substance in some interesting way affects the kisser's relation to the substance (perhaps allergies can be so transmitted, so flimping puts the flimper in jeopardy of acquiring an allergy). What is interesting, though, is McCawley's attempt to offer a formal account of alleged material universals, such as the absence from all languages of words like *flimp*.[3] We discuss lexical meanings in somewhat more detail in chapter 8.

Even if this particular approach to the kinds of words languages will have may now seem inadequate, the general idea of attempting to find explanations in terms of general linguistic principles for what can and cannot be lexicalized is of considerable interest. For instance, we do not know of any languages that lack a word that is more or less synonymous with *and*, joining expressions from different syntactic (and semantic) categories—sentences, noun phrases, or prepositional phrases—by using what can be seen as

the same semantic operation. Nor do we know of a language that uses a single word to mean what is meant by *not all* in English yet uses a syntactically complex expression to mean what *none* means. Although it is often said that comparatives (*taller*) are semantically simpler than the corresponding absolutes (*tall*), no language we know of expresses the comparative notion as a single morpheme and the absolute in a more complex way. Can semantic theory shed light on such observations (on the assumption that they are indeed correct)?

Certain quite abstract semantic notions seem to play an important role in many cross-linguistic generalizations. For example, agent, cause, change, goal, and source have been among the *thematic roles* proposed to link verb meanings with their arguments. Fillmore (1968) suggested a semantic case grammar in which predicates were universally specified in terms of the thematic roles associated with their arguments. Language-specific rules, along with some universal principles ranking the different thematic roles, then mapped the arguments of a verb into appropriate syntactic or morphological structures. The UCLA Syntax Project reported on in Stockwell, Schachter, and Partee (1973) adapted Fillmore's framework in developing a computational implementation of their grammar, and similar ideas have figured in other computational approaches to linguistic analysis. We discuss thematic roles in somewhat more detail in chapter 8.

Are such notions part of universal grammar, or is there another way to think about them? Are they connected more to general cognitive phenomena than to language as such? Perhaps, but in any case, certain empirical generalizations about linguistic phenomena seem linked to these semantic notions. For example, in language after language the words and constructions used to speak about space and spatial relations (including motion) are recycled to speak of more abstract domains, for example, possession. The precise details are not universal: Finnish uses the locative case in many instances where English would use the nonspatial verb *have* ("Minulla on kissa" literally glosses as "At me is a cat" but is equivalent to "I have a cat"). But English does use spatial verbs and prepositions to talk about changes in possession ("The silver tea set went to Mary"). The general claim, however, is that resources for describing perceptual experience and the principles that organize them are universally redeployed to speak of matters that are less concrete. As Jackendoff (1983, 188–189) puts it,

> In exploring the organization of concepts that ... lack perceptual counterparts, we do not have to start de novo. Rather, we can constrain the possible hypotheses about such concepts by adapting, insofar as possible, the independently motivated algebra of spatial concepts to our new purposes. The psychological claim behind this methodology is that the mind does not manufacture abstract concepts out of thin air, either. It adapts machinery that is already available, both in the development of the individual organism and in the evolutionary development of the species.

Investigations of the semantic value of words and grammatical particles, especially recurring general patterns of relationships, may help us understand more about human cognition generally.

One area where we find semantic universals is in combinatorial principles and relations; indeed, many investigators assume that it is only at the level of basic expressions that languages differ semantically, and it may well be true that the child need only learn lexical details. For example, languages are never limited to additive semantic principles like that of conjunction; predication, for example, seems to be universally manifested. Logical approaches to semantics have paid more explicit attention to composition than most other approaches and thus suggest more explicit hypotheses about how languages structure meaning. One question has to do with the different kinds of semantic values expressions can have: just as *to* and *number* are of different syntactic categories in English, they are associated with different semantic classes, or *types*, in any logical approach to semantics, and the semantic value associated with sentences is of yet another different type. Universally we need distinctions among types. Semantic theory should provide us with some account of these distinctions and allow us to investigate the empirical question of whether languages differ in the semantic types they encode.

Our discussion will focus primarily on English, since that is the language we and our readers share. Occasionally, however, we draw illustrations from other languages, and we intend our general approach to provide a framework in which to do semantics for human languages generally, not simply for English.

## 2.3 The significance of language: "aboutness" and representation

Meaning manifests itself in the systematic link between linguistic forms and things, what we speak of or talk about. This "aboutness"

of language is so familiar that it may not seem noteworthy. But the fact that our languages carry meaning enables us to use them to express messages, to convey information to one another. As Lewis Carroll observed, we can talk about shoes and ships and sealing wax and whether pigs have wings. We can also speak of South Africa, Ingrid Bergman, birthdays, wearing clothes well, fear of flying, and prime numbers. Were languages not to provide for significance in this sense, the question of meaning would hardly arise. Nonetheless, some semantic theorists have thought that such aboutness is not really part of the domain of semantics. They have focused instead on the cognitive structures that represent meaning, taking the fundamental significance of language to reside in relations between linguistic expressions and what are sometimes called "semantic representations."

On our view, the significance of language, its meaningfulness, can be thought of as involving both aboutness and representational components. Theorists differ in the emphasis they place on these components and in the view they hold of their connections. It will be convenient for the discussion that follows to have labels for these two aspects of significance. *Informational significance* is a matter of aboutness, of connections between language and the world(s) we talk about. Informational significance looks outward to a public world and underlies appraisal of messages in terms of objective nonlinguistic notions like truth. *Cognitive significance* involves the links between language and mental constructs that somehow represent or encode speakers' semantic knowledge. Cognitive significance looks inward to a speaker's mental apparatus and does not confront issues of the public reliability of linguistic communication.

**2.3.1 The informational significance of language** Language enables us to talk about the world, to convey information to one another about ourselves and our surroundings in a reliable fashion. What properties of language and its uses underlie this remarkable fact? What allows language to serve as a guide to the world and to enable us to learn from what others have perceived (seen, heard, felt, smelled) without having to duplicate their perceptual experience ourselves?

Informational significance does not require that language links to the world in ways that are predetermined by the physical structure of our environment. Nor does it require that environmental information is simply registered or received without active input from perceiving and thinking human minds. Yet it does probably require a regular and systematic correspondence between language and the shared environment, what is publicly accessible to many different human minds.

If you are skeptical about informational significance, consider the use of language in giving directions, warnings, recipes, planning joint activities, describing events. Things occasionally misfire, but by and large such uses of language are remarkably effective. Language could not work at all in such ways were it not imbued with some kind of informational significance, being about matters in a public world.

Let us make this more concrete with a couple of examples. Suppose we utter

(5) This is yellow.

Interpreting *this* and other demonstrative expressions is problematic if the interpreter does not have access to some contextually salient entity to which it refers—perhaps the drapes to which someone is pointing. Since we have provided no picture to accompany (5), readers do not know what *this* refers to and cannot fully understand what its use means. The important points here are (1) that certain expressions seem to be used to refer, to indicate certain nonlinguistic entities, and (2) that knowing how to grasp what such expressions refer to is part of knowing what they mean. Expressions like *this* provide particularly vivid illustrations, but the same point holds of expressions like *the man who is sitting in the third row* and many others.

Now let us consider another example.

(6) The door is closed.

This sentence would accurately describe the situation depicted on the right in (7) but not that depicted on the left.[4]

(7)
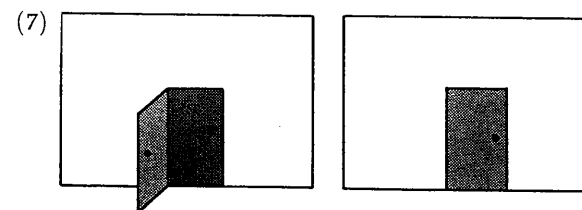
There are quite solid intuitions about the relation of sentence (6) to the two kinds of situations illustrated in (7). This fact is obvious yet nonetheless remarkable.

First, notice that the relation between the sentence and situations seems to be one that is independent of how those situations are presented. Instead of the drawings, we might have included photos or enclosed a videotape. We might even have issued you an invitation to come with us to a place where we could point out to you an open door and one that is closed. If you understand sentence (6), you can discriminate the two sorts of situation, no matter how we present them to you.

Second, observe that (6) can describe not just one or two, but a potential infinity of, different situations. In the picture on the right in (7), there is no cat in front of the closed door. But (6) would apply just as well to a situation like that depicted in (8), which is different from the right side of (7) only in that it contains a cat.

(8)
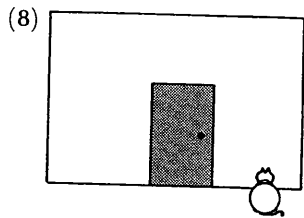
There is no need to stop with one cat or two or three, etc. We know how to keep going. The crucial point is that our knowledge of the relation between sentences and situations is not trivial and cannot consist in just remembering which particular situations are ones that a particular sentence can describe. Understanding what situations a sentence describes or, more generally, what information it conveys is crucial to grasping its meaning. It seems eminently reasonable to expect semantics to provide some account of this phenomenon.

Of course, language also enables us to talk with one another about more private internal worlds, to express our attitudes or mental states: hopes, beliefs, fears, wishes, dreams, fantasies. This too can be thought of as the conveying of information, but information in this case may seem less public or objective because the experiencing subject has some kind of privileged access to it. We cannot draw a picture to illustrate the situations described by sentence (9), but this does not mean that we do not know quite a lot about which situations it does, and which it does not, describe.

(9) Joan wants a tomato sandwich.

It is just that the differences among these situations are not apparent from purely visual signs. We would have equal difficulty using pictures to represent situations described or not described by sentence (10), yet what (10) is about is no less public than what (6) is about.

(10) Joan ate a tomato sandwich yesterday but not today.

