Quantity implicatures

Bart Geurts

December 21, 2009. To appear with Cambridge University Press.

Table of contents

Preface

Quantity implicatures had been haunting me for about a decade, when I finally decided it was time to put my mind at rest. Writing this book proved to be of considerable therapeutic value, but of course my hope is that it will be of value (therapeutic or otherwise) to its readers as well. Since I have been thinking about this topic for quite a few years, there was ample opportunity for changing my mind, which I did a number of times. I've had my heretical phases (most recently in my 2005 paper on disjunction), but they are behind me now, and I'm glad to be back in the Gricean fold.

I have tried to write a text that is accessible to a fairly broad audience, including students of linguistics, philosophy, and psychology. In recent years, the literature on my topic has become increasingly esoteric. I don't welcome this development, and I don't see it was inevitable, either. Be that as it may, my aim was to write a book that presupposes little more than general intelligence and common sense, though some familiarity with a handful of very basic linguistic and logical notions may prove to be useful. (A brief compendium of the latter will be found at the end of the book.)

For various kinds of help and support at various stages along the way, I would like to thank Patrick Blackburn, Richard Breheny, Vincent Homer, Janneke Huitink, Hans Kamp, Napoleon Katsos, Rick Nouwen, Ira Noveck, François Récanati, Philippe Schlenker, Katrin Schulz, Benjamin Spector, Dan Sperber, Rob van der Sandt, Tanja van Hummel, Noor van Leusen, Robert van Rooij, Bob van Tiel, and an anonymous reviewer. Special thanks is due to Corien Bary, Anne Hökberg, and Larry Horn, for their profuse comments on the first complete draft; to Emmanuel Chemla, who kindly demolished the first version of Chapter 6 for me; and to Nausicaa Pouscoulous, for joining me in my chase of embedded hippos. Finally, I'm grateful to the Netherlands Organisation for Scientific Research (NWO) for financial support in the last throes of the project.

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 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.
$\rightsquigarrow$ Julius may be in Amsterdam.
$\rightsquigarrow$ 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.

1 Grice

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 installments 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

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

there is contrast between being rich and being dull. Apparently, the demonstrative 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

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 isn'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.

- Quantity
  1. Make your contribution as informative as is required (for the current purposes of the exchange).
  2. Do not make your contribution more informative than is required.
- Quality: Try to make your contribution one that is true:
  1. Do not say what you believe to be false.
  2. Do not say that for which you lack adequate evidence.
- Relation: Be relevant.
- Manner: Be perspicuous:
  1. Avoid obscurity of expression.
  2. Avoid ambiguity.
  3. Be brief (avoid unnecessary prolixity).
  4. Be orderly.

It bears emphasising that these maxims are *not* to be thought of as based on convention, let alone linguistic convention. Rather, Grice sees them reasonable expectations flowing naturally from a universal theory of cooperative conduct, of which a theory of cooperative linguistic conduct is merely a special case. True, Grice never got round to working out such a theory, but the very banality of most of the maxims suggests rather strongly that his general conception is a feasible one. Surely, if a speaker wants to be cooperative, it would be a good idea for him to make his utterances sufficiently informative, speak the truth, and so on, and surely it is reasonable for hearers to expect speakers to behave in conformity with such patently sensible rules. In brief, the banality of the maxims strongly speaks in their favour.

It will be evident that Grice's maxims are a mixed bag. To begin with, the Quality maxims, especially the first one, are arguably more important than the others. Mendacity makes a speaker unreliable; verbosity, obscurity, or lack of orderliness are just a nuisance. Therefore, Grice suggests, it might be reasonable to accord a special status at least to the first Quality maxim: other maxims come into play only on the assumption that this maxim is satisfied. But the same goes for the maxim of Relation. After all, a speaker who makes no attempt at making his utterances relevant to the purpose of the discourse is every bit as uncooperative as one who is mendacious, and more obviously so.[5]

5. Cf. Sperber and Wilson (1995), who argue that Relevance (which is not the same Grice's maxim of Relation) is the most important, and indeed the only, maxim.

Concerning the Quantity maxims, Grice observes that the second one should perhaps be subsumed under the maxim of Relation: to the extent that an utterance is overinformative, it will be irrelevant to the discourse goals. Again, the same holds, mutatis mutandis, for the first maxim of Quantity. Therefore, it will not come as a surprise that many authors have tried to reduce Grice's maxims to a leaner set while of course criticising competing reductionist attempts (see, e.g., Atlas and Levinson 1981, Horn 1984, Sperber and Wilson 1995, and Levinson 2000).

It seems, then, that every maxim is special in one respect or another, and Grice's own classificatory scheme is likely to be redundant. Grice is clearly aware that this is so. In fact, he appears to take his own proposal with a considerable grain of salt. For one thing, as we have seen, Grice is at some pains to point out the defects in the classification he proposes. For another, the mere fact that he tries to shoehorn his maxims into a taxonomy devised by the philosopher Immanuel Kant for an entirely unrelated purpose suggests that Grice considers his proposal tentative at best. All of which is to say that we shouldn't set too much store by the precise form in which Grice casts his theory of discourse. The gist of Grice's theory is just that interlocutors expect each other to be cooperative and that this general expectation manifests itself in many different ways: we expect each other to be striving for relevance and truth, etc., and these expectations serve to make inferences about the speaker's intentional state: what she believes, wants, and so on. Any theory that accepts this general picture can be called "Gricean".

1.3 Conversational implicatures

A conversational implicature is "what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed" (Grice 1975/1989: 39-40). ψ is a conversational implicature licensed by the speaker's utterance of φ if (i) ψ is not conventionally associated with φ and (ii) the hearer has good reasons for believing that, assuming the speaker intends to be cooperative, he wouldn't have uttered φ unless ψ was the case. Let us have a look at some of examples given by Grice himself to illustrate how the maxims may be exploited for inducing conversational implicatures, starting with the petrol example of §1.1:

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

Evidently, the main goal of this mini-discourse is to help A obtain petrol for his car, and B's utterance wouldn't be relevant to this goal unless he believed the garage was open (or at least might be open), had petrol to sell, and so on. Hence, the implicature that B believes so is driven by the maxim of Relation.

The next example involves the Manner maxim:

(10) Miss X produced a series of sounds that corresponded closely with the score of "Home sweet home".

Here, the speaker goes of his way to avoid the verb "sing", using a considerable more prolix expression in its stead. He wouldn't do so unless he believed, and wanted to convey to his audience, that the predicate "sing" wasn't entirely appropriate, presumably because Miss X's singing wasn't very good.

Our third example involves the protagonist of this book, the first maxim of Quantity:

(11) A: Where does C live?
B: Somewhere in the South of France.

Grice glosses this case as follows:

> [B's answer is] less informative than is required to meet A's needs. This infringement of the first maxim of Quantity can be explained only by the supposition that B is aware that to be more informative would be to say something that infringed the second maxim of Quality, "Don't say what you lack adequate evidence for", so B implicates that he does not know in which town C lives. (Grice 1975/1989: 33)

The purpose of the exchange in (11) is that A should learn where C lives. A's question is a bid to establish this as the discourse goal, and B's response signals his acceptance. Grice's gloss states rather categorically that B's infringement of the Quantity maxim can be explained only by assuming that B lacks the evidence that would warrant a more specific claim, but that is clearly too strong. It could be, for example, that B considers his answer precise enough for A's purposes (Grice simply states that this isn't the case, but it might be), or A and B might be playing a guessing game in which the players aren't supposed to give fully specific answers, and there are many more possible situations in which (11B) would not implicate that the speaker doesn't know where C lives, though it may be that such situations tend to be rather special.

It is always possible to invent scenarios in which an implicature that seems natural and even compelling at first would fail to arise. But this is just to say that people don't arrive at implicatures by rigidly applying a set procedure,

and despite his momentary lapse in the cited passage, Grice is perfectly aware that this is not how things go. A conversational implicature is more like a stab at explaining a linguistic act, and the logic of this sort of stabbing (if indeed it has a logic) is quite unlike the classical logics in which valid conclusions are necessitated by their premisses. We will revisit this theme in the following chapter (§2.2).

Let me now give a somewhat rough general format for the derivation of conversational implicatures:

i. It starts with a speaker S performing a particular kind of action: S utters a sentence φ.
ii. In view of the current discourse goals, S's utterance wouldn't have been in accordance with the Cooperation Principle unless ψ was the case.
iii. Hence, ψ must be the case.

The implicature, i.e. ψ, will generally be of the form "S believes that …", "S doesn't know whether …", etc. But although ψ will often be about the speaker's knowledge state, it may be about other types of intentional state, too. For example, in a restaurant the statement "We haven't ordered yet" may be used to implicate that the speaker *wants* the waiter to take his order. Similarly, though φ will usually be a statement (Grice only discusses statements, and the formulation of some of the maxims is tailored to this type of speech act), it may also be a question or command, for example:[6]

(12) Please give me a pen or pencil.

This will usually imply that the addressee need not provide both a pen and a pencil, and this inference is naturally explained as a quantity implicature: if the request had been for a pen and a pencil, the speaker should have said, "Please give me a pen and a pencil", and since she chose to say (12) instead, that's probably not what she meant.

It is important to note that it is only speech acts that give rise to implicatures, where a speech act may be defined, somewhat loosely, as the utterance of a sentence that commits the speaker to having a certain belief, desire, etc. Since speech acts are not sentences, but rather actions of a specific kind, it would be a logical *faux pas* to say that this or that sentence carries such-and-

6. Here I'm assuming that the exclusive interpretation of "or", which licenses the inference that, according to the speaker, only one of the disjuncts is true, is due to a conversational implicature. See §3.3 for further discussion.

such an implicature. It is perhaps excusable to say that an implicature is associated with a sentence (as uttered on a given occasion), as long as it is firmly kept in mind that this is always to be thought of as informal shorthand: it is the speech act, not the sentence, that carries the implicature. However, it is wholly inexcusable to say that an *embedded* sentence, or any other expression smaller than a full sentence, may have conversational implicatures associated with it. For example, the embedded sentences in (13a,b) (marked by square brackets) could never give rise to conversational implicatures:

(13) a. If [C lives somewhere in the South of France], she must be bored.
b. There was a rumour that [Miss X produced a series of sounds that corresponded closely with the score of "Home sweet home"].

A speaker who utters (13a) doesn't state that C lives somewhere in the South of France, and therefore in this context it will not have the implicature it has in Grice's example (11). In fact, no matter how loosely we allow ourselves to speak of the connection between sentences and implicatures, there is no sense in which this occurrence of "C lives somewhere in the South of France" can be said to have implicatures at all. The same, mutatis mutandis, for (13b).

If embedded sentences cannot have implicatures associated with them, then a fortiori words can't have them, either. This may seem obvious enough, but in view of later developments it is worthwhile to belabour this point a bit.

(14) Harry is either in Antwerp or in Brussels.

Typically, an utterance of (14) would convey that the speaker neither knows whether Harry is in Antwerp nor whether he is in Brussels, and it has been argued that these inferences are encoded in the lexical content of "or" (Zimmermann 2000, Geurts 2005). However, the standard view is Grice's, who explains these inferences as quantity implicatures, along the following lines. If the speaker knew that Harry was in Antwerp, then he could have made a more informative claim by saying, "Harry is in Antwerp", but since he chose to make a less informative and more prolix statement instead, he probably does not know whether Harry is in Antwerp. By the same line of reasoning, the hearer can infer that the speaker doesn't know whether Harry is in Brussels.

Although I believe that Grice's pragmatic analysis of "or" is right, and I have my doubts about the lexicalist alternative, my main point, for now, is just that there is no middle ground between these two views, given the elementary fact that conversational implicatures are not carried by words. To claim otherwise (as has been done by, e.g., Levinson 2000 and Chierchia

2004) is quite simply an abuse of terminology. True, we have to allow for the possibility that, as a language evolves, conversational implicatures *turn into* conventional elements of content, and it might be argued that, in English, what started out as a conversational implicature has become part of the lexical content of "or". But if this is true (which I doubt), then this inference has ceased to be a conversational implicature.

Conversational implicatures are licensed not by words or sentences but by actions of a certain kind, i.e. speech acts. This already suffices to ensure that the etiology of conversational implicatures is far more general than that of, say, lexical entailments or conventional implicatures. It is a fact about English "or" that any English sentence of the form "φ or ψ" entails that at least one of φ and ψ is true. By contrast, that utterances of disjunctive sentences may give rise to ignorance inferences ("the speaker doesn't know whether φ and doesn't know whether ψ") is a fact about declarative statements with disjunctive content in general, be they expressed in English, Basque, or Mandarin Chinese. Indeed, Grice goes further even than this, in that he sees his theory of discourse as just a special instance of a general theory of cooperative action.

This view has a very important methodological corollary: it entails that if an inference can be explained in terms of conversational implicature, then ceteris paribus such an explanation is to be preferred. This is what Grice calls "Modified Occam's Razor": "Senses are not to be multiplied beyond necessity." (Grice 1978/1989: 47) Suppose, for example, that we had to adjudicate between the two theories of "or" discussed above: the lexicalist account which says that the ignorance inferences associated with "or" are part of its lexical content and the Gricean account according to which these inferences are conversational implicatures. Suppose, furthermore, that these two accounts described the relevant facts equally well. Then we would still have to go with the Gricean account, because it derives from general principles what according to its lexicalist competitor has to be stipulated.

1.4 Generalised vs. particularised

One of the best-known distinctions introduced by Grice is one he isn't too forthcoming about: it is the distinction between particularised and generalised conversational implicatures:

> I have so far considered only cases of what I might call "particularized conversational implicature"—that is to say, cases in which an implicature is carried

> by saying that p on a particular occasion *in virtue of special features of the context*, cases in which there is no room for the the idea that an implicature of this sort is normally carried by saying that p. But there are cases of generalized conversational implicature. Sometimes one can say that the use of a certain form of words in an utterance would normally (*in the absence of special circumstances*) carry such-and-such implicature or type of implicature. (Grice 1975/1989: 37; emphasis added)

Grice uses two cases of quantity implicature to illustrate the distinction. One of his examples is:

(15) X is meeting a woman this evening.

Grice observes that an utterance of this form "would normally implicate that the person to be met was someone other than X's wife, mother, sister, or perhaps even close platonic friend" (ibid.), adding that even when they are generalised, conversational implicatures do not thereby become conventional. And that is more or less what Grice has to say about the generalised/particularised distinction, leaving ample room for exegesis.

