

OXFORD COGNITIVE SCIENCE SERIES

*General Editors*

MARTIN DAVIES, JAMES HIGGINBOTHAM, JOHN O'KEEFE,  
CHRISTOPHER PEACOCKE, KIM PLUNKETT

Forthcoming in the series

*Mindreading*

Stephen Stich and Shaun Nichols

*Face and Mind: The Science of Face Perception*  
Andy Young

# CONTEXT AND CONTENT

*Essays on Intentionality in  
Speech and Thought*

ROBERT C. STALNAKER

OXFORD  
UNIVERSITY PRESS

relied on the heuristic accounts of others, but I have made no attempt to reduce them to each other, or to anything else. It may be charged that these concepts are too unclear to be the basic concepts of a theory, but I think that this objection mistakes the role of basic concepts. It is not assumed that these notions are clear. In fact, one of the points of the theory is to clarify them. So long as certain concepts all have *some* intuitive content, then we can help to explicate them all by relating them to each other. The success of the theory should depend not on whether the concepts can be defined, but on whether or not it provides the machinery to define linguistic acts that seem interesting and to make conceptual distinctions that seem important. With philosophical as well as scientific theories, one may explain one's theoretical concepts, not by defining them, but by using them to account for the phenomena.

## 2

## Pragmatic Presuppositions

There is a familiar intuitive distinction between what is *asserted* and what is *presupposed* in the making of a statement. If I say that the Queen of England is bald, I presuppose that England has a unique queen, and assert that she is bald. If I say that Sam regrets that he voted for Nixon, I presuppose that Sam voted for Nixon, and assert that he feels bad about it. If I say that Ted Kennedy is the only person who could have defeated Nixon in 1972, I presuppose that Ted Kennedy could have defeated Nixon in 1972, and assert that no one else could have done so. Philosophers have discussed this distinction mainly in the context of problems of reference. Linguists have discussed it in many other contexts as well. They have argued that the phenomenon of presupposition is a pervasive feature of the use of natural language, one that must play a role in the semantic analysis of many words and phrases.

The principle criterion that has been used to identify presuppositions can be stated in the following way: *Q* is presupposed by an assertion that *P* just in case under normal conditions one can reasonably infer that a speaker believes that *Q* from either his assertion or his denial that *P*. One who denies the example statements listed above—who says that the Queen of England is *not* bald, that Sam does *not* regret that he voted for Nixon, or that Ted Kennedy is *not* the only person who could have defeated Nixon in 1972, normally makes the same presuppositions as the person who makes the affirmative statements. Linguists have used this criterion to identify many examples of the phenomenon. The criterion, and many of the examples, are relatively clear and uncontroversial; it is clear that there is a phenomenon to be explained. But it is much less clear what kind of explanation of it should be given. Granted that either the statement that the Queen of England is bald, or the speaker who makes it, presupposes that England has a unique queen. But what is it about the statement, or

This paper was read at the University of Texas conference on performatives, conversational implicature, and presupposition in March, 1973, as well as at New York University. I, and I hope the paper, benefited from stimulating comments by linguists and philosophers at both places.

the speaker, which constitutes this fact? There are two very different kinds of answers to this question.

The first answer is that presupposition is a semantic relation holding between sentences or propositions. This kind of account draws the distinction between presupposition and assertion in terms of the content or truth-conditions of the sentence uttered or the proposition expressed. Here is an example of such a definition: a proposition that *P* presupposes that *Q* if and only if *Q* must be true in order that *P* have a truth-value at all. The presuppositions of a proposition, according to this definition, are necessitated by the truth, and by the falsity, of the proposition. When any presupposition is false, the assertion lacks a truth-value.

The second answer is that presupposition should be given a pragmatic analysis. The distinction between presupposition and assertion should be drawn, not in terms of the content of the propositions expressed, but in terms of the situations in which the statement is made—the attitudes and intentions of the speaker and his audience. Presuppositions, on this account, are something like the background beliefs of the speaker: propositions whose truth he takes for granted, or seems to take for granted, in making his statement.

The pragmatic account is closer to the ordinary notion of presupposition, but it has frequently been assumed that the semantic account is the one that is relevant to giving a rigorous theoretical explanation of the linguistic phenomena. I want to argue that this assumption is wrong. I will suggest that it is important for correctly understanding the phenomena identified by linguists to give the second kind of analysis rather than the first. In terms of the pragmatic account, one can give intuitively natural explanations of some facts that seem puzzling when presupposition is viewed as a semantic relation. The pragmatic account makes it possible to explain some particular facts about presuppositions in terms of general maxims of rational communication rather than in terms of complicated and *ad hoc* hypotheses about the semantics of particular words and particular kinds of constructions. To argue this, I will sketch an account of the kind I want to defend, and then discuss some of the facts identified by linguists in terms of it.

