# PRAGMATICS

R A READER

EDITED BY

## STEVEN DAVIS

# Pragmatics

## A Reader

*Edited by* Steven Davis

# Oxford University Press

Oxford   New York   Toronto
Delhi   Bombay   Calcutta   Madras   Karachi
Petaling Jaya   Singapore   Hong Kong   Tokyo
Nairobi   Dar es Salaam   Cape Town
Melbourne   Auckland

and associated companies in
Berlin   Ibadan

## Copyright © 1991 by Oxford University Press, Inc.

# Contents

## III Indexicals

## IV Direct and Indirect Speech Acts

## V Conversational Implicature and Relevance

# VIII  Psychology and Pragmatics

# Introduction

The term 'pragmatics' was first introduced in *Foundations of the Theory of Signs* by Charles W. Morris, who contrasts it with semantics and syntax. For Morris

> syntax [is] the study of the syntactical relations of signs to one another in abstraction from the relations of signs to objects or to interpreters; . . . semantics deals with the relation of signs to designata and so to objects which they may or do denote [and] 'pragmatics' is designated the science of the relation of signs to their interpreters. (1971, pp. 28, 35, 43)

Morris's use of 'sign' here is a bit confusing, since syntax studies the grammatical relations of morphemes[1] (for example, 'the'), which are not signs. Let us replace Morris's 'sign' with 'linguistic unit,' which applies to morphemes, phrases, and sentences. With this change we can take Morris to claim that syntax is the study of the grammatical relations of linguistic units to one another, and the grammatical structures of phrases and sentences that result from these grammatical relations; that semantics is the study of the relation of linguistic units to the world; and that pragmatics is the study of the relation of linguistic units to their users. Expanding his definition, Morris tells us that pragmatics is the study of "the biotic aspects of semiosis, that is, . . . [the study of] all the psychological, biological, and sociological phenomena which occur in the functioning of signs" (1971, p. 43). The problem with this broad view of pragmatics is that it is too inclusive to be of much use. Using this definition, pragmatics has as its domain any human activity involving language, and thus includes almost all human activity, from baseball to the stock market. The consequence is that all the human sciences become part of pragmatics. On this view, then, pragmatics is not on the same level as semantics and syntax, when these are construed as theories constructed to account for various aspects of a speaker's linguistic ability. Nor can pragmatics be regarded as a field of

study, like linguistics or sociology. What groups various activities and theories together in one field of study is that they share a set of questions or a methodology. But there is no common methodology or set of questions that groups together in a natural class the full range of the human sciences in which language is involved. Economics and socio-linguistics, for example, have very little in common to justify their inclusion in the same field of study. For 'pragmatics' to be a useful term, its domain must be restricted. Let us take a fresh look at the problem.

Pragmatic theories can be considered in two ways. On the first view, a pragmatic theory is part of a theory of a speaker's linguistic competence. As such, it is part of a psychological theory that plays a role in accounting for what speakers tacitly know which enables them to understand and to use sentences of their language. On the second view, pragmatics is not part of a theory of competence. It is a theory that attempts to account for a range of pragmatic facts[2] without making any commitment to whether the theory is psychologically realized. A clear view of this sort is held by Richard Montague, who claims that pragmatics is a branch of "mathematics, not of psychology" (1974, p. 2). Space does not allow me to argue for the first view, which I shall adopt here. I shall regard pragmatics as part of a theory of competence and, as such, take it to be psychologically realized.

We carry on conversations, and most of the time we are able to understand one another. How do we do this? What capacities, faculties, and knowledge enable us to communicate with one another with relative ease? According to Noam Chomsky, one of the capacities we have is our knowledge of the syntactic rules of our language. These rules account both for our intuitions about whether a string of words constitutes a well-formed sentence of our language and for our knowledge of the syntactic facts of our language. For example, our intuition of the following strings of words

(1) Boy the street the down ran.
(2) Flying planes can be dangerous.

is that (1) is not well formed, and (2) is ambiguous where the ambiguity arises from the syntactic structure of the sentence. Knowing syntactic rules is hardly sufficient to enable us to carry on a conversation. We must know more; we must understand the sentences of our language. To understand our language, we also must know its semantics. Our semantic knowledge is of two sorts. We know the meanings of the words of our language and the rules that enable us to combine these meanings to form meanings of phrases and sentences. We know, for example, that if a speaker, let us call her 'Alice,' were to say,

(3) Sam threw Sally a ball.