In my view, when Grice calls an implicature "generalised", it is merely to indicate that it tends to arise in such a broad range of contexts that it will not seem to depend on the details of the context too much. But this is just to say that the factors such conversational implicatures depend upon happen to be shared by many contexts. For example, it seems highly likely that in just about any context in which (15) is felicitously uttered, the interlocutors will be interested in the identity of the woman in question, regardless what other goals (if any) their discourse may serve, and in any such context, an utterance of (15) will be liable to license the inference that the woman is not X's wife, mother, etc.

If my reading of Grice is correct, then the distinction between generalised and particularised implicatures is "theoretically inert", as Neale (1992: 524, n. 18) puts it, which is in accordance with the fact that Grice has so little to say about it. The reason why Grice introduces this distinction, I believe, is that it is of *strategic* importance to him. One of the recurrent themes in Grice's work is that, in thinking about language, we have a tendency to overload our words with meaning, in the sense that, whenever we observe that a word is generally used in such-and-such a way, we are apt to put it down to the meaning of the word (cf. §1.1). It is this tendency that has led some authors to suppose that it is due to the lexical content of "or" that "φ or ψ" usually implies that the speaker doesn't know whether φ and doesn't know whether ψ. One of

Grice's main concerns is to show that this kind of reasoning is precipitate, and that such inferences are better treated as conversational implicatures. And to acknowledge the fact that such inferences can be pretty much standard, he dubs them "generalised conversational implicatures".

If this exegesis is on the right track, it is futile to argue against the generalised/particularised distinction, as some authors have done (e.g., Hirschberg 1985, Sperber and Wilson 1995), and it is mistaken to believe that generalised implicatures should be allocated an entire level of their own in a theory of interpretation (Levinson 2000). In general, it seems that Grice's notion of generalised conversational implicature has become undeservedly embroiled in discussions it wasn't designed for; more about which in Chapter 5.

1.5 Cancellability

Towards the end of "Logic and conversation", Grice discusses a number of "features" that are characteristic of conversational implicatures. For instance, he says that conversational implicatures have to be "calculable" and that at least some of them will be "non-detachable". "Calculability" just means that the hearer must be able to figure out how an implicature can be derived; which is such a basic requirement that it hardly needs mentioning. An implicature is "non-detachable" if it is dependent only on the content of the speaker's utterance, not its form. Hence, if an utterance of a sentence φ gives rise to an implicature ψ, then ψ is non-detachable from what is said by means of φ if alternative ways of expressing φ's meaning would have given rise to the same implicature. For instance, if we assume, if only for the purpose of illustration, that "some" is synonymous with "at least two", then we should expect that (16a) and (16b) should give rise to the same quantity implicatures:

(16) a. Some of the stewardesses were snoring.
b. At least two of the stewardesses were snoring.

And indeed, both sentences will tend to convey that, according to the speaker, not all the stewardesses were drunk. It should be obvious that not all conversational implicatures will be non-detachable in this sense, if only because Manner implicatures do depend on the form of the speaker's utterance. And there are more exceptions, for as we will see in Chapter 6, quantity implicatures, too, are sometimes contingent on linguistic form. Therefore, the diagnostic value of non-detachability is rather limited.

Another feature discussed by Grice is cancellability, and since this notion figures rather prominently in the literature and is much more problematic than it is usually given credit for, I will discuss it at greater length. According to Grice, there are two ways of cancelling a conversational implicature. If a speaker makes an utterance that gives rise (or might give rise) to an implicature φ, he may add a rider like "but not φ", in which case the cancellation is explicit, or φ may be cancelled implicitly, by the context. (17a,b) are instances of explicit cancellation:

(17) a. There is a garage round the corner, but it's closed.
b. X is meeting a woman this evening—his sister, in fact.

Good examples of implicit cancellation are harder to find, but we encountered one in the last section when discussing (18):

(18) Harry is either in Antwerp or in Brussels.

When uttered during an undoubtedly boring game in which the addressee has to guess where Harry is, (18) will not prompt the ignorance inferences otherwise associated with it. That is, it will not be construed as implying that the speaker doesn't know whether Harry is in Antwerp and doesn't know whether he is in Brussels, because the game presupposes that the speaker knows where Harry is.

Cancellability is a useful diagnostic, because it can be used to separate conversational implicatures from entailments and conventional implicatures. To illustrate, compare (17a), which is felicitous, with (19), which is not:

(19) *There is a garage round the corner, but it's not round the corner.

Since the propositional content of the first half of (19) is that there is a garage round the corner, this proposition can only be cancelled at the price of self-contradiction; which is why (19) is bad. Similarly, as we saw in §1.1, conventional implicatures cannot be cancelled with impunity, though the effect is not as strong:

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

If an inference is cancellable, it is neither a conventional implicature nor entailed by the sentence's propositional content, and therefore it might be a conversational implicature. However, it doesn't have to be, since other types of inference are cancellable, too:

(21) a. Harry doesn't realise October 2 was a Friday.

b. Harry doesn't realise October 2 was a Friday, because it wasn't: it was a Saturday.

Ceteris paribus, an utterance of (21a) will suggest that Harry got promoted, and it is widely agreed that, rather than a conversational implicature, this is a presupposition triggered by the factive verb "realise"; (21b) shows that this presupposition is cancellable. Hence, if an inference is cancellable, it may but need not be a conversational implicature.

Even though cancellability may an effective heuristic for ruling out the possibility that a given inference is either a conventional implicature or an entailment, it doesn't reveal the true nature of the inference in question. Hence, as a diagnostic test it is useful but less powerful than one might have hoped for. What is more worrisome is that the *word* "cancellability" has proved to be highly misleading. Grice's explanation of cancellability, which is echoed throughout the literature, is that conversational implicatures are cancellable because a speaker may decide, or be forced to decide, to opt out of the Cooperative Principle. On this view, a speaker who utters (17a) or (17b) is being uncooperative, if only in a small way. There are reasons for doubting that this is correct. An utterance of (17a) may not be of immediate help to the addressee, in the scenario envisaged by Grice, but arguably the speaker is still providing topical information, which may prove useful in the longer term. Hence, it is not so clear that someone who utters (17a) is being downright uncooperative. The same holds for (17b) and it holds a fortiori for (18). Indeed, in cases like (18), in which an implicature is preempted by the context, it is particularly doubtful that an inference is cancelled at all.[7]

Consider again the examples in (17). If these were instances of cancellation, we would have to say that in these cases the speaker first makes an utterance that gives rise to a implicature φ, and then goes on to cancel φ. This sounds somewhat perverse if not downright incoherent: a speaker who invites his audience to draw an inference only to cancel it immediately afterwards can hardly claim to be cooperative. Perhaps, then, it is better to say that the purpose of the second clause of (17a) is to *prevent* the hearer from deriving an implicature he might have derived otherwise; and the same,

7. To complicate things a bit further, there is an aspect of cancellability that many readers of Grice seem to have overlooked, which is that Grice considers it to be a feature of *generalised* conversational implicatures (Grice 1975/1989: 39). If this restriction is observed, then presumably (17a) will not count as an instance of cancellation, unless it can be argued that "There is a garage round the corner" standardly gives rise to the inference that the garage is open; which doesn't seem very plausible.

mutatis mutandis, for (17b). I believe that this makes more sense than the standard cancellation view, because conversational implicatures are uncertain in a way entailments and conventional implicatures are not. For one thing, conversational implicatures are generally dependent on relevance, which is an inherently fuzzy notion, not only because it is a matter of degree, but also because there will rarely be a perfect consensus, between interlocutors, on what is relevant and what is not. For another, conversational implicatures involve reasoning about the speaker's intentional state (his beliefs, desires, and so on), much of which is bound to be tentative. In contrast, given that it is part of the lexical content of "but" that "φ but ψ" implies a contrast between φ and ψ, there can hardly be any doubt that this is what the speaker means when he utters a sentence of this form.

What I would like to suggest, therefore, is that within the general framework outlined by Grice, the notion that conversational implicatures are literally cancellable is either misleading or misplaced. It is misplaced in cases like (18), where a conversational implicature doesn't arise, to begin with. It is misleading in cases like (17a) and (17b), which for reasons given in the last paragraph I would prefer to analyse in terms of prevention rather than cancellation. Needless to say, none of this is to suggest that we should give up the cancellability diagnostic, even if its value is somewhat limited, as I have argued. It is just that the name of the diagnostic shouldn't be taken too literally.

1.6 Gricean reasoning and the pragmatics of what is said[8]

Implicatures are suppositions made by the discourse participants in order to make sense of each other's speech acts. A speech act is a cooperative act in much the same way as giving someone a light, assisting at an operation, or playing the part of Fortinbras are cooperative acts, but speech acts are special in that they carry propositional content: "what is said", as Grice calls it. How does this sort of content come about? Grice's suggestion (and it is really no more than that) is that, by and large, it is determined by the rules of the language. Per linguistic convention, words and grammatical constructions are assigned meanings, and the meaning of a complex expression is determined by the meanings of its parts and the grammatical constructions they enter into. Thus, the proposition expressed by "Snow is white" is built up from the lexical

8. The second half of this section title was borrowed from Recanati (1989). See also Recanati (2004).

meanings of the words "snow", "is", and "white", pieced together according to the conventions of English grammar.

If this was the whole story, we could say that what is said is conventional simpliciter, but as Grice notes there is more to it.

(22) The ship left the lock.

Before this sentence can be assigned a truth value, we first have to decide: (*i*) what is the time at which the purported event took place and (*ii*) what are referents of "the ship" and "the lock"; and in order to make that last decision, (*iii*) the word "lock" has to be disambiguated. That is to say, in order to determine what is said, the hearer has to resolve any ambiguities the sentence might contain, interpret tense, and assign referents to pronouns and definite descriptions, as well as deal with a host of other aspects of interpretation that Grice doesn't mention. Clearly, such tasks may require that the hearer take contextual information into account: what is said is not determined by linguistic conventions alone, and contrary to what Grice seems to suppose, it is well possible that what is said is radically underdetermined by linguistic conventions.

In order to resolve ambiguities, referential terms, tense, and so on, the hearer will have to ask himself what the speaker's communicative intentions are, and inevitably the Cooperative Principle will be involved in answering this question. For example, the intended meaning of "lock" in (22) is probably the one that is the most relevant, at the point in the discourse where this sentence is uttered, so it would seem that the hearer needs to appeal to the maxim of Relation in order to justify her choice. It appears, therefore, that disambiguation has to rely on conversational implicatures. But how can that be? Given that conversational implicatures are *based on* what is said, how could they help to determine what is said? The Gricean project threatens to become circular, or so it seems; Levinson (2000) calls this "Grice's circle".

Fortunately, however, it is easy enough to see that no such circle exists. Consider example (22) again. How can the hearer determine the intended meaning of "lock" before she has determined what the sentence means? Easy: by asking herself what the sentence *would* mean if she chose one lexical meaning rather than another. If "lock" referred to a mechanism for locking a door, the propositional content of the sentence wouldn't make much sense, whereas if the word referred to a section of a canal in which the water level can be raised and lowered, what is said would make perfect sense, and therefore the latter meaning is preferred. The same, mutatis mutandis, for reference resolution,

the interpretation of tense, and so on. Hence, "Grice's circle" evaporates once it is realised that sub-propositional interpretation issues can be resolved by Gricean reasoning which is based not on what *is* said but on what *would be* said if this or that interpretive choice were made.

Though "Grice's circle" is a mirage, it brings us to a topic that deserves some discussion. It so happens that Grice concentrates his attention on pragmatic inferences associated with utterances of single sentences.[9] Due to this restriction, which is by no means intrinsic to his programme, Grice's claims about conversational implicatures do not always generalise to all kinds of inferences driven by the Cooperation Principle. Put otherwise, Grice's analysis sometimes fails to carry over to all instances of what I will call "Gricean reasoning". On the one hand, as we have already seen, some Gricean reasoning is sub-propositional, which is to say say that it is concerned with the disambiguation of words, reference resolution, and so on. On the other hand, some Gricean reasoning is supra-propositional: it is based on multi-sentence discourse segments. In the remainder of this section, I will say a bit more about each of these levels.

Not all linguistic acts are sentence-sized. If we reserve the term "speech acts" for those linguistic acts that are performed by means of sentences,[10] there are also linguistic acts which are expressed in units that are larger and others that are expressed in smaller units. Beginning with the former, it is natural enough to say that telling a story, giving a proof, pleading a case, exchanging gossip, etc., are linguistic acts made up of several speech acts; and such large-sized linguistic acts give rise to implicatures in much the same way as simple speech acts do. To illustrate, consider the following recipe for negotiating stairs with crutches, as provided by the American College of Foot and Ankle Surgeons (www.footphysicians.com):

1. Seat yourself on a low step.
2. Move your crutches upstairs by one of these methods:
   - If distance and reach allow, place the crutches at the top of the staircase.
   - If this isn't possible, place crutches as far up the stairs as you can—then move them to the top as you progress up the stairs.

9. As a matter of historical interest, it may be noted, though, that the first passage in Grice's work hinting at the Cooperative Principle concerns the resolution of ambiguities (Grice 1957). Cf. §8.5.

10. Though this usage is widely spread, some authors apply the term "speech act" more liberally. A prominent case in point is Searle (1969), who treats referring as a speech act. I don't have any principled quarrels with this, but I do believe that such cases are sufficiently special to deserve an appellation of their own.

3. In the seated position, reach behind you with both arms.
4. Use your arms and weight-bearing foot/leg to lift yourself up one step.
5. Repeat this process one step at a time. (Remember to move the crutches to the top of the staircase if you haven't already done so.)

Surely, this recipe would normally license the inference that the five steps have to be executed in the order given (Manner), that there isn't an extra step that isn't mentioned here (Quantity), and that the procedure will work (Quality). It is clear, then, that extended stretches of discourse give rise to conversational implicatures just like individual utterances do. There is nothing remarkable about this, since the distinction between sentences and larger discourse units is not very strict, to begin with. For example, stylistic considerations apart, there doesn't seem to be much of a difference between the sentences in (23a) and (24a) and the mini-discourses in (23b) and (24b):

(23) a. Harry is dull but rich.
b. Harry is admittedly dull. However, he is rich, too.

(24) a. Tobago is primarily hilly and of volcanic origin.
b. Tobago is primarily hilly. It is of volcanic origin.