Let me begin by rehearsing some truisms about communication. Communication, whether linguistic or not, normally takes place against a background of beliefs or assumptions which are shared by the speaker and his audience, and which are recognized by them to be so shared. When I discuss politics with my barber, we each take the elementary facts of the current political situation for granted, and we each assume that the other does. We assume that Richard Nixon is the President, that he recently defeated George McGovern by a large margin, that the United States has

recently been involved in a war in Vietnam, which is a small country in Southeast Asia, and so forth. That we can reasonably take these facts for granted obviously makes our communication more efficient. The more common ground we can take for granted, the more efficient our communication will be. And unless we could reasonably treat *some* facts in this way, we probably could not communicate at all.

Which facts or opinions we can reasonably take for granted in this way, as much as what further information either of us wants to convey, will guide the direction of our conversation—will determine what is said. I will not say things that are already taken for granted, since that would be redundant. Nor will I assert things incompatible with the common background, since that would be self-defeating. My aim in making assertions is to distinguish among the possible situations which are compatible with all the beliefs or assumptions that I assume that we share. Or it could be put the other way around: the common background is defined by the possible situations which I intend to distinguish among with my assertions, and other speech acts. Propositions true in all of them are propositions whose truth is taken for granted.

Although it is normally inappropriate because unnecessary for me to assert something that each of us assumes the other already believes, my assertions will of course always have consequences which are part of the common background. For example, in a context where we both know that my neighbor is an adult male, I say "My neighbor is a bachelor," which, let us suppose, entails that he is adult and male. I might just as well have said "my neighbor is unmarried." The same information would have been conveyed (although the nuances might not have been exactly the same). That is, the *increment of information*, or of content, conveyed by the first statement is the same as that conveyed by the second. If the asserted proposition were accepted, and added to the common background, the resulting situation would be the same as if the second assertion were accepted and added to the background.

This notion of common background belief is the first approximation to the notion of pragmatic presupposition that I want to use. A proposition *P* is a pragmatic presupposition of a speaker in a given context just in case the speaker assumes or believes that *P*, assumes or believes that his addressee assumes or believes that *P*, and assumes or believes that his addressee recognizes that he is making these assumptions, or has these beliefs.

I do not propose this as a definition or analysis, first since it is far from clear what it is to believe or assume something, in the relevant way and second since even assuming these notions to be clear, the definition would need further qualification. My aim is not to give an analysis but rather to

point to a familiar feature of linguistic contexts which, I shall argue, is the feature in terms of which a certain range of linguistic phenomena should be explained. The notion has, I think, enough intuitive content to enable us to identify a lot of particular cases, and the general outlines of the definition are clear enough to justify some generalizations about presuppositions which help to explain the facts. Before defending this claim by discussing some of the facts, I will make two remarks about the general notion.

First, note that it is persons rather than sentences, propositions or speech acts that have or make presuppositions. This goes against the prevailing technical use of the term, according to which presuppositions, whether semantic or pragmatic, are normally taken to relate two linguistic things. One might define such a relation in terms of the pragmatic notion in something like one of the following ways: (a) One might say that a sentence *x* presupposes that *Q* just in case the use of *x* to make a statement is appropriate (or normal, or conversationally acceptable) only in contexts where *Q* is presupposed by the speaker; or (b) one might say that the statement that *P* (made in a given context) presupposes that *Q* just in case one can reasonably infer that the speaker is presupposing that *Q* from the fact that the statement was made; or (c) one might say that the statement that *P* (made in a given context) presupposes that *Q* just in case it is necessary to assume that the speaker is presupposing that *Q* in order to understand or interpret correctly the statement. As stated, these suggested definitions are vague, and each is different from the others. But I do not think it would be fruitful to refine them, or to choose one over the others. It is true that the linguistic facts to be explained by a theory of presupposition are for the most part relations between linguistic items, or between a linguistic expression and a proposition. They are, as I interpret them, facts about the constraints, of one kind or another, imposed by what is said on what is appropriately presupposed by the speaker, according to various different standards of appropriateness. But I think all the facts can be stated and explained directly in terms of the underlying notion of speaker presupposition, and without introducing an intermediate notion of presupposition as a relation holding between sentences (or statements) and propositions.

This last point is a strategic recommendation, and not a substantive claim. As I said, one *could* define such a notion in various ways; I just doubt the theoretical utility of doing so. My purely strategic motive for emphasizing this point is that I want to avoid what I think would be a fruitless debate over which of various explications of the notion of pragmatic sentence presupposition best accords with the use of the term "presupposition" by linguists. I do not want to deny that, in an adequate

theory of conversation, one will need a notion or notions of conversational acceptability, and that once one has such a notion one has all the material for a definition of pragmatic sentence presupposition. A rough definition of "conversational acceptability" might be something like this: a speech act is conversationally acceptable in the relevant sense just in case it can reasonably be expected to accomplish its purpose in the normal way in which the normal purposes of such speech acts are accomplished. But such a notion would get its content from an account of the mechanisms by which the normal purposes of speech acts are accomplished, and the notion of speaker presupposition is intended to be one theoretical concept useful for giving such an account. It is in this way that it is a more basic concept than the concept of conversational acceptability.