What is noteworthy here is that language serves to bring private mental states into the public eye. Joan can speak about her desire to have a tomato sandwich today with the same ease that she speaks about the tomato sandwich that she actually consumed yesterday. Through language we not only inform one another about our external environment; we also manage to inform others of certain aspects of what our internal environment is like, thus externalizing or objectifying that internal experience to some extent. We can (sometimes) tell one another what is on our minds and we can use language to share what we imagine, suppose, or pretend.

Thus, when we speak of informational significance, we include not only links to physical or concrete phenomena but also to mental or abstract phenomena. There are deep philosophical questions that can be raised about the ontological status of different kinds of phenomena, but the important empirical fact for linguistic semantics is that for all of them we do indeed succeed in conveying information to one another by talking about them. It is in this sense that meaning always involves informational significance.

Semantic theories of informational significance are often called *referential* theories. *Truth-conditional semantics* is a particular kind of referential theory, which we will introduce in the next chapter and illustrate in more detail in succeeding chapters.

**2.3.2 The cognitive significance of language** The whole question of the meaningfulness of language has been approached from the inward-looking perspective of cognitive significance. The general idea is that we have ways of representing mentally what is meant by what we and others say. Perhaps, the suggestion seems to go, your understanding sentence (6), "The door is closed," is a matter of your recovering some internal representation of its meaning. Proponents of representational theories of meaning have usually not paid much attention to informational significance or even more generally to the capacity of people to judge with remarkable uniformity

relations between sentences and nonlinguistic situations. Rather, they have focused on understanding as a matter of what interpreters can infer about the cognitive states and processes, the semantic representations, of utterers. You understand us, on this view, to the extent that you are able to reconstruct semantic representations like the ones on which we have based what we say. Communicative success depends only on matching representations and not on making the same links to situations. As we will see in the next chapter, it is not impossible to connect a representational account with a referential one; nonetheless, most representationalists have simply ignored the question of objective significance, of how we manage to judge which of the situations depicted in (7) is described by sentence (6). They have seldom worried about the fact that there is an everyday sense of aboutness in which we take ourselves to be talking about our friends, the weather, or what we just ate for dinner, and not about our representations of them. Even if our impression that we are not just conveying representations but are talking about what is represented might ultimately be illusory, it does deserve explanation.

Some outward looking approaches view the cognitive significance of language as ultimately understood in terms of its informational significance. In such approaches, people may construct representations of what sentences mean, but the question of whether such representations are essentially identical need not arise. Understanding is a matter not of retrieving representations but of achieving consensus on informational significance.

It is almost certainly true that our talk about the world works so well because of fundamental similarities in our mental representations of it. Yet the similar representations required might not be semantic as such but connected to our perceptual experience. Nonetheless, that similar perceptual experience would depend on similar contact with a shared external environment. In this sense, a connection to the represented world is still basic, since it provides the basis for the similarities in perceptual experience, which in turn are somehow linked to linguistic expressions.

The semantic framework developed here emphasizes objective significance and referential connections but does not assume that the meaningfulness of language, its full significance, is exhausted by its informational significance. Indeed, we think that some aspects of how meanings are represented are meaningful even though they do not directly affect informational significance. Our guess is

that the aspect of meaningfulness that we have called cognitive significance has important implications for how conveyed information is processed. Chapter 6 discusses approaches to semantics that relate the informational significance of sentences to contextual factors and to the functioning of sentences in discourse, and in chapter 7 and part of chapter 8 we discuss some interesting proposals about the form of semantic representations.

## 3 Implication Relations

As we noted earlier, native speakers of a language have certain intuitions about what sentences or utterances convey, about the content and wider import of what is said, about what can be inferred on the basis of the sentence uttered, and about what is suggested. We often say that a sentence or utterance *implies* something. What is implied can be expressed by a sentence. For present purposes, we can think of *implication relations* as inferential relations between sentences. If $A$ implies $B$, we often say that $A$ suggests or conveys $B$ or that $B$ can be inferred from an utterance of $A$.

Implication relations can be classified on two axes. The first is what *licenses* or underwrites the implication. Where the basis for judging that $A$ implies $B$ is the informational or truth-conditional content of $A$, we say that $A$ *entails* $B$. Where what licenses the implication has to do with expectations about the reasons people talk and about their typical strategies in using language, we say that $A$ *implicates* (or *conversationally implicates*) $B$. Philosopher Paul Grice first argued for this distinction and proposed an account of how conversational implicatures work. Although there is still considerable disagreement on the theory of implicature, the need for such a distinction is now widely acknowledged.[5] We will discuss entailments in 3.1 and distinguish them from implicatures, which we discuss briefly in 3.2 (and in somewhat more detail in chapter 4). Formal semantic theories of the kind we develop in this book allow us to characterize entailment relations quite precisely. Distinguishing entailments from implicatures is important in developing semantic analyses, although it is by no means easy to do so (and there are often disagreements on where to draw the line).

The second axis of classification is the *discourse status* of the implication. The primary distinction here is between *assertions* (and various other things we might intend to accomplish when we say something: *questions*, *suppositions*, *orders*) and *presupposi-*

*tions*. An assertion aims to add content to the ongoing discourse, to effect some kind of change in what the conversationalists assume, whereas a presupposition presents its content as already assumed or taken for granted. Section 3.3 introduces presupposition and empirical tests to distinguish it from assertion and assertion-based implications. We will discuss assertion along with other kinds of speech acts in more detail in chapter 4 and again in chapter 6, where we return to presupposition. On this way of thinking of things, classifying an implication as a presupposition is neutral as to whether the implication might also be an entailment or some kind of conversational implicature (or licensed in some other way); *A* can, e.g., both entail *B* and presuppose *B*.

## 3.1 Entailment
Consider the following examples.

(11) *a.* This is yellow.
    *b.* This is a fountain pen.
    *c.* This is a yellow fountain pen.

(12) *a.* This is big.
    *b.* This is a sperm whale.
    *c.* This is a big sperm whale.

Imagine yourself uttering the sentences in (11) with reference to a particular object, perhaps a pen, perhaps something else. In such a situation you know that if your assertions of (11*a*) and (11*b*) are true (if the object is indeed yellow and indeed a fountain pen), then your assertion of (11*c*) is also true. It would be contradictory to assert the first two sentences and then deny the third; we discuss contradiction below. Any native speaker of English knows that the information conveyed by uttering (11*c*) is somehow already included in the information conveyed by uttering (11*a*) and (11*b*). This knowledge seems to be part of knowing what these sentences mean: we need know nothing about the object indicated by *this* beyond the fact that it is the same object for all three utterances. We say that the pair of sentences (11*a*) and (11*b*) *entails* sentence (11*c*).

Now imagine yourself uttering the sentences in (12), again keeping fixed what *this* refers to in all three utterances. Matters become very different. Suppose you take yourself to be pointing at a sperm whale. Sperm whales are pretty big creatures, so you might well

assert that (12*a*) and (12*b*) are true. Suppose in addition that you judge that this particular specimen is not especially distinguished in size among its fellow sperm whales, that it's one of the smaller ones. In such circumstances it would be quite reasonable to deny (12*c*). In this case the *a* and *b* sentences do not entail the *c* sentence.

We would find the same difference in the two sets of sentences if we used *automobile* instead of *fountain pen* and used *galaxy* instead of *sperm whale*. *Yellow* (along with other adjectives like *round*, *featherless*, *dead*) behaves differently from *big* (and other adjectives like *strong*, *good*, *intelligent*), and this difference seems semantic in nature. (See chapter 8, section 3.1, for discussion of this difference.)

As we have noted, the relation between the pair (11*a*, *b*) and (11*c*) is usually called *entailment*. Together (11*a*) and (11*b*) entail (11*c*), whereas (12*a*) and (12*b*) do not entail (12*c*).

An entailment can be thought of as a relation between one sentence or set of sentences, the entailing expressions, and another sentence, what is entailed. For simplicity we equate a set of entailing sentences with a single sentence, their conjunction, which we get by joining the sentences using *and*. The conjunction is true just in case each individual sentence in the set is true, and it describes exactly those situations that can also be described by each one of the individual sentences. We could, for example, simply look at the English sentences "This is yellow, and this is a fountain pen" and "This is big, and this is a sperm whale" in cases (11) and (12) above.

Theoretically, entailment relations might depend solely on the syntactic structure of sentences. However, the contrast between (11) and (12) (and a host of other such sentences) demonstrates that they cannot be simply a matter of surface syntax. Entailments seem to involve the information conveyed by sentences: if English sentence *A* entails English sentence *B*, then translating *A* and *B* into Finnish sentences *A'* and *B'* with the same informational significance will preserve the entailment relation.

Asked to define entailment, you might come up with any of the following:

(13) *A entails B* $=_{df}$
    • whenever *A* is true, *B* is true
    • the information that *B* conveys is contained in the
      information that *A* conveys

• a situation describable by *A* must also be a situation describable by *B*
• *A and not B* is contradictory (can't be true in any situation)

We will later discuss more formal characterizations of the entailment relation, but for the time being you can adopt any of the preceding definitions.

We can find countless examples where entailment relations hold between sentences and countless where they do not. The English sentence (14) is normally interpreted so that it entails the sentences in (15) but does not entail those in (16).

(14) Lee kissed Kim passionately.

(15) *a*. Lee kissed Kim.
    *b*. Kim was kissed by Lee.
    *c*. Kim was kissed.
    *d*. Lee touched Kim with her lips.

(16) *a*. Lee married Kim.
    *b*. Kim kissed Lee.
    *c*. Lee kissed Kim many times.
    *d*. Lee did not kiss Kim.

Looking at entailments shows, by the way, that what are conventionally treated as translation equivalents are not always informationally equivalent. The English sentence (17*a*) entails (17*b*), but the Finnish sentence (18), which most texts would offer as a translation of (17*a*), does not entail anything about the femaleness of the person or animal said to be big, the Finnish third-person pronoun *hän* being completely neutral as to the sex of its referent.

(17) *a*. She is big.
    *b*. Some female is big.

(18) Hän on iso.