Just as there are linguistic acts beyond the sentence level, there are linguistic acts below that level. In using a definite description like "the ship", the speaker performs a referential act targeted at a particular individual; in using a word like "lock", he intends to convey a particular meaning; and so on. It would be wholly remarkable if the hearer's interpretation of such acts did not proceed on the assumption that the speaker is trying to be cooperative. This is not to imply that at the sub-propositional level the maxims are operative in precisely the same way as they are on higher levels; since the maxims were tailor-made for (propositional) speech acts, this is hardly to be expected. But it is to say that we should expect to find implicature analogues below the propositional level, and we do. For example, the use of a definite description like "the ship" will normally be expected to be non-vacous (∼Quality) and to refer to a relevant object (∼Relation). Furthermore, in general referential expressions are expected to neither too general (∼Quantity 1) nor too specific (∼Quantity 2), since lack of specificity will render it impossible to identify the intended referent, while excessive specificity will result in moronic discourses like the following:

(25) Once Little Red Riding Hood's grandmother gave Little Red Riding Hood a little cap made of red velvet. Because it suited Little Red Riding

> Hood so well, and Little Red Riding Hood wanted to wear it all the time, Little Red Riding Hood came to be known as Little Red Riding Hood.

It is important to observe that there is a quite fundamental distinction between the propositional and supra-propositional levels, on the one hand, and the sub-propositional level, on the other, which lies in the fact that the linguistic acts that make up the former are quite different from the ones that define the latter. At the propositional level, the basic unit is the speech act, and it is natural enough to extend this to the supra-propositional level: discourses and discourse segments (paragraphs) are plausibly viewed as means for making compound speech acts. However, linguistic acts at the sub-propositional level are something different altogether. Referring to a ship and making a statement about a ship are patently different things, and we would be ill-advised to ignore the difference. If a referential expression α is insufficiently specific for the hearer to determine the intended referent, we might say that the speaker's use of α violates an *analogue* of the first Quantity maxim, but not the maxim itself, which was custom-made for speech acts. For these reasons, whereas I wouldn't mind extending Grice's notion of implicature to the supra-propositional level, I believe it would be a bad idea to apply this notion at the sub-propositional level, as well.

To sum up, Gricean reasoning is employed not only to draw inferences from what the speaker has said, but also to determine what he wants to say, and contrary to what Levinson and others have claimed, this is not particularly problematic, though it usefully serves to highlight the fact that Gricean reasoning may apply at various levels, which it is important to keep distinct. In the following chapters, our attention will be focused on the more familiar use of Gricean reasoning, i.e. on conversational implicatures proper, but the sub-propositional level will loom large at the end of the book (§8.5).

The Standard Recipe for Q-implicatures

2

Having surveyed the general framework of Gricean pragmatics, this chapter homes in on the variety of conversational implicatures we will be concerned with in the remainder of the book: quantity implicatures, or "Q-implicatures" for short (Horn 1984). Our starting point is what I call the "Standard Recipe" for deriving Q-implicatures. There are two reasons for calling it thus. First, the literature has by and large confined its attention to inferences that follow its pattern. Secondly, I suspect that in practice the Standard Recipe accounts for many if not most Q-implicatures associated with declarative sentences. However, standard though it may be, it doesn't cover everything that might justifiably be called a Q-implicature, even for declaratives.

The bulk of this chapter is devoted to fleshing out and amending the Standard Recipe, and this process will continue in the following chapters, where the amendments become more substantial. In fact, I will end up arguing that when push comes to shove the Standard Recipe gets things backwards: it has all the right ingredients, but fails to combine them in just the right way (Chapter 6). For the time being, however, we will be walking well-trodden paths.

2.1 The Standard Recipe

I will introduce the Standard Recipe by way of an example. Clyde says:

(1) Bonnie stole some of the pears.

On the classical Gricean account, (1) says that Bonnie stole at least some of the pears, and may implicate that she didn't steal all of them. This implicature

is explained by assuming that the hearer reasons, and is entitled to reason, as follows:[1]

i. Rather than saying (1), Clyde could have made a stronger statement:
(1*) Bonnie stole all the pears.
Why didn't he do so?

ii. The most likely explanation is that Clyde doesn't believe that (1*) is true: $\neg\textsc{Bel}_{\textsc{C}}(1^*)$.

iii. Clyde is likely to have an opinion as to whether (1*) is true: $\textsc{Bel}_{\textsc{C}}(1^*) \vee \textsc{Bel}_{\textsc{C}}(\neg(1^*))$.

iv. Between them, (*ii*) and (*iii*) entail $\textsc{Bel}_{\textsc{C}}(\neg(1^*))$: Clyde believes that Bonnie didn't steal all the pears.

Note, to begin with, that this train of reasoning is conducted in terms of the speaker's beliefs. In the more recent literature on implicature, there is a tacit consensus that this is right, but in the older literature, the same style of analysis was often cast in terms of speaker's knowledge (Gazdar 1979 and Soames 1982 are prominent examples). I don't see what could justify the latter choice, and in fact I believe it's plainly wrong, so I'm siding with the majority view. It should be noted, however, that for our purposes nothing hinges on the distinction between knowledge and belief (e.g., the logical form of the reasoning displayed above is the same in either case), and since it is sometimes more convenient to use the verb "know" rather than the weaker "believe", that's what I will sometimes do.

The Standard Recipe starts with the observation that the speaker could have made a stronger utterance than he did. Of course, since for any actual utterance there are always indefinitely many stronger utterances that in principle could have been made instead, there will have to be substantial constraints on the alternatives available to the speaker at any given point in the discourse. For example, if the speaker says φ and ψ is stronger than φ but also considerably longer, then one good reason why the speaker didn't bother to say ψ is that the extra effort required wouldn't have been worth the trouble. Another, and much more important, constraint is that the alternatives to be considered should be *relevant* to the interlocutors. Thus, in the case of (1), "Bonnie stole some of the pears and today is Saint Genebald's day", though

1. Readers who are unfamiliar with some of the notation appearing below are referred to the overview of notation and abbreviations used in this book (p. 192).

more informative than (1), would not count as an alternative.

Terminological note: Unless clearly indicated otherwise, I will use "alternative" as a theoretical term. An alternative is a sentence the speaker might have uttered instead of the one he did utter. The exact import of "alternative" will vary a bit, but its core meaning will remain constant: "the alternatives" associated with a sentence, as uttered on a given occasion, will always be a set of sentences.

Implicatures weak and strong Actually, the Standard Recipe involves not one but two inferences that we might want to call "Q-implicatures". The first one, $\neg\textsc{Bel}_{\textsc{C}}(1^*)$, is the weaker of the two, and it is important to realise precisely what it says:

It is not the case that Clyde believes that Bonnie stole all the pears.

This does *not* entail that Clyde believes that Bonnie didn't steal all the pears, which is what the strong implicature says. $\neg\textsc{Bel}_{\textsc{C}}(1^*)$ allows for the possibility that, as far as Clyde knows, Bonnie may or may not have stolen all the pears. According to the stronger implicature, i.e. $\textsc{Bel}_{\textsc{C}}(\neg(1^*))$, Clyde has made up his mind on this issue.

There is no established nomenclature for these two inferences. According to Soames, the weak one is a generalised implicature, while the second is particularised, but as I have my doubts about the generalised/particularised distinction (§1.4), I will not follow him in this point. Other options are to say that only the first inference is a genuine implicature (the second inference derives from the first plus an auxiliary assumption), or to use Sauerland's (2004) distinction between "primary" and "secondary" implicatures. Since this is just a terminological matter, and I happen to like the weak/strong distinction, I will stick with that.

The strong implicature follows from the weak one plus the assumption that Clyde either believes that Bonnie stole all the pears or believes that she didn't.[2] This assumption, too, has been given a variety of names, including "Experthood Assumption" (Sauerland 2004), "Authority Assumption" (Zimmermann 2000, Geurts 2005), and "Competence Assumption" (van Rooij and Schulz 2004). I will settle for the last one, but whatever its name, the underlying idea is the same in any case; it is that, in general, implicature strengthening is

2. As pointed out to me by Larry Horn, this style of disjunctive reasoning was already used by Bartsch (1973) to explain "neg-raising" in attitude reports. See Horn (1989) for further discussion.

licensed by the assumption that the speaker knows what he is talking about. In our example, this is to say that Clyde knows whether or not Bonnie stole all the pears. Note, however, that this is slightly stronger than what is required for obtaining a strong implicature: for that, we merely need to assume that Clyde believes either that Bonnie stole all the pears or that she didn't. That is to say, the Competence Assumption only has to rule out the possibility that Clyde hasn't made up his mind as to whether (1*) is true: it has to eliminate the dithering space between Clyde's believing that Bonnie stole all the pears and his believing that she didn't. If Clyde was dithering on the matter, (1*) and ¬(1*) would both be possible for him, and (2) would be true:

(2) $\textsc{Poss}_C(1^*) \wedge \textsc{Poss}_C(\neg(1^*))$, or equivalently: $\neg\textsc{Bel}_C(\neg(1^*)) \wedge \neg\textsc{Bel}_C(1^*)$

The Competence Assumption is equivalent to the negation of (2). Hence, our official definition of the Competence Assumption is somewhat weaker than the intuitive notion. In general, this mild discrepancy will not matter; we will briefly return to it in §2.3.

In a context that doesn't support the Competence Assumption, the strong implicature cannot be derived, and the Standard Recipe gives us only the weak one. To illustrate how that might happen, consider (1) in two slightly different scenarios. In both cases Clyde has witnessed Bonnie stealing some pears, but while in scenario S_1 he saw the platter from which she took the pears, in S_2 his view of the platter was obstructed by (say) a fat cat. Consequently, in S_1 but not in S_2 Clyde knows whether any pears are left; or in other words, while the Competence Assumption holds in S_1, it doesn't hold in S_2, and the Standard Recipe forecasts a strong implicature in the former case but not in the latter. This seems to be correct: while in S_1 we would infer from Clyde's uttering (1) that $\textsc{Bel}_C(\neg(1^*))$, S_2 only allows for the weaker inference that $\neg\textsc{Bel}_C(1^*)$.

In my opinion, one of the main virtues of the Standard Recipe is that it distinguishes between weak and strong implicatures, and connects them via the Competence Assumption. Generally speaking, non-Gricean theories are designed to account only for strong implicatures, which means that weak implicatures are relegated elsewhere and the link between weak and strong implicatures is severed (e.g., Gazdar 1979, Levinson 2000, Chierchia 2004, 2006). This is a less attractive solution than what the Gricean approach has to offer. We will take up this issue in §7.3; the weak/strong distinction and the Competence Assumption will be discussed at greater length in §2.3.

The Standard Recipe Example (1) is a somewhat special case, in the sense that its Q-implicatures are closely associated with one particular word, viz. the scalar expression "some"; which is why it is called a "scalar implicature". In cases like (1), we can say that Q-implicatures arise because the speaker uses a less informative *word* than he could have, and consequently his utterance could have been more informative than it is. However, as Hirschberg (1985) was the first to emphasise, Q-implicatures aren't always like this, as the following discourse illustrates:

(3) A: Has Clyde's book come out yet?
B: He has corrected the proofs.

B's answer clearly implies that the answer to A's question is no. This, too, is a Q-implicature, which is obtained by applying the Standard Recipe:

i. Rather than giving the answer in (3), B could have made a stronger statement:

(3*) Yes, it has.

Why didn't she do so?

ii. The most likely explanation is that B doesn't believe that (3*) is true: $\neg \textsc{Bel}_{\textsc{B}}(3^*)$.

iii. B is likely to have an opinion as to whether (3*) is true: $\textsc{Bel}_{\textsc{B}}(3^*) \vee \textsc{Bel}_{\textsc{B}}(\neg(3^*))$.

iv. Between them, (*ii*) and (*iii*) entail $\textsc{Bel}_{\textsc{B}}(\neg(3^*))$: B believes that Clyde's book hasn't come out yet.

The only difference between this derivation and the previous one is that the relevant alternative, (3*), cannot be generated from B's statement by substituting one word for another. It is important to see how small the difference is, because it is often overrated, if only implicitly: since a very large part of the literature on Q-implicatures confines its attention to the scalar variety, other varieties have received less attention than they deserve, and there seems to be a widespread tacit view that scalar implicature is so special that it is entitled to a treatment of its own, which it has duly received at the hands of a long list of distinguished authors, starting with Horn (1972). I consider this an unfortunate development, and take it to be obvious that we should be aiming for a more general account of Q-implicatures, in which scalar inferences come out as just a garden variety. We will revisit this issue in the next chapter.

Abstracting away from the peculiarities of the examples discussed in the foregoing, here is the general form of the Standard Recipe for deriving Q-implicatures:

S has said φ.
i. S could have made a stronger claim by saying ψ. Why didn't he do so?
ii. Presumably, it's because S doesn't believe that ψ is true: $\neg\textsc{Bel}_S(\psi)$.
iii. S has an opinion as to whether ψ is true: $\textsc{Bel}_S(\psi) \vee \textsc{Bel}_S(\neg\psi)$.
iv. Between them, (*ii*) and (*iii*) entail $\textsc{Bel}_S(\neg\psi)$: S believes that ψ is false.

A condensed form of the Standard Recipe was already given by John Stuart Mill in the 19th century (and quoted on p. 2). More developed versions can be found in Grice's own work (e.g., Grice 1975, 1978) and the argument became fully explicit at the hands of Soames (1982) and Horn (1989), and throughout the more recent Gricean literature. Its formal properties have been studied by van Rooij and Schulz (2004) and Spector (2006).

Non-declaratives One rather obvious limitation of the Standard Recipe is that, like the maxims proposed by Grice, it is tailored to assertions. It is not designed to work for questions, requests, commands, and so on. And despite the fact that the literature ignores them so completely that one might easily come away with the impression that non-declarative speech acts don't have Q-implicatures, it's not true that they don't. In the previous chapter, we saw that requests can have Q-implicatures (§1.3, example (12)). Here is an example involving a question. Suppose that Clyde is collecting the ingredients for a cake he wants to bake. Bonnie enters the scene, and the following exchange ensues:

(4) B: Do you have all the ingredients?
C: Where's the butter?

Clyde's question may plausibly be taken to implicate that he did manage to find the other ingredients, and this is arguably a Q-implicature, derivable as follows:

i. Clyde could have asked, for example: "Where is the butter and where are the eggs?" Why didn't he do so?
ii. One possible explanation is that Clyde doesn't need to be told where the eggs are.
iii. And the most likely explanation for (*ii*) is that he found them.

(And the same, mutatis mutandis, for the sugar, vanilla, and so on.) This is a bit sketchy, admittedly, and fleshing out this story is not going to be trivial. But I take it that something along these lines will be needed to explain the inference Clyde's question gives rise to, and that the skeletal form of this story is very similar to that of the Standard Recipe. In particular, both start out with the question why the speaker failed to produce a more specific utterance, though the way this question is answered is not exactly the same for interrogative and assertive speech acts.