she could mean that Sam threw Sally a spherical object, or that Sam arranged a party for Sally at which people danced and dressed in formal clothes. Therefore, we can add to our syntactic knowledge our knowledge of the meanings of the words of our language

and the rules by which we can combine them to form meanings for larger units. In some cases, our knowledge is the sort illustrated by (3), where we are in a position to give 'definitions' for a word. In other cases, our knowledge of the meaning of a word may consist in our being able to use the word appropriately in a wide range of circumstances. We are hard pressed to give the meaning of 'the,' but there is no doubt that we are able to use it correctly and, thus, we can be said to know the meaning of the word. We may argue that despite our inability to give a definition for 'the' and similar words, we tacitly know their meanings. Let us call the theory that describes the knowledge we have—tacit or explicit—of the meanings of the words of our language and the rules that combine them into meanings of larger units *a theory of meaning.*

A theory of meaning does not exhaust our semantic knowledge. We also know what the terms of our language refer to and what the *truth conditions* of the declarative sentences are. For example, in saying (3), Alice knows to whom 'Sam' and 'Sally' refer, what 'throw' is true of, and the conditions under which the sentence she utters is true. In some cases, Alice's knowledge of what the terms of her language refer to consists in her ability to identify the objects to which the terms refer; in other cases, she may know no more than that a name like 'Einstein' refers to Einstein, without knowing who he is. In fact, Alice may even be mistaken about who he is; she may confuse him with Bohr. Despite this, she can use 'Einstein' to refer to Einstein. What enables her to do so is that she is part of a linguistic community extending back in time to speakers who are causally connected to Einstein (Kripke 1980, pp. 96–98).

There is a relationship between reference and truth conditions. To illustrate this, let us consider an example that is simpler than (3). Suppose Alice says,

(4) Sally is a woman.

What Alice says in uttering (4) is true just in case 'a woman' is true of the person to whom Alice refers in using 'Sally.' We shall assume that Alice knows the truth conditions of the declarative sentences of her language. Her knowledge of these truth conditions cannot be given by a list of the sentences paired with their truth conditions. She can have no such knowledge, since the declarative sentences of her language are an infinite set. Rather, her knowledge consists in knowing a finite set of rules that specify both the referent for each term of her language and the truth conditions for each declarative sentence of her language.

Truth conditions are available for only the declarative sentences of a language. Analogous conditions apply to interrogatives and imperatives. For interrogatives, *answer conditions* specify the set of possible answers to the questions they can be used to ask. For imperatives, *compliance conditions* specify the set of actions that carry out the requests and commands which they can be used to make. Let us call truth conditions, answer conditions, and compliance conditions, *satisfaction conditions,* and a theory that specifies these conditions, a *theory of satisfaction.* We shall regard a theory of satisfaction to be a theory of Alice's knowledge of the satisfaction conditions of her language. Once again, this should not be taken to imply that Alice consciously knows these

conditions. If she were asked, she could not tell us what the satisfaction conditions are for certain sentences of her language. In the same way that speakers have tacit knowledge of the syntactic rules of their language, and of the rules that assign meanings to the sentences of their language, we can take it that Alice has tacit knowledge of the rules that assign satisfaction conditions to the sentences of her language. Taken together, a theory of meaning and a theory of satisfaction constitute a semantic theory for Alice's language and are part of a theory of her linguistic competence that enables her to communicate with others.

Let us digress for a moment to try to give an initial characterization of Alice's syntactic and semantic knowledge (a characterization that later will be revised), so that we can make a first attempt to distinguish it from the knowledge she has for which a pragmatic theory should give an account.

Three characteristics are often taken to be features of Alice's syntactic and semantic knowledge. First, it is held to be linguistic knowledge about the linguistic units of her language.[3] Second, it is regarded as independent of context. It is knowledge that does not depend on her knowing anything about the context in which the sentences of her language are used. Last, it is claimed that it is knowledge which is independent of her knowing her own intentions, wants, desires—her *intentional states*—which she has in saying what she does. For example, if Alice knows that 'run' is a verb and that 'ball' means *a spherical object used in games* or *a formal dance,* the knowledge she has is knowledge about English; it is not dependent on her knowing anything about the particular context of any utterance, and no reference is made to any of her intentional states. Let us call such knowledge *context-independent linguistic knowledge.* This is the sort of knowledge that Alice can be said to have in all contexts and that she can bring to bear in any context in which it is relevant in communicating with others and in understanding what they communicate to her.

Alice's context-independent linguistic knowledge does not exhaust the knowledge that enables her to communicate with others. She also knows how to use her language. As we shall see, this knowledge includes knowing the rules or maxims that govern conversation and the intentional states she must have in particular contexts which are necessary for various uses of her language. Let us call this knowledge of pragmatic facts *conversational knowledge* and take as an initial hypothesis (which we shall modify) that pragmatics is the study of conversational knowledge. It will be helpful to have before us examples of pragmatic facts. The examples clearly are not within the domain of syntax, and most of them do not appear to fall within semantics. The last example we shall consider, however, raises the question of whether a boundary can be drawn between conversational knowledge, on the one hand, and semantic knowledge, on the other. This boundary is important because if we were not able to draw a line between the two, we would not have separate sets of facts that constitute the domains of semantic theory and of pragmatic theory. Before considering the question of whether a boundary can be drawn between conversational knowledge and semantic knowledge, let us look at some examples of pragmatic facts.