Thus, although sentence (18) can be used to describe any situation (17*a*) describes, the Finnish can also be used to describe situations not describable by (17*a*), for example, to say of some man that he is big. That is, (18) is also a translation of (19*a*), but unlike (19*a*) it does not entail the information conveyed by (19*b*).

(19) *a*. He is big.
    *b*. Some male is big.

In particular contexts, the use of translations that are not informationally equivalent, translations where entailments are not preserved, may be unproblematic, since other information is available to ensure that only the desired information is actually conveyed. But neither (17*a*) nor (19*a*) is an informationally equivalent translation of the Finnish sentence (18), which is informationally equivalent to something like (20).

(20) She or he is big.

You might object to our claim that (14), "Lee kissed Kim passionately," entails (15*d*), "Lee touched Kim with her lips," by pointing out that sentence (21) can be true in a situation where (15*d*) is false.

(21) In her imagination Lee kissed Kim passionately.

Does your example defeat the claim that (14) entails (15*d*)? No. We could counter by claiming that if (15*d*) is false in the situation in which (21) is true then (14) is false in that same situation, and we might further claim that (21) entails (22).

(22) In her imagination Lee touched Kim with her lips.

On the other hand, if you manage to persuade us that Lee's mouthing of a kiss in Kim's direction from a distance of ten feet counts as her kissing him, then we have no good defense of our claim that (14) entails (15*d*) (since we agree that she is unable actually to touch him from that distance). Or your scenario might be romance via computer where Lee types in "I am kissing you passionately," addressing herself to Kim's computer. If we agree to accept either of your cases as real kissing, then our only possible line of defense is that there are different interpretations of *kiss* involved, only one of which requires that the kisser touch the kissee with her lips. In other words, we could accept one of your cases and continue to maintain that (14) entails (15*d*) only if we also argue that (14) is ambiguous, that it has more than one meaning. In this case, the string (14) could entail (15*d*) on one interpretation of *kiss* but not have that entailment on the interpretation your cases involve. We discuss later what considerations support claims of ambiguity.

Similarly, we claim that (14), "Lee kissed Kim passionately," does not entail (16*c*), "Lee kissed Kim many times." You might deny this by noting that the passionate kisser is unlikely to stop

with a single kiss. We can agree with that observation and may even agree with you that assertion of (14) does strongly suggest or imply the truth of (16c) but nonetheless disagree that the implication is an entailment. For example, we might want to maintain that a situation with one or a few kisses can nonetheless involve passionate kissing, perhaps persuading you by showing a film of a single kiss which you will agree is a passionate one. You might still maintain that Lee herself would never stop short of many kisses once she succumbs to passion, and thus that (14) would never be true without (16c) also being true. We must now take a slightly different tack, noting that this is a matter of what Lee happens to be like rather than a matter of what the sentences mean. Or perhaps we would remind you of the possibility that Lee could begin her round of passionate kissing but be allowed only one passionate kiss before Kim breaks free and runs away.

What we should not do in the face of your objections is simply to reiterate our initial claims. Judgments about entailment relations can be defended and supported by evidence. As in the case of any linguistic phenomenon, there may be areas of real diversity within the community of language users, dialectal and even idiolectal differences. This complication must not, however, obscure the important fact that judgments about semantic phenomena are interconnected, and thus that there is relevant evidence to be offered in support of such judgments. In learning to do semantics as a linguist, one must learn to develop semantic arguments and explore semantic intuitions systematically. And one must learn to discriminate between the strict notion of the entailment relation and looser varieties of implication. Test yourself on the following examples. Sentences (23a) and (24a) imply (23b) and (24b) respectively, but only one of the implications is an entailment. Try to discover for yourself which is which and why before reading the discussion that follows the examples.

(23) a. Mary used to swim a mile daily.
　　 b. Mary no longer swims a mile daily.

(24) a. After Hans painted the walls, Pete installed the cabinets.
　　 b. Hans painted the walls.

Sentence (23a) implies but does not entail (23b). Although in many contexts we would infer from an utterance of (23a) that (23b) is true, notice that (23a) could be used by someone familiar with Mary's routine last year but no longer in contact with her. It might

be true that Mary still swims a mile daily, and the speaker we've imagined could make clear that (23b) should not be inferred by continuing with something like (25).

(25) I wonder whether she still does [swim a mile daily].

In contrast, (24a) not only implies but entails (24b). Suppose that Hans did not paint the walls. Then even if Pete did install the cabinets, he did not do so after Hans painted the walls. That is, sentence (26) is contradictory.

(26) After Hans painted the walls, Pete installed the cabinets, but Hans did not paint the walls.

There is one further preliminary point that it is important to make about entailments; namely, that there are infinitely many of them. That is, there are infinitely many pairs of sentences $A$, $B$ such that $A$ entails $B$. Here are a couple of ways to construct indefinitely many such pairs. Intuitions are fairly sharp, for example, that (27a) entails (27c) and also that (27b) entails (27c).

(27) a. Lee and Kim smoke.
　　 b. Lee smokes and drinks.
　　 c. Lee smokes.

We can easily keep conjoining noun phrases (*Lee and Kim and Mike and Susan and* ...), adding descriptions like *the other Lee* or *the woman I love* should our stock of distinct proper names be exhausted. We can also, of course, just keep conjoining verb phrases: *smokes and drinks and has bad breath and lives in Dubuque and* ...). Either way we get more sentences that entail (27c), and we need never stop. That is, we have intuitions that seem to involve the meanings of indefinitely many sentences, a potential infinity. Only finitely many such intuitions could possibly be stored in memory. How, then, are such judgments possible? Here we see again the general issue of the productivity of meaning, which we introduced in 2.1.

---

**Exercise 1**　For each pair of sentences, say whether the $a$ sentence entails the $b$ sentence and justify your answers as well as you can. Where proper names or pronouns or similar expressions are repeated in $a$ and $b$, assume that the same individual is referred to in each case; assume also that temporal expressions (like *today* and the present tense) receive a constant interpretation.

(1)  *a.* Today is sunny.
 *b.* Today is warm.

(2)  *a.* Jane ate oatmeal for breakfast this morning.
 *b.* Jane ate breakfast this morning.

(3)  *a.* Jane ate oatmeal for breakfast this morning.
 *b.* Jane ate something hot for breakfast this morning.

(4)  *a.* Juan is not aware that Mindy is pregnant.
 *b.* Mindy is pregnant.

(5)  *a.* Every second-year student who knows Latin will get credit
 for it.
 *b.* If John is a second-year student and knows Latin, he will
 get credit for it.

(6)  *a.* If Alice wins a fellowship, she can finish her thesis.
 *b.* If Alice doesn't win a fellowship, she can't finish her thesis.

(7)  *a.* Maria and Marco are married.
 *b.* Maria and Marco are married to each other.

(8)  *a.* Only Amy knows the answer.
 *b.* Amy knows the answer.

(9)  *a.* Mary is an Italian violinist.
 *b.* Some Italian is a violinist.

(10) *a.* Some student will not go to the party.
 *b.* Not every student will go to the party.

(11) *a.* Allegedly, John is a good player.
 *b.* John is a good player.

(12) *a.* John knows that pigs do not have wings.
 *b.* Pigs do not have wings.

(13) *a.* John believes that pigs do not have wings.
 *b.* Pigs do not have wings.

(14) *a.* Oscar and Jenny are rich.
 *b.* Jenny is rich.

(15) *a.* Oscar and Jenny are middle-aged.
 *b.* Jenny is middle-aged.

(16) *a.* Not everyone will get the correct answer.
 *b.* Someone will get the correct answer.

## 3.2  Implicature

As we have set things up, it might look as if implicature is simply implication minus entailment. Implicature, however, is characterized more positively: we say that an utterance *A* implicates *B* only if we take *B* to be (part of what) the utterer of *A* meant by that utterance. An implicature must be something that the utterer might reasonably mean by making the utterance, something she expects to convey. And, critically, if *A* implicates *B*, there is a certain kind of explanatory account of that relation, one that invokes general principles of conversation, as well as (perhaps) certain specific assumptions about the particular context in which *A* happens to have been uttered. Grice says that implicatures must be *calculable*: there should be an argument that *A* implicates *B*, an argument that draws on the linguistic meaning of *A* and on expectations that speakers generally have of one another (e.g., that what is said will be "relevant" and "informative") and, in some cases, on particular features of the utterance context.

Suppose, e.g., that we have the dialogue in (28).

(28) *A:* Did you enjoy the dinner?
 *B:* We had mushroom salad and mushroom sauce on the
 pasta.

What might speaker *B* be implicating? Given a question like that asked by *A*, what becomes immediately relevant is for *B* to choose one of the possibilities in (29).

(29) *a.* I (namely *B*) enjoyed the dinner.
 *b.* I (namely *B*) didn't enjoy the dinner.

Thus, unless there's some reason to think that *B* is dodging the question, we will generally take *B*'s utterance to implicate either (29*a*) or (29*b*). But no general principles allow us to decide whether the implicature is positive or negative: to do that, we have to know more. Perhaps it is common knowledge that *B* hates mushrooms with a passion or, conversely, that *B* absolutely adores mushrooms in virtually any dish. In the first case, (29*b*) is implicated, whereas (29*a*) is implicated in the other case. If *A* knows nothing about *B*'s opinions of mushrooms, *A* will likely interpret *B*'s response as evasive. (An evasive answer might be in order if, e.g., *B* fears that *A* will report the evaluation to the person who hosted the dinner.) When the implicature to one of (29) works, then we are dealing with a *particularized* conversational implicature. Linguistic

theories cannot really predict such implicatures (except insofar as they can shed light on such issues as how questions make certain next contributions relevant). Not surprisingly, no one is likely to think that the relation between (28*B*) and either of the sentences in (29) is entailment, that it is the semantic content of the sentence in (28*B*) that licenses the inference to either (29*a*) or (29*b*).