Second, let me suggest one way that the definition given above needs to be qualified. In normal, straightforward serious conversational contexts where the overriding purpose of the conversation is to exchange information, or to conduct a rational argument, what is presupposed by the speaker, in the sense intended, is relatively unproblematic. The presuppositions coincide with the shared beliefs, or the presumed common knowledge. The difficulties in applying the notion come with contexts in which other interests besides communication are being served by the conversation. If one is talking for some other purpose than to exchange information, or if one must be polite, discreet, diplomatic, kind, or entertaining as well as informative, then one may have reason to act as if the common background were different than one in fact knows it to be. For example, when I talk to my barber, neither of us expects to learn anything; we are talking just to be civil, and to pass the time. If we haven't much to say, we may act as if the background of common knowledge is smaller than it really is. "Cold today, isn't it?" "Sure is, windy too." "Well, spring will be here before long." Although there is little actual communication going on here, it is clear that what is going on is to be understood in terms of genuine communication. We are pretending to communicate, and our pretense can be explained in terms of the same categories as a serious exchange of information.

In other cases, a speaker may act as if certain propositions are part of the common background when he knows that they are not. He may want to communicate a proposition indirectly, and do this by presupposing it in such a way that the auditor will be able to infer that it is presupposed. In such a case, a speaker tells his auditor something in part by pretending that his auditor already knows it. The pretense need not be an attempt at deception. It might be tacitly recognized by everyone concerned that this is what is going on, and recognized that everyone else recognizes it. In some cases, it is just that it would be indiscreet, or insulting, or tedious, or

unnecessarily blunt, or rhetorically less effective to assert openly a proposition that one wants to communicate.<sup>1</sup>

Where a conversation involves this kind of pretense, the speaker's presuppositions, in the sense of the term I shall use, will not fit the definition sketched above. That is why the definition is only an approximation. I shall say that one actually does make the presuppositions that one seems to make even when one is only pretending to have the beliefs that one normally has when one makes presuppositions. Presupposing is thus not a mental attitude like believing, but is rather a linguistic disposition—a disposition to behave in one's use of language as if one had certain beliefs, or were making certain assumptions.<sup>2</sup>

<sup>1</sup> This is a special case of what Grice has called *exploitation*, since the speaker exploits the rules governing normal conversation in order to communicate something which is not exactly said. See Grice (1989).

<sup>2</sup> It was suggested by Jerry Sadock (personal communication) that the definition should be modified in another way to account for examples of the following kind: I am asked by someone who I have just met, "Are you going to lunch?" I reply, "No, I've got to pick up my sister." Here I seem to *presuppose* that I have a sister, even though I do not assume that the addressee knows this. Yet the statement is clearly acceptable, and it does not seem right to explain this in terms of pretense, or exploitation. To meet this problem, Sadock suggests replacing the clause in the definition, "speaker assumes or believes that the addressee assumes or believes that *P*" with the clause, "speaker assumes or believes that the addressee *has no reason to doubt that P*."

The reason that I resist this suggestion, even though I recognize the force of the example, is that some basic generalizations about speaker presuppositions would fail if it were adopted. For example, one important generalization, alluded to above, is that it is unnecessary, in fact inappropriate, to assert what is presupposed. But consider a routine lecture or briefing by an acknowledged expert. It may be that everything he says is something that the audience has no reason to doubt, but this does not make it inappropriate for him to speak. The problem is that the modification would work only for cases where the addressee could infer what was being presupposed from the overt speech act. But this is not the only case where speaker presuppositions are important.

Two alternative responses to the example are possible: (a) one can explain it in terms of exploitation; (b) one can deny that there is a presupposition made at all in this kind of example.

To respond in the first way is, I admit, to stretch the notion of exploitation, first because the example lacks the flavor of innuendo or diplomatic indirection which characterizes the clearest cases of communication by pretense, and second because in the best cases of exploitation, it is the main point of the speech act to communicate what is only implied, whereas in this example, the indirectly communicated material is at best only a minor piece of required background information. Nevertheless, the explanation of how communication takes place in this example may be thought to be similar in form to explanations of how it takes place in the more familiar cases: the addressee infers that the speaker accepts that *Q* from the fact that he says that *P* because normally one says that *P* only when it is common background knowledge that *Q*.