Alice can play a dual role in a conversation. She can be at once speaker and hearer. Let us concentrate on her role as speaker. Suppose that in uttering

(5) I'll be in New York.

Alice is speaking literally. The sentence that Alice uttered can be used to make a promise, to declare her intention, or even to make a prediction about her future behavior. The acts that Alice can perform in uttering (5)—promising, declaring her intention, or predicting her own behavior—are called *speech acts.*

It may be thought that it is a semantic fact about (5) that it can be used to perform these speech acts. One proposal is that (5) is ambiguous and means the same as either

(6) (a) I promise that I'll be in New York.
   (b) I predict that I'll be in New York.
   (c) I declare my intention to be in New York.

Which of these (5) means can be taken to vary from context to context, but in any context it must mean one of these, if the speaker uses the sentence with one of its standard meanings. I believe that the above proposal is mistaken. The hypothesis is that (5) means the same as (6) (a), (b), and (c), and in a particular context which of these (5) means depends on which of these the speaker means by it. Let us suppose that the speaker uses (5) to mean the same as (6) (a). If (5) and (6) (a) have the same meaning, they should be true under the same satisfaction conditions. But (5) is true just in case Alice is in New York at some time after she says (5), while (6) (a) is true just in case Alice promises that she will be in New York. Since the sentences are true under different conditions, they cannot be synonymous. Thus the proposal is false.[4] This argument does not show that there is no semantic relation between (5) and the range of speech acts which it can be used to perform. But in the absence of such a proposal, I shall assume that the relationship is a pragmatic fact. Which speech act Alice performs in uttering (5) on a particular occasion depends on her knowing the rules that govern the speech acts that (5) can be used to perform, and on whether she intends in uttering (5) to be making a promise or prediction or to be declaring her intention (see Part IV).

This is not the only case in which Alice's knowledge of rules governing the use of language or her intentions on a particular occasion of utterance play a role in determining what she intends to communicate. Consider the following sentence:

(7) Sam really had a ball.

Since 'ball' is ambiguous, in uttering (7) Alice could be saying at least two things. The ambiguity of (7) is determined by the semantics of Alice's language. But her language does not determine which of these she means on a particular occasion. It is determined by what she intends the word 'ball' to mean on this occasion. In turn, her intention to

mean what she does by 'ball' plays a role in determining what she says in uttering (7) (see Part II).

In the preceding examples I have assumed that Alice is being straightforward. But this need not be so. Suppose that Alice is a professor of philosophy, and she is asked to write a letter of recommendation for Sam, one of her students. She writes,

(8) Sam has good handwriting.

In so writing, Alice implies that Sam is not a very good student of philosophy. There is no semantic or syntactic connection between (8) and what Alice implies. How, then, is she able to imply what she does by including in her letter of recommendation her comments about Sam's handwriting? She does so by exploiting certain maxims of conversation that she assumes her audience takes her to be following—one of which is 'Be relevant.' In violating the maxim of relevance, as she manifestly does in writing (8), she intends that her colleague to whom she wrote the letter recognize that she has violated the maxim and that she intended to do so. Moreover, Alice intends that her colleague, in recognizing that she has manifestly violated the maxim of relevance in commenting on Sam's handwriting (which clearly is not relevant to his philosophical ability), will draw the inference that she (Alice) does not think very highly of Sam's philosophical ability. Hence, in writing (8) Alice implies that Sam is not very good at philosophy by virtue of her knowledge of the maxims of conversation, her assumption that her audience has such knowledge, and her intention that her audience recognize her intention to exploit these maxims in saying what she does[5] (see Parts II and V).

Let us change the example to illustrate another contextual feature that plays a role in determining what Alice is communicating. Suppose that Alice says,

(9) She is tired.

Alice can use instances of (9) to say different things. In part, what she says in uttering (9) depends on the person to whom she refers in using 'she.' For example, if she were to refer to Betsy, she would be saying that Betsy is tired; if she were to refer to Ellen, she would be saying that Ellen is tired. Thus Alice's ability to use expressions to refer and her intention on the particular occasion in which she utters (9) to refer to some particular person play a role in determining both that she says something and what she says on that occasion (see Part II).

We distinguished between two sorts of semantic theories: a theory of meaning and a theory of satisfaction. Part of the theory of satisfaction is a theory of truth that is supposed to specify the truth conditions for each declarative sentence of the language. The question that arises here is whether a truth theoretic semantics can specify truth conditions without in some cases making reference to the intentions of speakers? The problem is that if truth theoretic semantics were to refer to speakers' intentions, it would cross the boundary between pragmatics and semantics, since we have supposed that speakers' intentions fall within the domain of pragmatics. The result would be that the

distinction between semantics and pragmatics would be blurred, and we would not have a separate domain of facts for semantics and pragmatics.

