

Introduction

This book is all about one simple idea: that speakers convey information not only by what they say, but also by what they *don't* say. In fact, the idea is so obvious that we may never know who had it first. What is known is that around the mid-19th century, John Stuart Mill thought it obvious enough to be mentioned almost in passing:

If I say to any one, “I saw some of your children to-day”, he might be justified in inferring that I did not see them all, not because the words mean it, but because, if I had seen them all, it is most likely that I should have said so: though even this cannot be presumed unless it is presupposed that I must have known whether the children I saw were all or not. (Mill 1865: 442, also cited by Horn 2009)

One century later, it was H. Paul Grice who saw that this simple idea contains the gist of a general framework for pragmatics, based on the premiss that discourse is a joint project undertaken by speakers who expect each other to be cooperative. It is this expectation, according to Grice, which gives rise to the pragmatic inferences he calls “conversational implicatures”. The inference Mill refers to is such an implicature; more specifically, it is a quantity implicature.

If it is so simple, why does it take a book to explain? There are several reasons. To begin with, it turns out that quantity implicatures are somewhat more complex than meets the eye (Chapter 2), but that is not the main reason. More important is that the extensive literature on quantity implicatures contains a number of distractors that we will have to clear out of the way. One is that the bulk of that literature has focused its attention on one species of quantity implicatures, i.e. scalar implicatures, which has led to a distorted view on the genus (Chapter 3). Mill's example is a case in point. If a speaker says,

(1) I saw some of your children today.

we can generate an alternative statement he could have made simply by supplanting “some” in (1) with a stronger expression of about the same complexity:

(2) I saw all of your children today.

Then the implicature is that, for all the speaker knows, this alternative statement doesn’t hold. There has been a tendency in the literature to treat all quantity implicatures on this model, but the trouble is that it isn’t general enough. For instance, it already fails in simple cases like the following:

(3) A: Who is coming to the meeting?
 B: Albinoni, Boccherini, and Corelli.

B’s answer may well implicate that apart from Albinoni, Boccherini, and Corelli, nobody else is coming to the meeting, and if this is inferred it must be a quantity implicature. However, it is evident that this inference cannot be obtained by replacing an expression in (3B) by one that is stronger and about equally complex, and then denying the result.

What is needed, I will argue, is a shift of perspective that may seem suspiciously subtle at first, but is entirely in the Gricean spirit and a prerequisite for a unified theory of quantity implicatures (Chapter 6). Instead of asking why the speaker didn’t say (2), the question should be whether the speaker might believe that (2) is true. In this case, the answer is likely to be “no”, for if the speaker believed that (2) is true, he would have uttered this sentence rather than (1), and thus we arrive at the conclusion that the speaker does not believe (2). I call this the “intention-based” approach. Unlike the standard scheme, it carries over to examples like (3) without a hitch. For A may now reason as follows: “Could it be that, according to B, more than three people are coming to meeting? Presumably not, because then B would have said so. Hence, I’m entitled to conclude that, for all B knows, only Albinoni, Boccherini, and Corelli are coming to the meeting.” Furthermore, this line of explanation extends to free choice inferences, like the following:

(4) Julius may be in Amsterdam or Berlin.
 ~ Julius may be in Amsterdam.
 ~ Julius may be in Berlin.

Despite the fact that they are obviously licit in some sense, these inferences have proved hard to account for. As I will argue in Chapter 6, one of the main attractions of the intention-based approach is that it allows us to explain them as run-of-the-mill quantity implicatures.

Another distracting factor we will have to engage with is related to the first. It is the notion that some quantity implicatures hold by default (Chapter 5). Again, scalar implicatures are the prime examples. It is tempting to suppose that an utterance of (1) will *normally* give rise to the inference that, for all the speaker knows, (2) isn't true. Apparently, this intuition is so compelling that many scholars have accepted it at face value. Nonetheless, I will argue that it is mistaken, and has done more than its share to keep implicature theorists off the right track.

Perhaps the main challenge for a Gricean theory of quantity implicatures, and a further reason why it takes more than thirty pages to develop, is that every once in a while it appears as if a quantity implicature occurs within the scope of an attitude verb, for example:

(5) Cleo believes that some of her shoes have been stolen.

This may be understood as implying that Cleo believes that not all her shoes were stolen, and this interpretation seems to require that a “not all” inference is drawn within the scope of “believes”. If that is the case, the inference in question cannot be a quantity implicature, because quantity implicatures, like all conversational implicatures, are derived on the basis of a speech act made by the speaker, and speech acts are made by way of complete sentences. Hence, there is no way a bona fide implicature could be derived within the scope of any expression.