To take the second option is to deny the generalization that the speaker *always* presupposes the existence of a unique referent (in the relevant domain of discourse) fitting any definite description (like "my sister") which he uses. To make this plausible, one would have to give an explanation of why one is *usually* expected to presuppose the existence of a unique referent when one uses a definite description—an explanation which also explains the exceptions to the rule.

The presumed background information—the set of presuppositions which in part defines a linguistic context—naturally imposes constraints on what can reasonably or appropriately be said in that context. Where the constraints relate to a particular kind of grammatical construction, or to a particular expression or category of expressions, one has a linguistic fact to be explained. This is the case with the sample sentences with which I began. One of the facts could be stated like this: it is inappropriate to say "The Queen of England is bald" (or to say "the Queen of England is not bald") except in a context in which it is part of the presumed background information that England has a queen. Compare this with a description that interprets the phenomena in terms of a semantic concept of presupposition: the proposition expressed by "the Queen of England is bald" has a truth-value only if England has a unique queen. The first description, in contrast to the second, makes no claim at all about the content of the statement—about the truth-conditions of what is said. The description in terms of the pragmatic notion does not rule out a semantic *explanation* for the fact that a certain presupposition is required when a certain statement is made, but neither does it demand such an explanation. That is, one *might* explain why it is appropriate for a speaker to say "the Queen of England is bald" only if he presupposes that England has a queen in terms of the following two assumptions: first, that the statement lacks a truth-value unless England has a queen, and second, that one normally presupposes that one's statements have a truth-value. But one also might explain the fact in a different way. The *facts* about presuppositions, I am suggesting, can be separated from a particular kind of semantic explanation of those facts. This separation of the account of presupposition from the account of the content of what is said will allow for more diversity among presupposition phenomena than would be possible if they all had to be forced into the semantic mold. Let me suggest, more specifically, four of the advantages of making this move.

First, if presupposition is defined independently of truth-conditions, then it is possible for the constraints on presuppositions to vary from context to context, or with changes in stress or shifts in word order, without those changes requiring variation in the semantic interpretation of what is said. This should make possible a simpler semantic theory; at the very least, it should allow for more flexibility in the construction of semantic theories. For example, D. T. Langendoen points out in a paper on presupposition and assertion that normally, if one said "my cousin isn't a boy anymore" he would be asserting that his cousin had grown up, presupposing that he is male. But one might, in a less common context, use the same sentence to assert that one's cousin had changed sexes, presupposing

that she is young.<sup>3</sup> If a semantic account of presupposition is given of this case, then one must say that the sentence is therefore ambiguous. On the pragmatic account, one just points to two different kinds of situations in which a univocal sentence could be used.

Second, if presupposition is defined independently of truth-conditions, then one can separate the question of entailment relations from the question of presupposition. On the semantic account, presupposition and entailment are parallel and incompatible semantic relations. *A* presupposes that *B* if and only if *B* is necessitated by *A* and its denial. *A* entails *B* if and only if *B* is necessitated by *A* but *not* by its denial. Thus the claim that the sentence, "Sam realizes that *P*" entails that *P* conflicts with the claim that that sentence presupposes, in the semantic sense, that *P*. But using the pragmatic account, one may say that sometimes when a presupposition is required by the making of a statement, what is presupposed is also entailed, and sometimes it is not. One can say that "Sam realizes that *P*" entails that *P*—the claim is false unless *P* is true. "Sam does not realize that *P*," however, does not entail that *P*. That proposition may be true even when *P* is false. All this is compatible with the claim that one is required to presuppose that *P* whenever one asserts or denies that Sam realizes it.

Third, the constraints imposed by a statement on what is presupposed seem to be a matter of degree, and this is hard to explain on the semantic account. Sometimes no sense at all can be made of a statement unless one assumes that the speaker is making a certain presupposition. In other cases, it is mildly suggested by a speech act that the speaker is taking a certain assumption for granted, but the suggestion is easily defeated by countervailing evidence. If a speaker says to me, "Sam was surprised that Nixon lost the election," then I have no choice but to assume that he takes it for granted that Nixon lost. But if he says, "If Eagleton hadn't been dropped from the Democratic ticket, Nixon would have won the election" (without an "even" before the "if" or a "still" after the "Nixon"), there is a suggestion that the speaker presupposes that Nixon in fact did not win, but if the statement is made in the right context, or with the right intonation, the suggestion is overruled. This difference in degree, and variation with context is to be expected on the pragmatic account, since it is a matter of the strength of an inductive inference from the fact that a statement was made to the existence of a background assumption or belief.