The problem is neatly illustrated by (9). Before considering this example further, it will be useful to distinguish between *speaker referent* and *semantic referent*. The semantic referent of an expression is the referent that the expression has in the language (e.g., 'dog' in English refers to dogs). The semantic referent of an expression is fixed by the conventions of the language. The speaker referent of an expression is fixed by what a speaker refers to in using the expression on a particular occasion, which is determined by what the speaker intends to refer to on that occasion. A speaker might use 'dog' on a particular occasion to refer, mistakenly, to some wolves (Kripke 1977, pp. 96–97).

Let us return to (9). Some pronouns—for example, 'she'—have no semantic referent. Their referents are not fixed by the conventions of the language but by what the speaker intends to refer to in using such terms on a particular occasion. In addition, the referent of 'she' plays a role in determining the truth conditions of (9), truth conditions that it has relative to the context in which it is used. (9) is true just in case the person to whom Alice refers is tired.

It may seem that the truth conditions for (9) can be specified in (10) independently of context and of the intentions of a particular speaker:

(10) 'She is tired' is true just in case the person to whom the speaker refers in uttering this sentence is tired.

This will not do. One of the purposes of a semantic theory for a language is to account for the logical relations among the sentences of the language. Consider the following argument:

(11) She is tired.
     She is Margaret Smith.
     _____
     Therefore, Margaret Smith is tired.

There is one interpretation of (11) in which it turns out to be valid and one in which it turns out to be invalid. If 'she' refers to the same person on both of its occurrences, then (11) is valid; if it does not, then (11) is invalid. Suppose we assume that (10) gives the truth conditions for the first sentence of the argument and that (12) gives the truth conditions for the second sentence.

(12) 'She is Margaret Smith' is true just in case the person to whom the speaker refers in using the sentence is Margaret Smith.

The problem is that (11) and (12) do not account for the interpretation of the argument in which it is valid. There is nothing in these truth conditions which guarantees that the speaker who uses the first and second sentences of the argument is the same or, even if this were the case, that the person to whom the speaker refers in using 'she' in

the two sentences is the same. Thus to give the truth conditions of certain sentences of natural languages requires that the truth conditions mention particular speakers and their intentions on particular occasions for those sentences that contain expressions whose referent is not determined by the conventions of the language. Consequently, semantic theory—at least the theory of satisfaction—must make reference to particular speakers and their referential intentions.

The problem is that these latter features seem to be the mark of the pragmatic, and thus, if the theory of satisfaction makes reference to particular speakers and their referential intentions, it should be considered to be part of pragmatic theory rather than semantic theory. This means in effect that the semantics of natural languages consists only of the theory of meaning.[6] I think, however, that there is a way to avoid this consequence by distinguishing among speakers' intentions—specifically, between their referential and communicative intentions. The referential intentions of a speaker determine the semantic referents of terms whose referents are not fixed by the conventions of the language. This includes demonstratives, such as 'this' and 'that'; indexicals, such as 'she' and 'it'; and incomplete definite descriptions, such as 'the table.' Speakers' communicative intentions are the intentions speakers have to get their audience to recognize what they are trying to communicate. In the case of terms that can be used to refer to something, the speakers' communicative intentions are to get their audience to understand to what they are referring. Let us illustrate the difference between referential and communicative intentions. Suppose that there is a room full of brown tables. Alice says,

(13)  That table is brown.

In saying this, Alice intends to refer to the only table that is under the window. Alice has referred to the table and said something true if the table is brown. Imagine further that Alice mutters (13) under her breath. In uttering (13) Alice need not intend that anyone understand to what she is referring and thus what she is saying. Given the situation described, it is difficult to see how anyone could understand either to what she referred or what she said. The utterance of 'that table' has no connection to what Alice is referring, except through her intention, which is not, however, manifest in what she says. Let us change the example a bit and suppose that Alice addresses (13) to an audience—for example, to Sam—and that in doing so she does not do anything to indicate to him to what table she is referring. In this context, it is natural for Sam to think that Alice wishes him to understand what she is saying and, therefore, intends that he understand to what she is referring. That is, because Alice addresses (13) to Sam, Sam has reason to think that Alice has communicative intentions. In part, for Alice to have communicative intentions is to have the intention that her audience recognize these intentions (Grice 1968, pp. 82–83). In this case, however, Sam has reason to be puzzled. The assumption that Alice has the communicative intention that Sam understand to which table she is referring is in conflict with the fact, which is mutually obvious to Sam and Alice, that on this occasion in uttering (13) Alice has not provided Sam with