2.2 Inference to the best explanation

Inferences and readings Implicatures are pragmatic inferences. Therefore, unless we are exceptionally liberal about what is to count as a "reading", they don't affect the reading of a sentence. To illustrate this point, consider the following sentences, all of which admit of several readings, and each in its own way:

(5) a. Clyde has a dog.
b. All the pears weren't ripe.
c. Bonnie believes that she is pretty.

(5a) is ambiguous between a reading according to which Clyde owns a male dog and one on which the animal in question may be any canis familiaris, male or female. Similarly, (5b) has two readings: one which allows for the possibility that some of the pears are ripe, and one which doesn't. Finally, (5c) is referentially ambiguous: the feminine pronoun may refer to Bonnie herself or to another woman. In all these cases, the hearer is expected to *select* the intended reading, using contextual cues as appropriate.

According to some authors, this is how (some) quantity inferences work, as well. For example, Fox (2007) and Chierchia et al. (to appear) maintain that (6a) is syntactically ambiguous between a strong reading that rules out (6b), and a weak reading that doesn't:

(6) a. Bonnie stole some of the pears.
b. Bonnie stole all the pears.

The merits and demerits of this view (especially the latter) will be discussed later (Chapter 7); the point I want to make now is just that it is fundamentally different from the Gricean view. If the "not all" inference associated with "some" in (6a) is an implicature, then the meaning of the sentence remains

the same no matter whether the implicature is derived or not. The hearer doesn't select between a reading that includes the scalar inference and one that doesn't. Rather, he first selects a reading, and then decides whether or not to *derive* the implicature.

Implicatures are abductive inferences Conversational implicatures are inferences of a special kind. To begin with, they are non-deductive. Clearly, Clyde's saying (6a) does not necessitate the conclusion that $\neg\textsc{Bel}_C$(6b). The inference may be licit on many and perhaps even most occasions, but it isn't valid in the logicians' sense of the word. What is it then? Implicatures are, in Peirce's (1931) terminology, abductive inferences, or inferences to the best explanation, in Harman's (1965) (I will use these terms interchangeably).[3] If Clyde utters (6a), and his interlocutors infer that $\neg\textsc{Bel}_C$(6b), then they do so because they judge it is the best way of explaining why Clyde failed to make a stronger claim.

A somewhat crude way of characterising abductive reasoning is that it uses conditionals backwards: starting out with an explanandum ψ, you observe that if any of $\varphi_1, \ldots, \varphi_n$ *were* true, then ψ would follow, and then proceed to select one among $\varphi_1, \ldots, \varphi_n$ as the best explanation. This sort of diagnostic reasoning is quite common in daily life. If I see that the milk bottle is empty, I'm likely to conclude that somebody (a fellow member of my household, presumably) drank it, although in principle it could have disappeared in countless other ways: it could have been used for watering the plants, its contents could have been transferred to another galaxy by aliens, and so on. Similarly, if a patient exhibits all the symptoms associated with measles, his physician is likely to infer that the best way of explaining the symptoms is by assuming that the patient has measles (rather than, say, dengue fever). Incidentally, in medicine, this mode of reasoning has been enshrined in the illustrious Zebra Principle: "When you hear hoofbeats, think horses, not zebras."

It is important to realise that abductive reasoning can go wrong without therefore being wrong: it is a form of reasoning that tolerates exceptions. To illustrate this characteristic feature, consider the following example. Having observed the usual symptoms (a rough engine, steam escaping from the exhaust, etc.), a mechanic suspects that the Volkswagen he is looking at must

3. See, e.g., Bach and Harnish (1979), Green (1996), Levinson (2000), and Chien (2008). Thagard and Shelley (1997) criticise various attempts at defining abduction in terms of deduction. An excellent book-length treatment of inference to the best explanation is Lipton's (2004).

have a blown head gasket. But then when he looks under the bonnet, he discovers various pieces of equipment expressly designed to produce these symptoms; there is nothing wrong with the engine. It turns out to be an elaborate prank. Clearly, the mechanic's diagnosis was incorrect: he thought the car had a blown head gasket, and it hadn't. Nevertheless, the man's reasoning was impeccable: this Volkswagen was merely an exception, not a counterexample.

The Gricean reasoning formalised in the Standard Recipe is very much like that of the mechanic. With some ingenuity, we can concoct all sorts of scenarios in which the Standard Recipe produces the wrong results. But they are merely exceptions, not counterexamples. For instance, there are cases in which Q-implicatures simply fail to arise. Consider (6a) once again, but now in a context in which it has been discovered that the pears were poisoned, and the only thing that matters is whether or not Bonnie took any of them. In such a context, the truth value of (6b) is *irrelevant*, and it seems unlikely that an implicature is intended by the speaker or inferred by the hearer (more about relevance in §2.4).

Hence, in some cases, Q-implicatures don't arise at all. More important for our current purposes are cases where Q-implicatures do arise, but not in the standard way. Consider the following example:

(7) A: Clyde seems to have become such a penny-pincher. What's the matter with him?
B: He had to borrow money for some expensive pet he wanted to buy, and now he's saving wherever he can.

B's use of the phrase "some expensive pet" may be construed as implicating that she doesn't know what species of animal Clyde bought. However, it may also convey that she doesn't care or even disapproves, perhaps because she dislikes pets or Clyde's infatuation with them. For reasons that aren't quite clear to me, being more general than might be expected may serve to express a negative attitude, as in:

(8) What's this thing doing on the mantlepiece?

—where it is obvious to all concerned that the thing in question is a priceless antique clock (cf. Cruse 1977). Perhaps the dismissive effect is caused by the speaker's studied indifference, but whatever the details of the explanation, it is clear that these are Q-implicatures that don't follow the Standard Recipe: the question why the speaker chose to be uninformative is answered not by assuming that there is something she doesn't know but by assuming that

she doesn't care, or worse. Put otherwise, while in the Standard Recipe the speaker's key attitude is epistemic, in these cases it is evaluative.

Another type of case in which the Standard Recipe doesn't apply is the following. There is a popular stereotype according to which mothers are liable to say things like:

(9) I see that somebody hasn't eaten his porridge yet.

In the kind of situation I'm alluding to, the mother is referring obliquely to little Clyde (say), so she could have been more explicit by addressing him directly:

(10) I see you haven't eaten your porridge yet.

Still, although there is a stronger statement at hand, there is no implicature to the effect that Clyde's mother doesn't believe that it is he who failed to eat his porridge. I don't know why mothers speak like this, but it is obvious enough that it is a case of Q-implicature: the speaker deliberately avoids using a (more informative) referential expression, and intends her audience to *infer* whom she has in mind.

Another example in the same vein involves a stereotype of a more recent vintage: it features two international terrorists calling each other on the phone. One of them suspects that an "intelligence agency" may be tapping the line, and says:

(11) A mutual acquaintance of ours will not be able to keep his appointment.

Like (9), the manifestly underinformative statement in (11) may serve to convey information about a particular individual, while at the same time the speaker is communicating that it is unsafe to be more informative. The latter is clearly a Q-implicature, since it is the answer to the question why the speaker chose to violate be less informative than he could have been; but the Standard Recipe doesn't apply here.

In examples like (9) and (11), it is clear that the Standard Recipe for deriving Q-implicatures would produce the wrong results, for in both cases it would be inappropriate to derive inferences of the form "The speaker doesn't know whether ..." Still, it is evident enough that these statements violate the Quantity maxim, and thereby give rise to Q-implicatures. It is just that in cases like these the best explanations are non-epistemic: they do not pertain to what the speaker knows or doesn't know.

Although some of the cases discussed in the foregoing are rather far-fetched, that doesn't affect my main point. Far-fetched or not, these examples tell us

something important about Q-implicatures, and indeed about conversational implicatures in general, namely that they heavily depend on world knowledge in a way deductive inferences utterly fail to do. Though in many cases, the Standard Recipe will give the best explanation for apparent violations of the Quantity maxim, there are bound to be exceptions, too: cases in which other explanations are better.

A Q-implicature is an answer to a particular type of why-question: "Why didn't the speaker say so-and-so?", where "so-and-so" is a more informative, or more specific, utterance than the one the speaker actually made. Nearly all the work on Q-implicatures has confined its attention to one particular type of answer to such questions, and I agree this is sensible enough: the cases covered by the Standard Recipe are of central importance, and probably the best place to start, so that's what I will do, as well. Still, we shouldn't lose sight of the fact that there are other ways of deriving Q-implicatures. Rather than an algorithm for producing Q-implicatures, the Standard Recipe is a heuristic form of reasoning that is interesting and common enough to deserve careful analysis. But there is only so much this analysis will be able to do. In particular, we cannot hope to explain why, on any given occasion, the hearer decides to adopt an explanation that follows the Standard Recipe, rather than any of infinitely many alternative explanations that are always available in principle.

2.3 Weak implicatures and competence

Competence and competence$^+$ To the best of my knowledge, the first fully developed version of the Standard Recipe was proposed by Soames (1982), whose analysis differs from the one presented in §2.1 in two ways. One is that, following Gazdar (1979), Soames casts his analysis in terms of knowledge rather than belief; for reasons given before I will ignore this part of his proposal. The second difference is that Soames uses a stronger version of the Competence Assumption, which I will adorn with a plus to distinguish it from the one I favour, and which (modulo the first difference) goes like this:

(12) *Competence*$^+$ (Soames 1982)
If ψ then $\textsc{Bel}_S(\psi)$, and if $\neg\psi$ then $\textsc{Bel}_S(\neg\psi)$.

S is competent$^+$ with respect to ψ if his beliefs agree with the facts; that is, if S knows whether or not ψ is true. In Soames's analysis, the Competence$^+$

Assumption replaces the Competence Assumption in the Standard Recipe as presented in §2.1. Clearly, Soames's notion of competence$^+$ is stronger than the unadorned one: assuming that ψ is either true or false, (12) entails that $\textsc{Bel}_S(\psi) \vee \textsc{Bel}_S(\neg\psi)$, but not vice versa. Therefore, Soames-style competence$^+$ is stronger than it need be for strengthening weak implicatures. By the same token, Soames's analysis is less general. For consider Travis, who has strong opinions on everything: give him any proposition ψ, and he will either believe ψ or $\neg\psi$, regardless what the facts of the matter are. According to the Standard Recipe, Travis's implicatures will always be strengthened, which seems to be right; Soames doesn't predict this.

Of course, Travis's case is exceptional, if not pathological, and I wouldn't want to argue on this basis that the Competence$^+$ Assumption is too strong. In fact, I suspect that Soames is basically right in that generally speaking it is competence$^+$, and not just competence, that is involved in the derivation of strong implicatures. If I stick with the weaker version nonetheless, it is mainly because it is simpler and more general: since it is defined in terms of knowledge rather than (mere) belief, Soames's notion imposes more stringent constraints on the derivation of strong Q-implicatures than ours does, and although these constraints may be satisfied in most cases, they are not needed in principle, as the case of Travis illustrates.

What are weak implicatures? According to the Standard Recipe, Q-implicatures are derived in a number of steps, resulting initially in a weak implicature of the form $\neg\textsc{Bel}_S(\psi)$, which may be strengthened to $\textsc{Bel}_S(\neg\psi)$, if the Competence Assumption holds for ψ. For ease of reference, I give here a compressed version of the argument:

i. Weak implicature: $\neg\textsc{Bel}_S(\psi)$
ii. Competence: $\textsc{Bel}_S(\psi) \vee \textsc{Bel}_S(\neg\psi)$
iii. Strong implicature: $\textsc{Bel}_S(\neg\psi)$

The weak implicature is usually paraphrased as "The speaker doesn't know whether or not ψ is the case." As before, I would prefer to talk in terms of belief rather than knowledge: "The speaker has no opinion as to whether ψ is true or false." But the difference doesn't really matter, because either way the paraphrase is *stronger* than the inference licensed by the Standard Recipe. For, the paraphrase is captured by the conjunction $\neg\textsc{Bel}_S(\psi) \wedge \neg\textsc{Bel}_S(\neg\psi)$, which says that the speaker believes neither ψ nor $\neg\psi$, and therefore entails, but isn't entailed by, $\neg\textsc{Bel}_S(\psi)$.

Before we continue, it may be helpful to give the following diagram, which displays how Competence and the various candidate implicatures relate to each other:

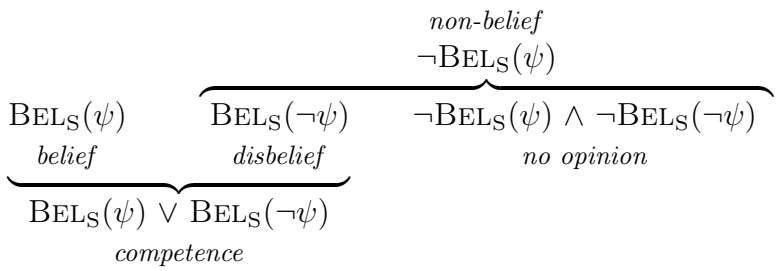

This is to be understood as saying that "non-belief" is the disjunction of "disbelief" and "no opinion", while "competence" is the disjunction of "belief" and "disbelief". Note, furthermore, that "no opinion", i.e. the intuitive paraphrase of the weak implicature, is the negation of "competence":

"$\neg\text{BEL}_\text{S}(\psi) \wedge \neg\text{BEL}_\text{S}(\neg\psi)$" is true iff "$\text{BEL}_\text{S}(\psi) \vee \text{BEL}_\text{S}(\neg\psi)$" is false.

Now, which of the two is the true weak implicature: $\neg\text{BEL}_\text{S}(\psi)$, which is what the Standard Recipe gives us, or the stronger $\neg\text{BEL}_\text{S}(\psi) \wedge \neg\text{BEL}_\text{S}(\neg\psi)$, which corresponds to the standard way of paraphrasing the weak implicature? According to Pouscoulous (2006), who was the first to broach this issue, it is the latter. To explain why, suppose hearer H has executed the first step of the Standard Recipe, deriving $\neg\text{BEL}_\text{S}(\psi)$. Then, in principle, there are three options for H to choose from:

i. H may decide not to adopt the Competence Assumption, and is then left with just $\neg\text{BEL}_\text{S}(\psi)$.
ii. H may decide that the Competence Assumption is true, and then $\text{BEL}_\text{S}(\neg\psi)$.
iii. H may decide that the Competence Assumption is false, and then $\neg\text{BEL}_\text{S}(\psi) \wedge \neg\text{BEL}_\text{S}(\neg\psi)$.