This is commonly known as the problem of “embedded implicatures” (which is strictly speaking a contradiction in terms, since by their very nature implicatures cannot be embedded). Embedded implicatures pose a problem for Gricean pragmatics in two ways. On the one hand, they have prompted a rash of proposals for analysing quantity implicatures, and especially scalar implicatures, with non-pragmatic methods. For instance, it has been suggested that “some” actually *means* “some but not all”, and if this were true, examples like (5) wouldn't be much of a problem. There is quite a variety of such conventionalist proposals, which are to be discussed and criticised in Chapter 7. On the other hand, even if this type of approach is on the wrong track, that doesn't give us a pragmatic solution to the problem of embedded implicatures. Chapter 8 takes up the challenge by arguing that there isn't one big problem but rather several smaller ones, most of which can be solved by showing that there are independent factors that interact with standard Gricean reasoning in such a way that embedded implicatures *seem* to arise. Put otherwise, I will argue that, by and large, embedded implicatures are a mirage. Still, there are exceptional cases in which, e.g., “some” does mean “some but not all”, and

these cannot be construed as conversational implicatures. However, I will try to show that they, too, can be accommodated within the framework of Gricean pragmatics.

In the meantime, it will have become evident that this is an essay in pure pragmatics; Gricean pragmatics, to be exact. However, although the theory to be presented in the following pages is squarely in the spirit of Grice's philosophy, it does not slavishly adhere to the letter of his writings. There are three points, in particular, in which I will deviate from the official Gricean party line. To begin with, whereas Grice usually (though by no means always) studied discourse from the speaker's perspective, I will resolutely adopt the hearer's point of view: my main interest lies in the interpretation of linguistic utterances. As a consequence, I will use some of the standard Gricean terminology in a way that is strictly speaking deviant (though it may be noted that this sort of terminological abuse is rife in the literature, even if it usually goes unacknowledged). In particular, this holds for the term "implicature", which Grice uses for certain commitments the speaker incurs by virtue of his utterances, whereas I will treat implicatures as hearers' inferences about such commitments. This is merely a difference in perspective, not opinion.

The second point is related to the first. In my view, a theory of interpretation—*any* theory of interpretation—should be of a piece with a psychological theory of language processing. This is not to say that a theory of interpretation should *be* a processing theory, but it is to say that pragmatics and psychology should mesh and that experimental data can and must be brought to bear on theories of interpretation. In this respect, too, I deviate from Grice, who (at least in his writings) never showed much of an interest in experimental psychology.

Thirdly, and most importantly, there is the much-debated issue of "what is said": Grice's term for what is also known as the "proposition", "literal meaning", or "truth-conditional content" carried by an utterance. Since conversational implicatures are calculated on the basis of what is said, it is of considerable importance to establish what *is* said by any given utterance and it is especially important to determine if and how pragmatic factors may affect truth-conditional content. It is widely agreed that Grice's own views on the matter were not entirely satisfactory, and I will go with the general consensus in this point, and argue, moreover, that truth-conditional content is dependent on pragmatic reasoning to a much larger extent than Grice was willing to allow for. But nevertheless, I will also argue, against the consensus, that this amendment is entirely consistent with the spirit of Grice's general programme.

Gricean pragmatics

1

The concept of conversational implicature was developed by Grice in his groundbreaking work of the 1960s, which culminated in the William James Lectures he delivered at Harvard University in 1967. These lectures had a formidable impact from the very start, but were slow to appear in print, the key instalments for our topic being “Logic and conversation” (1975) and “Further notes on logic and conversation” (1978). It wasn’t until 1989 that the William James Lectures were published collectively, as part of a selection of Grice’s papers, “Studies in the way of words”.

The theory of quantity implicatures that will unfold in the following chapters is thoroughly Gricean in spirit, although occasional departures from the letter of Grice’s writings will be hard to avoid. In this chapter, I will survey Grice’s views on conversational implicatures, and mark the chief points at which I deviate from his opinion. The discussion will be confined to matters that are relevant to the purposes of this book, which is to say that important parts of Grice’s work will be treated lightly, if that. For a more comprehensive introduction to the Gricean philosophy of language, I refer to Neale’s (1992) excellent review article on “Studies in the way of words”.

1.1 Saying vs. implicating

Consider the following sentences:

- (1) a. Harry is rich but dull.
- b. Harry is rich and dull.

Despite the undeniable fact that there is considerable semantic overlap between these sentences, it would seem that (1a) has a richer meaning than (1b). Both sentences entail that Harry is rich as well as dull, but while this is all that

is conveyed by (1b), (1a) conveys a bit more, viz. that one way or another there is a contrast between being rich and being dull.¹ Hence, we are inclined to say, pre-theoretically, that (1a) and (1b) have different meanings. More generally, in everyday practice we are liable to conclude that two sentences have different meanings whenever we observe that they don't communicate exactly the same message. One of the the most important insights to have emerged from the modern philosophy of language is that this practice is mistaken: our pre-theoretical notion of meaning covers an amalgam of very different species of content, which must be differentiated phenomenologically (they don't behave alike) as well as etiologically (they have different causes). Consequently, one of the main concerns of theories of interpretation is to isolate species of content and study their properties as well as the ways in which they interact. Small wonder, therefore, that in semantics and pragmatics demarcation issues are the order of the day. One theory's entailment is another's presupposition or implicature.