Fourth, and perhaps most important, the pragmatic analysis of presupposition, because it relates the phenomena to the general communication situation, may make it possible to explain some of the facts in terms of

general assumptions about rational strategy in situations where people exchange information or conduct argument. One way to explain the fact that a particular assertion requires or suggests a certain presupposition is to hypothesize that it is simply a fact about some word or construction used in making the assertion. In such a case, the fact about the presupposition requirement must be written into the dictionary, or into the semantics. But since we have an account of the *function* of presuppositions in conversations, we may sometimes be able to explain facts about them without such hypotheses. The propositions that *P* and that *Q* may be related to each other, and to common beliefs and intentions, in such a way that it is hard to think of a reason that anyone would raise the question whether *P*, or care about its answer, unless he already believed that *Q*. More generally, it might be that one can make sense of a conversation as a sequence of rational actions only on the assumption that the speaker and his audience share certain presuppositions. If this kind of explanation can be given for the fact that a certain statement tends to require a certain presupposition, then there will be no need to complicate the semantics or the lexicon.

For example, consider the word "know." It is clear that "x knows that *P*" entails that *P*. It is also clear that in most cases when anyone asserts or denies that x knows that *P*, he presupposes that *P*. Can this latter fact be explained without building it into the semantics of the word? I think it can. Suppose a speaker were to assert that x knows that *P* in a context where the truth of *P* is in doubt or dispute. He would be saying in one breath something that could be challenged in two different ways. He would be leaving unclear whether his main point was to make a claim about the truth of *P*, or to make a claim about the epistemic situation of x (the knower), and thus leaving unclear what direction he intended or expected the conversation to take. Thus, given what "x knows that *P*" means, and given that people normally want to communicate in an orderly way, and normally have some purpose in mind, it would be unreasonable to assert that x knows that *P* in such a context. One could communicate more efficiently by saying something else. For similar reasons, it would normally be inappropriate to say that x does not know that *P* in a context where the truth of *P* was in question. If the speaker's reason for believing his assertion were that he thought that *P* was false, or that he thought that x didn't believe that *P*, or didn't have reason to believe that *P*, then his statement would be gratuitously weak. And it would be unusual for a speaker to be in a position to know that one of these situations obtained, without knowing which.

This is a tentative and incomplete sketch of an explanation. Much more would have to be said to make it convincing. My point is to make it plausible that, in some cases at least, such explanations might be given, and to

<sup>3</sup> Langendoen (1971).

argue that where they can be given, there is no reason to build specific rules about presuppositions into the semantics.

I want now to illustrate these advantages of the pragmatic account by looking at some linguistic facts in terms of it. The two sets of facts I will consider are taken from two recent papers by Lauri Karttunen.<sup>4</sup>

First, on a distinction between two kinds of factive verbs. It is well known that among verbs that take a nominalized sentence complement (for example *believe*, *know*, *intend*, *see*) one can distinguish a subclass known as factive verbs (*know*, *regret*, *discover*, *see*, as contrasted with *believe*, *intend*, *assert*, *claim*). A number of syntactic and semantic criteria are used to draw the distinction, but the distinguishing mark that is relevant here is the following: if *V* is a factive verb, then *x V's that P* presupposes (and, I would say, entails as well) that *P*. If I assert or deny that Jones regrets, realizes, or discovers that Nixon won the election, then I presuppose that Nixon did in fact win. Karttunen has drawn a further distinction among two kinds of factive verbs which, he argues, requires a distinction between two kinds of presupposition relations. One kind of factive verb (labeled the *full factives*) includes *regret*, *forget* and *resent*. The basis for the distinction is as follows: with full factives, it is not only an assertion or denial of the proposition *x V's that P* that requires the presupposition that *P*, but also the *supposition* that *x V's that P* in the antecedent of a conditional, or the claim that the proposition *might* be true. With semi-factives, it is only the assertion or denial that require the presupposition. For example, consider the two statements

Sam may regret that he voted for Nixon.

If Sam regrets that he voted for Nixon, then he is a fool.

Because these two statements clearly require the presupposition that Sam voted for Nixon, *regret* is seen to be a full factive.

The following is Karttunen's example to illustrate the contrast between full factives and semi-factives. Compare

If I { regret  
realize  
discover } later that I have not told the truth,  
I will confess it to everyone.

In the first statement, the speaker clearly presupposes that he has not told the truth. In the other two cases, he clearly does not presuppose this. Thus *realize* and *discover* are seen to be semi-factives.

To explain the difference, Karttunen postulates a distinction between a strong and a weak kind of semantic presupposition. If *P* is necessitated by

<sup>4</sup> Karttunen (1971, 1973).

*Possibly Q*, and by *Possibly not-Q*, then *Q* strongly presupposes that *P*. Weak semantic presuppositions are defined in the usual way.

In discussing this example, I want to dispute both the data, and the theoretical account of them. I agree that there is a sharp contrast in the particular example given, but the matter is less clear if one looks at other examples. Consider:

If Harry discovers that his wife is playing around, he will be upset.

If Harry had discovered that his wife was playing around, he would have been upset.

If Harry had realized that his wife was playing around, he would have been upset.