What linguists have studied most systematically are what Grice called *generalized* conversational implicatures. These are the cases that often seem close to entailments. Take example (23) from the preceding section, repeated here.

(23) *a.* Mary used to swim a mile daily.
     *b.* Mary no longer swims a mile daily.

We argued that the relation between these sentences was not entailment, because we could follow an utterance of (23*a*) with an utterance of (25), also repeated here.

(25) I wonder whether she still does.

What (25) does is *defeat* the inference from (23*a*) to (23*b*): an empirical hallmark of conversational implicatures is that they are, in Grice's words, *defeasible*. An implication that can be defeated just by saying something that warns the hearer not to infer what might ordinarily be implied is not an entailment but something different. Notice that if we try to "defeat" entailments, we end up with something contradictory:

(30) #Lee kissed Kim passionately, but Lee didn't kiss Kim.

But even though the implication from (23*a*) to (23*b*) is defeasible, that implication is a very general one that holds unless it is specifically defeated. In contrast to the implication from (28*B*) to one of the sentences in (29), the implication from (23*a*) to (23*b*) does not depend on any special features of the contexts in which sentences like (23*a*) might be uttered. What, then, is the general argument, the calculation, that takes us from (23*a*) to (23*b*)?

Roughly, the argument goes like this. Hearers expect speakers to be adequately informative on the topic being discussed, and speakers know that hearers have this expectation. Sentence (23*a*) reports a past habit, in contrast to (31), which reports a present habit.

(31) Mary swims a mile daily.

Present habits, however, began earlier, and thus (31) might well be true of the same situation as (32*a*), which cancels implicature (23*b*).

(32) *a.* Mary used to swim a mile daily, and she still does.
     *b.* Mary used to swim a mile daily, but she no longer does.

Unless there is some special reason that the conversationalists are interested only in Mary's past habits, if the speaker is in a position to inform the hearer by uttering (31) rather than (23*a*), then she should do so (and, furthermore, she knows that the hearer expects her to do so). Thus to utter (23*a*) suggests one is not in a position to make the stronger claim (31), and in many circumstances it suggests that the stronger claim is false, i.e., (23*a*) conveys that (23*b*) holds. Indeed, it is normal to make the move from (23*a*) to (23*b*)— and, more generally, from *used to* to *does no longer*—unless there are explicit indicators to the contrary as in (32*a*). The strength of the implication is one reason why it is so often confused with entailment. Sentence (32*b*) illustrates another empirical test that distinguishes implicatures from entailments: they are typically *reinforceable*, without any flavor of the redundancy that generally accompanies similar reinforcement of entailments. Although (32*b*) sounds fine, (33), where an entailment is reinforced, sounds quite strange.

(33) #Lee smokes and drinks, but/and she smokes.

Reinforceability is the flip side of defeasibility. Because generalized implicatures are not part of the linguistic meaning of expressions in the same sense that entailments are, they can readily be explicitly set aside or explicitly underscored. However, they are strongly recurrent patterns, most of them found in similar form crosslinguistically.

Here are some more examples where a generalized implicature seems to hold between the (*a*) and the (*b*) sentences.

(34) *a.* Joan likes some of her presents.
     *b.* Joan doesn't like all of her presents.

(35) *a.* Mary doesn't believe that John will come.
     *b.* Mary believes that John won't come.

(36) *a.* If you finish your vegetables, I'll give you dessert.
     *b.* If you don't finish your vegetables, I won't give you dessert.

---

**Exercise 2** Choose one of the pairs of sentences in (34) to (36) and show that the relation between (*a*) and (*b*) is both defeasible and reinforceable.

---

We return to the topic of conversational implicature in chapter 4, where we say more about Grice's account of the conversational principles that underlie these relations.

## 3.3  Presupposition

Many expressions seem to "trigger" certain presuppositions; i.e., they signal that the speaker is taking something for granted. Utterances of sentences containing such expressions typically have two kinds of implications: those that are asserted (or denied or questioned or otherwise actively entertained) and those that are presupposed. As we noted above, presupposition is more than a species of implication: it is a matter of the discourse status of what is implied. If A presupposes B, then A not only implies B but also implies that the truth of B is somehow taken for granted, treated as uncontroversial. If A entails B, then asserting that A is true commits us to the truth of B. If A presupposes B, then to assert A, deny A, wonder whether A, or suppose A—to express *any* of these attitudes toward A is generally to imply B, to suggest that B is true and, moreover, uncontroversially so. That is, considering A from almost any standpoint seems already to assume or presuppose the truth of B; B is part of the background against was we (typically) consider A.

Consider, for example, the sentences in (37). Any one of (a–d) seems to imply (e) as a background truth. These implications are triggered by the occurrence of the phrase *the present queen of France*, a definite description. It is generally true of definite descriptions that they license such implications.

(37)  a. The present queen of France lives in Ithaca.
   b. It is not the case that the present queen of France lives in Ithaca (or more colloquially, the present queen of France does not live in Ithaca).
   c. Does the present queen of France live in Ithaca?
   d. If the present queen of France lives in Ithaca, she has probably met Nelly.
   e. There is a unique present queen of France.

Or consider (38). Again (using) any of (a–d) will generally imply (e). In this case, the implications are attributable to *regret*, which is a so-called factive verb. Factive verbs generally signal that their complements are presupposed. Other examples are *realize* and *know*.

(38)  a. Joan regrets getting her Ph.D. in linguistics.
   b. Joan doesn't regret getting her Ph.D. in linguistics.
   c. Does Joan regret getting her Ph.D. in linguistics?
   d. If Joan regrets getting her Ph.D. and linguistics, she should consider going back to graduate school in computer science.
   e. Joan got her Ph.D. in linguistics.

Look next at (39). Once again, each of the quartet (a–d) implies (e). In this case it is the quantifying determine *all* that is responsible. A number of quantificational expressions serve to trigger presuppositions.

(39)  a. All Mary's lovers are French.
   b. It isn't the case that all Mary's lovers are French.
   c. Are all Mary's lovers French?
   d. If all Mary's lovers are French, she should study the language.
   e. Mary has (three or more?) lovers.

Finally, look at (40), where we find the same pattern. In this case it is the cleft construction that is responsible.

(40)  a. It was Lee who got a perfect score on the semantics quiz.
   b. It wasn't Lee who got a perfect score on the semantics quiz.
   c. Was it Lee who got a perfect score on the semantics quiz?
   d. If it was Lee who got a perfect score on the semantics quiz, why does she look so depressed?
   e. Someone got a perfect score on the semantics quiz.

A distinguishing empirical feature of presupposition, then, is that it involves not just a single implication but a family of implications. By this we mean that not only assertive uses of sentence A (the affirmative declarative) imply B but also other uses of A where something is, for example, denied, supposed, or questioned. That we are dealing with a family of implications derives from the fact that the presupposition is background. Each of (a–d), what we will call the P family, is said to presuppose (e) because uttering each (typically) implies (e) and also implies that (e) is being taken for granted. It is convenient for testing purposes to identify the P family in syntactic terms: an affirmative declarative, the negative of that declarative, the interrogative, and the conditional antecedent. In semantic/pragmatic terms, these represent a family of different

sorts of attitudes expressed towards $A$. We can thus informally characterize when $A$ presupposes $B$ as follows:

(41) $A$ presupposes $B$ if and only if not only $A$ but also other members of the $P$ family imply (and assume as background) $B$.

Presuppositions come in families, even if sometimes certain members of the family may be stylistically odd.

Notice that we have said that $A$ and other members of its $P$ family *imply* $B$ when $A$ presupposes $B$. We do *not* require that these implications be entailments. As we have defined entailment, it is not even possible for all these relations to be entailments. However, it is possible that some member of the family entails $B$. Sentence (40$a$), for example, not only presupposes (40$e$); it also entails (40$e$). If (40$a$) is true, then (40$e$) must also be true. The negation, (40$b$), also presupposes (40$e$) but does not entail it. The implication to (40$e$) is *defeasible*; that is, there are contexts in which it can be defeated, contexts in which (40$b$) is asserted yet (40$e$) is not assumed to be true. We might take (42) as a discourse context that defeats the implication from (40$b$) to (40$e$).

(42) *Speaker 1*: I wonder whether it was Lee or someone else who got a perfect score on the semantics quiz.
*Speaker 2*: It wasn't Lee who got a perfect score [on the semantics quiz]. I happen to know that Lee scored only 70 percent. I wonder if anyone managed to get a perfect score.

Speaker 2 has taken issue with speaker 1's presupposing that someone got a perfect score by suggesting that (40$e$) may be false and asserting that (40$b$) is indeed true. Of course, speaker 2 chooses this way of conveying the information that Lee did not get a perfect score because speaker 1 has already implied that someone did do that.

We need only look at noncleft counterparts of the sentences in (40) to see that $A$ may entail $B$ yet not presuppose $B$.

(43) *a.* Lee got a perfect score on the semantics quiz.
*b.* Lee didn't get a perfect score on the semantics quiz.
*c.* Did Lee get a perfect score on the semantics quiz?
*d.* If Lee got a perfect score on the semantics quiz, why does she look so depressed?
*e.* Someone got a perfect score on the semantics quiz.

If focal stress is not placed on *Lee*, then none of (43$b$–$d$) typically imply (43$e$), even though (43$a$) entails (43$e$). Someone's getting a

perfect score on the semantics quiz is not part of the usual background for talking about Lee's achieving the feat in question, as stated by (43$a$). Indeed, it seems reasonable to say that a major semantic difference between the subject-verb-object (S-V-O) sentence (43$a$) and its cleft correlate (40$a$), "It was Lee who got a perfect score on the semantics quiz," is that the latter but not the former carries a presupposition that someone got a perfect score. Whether this difference can ultimately be explained in terms of some other difference between the two is an issue we cannot answer here.