To make this more concrete, we return to our original example: Clyde says:

(13) Bonnie stole some of the pears.

Given that Clyde could have said,

(13*) Bonnie stole all the pears.

the three options distinguished above amount to the following:

i. $\neg\textsc{Bel}_{\textsc{C}}(13^*)$: Clyde either believes (13*) is false or he believes (13*) may or may not be true.
ii. $\textsc{Bel}_{\textsc{C}}(\neg(13^*))$: Clyde believes (13*) is false.
iii. $\neg\textsc{Bel}_{\textsc{C}}(13^*) \wedge \neg\textsc{Bel}_{\textsc{C}}(\neg(13^*))$: Clyde believes that (13*) may or may not be true.

Pouscoulous considers it highly unlikely that by saying (13) Clyde should have wanted to *communicate* merely that (*i*) is the case, and I believe she is right about this: presumably, what Clyde intended to communicate to his audience was either (*ii*) or (*iii*). However, taking the hearer's point of view, it may be doubtful what the status of the Competence Assumption is, so in at least some contexts (*i*) may be all the hearer is able to infer with reasonable confidence, in which case he may want to ask for clarification:

(14) What do you mean: she didn't steal all of them or you don't know?

So, although hearers may prefer to either adopt or reject the Competence Assumption, they may sometimes remain undecided, and be left with a rather weak construal of what the speaker meant to convey.

The upshot of the foregoing discussion is as follows. Although the Standard Recipe isn't wrong, it is incomplete, and the common way of presenting it is not entirely correct. Here is how it should go. Having derived an inference of the form

(15) $\neg\textsc{Bel}_{\textsc{S}}(\psi)$

the hearer may go on to assume either that the Competence Assumption holds or that it doesn't hold. In the former case, it follows that

(16) $\textsc{Bel}_{\textsc{S}}(\neg\psi)$

In the latter case, the inference is that

(17) $\neg\textsc{Bel}_{\textsc{S}}(\psi) \wedge \neg\textsc{Bel}_{\textsc{S}}(\neg\psi)$

The problem with the standard way of presenting the Standard Recipe is that (15) is offered as meaning (17), which it doesn't, since the former is weaker than the latter. Note that (16) and (17) are logically independent (neither entails the other), which is to say that although (16) may seem to be stronger than (17), it isn't. Therefore, we can rule out the possibility that (16) is derived by strengthening (17): it can only derive from (15).

Having made the requisite distinctions, I will in the following take the liberty of not being too strict about them, the reason being simply that for most of the discussion the difference between (15) and (17) is immaterial.

Disjunction and ignorance Disjunctive sentences of the form "φ or ψ" generally give rise to the implicature that the speaker doesn't know which of φ and ψ is true. For example, if Clyde says (18), his utterance is likely to be construed as implying that he is agnostic about the truth values of its disjuncts, (19a,b):

(18) Bonnie stole an apple or a pear.

(19) a. Bonnie stole an apple.
b. Bonnie stole a pear.

Grice's (1978/1989: 46-47) explanation for these ignorance inferences is seemingly straightforward: had Clyde known that (19a) is true, he should have said so, and since he didn't he is implicating that $\neg\textsc{Bel}_C$(19a); the same for the (19b). The reason why this is only seemingly straightforward is that, according to the Standard Recipe, the next step should be to strengthen these implicatures, which would lead us to conclude that $\textsc{Bel}_C(\neg$(19a)) and $\textsc{Bel}_C(\neg$(19b)): Clyde believes both disjuncts to be false. Assuming that Clyde is not violating the Quality maxim, he believes what he says, so we cannot infer both $\textsc{Bel}_C(\neg$(19a)) and $\textsc{Bel}_C(\neg$(19b)); but why not infer one of the two? The answer is that we cannot make the Competence Assumption for either disjunct, on pains of contradiction (Sauerland 2004). Still assuming that Clyde is not violating the Quality maxim, he believes what he says, and therefore $\textsc{Bel}_C$((19a) $\vee$ (19b)); so, taken together with the weak implicatures we have already, this yields:

i.	$\textsc{Bel}_C$((19a) $\vee$ (19b))	[*Quality*]
ii.	$\neg\textsc{Bel}_C$(19a)	[*Quantity*]
iii.	$\neg\textsc{Bel}_C$(19b)	[*Quantity*]

If, in addition, we adopt the Competence Assumption for (19a), we get a contradiction, as follows:

iv.	$\textsc{Bel}_C$(19a) $\vee$ $\textsc{Bel}_C(\neg$(19a))	[*Competence*]
v.	$\textsc{Bel}_C(\neg$(19a))	[*from* (*ii*) *and* (*iv*)]
vi.	$\textsc{Bel}_C$(19b)	[*from* (*i*) *and* (*v*)]
vii.	contradiction	[*between* (*iii*) *and* (*vi*)]

Of course, assuming competence for (19b) will yield the same result. In short: provided Clyde believes what he says, the weak implicatures of his utterance imply that the Competence Assumption fails for both disjuncts, and therefore no strong implicatures can be inferred. We might say, therefore, that in the case of disjunction, weak implicatures are non-competence inferences.

Competence by default? It is sometimes suggested that hearers adopt the Competence Assumption by default. Very roughly, the idea is that, in the absence of evidence to the contrary, hearers will assume that the speaker knows what he is talking about. Thus formulated, the claim is plausible enough, though it is less than self-evident what it means for Q-implicatures. I propose to read it as follows:

(20) *Competence by default*
If the hearer applies the Standard Recipe, and asks himself why the speaker uttered φ instead of the stronger ψ, he will assume by default that the speaker is competent with respect to ψ.

If this is correct, we predict that, *if* the Standard Recipe is deployed, there will be an inference of the form $\textsc{Bel}_S(\neg\psi)$ by default. Note that this is superficially similar to, but really *very* different from, the claim that certain expressions trigger default inferences of the form $\textsc{Bel}_S(\neg\psi)$, a claim that will be discussed at length in Chapter 5.

Another way of formulating (20) is that, generally speaking, a speaker who has sufficient evidence for making a claim φ will be taken to also have sufficient evidence for deciding on the truth value of ψ, where ψ is a stronger alternative to φ. This seems plausible to me, though of course it may turn out to be false. The only way to settle this issue, I suspect, is by collecting quantitative data, but unfortunately I don't see how this might be done. In the meantime, I suggest that we accept (20) as a working hypothesis.

2.4 Relevance

The derivation of a Q-implicature starts with the observation that the speaker could have made a stronger statement than he did, and then proceeds to explain why he didn't do so. It is easy enough to see that it would be a bad idea to raise this question for every stronger statement the speaker could have delivered, not only because doing so would cause a dramatic slowdown of the interpretation process, but also because it would produce absurd results like the following. Bonnie says:

(21) Clyde is a dope.

Instead of saying (21), Bonnie could have a made the following statement, which is patently stronger:

(22) Clyde is a dope and Jupiter is larger than Saturn.

So, apparently Bonnie doesn't believe that (22) is true, and since she has committed herself to the truth of the first conjunct, we are forced to conclude that she doesn't believe that Jupiter is larger than Saturn.

Something must have gone wrong, and it is fairly evident what it is: not any old statement that happens to be stronger than (21) qualifies as a *relevant* alternative. Indeed, for practically any given statement and context, the overwhelming majority of stronger statements fail to be relevant. Thus the question arises: how to separate the relevant alternatives from the irrelevant ones? This question is … uhm … difficult, and since I have little hope of advancing the state of the art in this respect, the following remarks serve to outline the problem more than anything else.

Discourse purposes Let's start with Grice's own formulation of the first Quantity maxim:

> Make your contribution as informative as is required (for the current purposes of the exchange). (Grice 1975/1989: 26)

This formulation suggests that according to Grice the only relevant alternatives are those that would have contributed to the purposes of the discourse (though elsewhere he seems to adopt a more liberal view on relevance; e.g., Grice 1975/1989: 38). This is an appealing idea, and to illustrate how it might work, I will use one of Grice's own examples:

(23) The prize is either in the garden or in the attic. (Grice 1978/1989: 45)

As explained in the last section, an utterance of (23) will tend to implicate that, as Grice puts it, the speaker has "non-truth-functional grounds" for believing that his statement is true; which is Grice's way of saying that (23) will typically license the implicature that the speaker doesn't know which of the following is true:

(24) a. The prize is in the garden.
b. The prize is in the attic.

But now suppose (23) is uttered in the context of a treasure hunt, by a parent who knows full well where the prize is; then there will be no such implicature, obviously. In such a case, neither (24a) nor (24b) count as stronger claims a cooperative parent could have made, in view of the purpose of the exchange, because either statement would have spoiled the game.

What are discourse purposes? This, too, is a hard question, but part of the answer must be that discourse purposes are common goals, i.e. goals that are shared between speakers. All parties can come to the discourse with any number of private goals, but only some of these will become discourse goals, and not every goal the discourse helps to achieve is a discourse goal. For example, even if my main reason for steering the conversation towards minimalist syntax is that I want to impress my audience with my knowledge of the topic (the example is fictitious), that would not be a common objective, and therefore not a discourse goal. Rather, the purpose of the discourse would be to exchange information about minimalist syntax.

A common way of setting up discourse goals is by asking questions. For instance, by asking,

(25) Where's the Earl Grey?

Bonnie is making a bid for establishing a discourse goal of a very particular kind, viz. that she should learn where the Earl Grey is. If Clyde accepts this as the next discourse purpose on the agenda, then he should undertake to answer the question to the best of his ability; for instance, by offering one of the following:

(26) a. In the blue cupboard.
b. On the top shelf of the blue cupboard.

Now (26b) is more informative than (26a), and therefore an utterance of the former might well implicate that Clyde cannot be more precise. However, if it is agreed that (26a) is good enough for achieving the current discourse goal, then no Q-implicature will arise.

An illustration of how implicatures are constrained by questions is offered by the following minimal pair:

(27) A: Has your article been published yet?
B: I corrected the proofs.

(28) A: What was your contribution to that article?
B: I corrected the proofs.

Whereas B's statement in (27) may be interpreted as implicating that for all B knows, the article hasn't been published yet, the same statement in (28) is unlikely to give rise to such an implicature. What explains the contrast is that the discourse goals are different.

A question serves to set up a discourse goal which its answer should help to achieve, and thus helps to establish what is relevant in the following. Hence, questions are employed for negotiating relevance, and this observation provides us with a useful tool for studying the effects of relevance on implicatures, as well as other aspects of interpretation: by observing how the interpretive effects of a statement change depending on what questions it might be used to address (as we did in the foregoing examples), we are in effect varying the relevance factor.

Hearer's interest If it was determined entirely by the immediate purposes of the discourse, the notion of relevance would be relatively unproblematic. Unfortunately, however, it isn't.

(29) A: I tried to call you yesterday afternoon. Where were you?
B: I was playing tennis with a friend of mine.

B's utterance is likely to implicate that his tennis partner, Clyde, is not among A and B's mutual acquaintances. But this information is not relevant to the immediate purpose of the discourse, which is to answer A's question, and thus explain why B wasn't answering the phone. In cases like (29), we probably have to say that relevance is determined by the hearer's interests. B will be aware that Clyde's identity is of potential interest to A if Clyde happens to be a mutual friend, A's employer, the Dalai Lama, etc., and this is enough to make relevant some of the alternative statements that are stronger than (29B).[4]

General interest We have seen that relevance is determined not only by the immediate purposes of the discourse, but by the hearer's interests, as well. And that's not the end of it, as the following dialogue illustrates:

4. Although discourse purposes and hearer's interests will typically overlap, they are not the same. For example, even if it is common knowledge between a journalist and a politician she is interviewing that neither has any personal interest in the other, whatsoever, this will not prevent them from producing all manner of implicatures.

It is the hearer-centred notion of relevance that takes pride of place in Sperber and Wilson's work (e.g. Sperber and Wilson 1995, Wilson and Sperber 2004), though I should note that Sperber and Wilson's use of "relevance" doesn't align with any intuitive notion of relevance, as far as I can tell. Rather, what they call "relevance" is the *ratio* between hearer's relevance and processing effort; so the unit of measurement of relevance à la Sperber and Wilson is common relevance per kilojoule.

(30) A: Where is Bonnie?
C: She went out to buy a piece of furniture.

C's answer is very likely to implicate that he doesn't know what kind of furniture Bonnie went out to buy (or that he doesn't care, cf. the discussion of (7), §2.2). This information does not seem relevant to the immediate purpose of the discourse, and though it may be relevant to A's personal interests, it doesn't have to be. For suppose that it is mutual knowledge between A and C that A has no interest in furniture whatsoever, and couldn't care less whether Bonnie bought a sofa, a dinner table, or a rocking chair.[5] Then, still, the question is bound to arise why C chose to be so unspecific. Note that if C had said (31) instead,

(31) C: She went out to buy a sofa.

there wouldn't have been an implicature to the effect that, e.g., B didn't know if B went out to buy a blue sofa. In this respect, there is a sharp contrast between (30C) and (31C), as possible answers to A's question. The same contrast can be observed in the following pair:

(32) a. I saw an animal on the lawn this morning.
b. I saw a dog on the lawn this morning.

Whereas (32a) seems quite liable to implicate that the speaker isn't sure what species of animal he saw, it would require special circumstances for (32b) to implicate that the speaker isn't sure what breed of dog he saw (cf. Brown 1958, Cruse 1977, Rosch 1978, Hirschberg 1985).

The examples in (30)-(32) indicate that there are general expectations to the effect that, when introducing a new discourse entity, speakers should employ expressions of at least a minimum level of specificity: "sofa" and "dog" are sufficiently specific; "piece of furniture" and "animal" are not. These expectations are general in that they don't fluctuate wildly from context to context

5. On the other hand:

> Information, like money, is often given without the giver's knowing to just what use the recipient will want to put it. If someone to whom a transaction is mentioned gives it further consideration, he is likely to find himself wanting the answers to further questions that the speaker may not be able to identify in advance; if the appropriate specification will be likely to enable the hearer to answer a considerable variety of such questions for himself, then there is a presumption that the speaker should include it in his remark; if not, then there is no such presumption. (Grice 1975/1989: 38)

or from hearer to hearer, though they may be culture specific. This point was first made by Matsumoto (1995), on the basis of evidence from Japanese and Korean.