Harry may realize that his wife has been playing around.

Harry may never discover that his wife has been playing around.

There is, I think, in all these cases a presumption that the speaker presupposes that Harry's wife is, or has been, playing around. The presumption is stronger in some of the examples than in others, but it seems to me that in some of them it is as strong as with *regret*. Further, if we assume that with the so-called semi-factives like *discover* and *realize*, there is *always* a presumption that the speaker presupposes the truth of the proposition expressed in the complement, we can still explain why the presumption is defeated in Karttunen's particular example. The explanation goes like this: if a speaker explicitly supposes something, he thereby indicates that he is not presupposing it, or taking it for granted. So when the speaker says "if I realize later that *P*," he indicates that he is not presupposing that he will realize later that *P*. But if it is an open question for a speaker whether or not he will at some future time have come to realize that *P*, he can't be assuming that he already knows that *P*. And if he is not assuming that he himself knows that *P*, he can't be assuming that *P*. Hence *P* cannot be presupposed. A roughly parallel explanation will work for *discover*, but not for *regret*.

One can explain another of Karttunen's examples in a similar way. Consider the three questions:

Did you { regret  
realize  
discover } that you had not told the truth?

Here *realize* seems to go with *regret* and not with *discover*. The first two questions seem to require that the speaker presuppose that the auditor did not tell the truth, while the third does not. Again, we can explain the difference, even while assuming that there is a presumption that the presupposition is made in all three cases. The reason that the presumption is

defeated in the third case is that the speaker could not make that presupposition without assuming an affirmative answer to the question he is asking. But in general, by asking a question, one indicates that one is not presupposing a particular answer to it. This explanation depends on the particular semantic properties of *discover*, and will not work for *realize* or *regret*.<sup>5</sup> It also depends on the fact that the subject of the verb is the second-person pronoun. Hence if the explanation is right, one would expect the presupposition to reappear in the analogous third-person question, "Did Sam discover that he hadn't told the truth?" It seems that it does.

Since on the pragmatic account, the constraints on presuppositions can vary without the truth-conditions changing, we can allow presupposition differences between first- or second-person statements and questions and the corresponding third-person statements and questions without postulating separate semantic accounts of propositions expressed from different points of view. So, while we have noted differences in the presuppositions required or suggested by the following two statements,

If Harry discovers that his wife has been playing around, he will be upset.

If I discover that my wife has been playing around, I will be upset (said by Harry).

This difference does not prevent us from saying that the two statements both have the same semantic content—that the same proposition is expressed in both cases. It would not be possible to say this on a semantic account of presupposition.

If the explanations I have sketched are on the right track, then we can account for at least some of the differences between factive and semi-factive verbs without distinguishing between two different kinds of presupposition relations. We can also account for some differences among semi-factives, and differences between first- and third-person statements without complicating the semantics. The explanation depends on just two things: first, some simple and very general facts about the relation between pragmatic presuppositions and assertions, questions, and suppositions;

<sup>5</sup> The relevant difference between *realize* and *discover* is this: because *realize* is a stative verb, a past tense statement of the form *x didn't realize that P* must be about some particular time in the past (determined by the context), and not about *all* times in the past. This means that *x didn't realize that P* may be true, even though *x now* knows that *P*. In contrast, because *discover* is an inchoative verb, *x didn't discover that P* may be about *all* times in the past. For this reason, normally, *x didn't discover that P* implies that *x has not yet discovered that P*, and so does not know that *P*. Therefore, if a speaker presupposes that *P*, he assumes that *x has* discovered that *P*, and so assumes a particular answer to the question he is asking.

second, on the ordinary semantic properties of the particular verbs involved.<sup>6</sup>

The second set of facts that I will discuss concerns the presuppositions of compound sentences: How do the presuppositions required by a conditional or conjunctive statement relate to the presuppositions that would be required by the component parts, stated alone? In general, what is the relation between the presuppositions required by an assertion that *A* and the assertion that *B* on the one hand, and by an assertion that *A and B* or that *if A, then B* on the other? Karttunen defends the following answer to the question: let *S* be a sentence of the form *A and B* or *If A, then B*. *S* presupposes that *C* if and only if either *A* presupposes that *C*, or *B* presupposes that *C* and *A* does not semantically entail that *C*. In other words, the presuppositions of a conjunction are the presuppositions required by either of the conjuncts, *minus any required by the second conjunct which are entailed by the first*. The presuppositions of a conditional are the presuppositions of either antecedent or consequent minus those required by the consequent and entailed by the antecedent. So if I say "Harry is married, and Harry's wife is a great cook," I assert, and do not presuppose, that Harry is married. But the second conjunct, stated alone (*Harry's wife is a great cook*), would require the presupposition that Harry is married. The sentence with conjuncts in reverse order would be unacceptable in any normal context (*Harry's wife is a great cook, and Harry is married*).