What the sentences in (43) show is that $A$ can entail $B$ without other members of the $P$ family also implying $B$. Presupposition and entailment are thus quite distinct. $A$ may entail $B$ but not presuppose it, as in (34); conversely, $A$ may presuppose $B$ but not entail it, as in (40). And given the way we have defined entailment and presupposition, it is also possible for $A$ both to entail and to presuppose $B$. (Some accounts of presupposition do not admit this possibility; we will discuss this and related issues in more detail in chapter 6.)

Presupposition requires a family of implications, not all of which can be licensed by an entailment. Interrogatives, for example, would never entail other sentences, since they are not ordinarily valued as true or false; use of an interrogative may, however, imply something. Thus, one important question presupposition raises is about the nature of implications that are not backed by entailment relations. Some presuppositions, it has been argued, derive from quite general conversational principles and thus might be held to be licensed in much the same way as the conversational implicatures we briefly discussed in the preceding section. And there may be other mechanisms at work.

A related issue is the speaker's responsibilities with respect to what the utterance presupposes. What is presupposed in a discourse is what is taken for granted. Thus, a speaker who says $A$, presupposing $B$, in a context where $B$ is at issue has thereby spoken inappropriately in some sense. For example, suppose that Sandy is on trial for selling illicit drugs and the prosecuting attorney asks question (44).

(44) Sandy, have you stopped selling crack?

As we know, the question is unfairly loaded, since it presupposes (45), which is very much at issue.

(45) Sandy has sold crack.

If Sandy simply answers yes or no, the presupposition is un-challenged, and she appears to go along with the implication that (45) is true. A defensive answer must explicitly disavow that implication:

(46) Since I never did sell crack, I have not stopped selling crack.

In many contexts, however, it is perfectly appropriate for a speaker to say *A*, presupposing *B*, even though the speaker does not believe that *B* is taken for granted by other discourse participants. For example, (47) might be uttered by a passenger to the airline representative, who can hardly be thought to know anything about the passenger's personal habits. Although the last clause in (47) presupposes the clause that precedes it in square brackets, it would seem unduly verbose to express that presupposed information overtly.

(47) I don't want to be near the smoking section because [I used to smoke and] I've just stopped smoking.

An obvious difference between the airline passenger and the prosecuting attorney is that the latter knows full well that what the utterance presupposes is controversial, whereas the former can safely assume that the reservations clerk has no opinion about what is being presupposed (and no real interest in the matter). With no reason to suppose otherwise, the clerk can quite reasonably be expected to accept the passenger's presupposition as if it were already taken for granted and discourse should proceed unproblematically. What happens in such cases is called *accommodation*.

We have barely begun to explore the topic of presupposition, and we will consider some of these phenomena in more detail in chapter 6. But it is clear already that presupposition raises questions not just about individual sentences and their truth or falsity but also about the uses of sentences in connected discourse (including uses of interrogatives, which are generally not said to be either true or false).

---

**ise 3**  Consider the following:

(1) *a.* That John was assaulted scared Mary.
   *b.* Mary is animate.
   *c.* John was assaulted.
   *d.* That John was assaulted caused fear in Mary.

(2) *a.* That John was assaulted didn't scare Mary.
   *b.* Mary is animate.
   *c.* John was assaulted.
   *d.* That John was assaulted didn't cause fear in Mary.

(3) *a.* John didn't manage to get the job.
   *b.* It was kind of hard for John to get the job.
   *c.* John didn't get the job.

In each of these examples, the *a* sentences presuppose and/or entail the other sentences. Specify which is a presupposition and which a simple entailment and which is both an entailment and a presupposition. Explain what test convinced you of your answer.

What relationship holds between the sentences in the following examples? Explain why you think that that relation holds.

(4) *a.* It is false that everyone tried to kill Templeton.
   *b.* Someone did not try to kill Templeton.

(5) *a.* That John left early didn't bother Mary.
   *b.* John left early.

(6) *a.* Someone cheated on the exam.
   *b.* John cheated on the exam.

(7) *a.* If John discovers that Mary is in New York, he will get angry.
   *b.* Mary is in New York.

(8) *a.* Seeing is believing.
   *b.* If John sees a riot, he will believe it.

---

## 4  More Semantic Relations and Properties

Implication relations are not the only kind of semantic relations speakers recognize. In this section we look at a number of other semantic relations and properties.

### 4.1  Referential connections and anaphoric relations

Consider the sentences in (48).

(48) *a.* *She* called me last night.
   *b.* Did you know that *he* is a Nobel Prize winner?
   *c.* I had a terrible fight with *that bastard* yesterday.

Each of the italicized expressions is used to *refer* to someone, to pick out an individual about whom something is being said, but a pointing gesture or a nod or some similar nonlinguistic means may be needed to indicate who this is. These same expressions, however, can be used in contexts where such pointing is unnecessary because they are linked to other *antecedent* expressions. In (49) speakers judge that the bracketed italicized expressions can be understood as *coreferential* with, having the same reference as, the bracketed unitalicized expressions that serve as their antecedents, and furthermore, they can be understood as dependent for their reference on the reference assigned to their antecedents. Intuitive judgments are quite clear-cut in these cases: the italic expressions are *referentially dependent* on the unitalicized expressions.

(49) *a.* If [*she*] calls, please tell [Teresa] I've gone to the pool.
   *b.* [The computer repairman] insists that [*he*] found nothing wrong.
   *c.* I talked to [Kim] for an hour, but [*that bastard*] never once mentioned the gift I sent him from Peru.

Expressions are said to be interpreted *anaphorically* when their reference is derived from that of antecedent expressions. The italicized expressions in (49) illustrate this. There are some expressions that can only be interpreted anaphorically and not through anything like pointing. The reflexive pronoun *herself* falls in this category; compare (50*a*), where *she* can serve as antecedent, with (50*b*), where there is no antecedent for *herself*.

(50) *a.* [She] is proud of [*herself*].
   *b.* *Be proud of herself.

In the syntactic literature, coindexing, as in (51), is the commonest device for indicating coreference.

(51) *a.* If [*she*]$_i$ calls, please tell [Teresa]$_i$ I've gone to the pool.
   *b.* [The computer repairman]$_j$ insists that [*he*]$_j$ found nothing wrong.
   *c.* I talked to [Kim]$_k$ for an hour but [*that bastard*]$_k$ never once mentioned the gift I sent [*him*]$_k$ from Peru.
   *d.* [She]$_l$ is proud of [*herself*]$_l$.

Chomsky (1981) discusses indexing as a formal process in some detail, but its informal use for this purpose far predates contem-

porary government-binding (GB) theory (see, for example, Postal (1971)).

What are called judgments of coreference in the literature typically involve judging not sameness of reference as such but dependence of reference of one expression upon that assigned to another.[6] Directed linking is another device sometimes used to show nonsymmetric dependence relations;[7] (52) shows a notation for linking.

(52) *a.* If [*she*] calls, please tell [Teresa] I've gone to the pool.
   *b.* [The computer repairman] insists that [*he*] found nothing wrong.
   *c.* I talked to [Kim] for an hour, but [*that bastard*] never once mentioned the gift I sent [*him*] from Peru.
   *d.* [She] is proud of [*herself*].

Referential connections may be somewhat more complex. Much of chapter 3 is devoted to making precise the nature of the dependencies speakers recognize as possible in (53), where the dependencies are indicated by coindexing, just as in the simpler cases above. In (53) the anaphorically interpreted NPs (*she, her, himself, his,* and *themselves*) are said to be *bound* by their antecedent NPs.

(53) *a.* [Every woman]$_i$ thinks that [*she*]$_i$ will do a better job of child rearing than [*her*]$_l$ mother did.
   *b.* [No man]$_i$ should blame [*himself*]$_i$ for [*his*]$_i$ children's mistakes.
   *c.* [Which candidates]$_i$ will vote for [*themselves*]$_i$?

In (53) repetition of an index does not indicate straightforward sameness of reference, as it did in (51). Expressions like *every woman, no man,* and *which candidates* do not refer in the intuitive sense, though their relations to anaphors are often called "coreference." Although *she* in (53*a*) is not used to refer to any individual, the interpretation of (53*a*) can be understood in terms of sentences in which NPs in the analogous positions both refer to the same individual. Roughly, (53*a*) says that if we point to any particular woman and say (54), where each of the indexed NPs refers to that woman, then what is said will be true, no matter which woman we pick.

(54) [She]$_i$ thinks that [*she*]$_i$ will do a better job of child rearing than [*her*]$_i$ mother did.

Linguistic questions about the nature of anaphoric relations provided a major impetus for exploration of how classical logical theories might shed light on natural language semantics. In exploring how syntactic structures affect the possibilities of interpreting expressions, linguists and philosophers have discovered other cases of so-called coreference where referential dependency may be somewhat different both from simple sameness of reference and from the standard binding relations elucidated by quantification theory.

(55) *a.* Kath caught [some fish]$_i$, and Mark cooked [*them*]$_i$.
 *b.* If [a farmer]$_j$ owns [a donkey]$_i$, [*he*]$_j$ beats [*it*]$_i$.
 *c.* [Gina]$_i$ told [Maria]$_j$ that [*they*]$_{i+j}$ had been assigned clean-up duty.

In (55*c*) the plural pronoun *they* has what have been called *split antecedents*; the index $i+j$ indicates referential dependence on both the distinct indexes $i$ and $j$. The notation $i, j$ is often used for indicating split antecedents, but we want to reserve this notation for cases where an expression may be linked either to something with index $i$ or to something with index $j$. In the rest of this section we ignore split antecedents.

These and many other examples have been widely discussed in the recent syntactic and semantic literature. Though there continues to be debate on the appropriate analysis of particular anaphoric relations, there is no question that speakers do recognize the possibility of some kind of interpretive dependencies in all these and indefinitely many other cases. Judgments of coreference possibilities (broadly understood) are fundamentally important semantic data.