Matsumoto observes that, in English, a sentence like (33) will not tend to implicate that the speaker doesn't know whether Michio is older or younger than Takashi:

(33) Michio is Takashi's brother.

However, if (33) is translated into Japanese, the result is a sentence that does have this implicature:

(34) Kochira wa Takashi-kun no kyoodai no Michio-kun desu.
this TOP Takashi-Mr. GEN brother GEN Michio-Mr. COP

The contrast between the two languages is explained by the fact that, in addition to words corresponding to English "brother" and "sister", Japanese also has a system of sibling terms that encode sex as well as seniority:

ani "older brother"	*otooto* "younger brother"
ane "older sister"	*imooto* "younger sister"

Moreover, and crucially, these are the preferred terms for introducing siblings, and therefore the use of "kyoodai" in (34) is marked in much the same way that "animal" is marked in (32a). Quite possibly, this preference is due to the fact that age differences play a more important role in Japanese culture than they do in other parts of world (see §6.5 for further discussion).

To recapitulate: If a speaker utters a sentence φ when a stronger sentence ψ could have been used instead, the difference doesn't matter (i.e., will not trigger Q-implicatures) unless it is relevant, and it can be relevant in at least three ways. The difference between φ and ψ may be (*i*) relevant to the purposes of the discourse, (*ii*) of potential interest to the hearer, or (*iii*) of general interest. These three options need not be disjoint, but they are not necessarily the same, either.

Relevance isn't one thing: there are a number of disparate factors that determine what is relevant and what is not, and while some of these factors are context dependent to a high degree, others are less so. Unfortunately, although relevance is of crucial importance to an adequate understanding of conversational implicatures of all stripes, I will have very little to add to the preliminary observations offered in this section. Relevance-related considerations will loom large at many points in the book, but I don't have anything

like a theory of relevance to offer; that is to say, a theory that is sufficiently precise, not overly simplistic, and that draws together the foregoing observations with many others that have been made in the literature. I'm not aware of any theories that meet these basic requirements, and, sadly, I doubt that I will see one in my lifetime.[6]

2.5 Conclusion

In this chapter, we have studied the Standard Recipe, a form of reasoning that gives rise to two types of Q-implicatures, weak and strong, which are connected by the Competence Assumption. It is important to bear this general pattern in mind, not only because it will recur again and again, but also because, as we will see in §7.3, alternative theories that seek to explain these inferences in non-pragmatic terms can only account for strong implicatures and have to let the weak ones fall by the wayside.

More importantly, however, the Standard Recipe is a specimen of inference to the best explanation, which is to say that it is not the only way of deriving Q-implicatures and, *a fortiori*, it is not an algorithm for deriving Q-implicatures. The fact that Gricean reasoning is essentially abductive in nature is often overlooked, by friend and foe alike, and all too often the Standard Recipe is treated as if it were a logical calculus; which it isn't. This trend is particularly strong when the focus is on scalar implicatures, a sub-class of Q-implicatures to which we now turn.

6. This may seem like a slur on theories in which relevance plays a central role, but that's not how it is intended. It's just that, as far as I can tell, all such theories *presuppose* some notion of relevance, which is left unaccounted for.

Scalar implicatures | 3

By far the most popular variety of Q-implicatures are scalar implicatures, which were put on the pragmatic map by Horn (1972). Most of the examples we discussed in the foregoing were of this sort; (1) is another one:

(1) Wilma read some of the papers.

(2) a. Wilma read many of the papers.
    b. Wilma read most of the papers.
    c. Wilma read all of the papers.

The distinctive feature of scalar implicatures is that we can use lexical substitution to generate the relevant alternatives from the sentence uttered. Thus, stronger alternatives for (1) would be (2a-c), which are obtained simply by replacing "some" with "many", "most", and "all", respectively. The idea is that there is a scale, say ⟨some, many, most, all⟩, which orders its elements from weaker to stronger, and which in effect generates the alternatives to a given sentence. Hence, using this substitution method, an alternative to a given sentence φ is generated by replacing an expression α in φ by a stronger scalemate of α's. This generative view, as I propose to call it, originated in Horn's (1972) seminal work, and was subsequently adopted by Gazdar (1979), Soames (1982), Hirschberg (1985), and many others. More recently, the view is still being taken for granted, e.g., by Sauerland (2004), Chierchia (2004), and Chemla (2007).

In this chapter, I will first expound and then criticise the generative approach; the latter on two main counts. First, I will argue that Horn scales are redundant: we don't need a scale to tell us that "all" is stronger than "some", for this much follows from the meanings of the words, and if we really want to use the ordering encoded in the scale, we get into trouble. Secondly, I will try to show, using the case of disjunction, that the class of scalar implicatures, as

staked out by the substitution method, does not consitute a natural kind: the generative view is compelled to draw boundaries that shouldn't be there.

Before getting down to business, I should like to stress that, though his name will occur rather frequently in the coming pages, this chapter is not a critique of Horn's views, past or present. In fact, I won't be criticising anybody in particular. The version of the substitution method that will serve as my scapegoat is extremely simple, and is not supported by anyone anymore. However, that doesn't really matter for the main points I want to make.

3.1 Horn scales and the generative view

Let's start by introducing a notational device that will make it easier to discuss sentences and their alternatives: $\varphi[\ldots]$ is a sentence frame in which the dots mark an open position, for example "Wilma read [...] of the papers." If $\varphi[\ldots]$ is a sentence frame, and α is an expression of the appropriate type, then $\varphi[\alpha]$ is that sentence in which α fills the open position in $\varphi[\ldots]$. Hence, $\varphi[\alpha]$ and $\varphi[\beta]$ are the same except that the latter has β where the former has α.

A Horn scale is simply a sequence of increasingly informative expressions. Examples are:

(3) a. ⟨some, many, most, all⟩
b. ⟨warm, hot, scalding⟩
c. ⟨clever, brilliant⟩
d. ⟨or, and⟩

Two expressions will be called "scalemates" iff there is a Horn scale that contains them both; β is a stronger scalemate of α's iff there is a Horn scale in which β occurs to the right of α.

The elements of a Horn scale are ordered according to how informative they are.[1] For example, everyone who is "brilliant" is "clever", as well, but not vice versa; or in other words, the set of brilliant people is a proper subset of the set of clever people.[2]

1. I will be assuming throughout that informativeness is defined in terms of entailment. There has been some discussion in the literature about so-called "non-entailment scales" (see especially Hirschberg 1985). For reasons that will transpire in due course, I believe this is a non-issue.
2. Note that this is not an empirical claim about the meanings of words like "clever" and "brilliant". More accurately, it is not a claim that can be verified directly, for to *call* somebody "clever" often is to imply that the person in question is not brilliant. However, since we

With these preliminaries out of the way, we proceed to define a very simple procedure for generating the alternatives for a given sentence, using Horn scales:

(4) *Using Horn scales to generate alternatives*
$\varphi[\beta]$ is an alternative to $\varphi[\alpha]$ iff β is a stronger scalemate of α's.

Assuming that ⟨some, many, most, all⟩ is the only available scale, this yields the following results for the sentences in (1) and (2):

(2a), (2b), and (2c) are the alternatives to (1).
(2b) and (2c) are the alternatives to (2a).
(2c) is the sole alternative to (2b).
There are no alternatives to (2c).

When the Standard Recipe is applied to the alternatives of sentence (1), for example, this yields the following weak implicatures, each of which may be strengthened, as usual, provided the Competence Assumption holds:

$\neg$BEL$_S$(Wilma read many of the papers)
$\neg$BEL$_S$(Wilma read most of the papers)
$\neg$BEL$_S$(Wilma read all of the papers)

The exclusive interpretation of "or" is explained along the same lines:

(5) a. Barney has bought a Manet or a Monet.
b. Barney has bought a Manet and a Monet.

On the classical, inclusive interpretation of "or", the truth-conditional meaning of (5a) merely rules out the possibility that Barney has bought neither a Manet nor a Monet: the sentence is true if Barney has bought a Manet or a Monet or both. Therefore, we can assume that "or" and "and" share a scale and that the latter is stronger than the former, and then the substitution method predicts that (5b) is an alternative to (5a). If we put this alternative through the Standard Recipe, we get two grades of exclusivity implicature: $\neg$BEL$_S$(5b) (weak) and BEL$_S$($\neg$(5b)) (strong). Thus the exclusive construal of (5a) is explained as an ordinary scalar implicature.

We now have an apparatus that has been viewed as something I have argued could not exist: a mechanical procedure for generating implicatures. On this view, there *is* an algorithm for churning out scalars implicatures, after all, and it consists of the following components:

---

want to treat this as a conversational implicature, it had better not be entailed by the lexical meaning of "clever". Hence, to say that "brilliant" is stronger than "clever" is to adopt an hypothesis that can only be tested as part of a larger theory of interpretation.

i. a set of Horn scales
ii. a procedure for generating alternatives
iii. the Standard Recipe

It is clear why the substitution method has proved to be so popular: it is intuitively appealing, easy to formalise, and simple. And it is even more appealing to suppose that this method is in fact an algorithm for producing a special class of Q-implicatures, viz. scalar ones. There are two reasons for doubting that this could be right. One is that implicatures are inferences to the best explanation, which means that even if a procedure like this were correct in many and perhaps even the majority of cases, it is bound to fail in some (cf. §2.2). Another is that implicatures are contingent on considerations of relevance, which, prima facie at least, don't seem to be taken into account by the generative view (cf. §2.4). Hence, two perfectly sound *a priori* reasons for rejecting this position out of hand.

However, although the abduction argument is incontrovertible, as far as I can see, the relevance argument may not be. For, it might be argued that, contrary to what I have alleged, the generative view already takes into account relevance: it is encoded in Horn scales. To explain how, let us turn to a question that has received remarkably little attention in the literature: *What are Horn scales, anyway?* I believe I can answer that question, and I suspect that my answer will satisfy most scale theorists (though I'm not sure). Briefly, what I would like to suggest is that Horn scales represent hearers' expectations in the sense that if β is a scalemate of α's, β is taken to be relevant in a context in which α is used.

To explain this a bit further, let's start by observing that, although Horn scales are ordered in terms of semantic strength, most sequences of increasingly informative expressions do not qualify as Horn scales. To see why, consider the following candidate scales:

(6) a. ⟨animal, dog⟩
    b. ⟨animal, dog, schnauzer⟩

If (6a) is a bona fide Horn scale, we predict that (7) should give rise to the implicature that the speaker doesn't know (or, perhaps, didn't know at the time) whether or not the sighted animal was a dog:

(7) I saw an animal on the lawn this morning.

This prediction is plausible enough. However, if (6b) was a Horn scale, as well, we would make an analogous prediction for (8): it should yield the implicature

that the speaker doesn't (or didn't) know whether or not the dog in question was a schnauzer:

(8) I saw a dog on the lawn this morning.

This prediction is not nearly as plausible, and if we want to avoid it, we will have to assume that (6b) is not a Horn scale.

On the interpretation I'm suggesting, if (6b) is a Horn scale while (6a) is not, then it is because hearers have different sorts of expectations about the words (or concepts) "animal" and "dog": if an individual x is introduced as an "animal", the hearer is likely to wonder what species of animal x is, whereas if x is introduced as a "dog", the question whether x is a schnauzer will not in general present itself with equal force. Of course, such expectations may be dependent on the context to a greater or lesser degree. For instance, if (8) occurred in a conversation between avid dog lovers, it is quite possible that the hearer would ask herself whether or not the dog in question was a schnauzer, and this could be represented by assuming that, in the given context, (6b) *is* a Horn scale.

If this is the right way of viewing Horn scales (and it's the best I can think of), such scales are partial representations of contextual relevance. Some of these representations, like (6a), will be fairly constant across contexts; others, like (6b), will be more dependent on specific contextual features. Either way, it will be clear that Horn scales don't explain what relevance is; they merely provide a format for representing it. It will also be clear that this format will not be suitable for representing relevance in general. For example, the way relevance expectations are affected by questions (§2.4) cannot be represented by scales, obviously. Which is to say that Horn scales have their limitations; but that is not necessarily a bad thing.

Assuming that this is the right way of viewing the connection between relevance and scales, I will proceed by criticising Horn scales on two points. In the next section (§3.2) we will see that Horn scales are in a sense too rich. A Horn scale collects a number of expressions and orders them in terms of informativeness. I will argue that the ordering is an idle wheel in the machine at the best of times. That is to say, the generative theory doesn't need scales, because sets will do. Though the change from scales to sets is a non-trivial one, it doesn't undercut the key tenets of the generative approach. That job is left to §3.3, where it will be argued that the generative view stands in the way of a unified theory of Q-implicatures.

3.2 Implicatures and downward entailing environments

Consider the following sentence:

(9) Fred didn't read many of the books Wilma gave him.

Assuming, as before, that "many" figures in the scale ⟨some, many, most, all⟩, the substitution method yields (10) as one of the alternatives to (9):

(10) Fred didn't read all the books Wilma gave him.

Applying the Standard Recipe to this alternative, we get $\neg\textsc{Bel}_S(10)$ (weak implicature) and possibly $\textsc{Bel}_S(\neg(10))$ (strong implicature). But this prediction is patently wrong: since (9) *entails* (10), it cannot be that $\neg\textsc{Bel}_S(10)$, let alone that $\textsc{Bel}_S(\neg(10))$. On the other hand, (9) seems to license a Q-implicature that the substitution method does not account for, viz. that $\textsc{Bel}_S(\neg(11))$; or in prose: according to the speaker, Fred read at least some of the books Wilma gave him.

(11) Fred didn't read any of the books Wilma gave him.

(To avoid irrelevant complications, we'll assume that "some" and "any" are just different forms of the same word: one for positive, the other for negative environments.) Unlike (10), (11) *is* stronger than (9), and therefore should qualify as a true alternative, which it isn't according to the definition in (4). Note, furthermore, that when (10) is uttered in its own right it may also implicate that the speaker doesn't believe (11), despite the fact that "all" is the strongest element of its scale, and therefore the generative procedure doesn't yield any alternatives for (10). Apparently, as things stand the substitution method produces the wrong alternatives for negative sentences like (9) and (10).

What causes the problem is that the substitution procedure only checks for the occurrence of certain telltale words, i.e. scalar expressions like "some", "or", or "clever", and ignores the remainder of the sentence that contains them. The procedure presupposes that when we take a statement and substitute a stronger term for a weaker one, the result is bound to be a stronger statement. This is not true for negative statements, because negation reverses the direction of entailment: if φ entails ψ, then $\neg\psi$ entails $\neg\varphi$. Negative environments are not the only ones to have this property: it is the defining property of

downward entailing environments. Downward entailment and its dual, upward entailment, are defined as follows:

Upward entailing environments
$\varphi[\ldots]$ is an upward entailing environment iff, for any two expressions α and β: if α entails β, then $\varphi[\alpha]$ entails $\varphi[\beta]$.