Consider again the contrast between (1a) and (1b). According to Frege (1892, 1918) both sentences have the same meaning; they express the same proposition. Meaning, in Frege's book, is what determines whether or not a sentence is true, in a given context, and this is the same in either case: Frege considers both sentences true if Harry is rich and dull; otherwise they are both false. Frege concedes that (1a) "intimates" a contrast between the property of being rich and that of being dull, but he denies that this is part of the sentence's truth conditions. Therefore, Frege would say that (2b), while surely infelicitous, does not have contradictory truth conditions, as (2a) does:

- (2) a. *Harry is rich but dull, but he isn't rich.
- b. ?Harry is rich but dull, though I wouldn't want to suggest that there is a contrast between these two properties.

If the contrast between (2a) and (2b) seems a bit fragile, there are other ways of bringing out the difference between truth-conditional and intimated content. Consider, for example, the following pair of discourses:

- (3) a. Harry is rich and dull, and that's a good thing.
- b. Harry is rich but dull, and that's a good thing.

In both cases, the intended referent of "that" is the fact that Harry is rich *and* dull: the second conjunct of (3b) does not imply that it is a good thing that there is contrast between being rich and being dull. Apparently, the demon-

1. This is a bit of a simplification, as pointed out by Dummett (1973), for example.

1.1 Saying vs. implicating

strative pronoun ignores the intimated content of the first conjunct, which supports the view that truth-conditional and intimated content are distinct. The following example makes the same point in a different way:

- (4) a. Most of these girls are rich and dull.
- b. Most of these girls are rich but dull.

(4a) says that for most x in a given set of girls G it is the case that (i) x is rich and (ii) x is dull. That is, the entire content of “rich and dull” is attributed to each of a majority in G . By contrast, (4b) does not say that for most x in G it is the case that (i) x is rich, (ii) x is dull, and (iii) there is a contrast between x being rich and x being dull. Rather, what this sentence expresses is that there is a contrast between being rich and being dull and that most of the girls are rich and dull. Hence, the intimated content associated with “but” does not behave as ordinary truth-conditional content would: it seems oblivious to the fact that it is within the scope of a quantifier.

Frege’s distinction between truth-conditional and intimated content foreshadows Grice’s distinction between “what is said” and “conventional implicature” (Horn 2007). What is said is truth-conditional content; conventional implicatures have no bearing on the truth value of a sentence. However, in another way these two notions are closely related, for in Grice’s view both are conventional in nature—or very nearly so. The contrasts in (1)-(4) are due to the fact that “but” comes with an implicature that “and” lacks, and since it is part of the conventions of English that “but” is used this way, Grice calls it a “conventional implicature”. To a first approximation, at least, what is said is conventional in much the same way: linguistic conventions assign truth-conditional meanings to words and grammatical constructions, and the meaning of a complex expression is determined by the meanings of its parts and the grammatical constructions they enter into; thus the truth-conditional meaning of a sentence, “what it says”, is conventional. Grice is aware of the fact that this picture is somewhat simplified, because truth-conditional meaning is affected by various contextual factors, and therefore not *purely* conventional. This context dependence raises some rather deep issues, which we will dwell on later in this chapter (§1.6). In the meantime, let’s ignore these complications and agree that what is said is primarily determined by linguistic conventions.

... *A digression on truth*

Like Frege, Grice considers conventional implicatures to be truth-conditionally inert. (In fact, “implicature” is Grice’s cover term for non-truth-conditional

content.) Considerable quantities of ink have been spilled on the question of whether or not this is right. To illustrate what the issue is about, suppose I utter the following:

(5) ?Harry is British but he speaks English.

In view of the well-known fact that English is spoken by the overwhelming majority of British citizens, there would not normally be a contrast between being British and speaking English, and therefore my utterance is odd. Now, the mooted question is this: When I said (5), did I speak falsely or was my utterance merely defective in some non-alethic sense of the word? In my opinion, for all the attention it has received, this question is not particularly interesting, for the following reason. Nobody will deny that my utterance of (5) *carried the information* that there is a contrast, and that this information is false because it fails to agree with the facts. I take it that this much is uncontroversial. It is also agreed upon, I believe, that a lack of contrast renders my utterance defective. The question is just whether its defect should be called “falsity” or something else, and that is a terminological issue.