There are also indefinitely many cases where the intuitive judgments are that such dependencies are not possible. These are usually called judgments of *disjoint reference*, a kind of independence of reference assignment. The terminology was introduced in Lasnik (1976), but as with "coreference," it must be understood somewhat loosely. The asterisks in (56) mean that the indicated referential dependencies are judged impermissible. The NPs in question are, according to speakers' judgments, necessarily interpretively independent of one another and are not anaphorically relatable.

(56) *a.* \*Behind [Teresa]$_i$, [*she*]$_i$ heard Mario.
 *b.* \*[He]$_i$ insists that [the computer repairman]$_i$ found nothing wrong.
 *c.* \*If [*that bastard*]$_i$ calls, tell [Kim]$_i$ I've gone to Peru.
 *d.* \*[*Herself*]$_i$ is proud of [her]$_i$.

Sentences (56*a*–*c*) are bad with the indicated coindexing; they can be used only if the italicized expressions are interpreted nonanaphorically (through pointing or something similar). Sentence (56*d*) is unusable because *herself* happens to be an expression that requires anaphoric interpretation.

Much interesting recent linguistic research in semantics has tried to elucidate and systematize judgments about referential relations, and such data have figured prominently in developing theories of the map between syntactic structures and their interpretation.

---

**Exercise 4** Each of the following sentences contains some nonpronominal NPs and a pronoun (in some cases, a possessive pronoun). Assign a distinct index to each nonpronominal NP. Copy all such indices on the pronoun in the sentence, and star those indices copied from NPs that cannot be antecedents for the pronoun. For example,

(1) *a.* John believes that few women think that they can be successful.
 *b.* John$_1$ believes that [few women]$_2$ think that they$_{2, *1}$ can be successful.

(2) *a.* They know few women.
 *b.* They$_{*1}$ know [few women]$_1$.

(3) She thinks that Barbara is sick.

(4) If she is sick, Barbara will stay home.

(5) When he is unhappy, no man works efficiently.

(6) Neither of Ann's parents thinks he is adequately paid.

(7) That jerk told Dick what Mary thinks of him.

(8) If she wants to, any girl in the class can jump farther than Mary.

(9) Her mother is proud of every woman.

(10) Her mother is proud of Lisa.

(11) My friends think that Joan's parents met each other in college.

(12) John promised Bill to help him.

(13) John persuaded Bill to help him.

(14) Every girl on the block jumps rope, but she knows few rhymes.

(15) The man who likes him will meet Bill tomorrow.

(16) John needs to talk to Bill about himself.

(17) John needs to talk to Bill about him.

(18) She does not realize that every girl is talented.

___

### 4.2  Ambiguity

*Ambiguity* arises when a single word or string of words is associated in the language system with more than one meaning. Each of the sentences in (57) illustrates a different way in which a single expression may be assigned multiple interpretations.

(57) a. You should have seen the bull we got from the pope.
   b. Competent women and men hold all the good jobs in the firm.
   c. Mary claims that John saw her duck.
   d. Someone loves everyone.

Sentence (57a) illustrates what is called lexical ambiguity: the form *bull* can be assigned at least three quite different interpretations (roughly, a papal communication, a male cow, or nonsense). The sentence is ambiguous because *bull* is ambiguous. To understand sentences containing that form, to identify their entailments, we need to know which of its three interpretations is being used. Lexical disambiguation is exactly like knowing which word has been used, like knowing, for example, that someone has uttered *cow* rather than *sow*. That is, an ambiguous lexical item can be thought of as several different lexical items that happen to be written and pronounced in the same way.

Sentence (57b) shows a simple kind of structural, or syntactic, ambiguity. We need not interpret any individual word as ambiguous but can attribute the ambiguity to distinct syntactic structures that give rise to distinct interpretations. Is *competent* modifying the

conjunction *women and men*, or is the NP *competent women* conjoined with the single-word NP *men*? One interpretation entails that the men holding the good jobs are competent, whereas the other does not. The English sentences in (58) unambiguously convey the two possible interpretations and thus allow us informally to disambiguate the original sentence.

(58) a. Women who are competent and men hold all the good jobs in the firm.
   b. Women who are competent and men who are competent hold all the good jobs in the firm.

Example (57c) illustrates both syntactic and lexical ambiguity. Is Mary claiming that John saw the bird she possesses or that he saw her lowering herself? These two interpretations are associated with radically different syntactic structures (*her duck* is in one case like *me jump* and in the other case like *my dog*) and also with distinct lexical meanings (the noun and the verb *duck* have the same spelling and pronunciation but quite distinct interpretations).

Sentence (57d) illustrates scope ambiguity. We can interpret the sentence as simply assigning some lover to each person (there is always the person's mother!) or as saying that someone is a universal lover (perhaps a divinity). The ambiguity here arises from the relation between *someone* and *everyone*: a scope ambiguity is not lexical but structural. But (57d) differs from (57b) and (57c) in having only a single surface syntactic structure. There have been arguments offered that sentences like (57d) do have multiple syntactic structures at some nonsurface level; we adopt such an approach in chapter 3. It is controversial, however, whether all scope ambiguities reflect syntactic ambiguities. If there are sentences whose ambiguity is nonlexical and that do not involve distinct syntactic structures, then structures or constructional principles that play no syntactic role are needed for semantic interpretation. We leave it as an open question whether there are any nonlexical, nonsyntactic ambiguities of this kind.

For linguistic purposes, ambiguity (multiplicity of interpretations assigned by the language system) is distinguished both from vagueness and from deixis or indexicality.

*Vagueness* is a matter of the relative looseness or of the nonspecificity of interpretation. For example, *many linguists* is noncommittal as to the precise number of linguists involved. It seems to be part of what we know about *many* that it is imprecise in this

sense. We discuss semantic imprecision in chapter 8. Virtually all expressions are general: *kiss* does not specify whether the kiss lands on the lips or cheek, etc., of the one kissed. But neither *many linguists* nor *kiss* would count as having multiple meanings on these grounds (that is, as synonymous with, for example, *350 linguists*, *400 linguists*, *379 linguists*, or again with *kiss on the lips*, *kiss on the cheek*).

*Deixis*, or *indexicality*, is involved when the significance of an expression is systematically related to features of the contexts in which the expression is used. For example, the first-person pronoun *I* is an indexical expression, but it is hardly ambiguous simply because it is sometimes interpreted as referring to Gennaro, sometimes to Sally, sometimes to you.

It is not always as easy to distinguish ambiguity from vagueness and indexicality as our examples might suggest, and we will return to these topics in later chapters. One test of ambiguity is the existence of distinct paraphrases for the expression in question, each of which conveys only one of the interpretations in question. An expression is a paraphrase of a declarative sentence for these purposes if it expresses exactly the same information as the original does on one way of understanding it; paraphrases will share all entailments with the given interpretation. Distinct paraphrases will usually have distinct entailments. The distinct interpretations must not be explicable in pragmatic terms; for example, "I'd like a glass of water" probably does not count as ambiguous, because how it is understood depends on pragmatic factors: on what an utterance of it is intended to accomplish. In general, expressions that are ambiguous can be used only with one of their meanings in any given situation. Exceptions are cases of punning and are clearly very special. There are many clear cases of lexical, structural, and scope ambiguities, and there are also some instances where intuitions do not settle the question of how different interpretations should be analyzed. For now, however, we simply want to emphasize that ambiguity is an important semantic phenomenon and that it is distinct from both vagueness and indexicality.

---

**Exercise 5** For each of the following sentences, state whether you judge it to be ambiguous, and for ambiguous sentences, disambiguate them by providing unambiguous distinct paraphrases of their possible interpretations.

(1) Everyone didn't like the movie.

(2) Someone came.

(3) Joan should be in New York.

(4) The missionaries are too hot to eat.

(5) The students are revolting.

(6) A lot of people came to Chomsky's lecture.

(7) Andrea is feared.

(8) Mindy likes Cynthia better than Jonathan.

(9) Visiting relatives can be tedious.

(10) Elizabeth didn't finish her thesis to please Jim.

(11) She was upset.

(12) John hit a boy with a book.

(13) John left early and Bill left early or Sue left early.

(14) Zelda ran the Boston marathon.

(15) Every faculty member was met by two student guides.

(16) Every student thinks that she is a genius.

---

## 4.3 Synonymy

In discussing ambiguity, we mentioned the notion of one expression's being a paraphrase of another, or synonymous with it. Judgments of *synonymy*, or semantic equivalence—that distinct expressions have the same meaning—turn out to be somewhat complex: they are relative to certain purposes or restricted to certain domains. If explicit content, that is, informational significance, is all that is at stake, then the sentences in (59) count as synonymous with one another: they share all their entailments, which is what we required of a disambiguating paraphrase.

(59) *a.* Those women at the corner table look ready to order.
     *b.* Those ladies at the corner table look ready to order.
     *c.* Those dames at the corner table look ready to order.

Suppose that one of these sentences is uttered by the head waiter to his underling. She doesn't quite catch what was said and asks an-

other one of the servers, who, to report what was said, might perfectly well reply,

(60) He said that ...

and choose any one of the sentences in (59) to complete (60). It is irrelevant to the server's immediate purposes how the other server identifies the customers that the head waiter says are ready to place an order, so long as she does so accurately. Even if the report is not the same sentence that the head waiter actually uttered, the reporter has not misrepresented the content of what he said. She has made a judgment of synonymy, or semantic equivalence, that fits with judgments of other native speakers.

The notion of synonymy involved here we call *content synonymy*, and we can define it in terms of mutual entailment.

(61) *A is (content) synonymous with B* $=_{df}$ *A entails B and B entails A.*

We could equally well have required that $A$ and $B$ share all their entailments, that is, that for any $C$, if $A$ entails $C$, then $B$ entails $C$, and vice versa. Two sentences will satisfy definition (61) if and only if they have all the same entailments. What content synonymy requires is just that $A$ and $B$ are true in exactly the same set of circumstances.