Downward entailing environments
$\varphi[\ldots]$ is a downward entailing environment iff, for any two expressions α and β: if α entails β, then $\varphi[\beta]$ entails $\varphi[\alpha]$.

Another way of saying this is that an upward entailing environment $\varphi[\ldots]$ *preserves* the direction of the entailment relations between the expressions that can fill the dots; a downward entailing environment, on the other hand, *reverses* entailment relations. To illustrate, assuming that "scarlet" entails "red", and using "⇒" to symbolise entailment:

(12) Fred is wearing scarlet socks ⇒ Fred is wearing red socks.
Fred isn't wearing red socks ⇒ Fred isn't wearing scarlet socks.

Hence, while "Fred is wearing [...] socks" is an upward entailing environment, "Fred isn't wearing [...] socks" is a downward entailing environment.

For examples of downward entailing environments that are not negative (at least not in any intuitive sense), we turn to the quantifiers. In (13)-(15), square brackets are used to demarcate the restrictor (first argument) and the nuclear scope (second argument) of the quantifiers "some", "at most five", and "all". These examples show that "some" is upward entailing in both arguments, "at most five" is downward entailing in both arguments, and "all" has a mixed profile; Table 1 gives a summary.

(13) Some of [the students with scarlet socks] [passed]
⇒ Some of [the students with red socks] [passed]

quantifier	*restrictor*	*nuclear scope*
some	UE	UE
at most five	DE	DE
all	DE	UE

Table 1: Entailment properties of three quantifiers.
UE = upward entailing; DE = downward entailing.

Some of [the students who passed] [wore scarlet socks]
⇒ Some of [the students who passed] [wore red socks]

(14) At most five of [the students with red socks] [passed]
⇒ At most five of [the students with scarlet socks] [passed]

At most five of [the students who passed] [wore red socks]
⇒ At most five of [the students who passed] [wore scarlet socks]

(15) All [the students with red socks] [passed]
⇒ All [the students with scarlet socks] [passed]

All [the students who passed] [wore scarlet socks]
⇒ All [the students who passed] [wore red socks]

One of the main reasons why entailment properties have received quite a lot of attention in the linguistic literature is that they affect the distribution of certain expressions. All languages have negative polarity items, which are so-called because they can occur within the scope of a negative expression, and are banned from positive environments. English "any" is a case in point:

(16) a. Wilma {*has/doesn't have} any luck.
b. {*Someone/No one} has any luck.

On closer inspection, it turns out that negative polarity items do not necessarily require a negative environment, though there certainly are constraints on where they may occur, as the following sentences illustrate:

(17) a. *Some nurses with any experience passed the exam.
b. *Some nurses who passed the exam had any experience.

(18) a. At most five nurses with any experience passed the exam.
b. At most five nurses who passed the exam had any experience.

(19) a. All nurses with any experience passed the exam.
b. *All nurses who passed the exam had any experience.

The generalisation is that negative polarity items may only occur in downward entailing positions. In a sense, a negative polarity item signals that the environment in which it occurs is downward entailing.[3]

To return to implicatures, the reason why, despite the fact that "all" is the strongest quantifier in its scale, (20) still manages to give rise to a Q-implicature is that it happens to occur in a downward entailing environment.

3. Actually the facts about polarity are rather more complex than this brief discussion may suggest. See van der Wouden (1997) for an overview of the main issues.

(20) Fred didn't read all the books Wilma gave him. (= (10))

Given that "all" is *stronger* than any of its scalemates, (20) is *weaker* than:

(21) Fred didn't read any of the books Wilma gave him.

Therefore, the weak implicature associated with (20) should be $\neg\textsc{Bel}_S(21)$, and the strong one $\textsc{Bel}_S(\neg(21))$.[4]

Since negation is not the only environment to be downward entailing, we should observe this phenomenon in some non-negative environments, as well; and we do. We have seen that the restrictor of a universal quantifier is a downward entailing environment. Now consider (22):

(22) Everyone who read all these books passed the exam.

Here the universal quantifier "all" occurs in the restrictor of an "every"-type quantifier, so we will get a *stronger* statement if we replace "all" by a *weaker* term, like "any", for example:

(23) Everyone who read any of these books passed the exam.

Therefore, we should expect that an utterance of (22) may implicate that $\neg\textsc{Bel}_S(23)$, and possibly $\textsc{Bel}_S(\neg(23))$. These predictions seem to be correct.

Let us now consider how we could amend the substitution method so as to accommodate the foregoing observations. It seems that what we have to do, if we want to adhere to the view that alternatives are generated by scales, is stipulate that scales are *reversed* in downward entailing environments: if "all" occurs in a downward entailing environment, its scale isn't ⟨some, many, most, all⟩, but ⟨all, most, many, some⟩, and this is what causes (20) to be weaker than its alternatives with "many" and "some" (Fauconnier 1978, Levinson 2000).

This explanation has one major drawback: it's wrong. Scale reversal implies that, in a downward entailing environment like (20), "all" becomes weaker than "some", "many", and "most". But that's absurd. Surely, the *word* "all" is always stronger than its scalemates. So the concept of scale reversal is untenable. True, this notion seems entirely plausible at first, but that's because it

4. One problem with this explanation is that, although the following alternatives are also stronger than (20), it seems that scalar implicatures are less likely to arise here than in the case of (21).

> Fred didn't read most of the books Wilma gave him.
> Fred didn't read many of the books Wilma gave him.

I don't know what causes this discrepancy, but it is presumably related to the fact, observed by Horn (1972), that there is a positive correlation between the robustness of a scalar implicature and the difference in strength between the scalar expressions involved.

is mistaken for something else. We have seen that there *is* a sense in which a downward entailing environment reverses entailment: given that "scarlet" entails "red", "not red" entails "not scarlet"; i.e. the direction of entailment between the complex expressions runs against that between their adjectival parts. But it patently does not follow from this that in a negative context "red" entails "scarlet".

If we really want to stick to the idea that alternatives are generated by lexical scales, then the best thing to do is give up on the idea that scales are scales. This must sound a bit Zen, so let me explain. The original proposal was that the quantifiers "some", "many", "most", and "all" line up in a scale ⟨some, many, most, all⟩. But that's what they do *anyway*: it follows from the lexical meanings of these quantifiers that "many" entails "some", and so on. We may as well leave the ordering implicit, and replace the scale with the unordered *set* {some, many, most, all}. So instead of assuming that words have scales associated with them, we now assume that they are associated with sets of related words, which we might call "Horn sets" (Sauerland 2004). Alternatives are now generated as follows:

(24) *Using Horn sets to generate alternatives*
$\varphi[\beta]$ is an alternative to $\varphi[\alpha]$ iff α and β share a Horn set and $\varphi[\beta]$ is stronger than $\varphi[\alpha]$.

The crucial change with respect to previous definition in (4) is that the notion of strength does not enter at the lexical level anymore; rather, it is full sentences that are compared for strength. A sentence has to be weaker than any of its alternatives, but we're no longer trying to derive this ordering from lexical scales; there was no point doing so, anyway, because that is already being taken care of by the meanings of the sentences in question. Table 2 illustrates the predictions yielded by the amended procedure.

A Horn scale does two things: it collects a number of expressions and orders them according to their strength. We have seen that there is no work to do for the ordering imposed by a Horn scale, since it is the entailment properties of entire sentences, rather than their parts, that is relevant for defining alternatives. If we want to adopt a generative approach to scalar implicatures, the only real work for a Horn scale to do is to select a number of alternative words; which is to say that it doesn't have to be a scale: a set will do.

The move from scales to sets may look like a relatively minor modification, and technically speaking it is. However, the change in spirit is considerable.

	sentence	*alternatives*
i.	Wilma read some of the papers	(*ii*), (*iii*), (*iv*)
ii.	Wilma read many of the papers	(*iii*), (*iv*)
iii.	Wilma read most of the papers	(*iv*)
iv.	Wilma read all of the papers	—
v.	Wilma didn't read any of the papers	—
vi.	Wilma didn't read many of the papers	(*v*)
vii.	Wilma didn't read most of the papers	(*v*), (*vi*)
viii.	Wilma didn't read all of the papers	(*v*), (*vi*), (*vii*)

Table 2: Predictions yielded by the amended procedure for generating alternatives.

The original intuition underlying the generative view is that the alternatives involved in scalar implicatures are generated by *words*, and the previous discussion shows that this intuition is wrong: in the final analysis, it is decided at sentence level what the alternatives are.

3.3 Disjunction: exclusivity and ignorance

Having argued that the classical Horn scales should be replaced with a notion that, though similar in one respect, is quite different in another, I will now try to show that the generative view on scalar implicatures runs into rather serious difficulties. The argument will proceed by way of the pragmatics of disjunction.

Disjunctions belong to the stock-in-trade examples of sentences that license Q-implicatures. In fact, they come with two kinds of Q-implicatures. On the one hand, disjunctive statements generally give rise to ignorance inferences (§2.3), while on the other hand, "or" may be given an exclusive construal, which is standardly analysed as a scalar implicature, and the main concern of this section. In the following I will try to show that, if the generative view were right, it would be a mystery how the exclusivity of "or" can be a Q-implicature. Here's the example we'll be working with:

(25) Wilma is dating Albinoni or Boccherini.

Let's say that the logical form of (25) is "A $\vee$ B", where "$\vee$" is the standard symbol for inclusive disjunction: "A $\vee$ B" is true iff Wilma is dating Albinoni

or Boccherini or both. Now we can summarise the various pragmatic inferences associated with (25) as follows:

(26) *Ignorance inferences*
$\neg \textsc{Bel}_S(A)$ and $\neg \textsc{Bel}_S(\neg A)$
$\neg \textsc{Bel}_S(B)$ and $\neg \textsc{Bel}_S(\neg B)$

(27) *Exclusivity inferences*
$\neg \textsc{Bel}_S(A \wedge B)$ [*weak*]
$\textsc{Bel}_S(\neg(A \wedge B))$ [*strong*]

In §2.3, we saw how ignorance inferences can be analysed as Q-implicatures. Here I would merely like to point out that there is no way the substitution method could account for these inferences, for the simple reason that the alternatives involved are shorter than the sentence uttered: it is obviously impossible to derive "Wilma is dating Albinoni" from (25) by substituting one word for another.[5] Hence, the generative view entails that there is a stark difference between ignorance inferences and exclusivity inferences: the substitution method is implicated in the latter but not the former. I believe this is a serious problem, and will return to it at the end of this section.

Exclusivity inferences are standardly explained as scalar implicatures, on the basis of the assumption that (28), whose logical form is "$A \wedge B$", is one of the alternatives to (25):

(28) Wilma is dating Albinoni and Boccherini.

Since the speaker chose not to make this statement, we infer that $\neg \textsc{Bel}_S(A \wedge B)$, and if the Competence Assumption holds, we get $\textsc{Bel}_S(\neg(A \wedge B))$, as well, which is to say that, according to the speaker, Wilma is dating Albinoni and Boccherini, but not both.

So much for the standard story. There are two problems with it: one general and one specific to the substitution method. The general problem is that, in the case of disjunction, it is less than obvious why the Competence Assumption should hold. For, given that (25) typically licenses the ignorance inferences in (26), it is anything but evident why the hearer should believe that the speaker believes either that $A \wedge B$ or that $\neg(A \wedge B)$, and therefore it isn't obvious at all that, in general, strong exclusivity can be explained in terms of Q-implicature. It may well be that the exclusive interpretation of "or" is often

5. This is not quite true: Sauerland (2004) has a trick for deriving alternatives that are, in effect, shorter than the sentence uttered. However, even by the author's own admission, it *is* a trick.

based on considerations of plausibility rather than implicature, and (25) may be a case in point: since dating is usually, though by no means invariably, a one-on-one affair, it may just be rather likely that Wilma is not dating two men at the same time. Whenever such considerations apply, they will obviously preempt any exclusivity implicatures, weak or strong, that might have arisen. It is an empirical question how frequently this happens, but I will not address that question here. In the following I will ignore this complication and accept, if only for argument's sake, the widely held premiss that exclusivity is due to a Q-implicature.

The more serious problem is specific to the generative treatment of scalar implicatures, and has to do with long disjunctions like the following:

(29) Wilma is dating Albinoni, Boccherini, or Corelli.

To see what the problem is, consider that, if (28) is the relevant alternative to (25), then presumably (30) is the relevant alternative to (29):

(30) Wilma is dating Albinoni, Boccherini, and Corelli.

If the speaker doesn't believe that (30) is true, then he doesn't believe that Wilma is dating each of Albinoni, Boccherini, and Corelli. But this is not enough to qualify as an exclusive construal of (29), for such a construal should entail that, according to the speaker, Wilma is dating just one man on the list.

In order to get a better view on the problem, let's have a closer look at the grammar of disjunction and conjunction. "And" and "or" are usually treated as binary operators, but the examples in (29) and (30) suggest that they can take any number of arguments. The most natural way of accommodating this observation is by generalising the semantics of the underlying connectives, as follows:

(31) $\vee\varphi_1 \ldots \varphi_n$ is true iff at least one of φ_1, …, φ_n is true.
$\wedge\varphi_1 \ldots \varphi_n$ is true iff all of φ_1, …, φ_n are true.

Both $\vee$ and $\wedge$ now combine with an arbitrary number of sentences, and are interpreted, in effect, as existential and universal quantifiers, respectively. It will be clear that the familiar binary treatment of the connectives comes out as a special case of this.

On this analysis, the logical form of (29) is $\vee$ABC, and (30) is parsed as $\wedge$ABC. Hence, the strongest Q-implicature we are going to have for (29) is that $\textsc{Bel}_S(\neg(\wedge ABC))$, which merely says that, according to the speaker, Wilma is not dating all the men mentioned in the sentence. This is correct as far as it goes, but weaker than a truly exclusive construal. The predictions made by

the substitution method may be correct for binary disjunctions, but they are too weak for longer ones.