Now if we regard Karttunen's generalization as a generalization about *semantic* presuppositions, then we will interpret it as a hypothesis about the way the truth-value (or lack of it) of a conjunction or conditional relates to the truth-values of the parts. The hypothesis has the consequence that the conjunction *and* is not truth-functional, since the truth-value of a conjunctive statement will in some cases depend on entailment relation between the conjuncts. It has the consequence that *and* is not symmetric. *A and B* may be false while *B and A* lacks a truth-value. Finally, it has the consequence that the simple conjunction *and* is governed by mysteriously complicated rules.

<sup>6</sup> Two disclaimers: First, I do not want to leave the impression that I think I have explained very much here. I have not made any attempt to explain the source of the presumption that the complements of both factive and semi-factive verbs are presupposed. I have tried to explain only how the presupposition is cancelled in certain cases. Also, the presumption is clearly harder to defeat in some cases than in others: harder with *realize* than with *discover*, and harder with full factives than with semi-factives. I have said nothing that would explain this. My hope, however, is that such explanations can be given using the general strategy which I am recommending. Second, I do not want to deny that there are systematic differences between factives and semi-factives. One difference is that full factives all require not only the presupposition that the proposition expressed in the complement is true, but also the presupposition that the subject of the verb knows or knew that it is. None of the semi-factives requires or suggests this second presupposition: in fact, they rule it out.



On the other hand, if we regard Karttunen's generalization as a generalization about *pragmatic* presuppositions, then we can reconcile it with the standard truth-functional account of *and*, and we can explain the generalization without postulating any *ad hoc* semantic or pragmatic rules. The explanation goes like this: first, once a proposition has been asserted in a conversation, then (unless or until it is challenged) the speaker can reasonably take it for granted for the rest of the conversation. In particular, when a speaker says something of the form *A and B*, he may take it for granted that *A* (or at least that his audience recognizes that *he* accepts that *A*) after he has said it. The proposition that *A* will be added to the background of common assumptions before the speaker asserts that *B*. Now suppose that *B* expresses a proposition that would, for some reason, be inappropriate to assert except in a context where *A*, or something entailed by *A*, is presupposed. Even if *A* is *not* presupposed initially, one may still assert *A and B* since by the time one gets to saying that *B*, the context has shifted, and it is by then presupposed that *A*.<sup>7</sup>

As with the explanation sketched in the earlier discussion, this explanation rests on just two things: first, a simple pragmatic assumption about the way presuppositions shift in the course of a conversation—an assumption that says, roughly, that a speaker may build on what has already been said; second, an uncontroversial assumption about the semantic properties of the word *and*—in particular, that when one asserts a conjunction, he asserts both conjuncts. If we interpret presupposition to mean *pragmatic* presupposition, then we can deduce Karttunen's generalization from these two almost trivial assumptions.

The analogous generalization about conditional statements is explainable on equally simple assumptions. Here we need first the assumption that what is explicitly *supposed* becomes (temporarily) a part of the background of common assumptions in subsequent conversation, and second that an *if* clause is an explicit supposition. Again, Karttunen's generalization can be derived from these obvious assumptions.

I have been arguing in this paper for the fruitfulness of separating semantic from pragmatic features of linguistic expressions and situations, and of explaining a certain range of phenomena in terms of pragmatic rather than semantic principles. This goes against the trend of the work of generative semanticists such as George Lakoff and John Ross, who have emphasized the difficulty of separating syntactic, semantic, and pragmatic problems, and who have sometimes suggested that such distinctions as

<sup>7</sup> In a paper at the Texas conference on performatives, conversational implicature and presuppositions, Karttunen put forward an explanation of his generalization which is very similar to this. Our accounts were developed independently.

between syntactic and semantic deviance or semantic and pragmatic regularities are of more use for avoiding problems than for solving them. Partly to respond to this concern, I will conclude with some general remarks about the distinction between semantics and pragmatics, and about what I am *not* recommending when I suggest that the distinction be taken seriously.

First remark: semantics, as contrasted with pragmatics, can mean either the study of *meaning* or the study of *content*. The contrast between semantic and pragmatic claims can be either of two things, depending on which notion of semantics one has in mind. First, it can be a contrast between claims about the particular conventional meaning of some word or phrase on the one hand, and claims about the general structure or strategy of conversation on the other. Grice's distinction between conventional implicatures and conversational implicatures is an instance of this contrast. Second, it can be a contrast between claims about the truth-conditions or *content* of what is said—the proposition expressed—on the one hand, and claims about the *context* in which a statement is made—the attitudes and interests of speaker and audience—on the other. It is the second contrast that I am using when I argue for a pragmatic rather than a semantic account of presuppositions. That is, my claim is that constraints on presuppositions are constraints on the contexts in which statements can be made, and not constraints on the truth-conditions of propositions expressed in making the statements. I also made use of the other contrast in arguing for this claim. I conjectured that one can explain many presupposition constraints in terms of general conversational rules without building anything about presuppositions into the meanings of particular words or constructions. But I make no general claim here. In some cases, one may just have to write presupposition constraints into the dictionary entry for a particular word. This would make certain presupposition requirements a matter of *meaning*, but it would not thereby make them a matter of *content*. There may be facts about the meaning of a word which play no role at all in determining the truth-conditions of propositions expressed using the word.