There is another sense in which speakers judge that the sentences in (59) have different meanings and thus are not (fully) synonymous. In choosing to utter one rather than another of these sentences to describe a situation, speakers can convey something important about their attitudes toward that situation and those involved in it. The differences involved are traditionally said to be connotations or a matter of tone; they may ultimately be a matter of presuppositions. In any case, they can be quite consequential. Suppose, for example, that the head waiter must later defend himself in a sex-discrimination suit filed by the server who was told what he had said. In this case how he said it does indeed matter.

Let us turn to some different examples. Speakers judge that the sentences in (62) share the same informational significance; they are content synonymous.

(62) *a.* The police searched Sarah.
    *b.* Sarah was searched by the police.

Again, this judgment seems to be grounded in the fact that (62*a*) entails (62*b*) and vice versa, that they share all their entailments. Yet doubt has been raised about the claim that passives are always content synonymous with the corresponding actives. Why? Precisely because in some cases it is not obvious that corresponding actives and passives do share all their entailments. For example, (63*a*), first discussed in Chomsky (1957), certainly does not entail (63*b*); (63*a*) is true, and (63*b*) false, if the circumstances are as described in (63*c*).

(63) *a.* Everyone in this room speaks two languages.
    *b.* There are two particular languages such that all the people in the room speak those languages.
    *c.* There are four people in the room, one of whom speaks only Italian and English, another only Finnish and Swedish, another only Hebrew and Russian, another only Yoruba and French.

The question is whether (64), the passive counterpart of (63*a*), is also true in the situation described by (63*c*) or in any other situation where (63*b*) fails to be true.

(64) Two languages are spoken by everyone in this room.

Here judgments are much less clear. What is clear is that the syntactic difference in (63*a*) and (64) leads to a difference in what an interpreter is likely to infer. From an utterance of (64) we are inclined to infer that the situation is not that described in (63*c*) but rather one where there are two particular languages that all speak, perhaps English and Japanese. Is this inclination a matter of entailment, or is it some less strong kind of implication? Here judgments are divided. The important point for our present purposes is not whether (63*a*) and (64) are content-synonymous, whether they express the same literal content. What matters for this discussion is the strong link between negative judgments on equivalence of content and negative judgments on identity of entailments.

No one is likely to deny, of course, that the difference between the active and passive can be important in interpretation. As we have just noted, (64) certainly suggests something that (63*a*) does not. And even where an active and passive clearly entail one another, as in (62) and many other pairs, substitution of one string for the other in certain contexts may fail to preserve mutual entail-

ments. The sentences in (65), for instance, clearly do not entail one another.

(65) *a.* Unwillingly the police searched Sarah. [The mayor forced them.]
   *b.* Unwillingly Sarah was searched by the police. [They had to tie her down.]

In other words, how a sentence structures the content it expresses can apparently matter to the contribution that sentence makes to the content of sentences in which it is embedded. Even if $A$ and $B$ have exactly the same entailments, it seems that two sentences $C(A)$ and $C(B)$ that differ from one another only in that $C(B)$ contains $B$ where $C(A)$ contains $A$ may differ in their entailments.

There are other ways in which sentences that express the same content can, in some sense, differ in meaning. For example, consider the different utterances in (66), the first of which places focus on *Mary*, the second of which places focus on *cake* (italics indicate focal stress). The sentences in (67), while structurally different, are identical in focal structure (and arguably also in entailments) to those in (66).

(66) *a.* *Mary* baked the cake.
   *b.* Mary baked the *cake*.

(67) *a.* It was Mary who baked the cake.
   *b.* It was the cake that Mary baked.

Sentences (66*a*) and (67*a*), which focus on *Mary*, might both be used, for example, to answer someone who uttered (68*a*), whereas (66*b*) and (67*b*), which focus on *cake*, strike us as badly suited for that job but just what is needed to answer someone who asks (68*b*).

(68) *a.* Who baked the cake?
   *b.* What did Mary bake?

It is sometimes claimed that perfect synonymy does not exist. What is usually meant by this is that formally distinct expressions are nearly always used in somewhat different ways, are appropriate in somewhat different contexts. This can involve their syntactic structure, their tone, what they suggest, the metaphoric possibilities they evoke, even matters of phonological and phonetic structure. If synonymy of distinct expressions means that we judge them appropriate in exactly the same range of contexts, effective for pre-

cisely the same purposes, then it is no surprise that plausible candidates are hard to find.

On the other hand, mutual entailment can be quite reliably judged, as can certain other properties relevant to semantic equivalence (for example, identity of focal structure). Mutual entailment, however, is basic; it generally provides the minimal basis for judgments of synonymy relied on in assessing accuracy of translations from one language to another and of second-party reports of what someone has said. Sometimes more is needed for an adequate translation or report, but mutual entailment is the necessary starting point.

### 4.4 Contradiction

*Contradiction* is intimately linked to entailment. When we said that (14), "Lee kissed Kim passionately," entails (15*d*), "Lee touched Kim with her lips," for example, we were guided by the judgment that (69), the conjunction of (14) with the negation of (15*d*), is contradictory.

(69) Lee kissed Kim passionately, but she [Lee] didn't touch him [Kim] with her lips.

What is meant by saying that (69) is contradictory? We can informally define contradiction in either of the following ways:

(70) $A$ is *contradictory* $=_{df}$
   • $A$ can never be true
   • there is no possible situation describable by $A$

That is, in judging (69) to be contradictory, we deem that it is false no matter what the facts might be, that it describes no possible situation. Contradiction can also be thought of as a relation between sentences; the informal definitions in (71) can get us started.

(71) $A$ and $B$ are *contradictory* $=_{df}$
   • $A$ and $B$ cannot both be true; whenever $A$ is true, $B$ is false, and whenever $B$ is true, $A$ is false.
   • a situation describable by $A$ cannot also be a situation describable by $B$

When we speak of one *person x* contradicting another *person y*, we mean that what $x$ has asserted contradicts what $y$ has asserted. Lois's response of no to her mother's assertion $A$ is tantamount to

an assertion by Lois of "not *A*," which contradicts her mother. *A* and *B* are said to be contradictories if each contradicts the other; *A* and *not A* are contradictories par excellence. If a sentence is contradictory, it will have entailments that are contradictories. More specifically, among its entailments will be a pair of sentences one of which is the negative of the other.

As with intuitions about entailments, initial judgments about contradictoriness can be subjected to further tests. We can defeat a claim that *A* and *B* are contradictory by showing a situation to which they both apply.

Sometimes sentences that overtly express contradictions are used for other purposes. For example, (72*a*) might receive as an answer (72*b*), which looks like a contradiction but is interpreted along the (noncontradictory) lines suggested in (72*c*). We do not simply interpret the speaker who utters (72*b*) as committed to an impossibility.

(72)  *a.* Is Andrea smart?
    *b.* She [Andrea] is [smart], and she [Andrea] isn't [smart].
    *c.* Andrea is smart in some respects but not smart in other respects.

We consider similar examples in more detail in chapter 8, section 5.

### 4.5  Anomaly

Contradictions are clearly incoherent; we might well say that (69) doesn't make sense because it entails contradictories. Few would be tempted to say that (69) is ungrammatical, however, or that it is completely meaningless. The problem seems to be that its meaning includes, in some sense, obviously incompatible parts, the two clauses that are conjoined. Each of the constituent clauses is, however, perfectly fine on its own; incoherence arises from combining them.

Incoherent sentences that are not surface conjunctions of contradictory sentences do not so blatantly generate contradictory entailments. Indeed, their incoherence is often such that we are hard pressed to see that they have any entailments at all. Linguists have spoken of *anomaly* in cases like those illustrated in (73).

(73)  *a.* The square root of Milly's desk drinks humanity.
    *b.* Colorless green ideas sleep furiously.

    *c.* To laugh is very humid.
    *d.* The fact that cheese is green skipped inadvertently.
    *e.* Being a theorem frightens consternation.
    *f.* My toothbrush is blonde and buxom.
    *g.* That rock thinks it's too good to hold the door open.

Chomsky (1965) introduced the notion of selectional restrictions to mark such sentences as ungrammatical. A verb like *drink*, he noticed, carries the information that its object designates something drinkable—a liquid or semiliquid substance perhaps, but at the very least something concrete rather than abstract—and that its subject designates something that might be a drinker, minimally, an animate being, we might suppose. The idea, then, was to provide a mechanism to ensure that *drink* selects only arguments satisfying such restrictions. From information given in its lexical entry, *drink* would be marked by something like the following "selectional feature":

(74)  [+[+animate] ____ [−abstract]]

This is a contextual feature indicating that *drink* must only be inserted where there is a preceding animate subject and a following nonabstract object. Subject and object NPs, it was assumed, would receive feature specifications from their head nouns; *humanity*, for example, would be marked [+abstract] and *square root* [−animate]. Violations of selectional restrictions would arise from mismatches between features and would be ungrammatical.[8]

Sentences like those in (73) do seem very strange, and their strangeness seems different from that of a simple contradiction like (69), "Lee kissed Kim passionately, but she didn't touch him with her lips." The constituent clauses in (69), "Lee kissed Kim passionately" and "she [Lee] didn't touch him [Kim] with her lips," are each semantically unproblematic; each describes a possible situation. The oddness of (69) is that passionate kissing and not touching with the lips are brought together in a single event. The anomalous sentences in (73) are not strange in precisely the same ways or to the same degree. Some of them even seem more susceptible to being put to good use than does (69). We can imagine ways of interpreting sentences like (73*f*) and (73*g*), for example (someone might, for instance, have a toothbrush that looks like a woman, or someone might pretend or even believe that rocks are thinking beings). Yet (73*a*) and (73*e*) seem virtually impossible to

make any sense of (the very notion of square root would seem to preclude a desk's having one or it's being something that might drink, and similarly, consternation seems incomparable to the kinds of things that can be frightened: people and other sentient beings).