Whether or not it is merely about terminology, it should be noted that speakers’ intuitions will not help to resolve the issue. If a professional philosopher were to use words like “true” and “false” as loosely as they are used in real life, his job would be on the line. The predicate “true” is generally used for expressing agreement, and not merely to indicate that a claim is factually correct. Conversely, calling a claim “not true” is a perfectly acceptable way of saying you disagree, even if there can hardly be a fact of the matter. Here are some examples from the internet:

- (6) a. Somebody was saying I am a funny and sweet girl. That’s true.
b. The churros with chocolate were lovely, it’s true.
c. It is easier to bottle-feed than to breast-feed. Not true!

In all likelihood, it is a matter of taste whether the speaker of (6a) is a funny and sweet girl, whether the churros were lovely, and whether bottle-feeding is easier than breast-feeding. Nonetheless the predicate “true” is freely used in these cases for expressing authors’ agreement or lack thereof. The upshot of these observations is that “truth-conditional content” is a *theoretical* notion which need not always align with speakers’ intuitions about truth.

The question of whether conventional implicatures have any bearing on truth is of much less interest than some scholars seem to believe. However, I suspect that one of the reasons why this question has been discussed so

1.1 Saying vs. implicating

9

zealously is that it has been mixed up with another one, viz. whether or not conventional implicatures should be distinguished from other types of content, specifically from what is said. That question *is* important, and can be answered in the affirmative, e.g., on the grounds that conventional implicatures interact in their own special way with pronouns and quantifiers, as we saw in our discussion of examples (3) and (4). *End of digression*

Whereas what is said and conventional implicatures are both conventional (or largely so), conversational implicatures are something different altogether (and it is unfortunate that Grice chose to use the same term, “implicature”, for two notions that are so far apart). Conversational implicatures are, first and foremost, *non-conventional*; they are not due to linguistic conventions of any kind. To explain how such inferences might arise, consider the following example. Grice (1975/1989: 32) invites us to imagine a person A “standing by an obviously immobilized car”,² addressing B as follows:

- (7) A: I am out of petrol.
 B: There is a garage round the corner.

Grice observes that B’s utterance allows A to infer that, as far as B knows, the garage in question is open. It will be clear that this is neither said nor conventionally implicated: the conventions of the English language aren’t such that (7B) encodes in any way that the garage is open. I take it that this much is obvious, but to underscore this point, note that B could have said (8) without fear of delivering an infelicitous statement:

- (8) There is a garage round the corner, but I’m afraid it is closed.

If it was part of the conventional content of (7B) that the garage is open, (8) should be infelicitous, which it isn’t. How, then, does the inference come about? Grice’s answer is really very simple: A assumes that B is trying to be helpful, and reasons that B’s utterance wouldn’t be helpful if he knew or suspected that the garage wasn’t open. Hence, A concludes, for all B knows the garage is open. Put otherwise, supposing that B is trying to be cooperative, it wouldn’t make sense for him to utter (7B) unless he believed the garage was open, and therefore he probably believes that it is open.

The inference that for all B knows, the garage is open, Grice calls a “conversational implicature”. It arises on the assumption that the speaker is doing

2. I’ve always been puzzled by Grice’s suggestion that one can *see* that a car with an empty tank is “obviously immobilized”.

his best to make a useful contribution to the discourse. Like their conventional namesakes, conversational implicatures are truth-conditionally inert, but apart from that the two types of inference couldn't be more different, the most important difference being that conversational implicatures aren't encoded in the utterances that give rise to them. We will have much more to say about this, but first let us consider how Grice unpacks the notion that utterances can be more or less useful.

1.2 Discourse and cooperation

Grice sees a discourse as a collaborative effort.³ A discourse is a joint project in which the interlocutors aim at achieving one or more common goals. In (7), the goal is to provide A with petrol. (True, in the first instance, this is A's goal, but B's response shows that he is making A's cause his own.) Obviously, there are indefinitely many goals a discourse may serve. People talk with each other to exchange information, to negotiate a deal, to settle disputes, and so on.⁴ But in any given case the number of discourse goals will be fairly small, and apparent to all interlocutors; they are, after all, common goals.

Given that a discourse is a joint project between interlocutors, it is natural enough to suppose that an utterance φ will normally be interpreted in the light of the current discourse goals, on the assumption that φ was designed by the speaker to further these goals. In other words, the hearer will assume that the speaker intends to abide by the

Cooperative Principle

Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. (Grice 1975/1989: 26)

This is Grice's "overarching principle" for an informal theory of cooperative linguistic behaviour. Grice presents that theory as a collection of more specific "maxims", loosely arranged into four rubrics:

3. Following common usage in pragmatics, I will generally employ the term "discourse" instead of Grice's "conversation".
4. I'm reminded by Larry Horn that language is also used in projects that are non-cooperative in one way or another: courtroom arguments, political debates, matrimonial fights, and so on. Note, however, that these are still *joint* projects which are cooperative at least in the sense that an exchange of views is supposed to take place.