Second remark: in recommending a separation of content and context I am not suggesting that there is no interaction between them. Far from it. The semantic rules which determine the content of a sentence may do so only relative to the context in which it is uttered. This is obviously the case with sentences using personal pronouns, demonstratives, quantifiers, definite descriptions, or proper names. I suspect it happens in less obvious cases as well. But this interaction does not prevent us from studying the features which define a linguistic context (such as a set of pragmatic presuppositions) in abstraction from the propositions expressed in such

contexts, or from studying the relationships among propositions in abstraction from the contexts in which they might be expressed.

A final remark: in some cases, distinctions such as that between semantic and pragmatic features may be used as a way to set problems aside. Some linguists have accused other linguists of using the distinction between syntax and semantics in this way. Deviant sentences which seem to conflict with syntactic generalizations are not treated as counter-examples, but instead are thrown into a "semantic wastebasket" to be explained away by some future semantic theory. In the same way, some may be suspicious that I am setting up a pragmatic wastebasket, and recommending that all the interesting problems be thrown away.

I do not think that this is always a bad procedure, but it is not what I am suggesting here. I am recommending instead the development and application of a pragmatic theory in which detailed explanations of phenomena relating to linguistic contexts can be given. It is true that traditionally the more well-developed and the more rigorous linguistic theories have focussed on questions of grammar and content, while the discussions which emphasized the role of conversational context have been more informal and less theoretical. But there is no necessity in this. Potentially at least, a theory of pragmatics, and the notion of pragmatic presupposition can be as precise as any of the concepts in syntax and semantics. Although the explanations I have sketched in this paper are informal and incomplete, I think they suggest a strategy for giving explanations of linguistic phenomena relating to contexts which are both rigorous and intuitively natural.<sup>8</sup>

<sup>8</sup> I have been accused, partly on the basis of this concluding paragraph, of being overly optimistic about the possibility of a formal theory of pragmatics which is both rigorous and sufficiently detailed to provide substantive explanations of linguistic phenomena. This accusation may be just, but my main point here is independent of this. However easy or difficult it proves to be to develop an adequate theory of conversation, one cannot simplify the task by building conversational rules into a semantic theory of the content of what is said.

### 3

## Indicative Conditionals

"Either the butler or the gardener did it. Therefore, if the butler didn't do it, the gardener did." This piece of reasoning—call it the *direct argument*—may seem tedious, but it is surely compelling. Yet if it is a valid inference, then the indicative conditional conclusion must be logically equivalent to the truth-functional material conditional,<sup>1</sup> and *this* conclusion has consequences that are notoriously paradoxical. The problem is that if one accepts the validity of the intuitively reasonable direct argument from the material conditional to the ordinary indicative conditional, then one must accept as well the validity of many arguments that are intuitively absurd. Consider, for example, "the butler did it; therefore, if he didn't, the gardener did." The premise of this argument entails the premise of the direct argument, and their conclusions are the same. Therefore, if the direct argument is valid, so is this one. But this argument has no trace of intuitive plausibility. Or consider what may be inferred from the *denial* of a conditional. Surely I may deny that if the butler didn't do it, the gardener did without affirming the butler's innocence. Yet if the conditional is material, its negation entails the truth of its antecedent. It is easy to multiply paradoxes of the material conditional in this way—paradoxes that must be explained away by anyone who wants to defend the thesis that the direct argument is valid. Yet anyone who denies the validity of that argument must explain how an invalid argument can be as compelling as this one seems to be.

There are thus two strategies that one may adopt to respond to this puzzle: defend the material conditional analysis and explain away the paradoxes of material implication, or reject the material conditional analysis and explain away the force of the direct argument.<sup>2</sup> H. P. Grice, in his

The ideas in this paper were developed over a number of years. During part of this time my research was supported by the National Science Foundation, grant number GS-2574; more recently it was supported by the John Simon Guggenheim Memorial Foundation.

<sup>1</sup> The argument in the opposite direction—from the indicative conditional to the material conditional—is uncontroversially valid.

<sup>2</sup> This does not exhaust the options. Three other possible strategies might be mentioned. (1) Defend the direct argument, not by accepting the truth-functional analysis of the conditional, but by rejecting the truth-functional analysis of the disjunction. (2) Give a