It has often been pointed out that poetic uses of language are sometimes anomalous if interpreted in the usual and most obvious ways. Personification, for example, is a familiar poetic device, and (73f) and (73g) might easily be interpreted from that perspective. But the very fact that interpretation of these sentences typically suggests that their utterers are presenting toothbrushes and rocks as personlike calls for some explanation. Sentence (73b), famous from Chomsky's use of it in *Syntactic Structures* to illustrate the possibility of divergence between grammatical and semantic well-formedness, is the final line of a poem by John Hollander, where it seems vaguely evocative.[9] Again, the question of how its appropriateness there is achieved needs to be addressed.

The point is not that the sentences in (73) are semantically acceptable (although some of them may be) but rather that they are semantically distinct from one another, and a theory that simply marks them all as meaningless does not reveal this. As in the case of straightforward contradictions, the individual words and the syntactic constructions are semantically unproblematic; what is odd are the combinations, and some are much odder than others.

In some cases the oddness seems linked more to the structure of the world than to facts about linguistic meaning: rocks just aren't the kind of thing that thinks, as it happens, but this seems less a matter of what *rock* and *think* mean than a matter of what rocks and thinking are like. People are inclined to say that someone might wonder or claim or wish that rocks think. The study of artificial intelligence has raised the possibility of machines' thinking, a possibility that might well have been deemed as strange a century or so ago as that of rocks' thinking. On the other hand, (73e) seems far more peculiar; because it is an abstract entity, consternation is completely outside the realm of things than might be frightened. We cannot begin to understand someone's wondering whether consternation has been frightened. Someone who utters (73e) with apparent seriousness will be thought to have made a slip of the tongue or some other linguistic mistake (perhaps not knowing the meanings of some of the words used), to be suffering from some form of aphasia, to be mentally disturbed in some way. It would be quite strange for another to report the event by saying,

(75) Lindsey wonders whether being a theorem frightens consternation.

Sentence (75) seems hardly easier to interpret than sentence (73e). Similarly, sentence (73a) seems to resist any kind of interpretation: a desk is not a number and therefore in some fundamental way not the sort of thing that could have a square root, and numbers are not the sort of things that drink.

The correct conclusion may be that judgments of anomaly pick out a somewhat heterogeneous set of expressions, some of which are simply contradictions (with the incompatible entailments perhaps less immediate than in the cases that are obvious contradictions), others of which describe situations that are bizarre because of how the world works, and others of which involve a kind of semantic incompatibility other than that of contradiction (perhaps a semantic analogue of the notion of a violation of selectional restrictions).

What might this special kind of semantic incompatibility be like? It might somehow be part of the meaning of *drink*, for example, that it is only predicable of a certain range or *sort* of object, a sort that does not (at least in normal or literal uses) include square roots. Though it might be difficult to decide for a particular sentence whether it is *sortally deviant* (what is often called a *category mistake* in the philosophical literature) or anomalous in some other way, semantic anomaly, as illustrated in (73), is quite pervasive, is apparently distinct from the other phenomena we have considered, and seems clearly to call for some kind of semantic account.

One proposal is that some kinds of anomaly involve incompatible presuppositions. This would make anomaly analogous to contradiction, which involves incompatible entailments. The problem of distinguishing (certain cases of) anomaly from contradiction would then reduce to the problem of distinguishing presupposition from entailment, a matter we have touched on already and will later take up in more detail.

### 4.6 Appropriateness

One characteristic of anomalous expressions is that they are inappropriate for use in most contexts. People seem able to judge that particular expressions are or are not *appropriate* for uttering in particular contexts, and some have tried to incorporate an account of appropriateness conditions into a theory of linguistic semantics.

As we noted above in section 3.2, sentences are often judged inappropriate for contexts where their presuppositions are at issue or somehow controversial. Appropriateness is sometimes held to be a more general and useful notion for semantic theory than that of truth, or descriptive applicability, which was central to our discussion of entailments and contradictions. Only declaratives are sensibly said to describe a situation, or to be true of certain circumstances; interrogatives and imperatives are susceptible to the defect of inappropriateness rather than that of falsity. It is sometimes thought that a theory of appropriateness might replace a semantic theory based on truth. Appropriateness is often appealed to in explaining how speech acts are performed, how we manage to "do things with words": assert, inquire, promise, entreat, and the like. Some examples will illustrate. It is inappropriate for us to promise you to do something that we do not believe ourselves capable of doing (teach you all there is to know about meaning) or to do something we have no intention of doing (resign our positions if you don't like our book). It is inappropriate to assert something that we do not ourselves believe or that we do not want to give you reason to believe. It is generally inappropriate to inquire whether pigs have wings if we know whether pigs have wings (though, of course, examiners in pig biology may put the question to their students, knowing full well its answer). In chapter 4, we discuss speech acts in some detail. To perform a certain speech act is, in part, to adopt a certain attitude toward the content of what one says and perhaps also sometimes to urge a certain attitude on the part of the hearer ("Is that a promise or a threat?").

A related but slightly different area where appropriateness is appealed to is in judgments of whether a particular expression fits in a particular discourse slot, whether the discourse itself is sensible, coherent. If you have just uttered (76a) to the instructor, then (76b) seems highly inappropriate as her response.

(76)  a. Can I have a copy of the answer sheet?
      b. Yes, and Joan is similar.

There are clearly many more factors involved in assessing discourse appropriateness than what linguistic expressions mean. For example, relevance is a factor in assessing discourse appropriateness, and knowing what is relevant may involve all kinds of nonlinguistic knowledge. It seems quite unlikely that we could explicitly specify for all sentences of the discourse all the contexts in

which they might be appropriate, though for some expressions we might be able to characterize at least partially the class of inappropriate contexts (see the discussion of presupposition in chapter 6).

Appropriateness is also invoked in dealing with matters of stylistic register: certain forms are reserved for church services, others are appropriate for the locker room, others for family dinners. It is generally inappropriate to mix registers, to use them in the wrong contexts, just as it is inappropriate to wear tennis shoes with a ball gown or to wear a ball gown to your linguistics class. Appropriateness here seems linked to cognitive significance: choosing a certain style signals a certain attitude toward the speech situation.

The notion of appropriateness is thus something of a mixed bag. Appropriateness does not seem to be structured like truth. There is no generally recognized relation of one expression's being dependent on another for its appropriateness parallel to the entailment relation, where one sentence must be true if another is. Nor does appropriateness seem to be readily amenable to a compositional treatment; certainly, no one has offered any general account of how to project appropriateness of (indefinitely many) complex expressions from appropriateness-related properties of their constituents. In other words, it does not seem that appropriateness will replace truth as a fundamental notion for semantic theory.

Nonetheless, recent work on such topics as presupposition has suggested that certain aspects of appropriateness may be characterizable in a much more rigorous way than was once thought possible. As we pointed out, the sentences "Lee got a perfect score on the semantics quiz" and "It was Lee who got a perfect score on the semantics quiz" entail one another; truth-based considerations do not distinguish them. The latter sentence, however, presupposes that someone got a perfect score, whereas the former does not. As we shall see in chapter 6, the presupposition of the cleft restricts the range of contexts in which its utterance is appropriate. It would be inappropriate to utter it in response to the question "Did anyone get a perfect score on the semantics quiz?" for example. Considerable progress is being made in developing empirically sound and theoretically sophisticated discourse theories that elucidate what is involved in such judgments for these and certain other kinds of cases. We will also see that something systematic can be said about how presuppositions of complex sentences relate to the presuppositions of constituent sentences.

**mmary**

We have given the reader an indication of the main aspects of language that a theory of meaning must deal with. Meanings form a productive system in which new meanings can always be expressed. There are aspects of meaning that may be constant across all human languages. Furthermore, meaning encodes information about the world and plays a role in giving a shape to our mental states. A theory of meaning must shed light on all these issues. We have also discussed the different types of semantic judgments in which what we know about meaning manifests itself, and we have provided a preliminary classification of such judgments. We are capable of assessing certain semantic properties of expressions and how two expressions are semantically related. These properties and relationships and the capacity that underlies our recognition of them constitute the empirical base of semantics.

In presenting a theory of semantics that tries to shed light on all these aspects of meaning, we are guided throughout by what Jackendoff (1983, 13) dubs the "grammatical constraint": "prefer a semantic theory that explains otherwise arbitrary generalizations about the syntax and the lexicon." The adherence to this constraint is what perhaps most sharply distinguishes our approach from that of philosophical logicians.

# 2 Denotation, Truth, and Meaning

## 1 Introduction

We have outlined what we think the empirical coverage of a theory of meaning should be. This will help us in directly addressing the question, What is meaning? Answers should be evaluated on the basis of how well they account for the phenomena singled out in chapter 1.

The question of what meaning is, is important to any discipline concerned, directly or indirectly, with cognition, that is, with how humans process information. To indicate where we stand with respect to some of the traditional views of meaning, it is convenient to classify approaches to meaning in three groups.

The first family of theories can be labeled "referential" or "denotational." This kind of theory is outward looking; its main emphasis is on the informational significance of language, its aboutness. Meaningfulness lies, according to this view, in the relations of symbols and configurations thereof to objects of various kinds. The study of meaning is the study of such relations. This tradition is the basis of the semantic techniques that have been developed within mathematical and philosophical logic.

It seems reasonable to maintain that the study of the relation of symbols to what they stand for must indeed be part of an account of meaning. For otherwise, how could we understand the fundamental fact that configurations of symbols carry information about all the diverse aspects of our experience?

A second family of theories of meaning might be labeled "psychologistic" or "mentalistic." Theories of this sort are inward looking and focus on the cognitive significance of language. The meaning of a configuration of symbols, according to this view, lies in what we grasp when we manipulate them; that is, it lies in the internalized representation of their retrievable content. The study of meaning is the study of how contents are mentally represented, the study of semantic representations. This tradition is the basis of much semantic work in psychology and artificial intelligence.

It seems reasonable to maintain that a given configuration of symbols has meaning for us only if we are able to grasp its content,