As far as I can see, this conclusion is pretty much inevitable, because it only requires the premiss that, in order to deal with long disjunctions and conjunctions, "and" and "or" will have be construed as in (31), and I find this analysis more appealing than any other I can think of. However, it may be instructive to consider what will happen if we revert to the logicians' binary orthodoxy, and postulate that underlying (29) and (30) are multiple occurrences of good-old binary "$\vee$" and "$\wedge$". On this account, the logical form of (29) is either "$(A \vee B) \vee C$" or "$A \vee (B \vee C)$". If this return to logical lore is to make any difference for the issue in hand, we will have to assume, furthermore, that alternatives are computed, not on the basis of the surface form of a sentence, but rather on the basis of its logical form. (This is quite a departure from the original intuition underlying the substitution method, but we'll let it pass.) Then, if (29) is parsed as "$(A \vee B) \vee C$", its alternatives are:

(32) a. $(A \vee B) \wedge C$
b. $(A \wedge B) \vee C$
c. $(A \wedge B) \wedge C$

Whereas if the same sentence is parsed as "$A \vee (B \vee C)$", we get the following alternatives:

(33) a. $A \vee (B \wedge C)$
b. $A \wedge (B \vee C)$
c. $A \wedge (B \wedge C)$

Consequently, we will have different sets of scalar implicatures, depending on how we resolve the scope ambiguity hypothesised for (29). To see where things go wrong, consider the following table, which gives all the possible combinations of truth values for A, B, C:

	a	b	c	d	e	f	g	h
A	1	1	1	1	0	0	0	0
B	1	1	0	0	1	1	0	0
C	1	0	1	0	1	0	1	0

Note that on an exclusive construal of "or", the interpretation of (29) should retain the columns in which just one of A, B, and C is true, and eliminate all others; that is to say, only the shaded columns should remain. The problem is that either set of alternatives will eliminate one column too many: on the first

parse, the alternative in (32b) will give rise to the implicature that $\textsc{Bel}_S(\neg((A \wedge B) \vee C))$, which entails $\textsc{Bel}_S(\neg C)$, thus ruling out column (g); whereas if we adopt the second parse, the alternative in (33a) will yield the implicature that $\textsc{Bel}_S(\neg(A \vee (B \wedge C)))$, which entails $\textsc{Bel}_S(\neg A)$, and therefore rules out column (d). In the former case, (29) is interpreted as "Wilma is dating Albinoni or Boccherini, but not both"; in the latter, as "Wilma is dating Boccherini or Corelli, but not both". Either way, the construal we end up with is too strong.

To summarise, it seems unlikely that the substitution method can provide us with an adequate account of the exclusive readings of long disjunctions. On the most promising analysis of long disjunctions, the readings it delivers are too weak, and if we revert to a binary analysis of "or", which is decidedly less plausible, the predicted readings are too strong.

On reflection, it isn't surprising that the substitution method should run into trouble with long disjunctions. For, in order to arrive at an exclusive interpretation of (29), the following alternatives will have to be excluded:

(34) a. Wilma is dating Albinoni and Boccherini.
b. Wilma is dating Albinoni and Corelli.
c. Wilma is dating Boccherini and Corelli.
d. Wilma is dating Albinoni, Boccherini, and Corelli.

Once we have these alternatives, the exclusive reading of (29) can be derived without further ado. The problem is that only (34d) is derivable from (29) by way of substitution: each of (34a-c) is shorter than (29), and therefore out of reach of the substitution method.

Having framed the issue thus, I would like to return to the ignorance implicatures associated with "or". In the case of (29) these would be of the form, "The speaker doesn't know whether or not Wilma is dating x", where x is either Albinoni, Boccherini, or Corelli. These implicatures are enabled by the following alternatives:

(35) a. Wilma is dating Albinoni.
b. Wilma is dating Boccherini.
c. Wilma is dating Corelli.

As I pointed out at the beginning of this section, the substitution method cannot account for ignorance implicatures because the relevant alternatives are shorter than the uttered sentence. But this is the same problem that long disjunctions give rise to, as we have just seen.

At this point, one begins to suspect that the generative approach is blocking a unified account of exclusivity and ignorance implicatures. And it is. Here is how such an account might go. The core alternatives associated with (29) are (35a-c), but these are not the only alternatives, if we assume (as seems plausible) that sets of alternatives are closed under conjunction: if φ and ψ are alternatives, then "φ and ψ" is an alternative, as well. Therefore, in addition to (35a-c), the sentence's alternatives also include (34a-d). Now the Standard Recipe applies to these alternatives, as always, yielding exclusivity inferences based on the alternatives in (34) and ignorance inferences based on the alternatives in (35). A unified explanation if there ever was one.

To be sure, this story is not quite without its problems. For example, it is natural to suppose that, if sets of alternatives are closed under conjunction, they should be closed under disjunction, as well: if φ and ψ are alternatives, then "φ or ψ" should be an alternative, too. But this we can't allow, because it will get us into the same kind of trouble the substitution method runs into if it adopts a binary analysis of the connectives. However, although it is certainly a cause for concern, I will not address this issue here, because it is only the tip of an iceberg that a large part of Chapter 6 will be devoted to.

I have argued that although the generative approach has no problems explaining the exclusive construal of binary "or", the proposed explanation fails for longer disjunctions. Moreover, even in the binary case, the generative view forces us to adopt a strict dichotomy between exclusivity and ignorance inferences, which become amenable to a unified treatment only by giving up the presupposition that exclusive "or" derives from alternatives generated by the ⟨or, and⟩ scale. More generally, the view that (some) alternatives are generated from Horn scales (or sets) yields a distorted picture of Q-implicatures.

3.4 Conclusion

Frankly, I'm a bit wary of the notion of scalar implicature, because it seems to me that all too often it is intimately connected with the generative view, which I've been criticising in the foregoing. If scalar implicatures are defined as those Q-implicatures which derive from alternatives generated by the substitution method, then I'm against them. However, this is not to say that I don't believe in scalar implicatures tout court. My point is just that from its very inception the notion of scalar implicature has been associated with the wrong sort of theory. My main objection against that theory is that it stands in

the way of a truly unified treatment of Q-implicatures. In the foregoing, I illustrated this with the two types of Q-implicatures associated with "or", but my point is more general: if we adopt the substitution method for dealing with scalar implicatures, we create an undesirable dichotomy between scalar and non-scalar Q-implicatures.

Rather than treating scales as primitives and using them to generate alternatives, as the generative approach tries to do, I believe it is better to take alternatives as primitive, and start from there. If we do so, we can still distinguish between Q-implicatures according to the alternatives they derive from, but such distinctions need not be as profound as what is entailed by the generative approach. A little reflection shows that, lengthwise, there are three ways a sentence can relate to any of its alternatives: it can be shorter, the same length, or longer:

(36) a. Betty read "Ulysses" or "Finnegans Wake".
b. Betty read "Ulysses".

(37) a. Betty read some of these novels.
b. Betty read all of these novels.

(38) a. Betty read "Finnegans Wake".
b. Betty read "Finnegans Wake" and "Ulysses".

Each of the (a) sentences can license the implicature that the speaker doesn't believe that the corresponding (b) sentence is true, and while (37a) and (37b) are roughly the same length, (36a) is longer than (36b), and (38a) is shorter than (38b). These differences may be important, for if an alternative φ is significantly longer than the sentence the speaker chose to utter, one of the reasons why φ was not uttered could be that it wouldn't have been worth the extra effort, in which case the Standard Recipe wouldn't apply (cf. Horn 2000). This being so, the Q-implicature associated with (38a) should be less robust than the ones associated with (36a) or (37a), which seems to be right. See §6.6 for further discussion of this topic.

If we had to define scalar implicatures in purely descriptive terms, we might say that a scalar implicature is a Q-implicature associated with a sentence $\varphi[\alpha]$, where the relevant alternative is of the form $\varphi[\beta]$ and α and β are of roughly the same length, as a consequence of which $\varphi[\alpha]$ and $\varphi[\beta]$ are similar in length, as well. It seems to me that, implicitly or explicitly, many scale theorists would agree with this description, and for the reason given in the last paragraph I believe that scalar implicatures, thus characterised, are an interesting special case of Q-implicature.

Psychological plausibility | 4

For Grice, conversational implicature is a chapter in what he calls "philosophical psychology". I will not attempt to define what he means by this, but it is obvious enough that Grice's theory is deeply imbued with such psychological notions as belief, intention, and so forth. It is somewhat remarkable, therefore, that the Gricean framework is often claimed not to be a psychological one, or not to make "psychological predictions". In a sense, this is true. When taken on its own, Grice's general theory of conversational implicature makes no predictions about what will be observed in everyday life or in the psychology lab, and the Gricean theory of Q-implicatures outlined in the foregoing is in the same boat. For example, though we predict that, under such-and-such circumstances, an utterance of sentence φ will have ψ as an implicature, this does not entail that, under said circumstances, hearers will interpret φ as implying ψ. A Gricean theory of implicature is about inferences hearers are *entitled* to make, not about the inferences they will make.

Nevertheless, it is a mistake to write off the Gricean approach as an idle exercise in philosophy devoid of any bearing on "real" psychology whatsoever. For, given its distinctively psychologistic signature, it is only a small step from a Gricean theory of pragmatics to a theory of experimental pragmatics. All it takes is minimal auxiliary assumptions like the following, for instance:

i. If, according to the Gricean theory, φ implicates that ψ, then hearers will tend to derive ψ as an implicature.

Or alternatively:

ii. If, according to the Gricean theory, φ implicates that ψ, and hearer H actually infers ψ from φ, then H has derived ψ as an implicature.

(*i*) is a much stronger claim than (*ii*), and therefore less likely to be true. But both illustrate my point equally well, which is that it doesn't take that much to make the Gricean theory yield an ample supply of predictions that can be tested in the lab.

So, it's easy enough to interpret Gricean pragmatics as a psychological theory. But is it a *good* theory? My answer to this question is an unqualified yes. It is often said that the Gricean theory of conversational implicature is implausible from a psychological point of view. Even self-confessed Griceans have been careful to dissociate themselves from claims to the effect that Grice's work is a sound basis for a psychological model of pragmatic inference. I'm not one of them. I believe the Gricean approach is eminently plausible in every sort of way, including the psychological.

The purpose of this chapter is twofold. To begin with, I will try to dispel the myth that the Gricean framework is inadequate from a psychological point of view. However, what starts out as defensive enterprise will gradually turn into a more positive argument, as we will be surveying all sorts of experimental data that favour a Gricean approach to experimental pragmatics.

4.1 Charges of psychological inadequacy

Unfortunately, the view that Gricean pragmatics is psychologically implausible has never been argued for in any detail. The following is rather typical: commenting on the Gricean way of deriving conversational implicatures, Wilson (2000: 416) writes:[1]

> It is hard to imagine even adults going through such lengthy chains of inference …

And that's all Wilson has to say about this topic. There may be the rudiments of an argument here (to which we will return shortly), but it is clear that Wilson is voicing a sentiment more than anything else. The following passage is atypical in that it spends more than two lines dismissing the Gricean account of Q-implicature:

> In particular, its implications for processing are less attractive. According to such an account, the inference from the utterance to its scalar implicature goes through a consideration not just of what the speaker said and the context but also of what the speaker might have said but did not. It is this type of

1. The following statement is reiterated (verbatim) by Sperber and Wilson (2002: 8-9).

> onerous inference that makes the Gricean account of implicature derivation seem implausible from a cognitive and developmental point of view. (Noveck and Sperber 2007)

Though less succinct than Wilson's statement, this is hardly more explicit. Like Wilson, Noveck and Sperber merely hint at a possible argument without developing it in any detail.

Basically, some scholars just seem to experience a powerful intuition that the Gricean story makes pragmatic inference too difficult for ordinary mortals to handle. What might cause this feeling? There are five possible reasons I can think of. Looking at what I have called the Standard Recipe for deriving Q-implicatures, I see the following ways in which it might be thought to be implausible from a psychological point of view:

i. *Logical complexity:* The logical structure of the argument is too complex. I take it that this is Wilson's worry.

ii. *Abduction:* Conversational implicatures are abductive inferences, and abduction is hard.

iii. *Incremental processing:* According to Grice, you have to know the meaning of a sentence before you can start deriving conversational implicatures. This premiss sits uneasily with all the experimental evidence showing that utterances are interpreted incrementally, i.e. word by word.

iv. *The intentional stance:* The derivation of a conversational implicature involves reasoning about the speaker's intentional state: what he wants, believes, and so on. Such reasoning is hard, especially for children, and it is improbable that hearers are doing this all the time, as they should if the Gricean view is to be upheld.

v. *Alternatives:* It is psychologically implausible that hearers are constantly monitoring what the speaker could have said, but didn't. This is Noveck and Sperber's point.

In the following, we will consider each of these arguments in turn. The first three will be disposed of rather swiftly, but I will dwell longer on the fourth one and longer still on the fifth. This differential attention reflects my belief that the last two objections are more serious than the others. It is the essence of Grice's view on interpretation that speakers take the intentional stance towards each other, and should this assumption prove to be wrong, then that's the end of the Gricean story. The last objection is directed specifically against the Gricean analysis of Q-implicatures: if there are good empirical reasons for

doubting that hearers take into consideration not only what the speaker says but also what he could have said, then that analysis is of philosophical interest, at best, and irrelevant to the psychology of interpretation.

4.2 Logical complexity

Logical reasoning is one of the oldest topics in cognitive psychology, and there is a wealth of experimental data on how people reason with conditionals, syllogisms, and so on. None of these data indicate that logical reasoning is intrinsically difficult: some argument forms are easy, some are hard. For example, subjects' performance on syllogistic reasoning ranges from near-perfect to abysmal, and the same goes for propositional reasoning: modus ponens is easy, modus tollens is harder, and so on.[2] To illustrate, one of the very easy argument forms is the so-called "Disjunctive Syllogism", of which the following is an instance:

(1)

Fred is not in Antwerp	$\neg\psi$
Fred is in Antwerp or Brussels	$\psi \vee \chi$
Fred is in Brussels	χ

Intuitively, it is pretty evident that this is easy, and psychological experiments confirm that intuition. For example, in a series of studies reported by Braine et al. (1995), participants' performance on Disjunctive Syllogism was essentially perfect, ranging as it did between 94% and 100% correct.

It so happens that the logical form of the Standard Recipe is the same as that of the argument in (1):

reasoning step	*logical form*	
i. Since S didn't say φ, $\neg\textsc{Bel}_S(\varphi)$	$\neg\psi$	(abduction)
ii. $\textsc{Bel}_S(\varphi) \vee \textsc{Bel}_S(\neg\varphi)$	$\psi \vee \chi$	(competence)
iii. $\textsc{Bel}_S(\neg\varphi)$	χ	(from (*i*) and (*ii*))

How hard can this be? It starts with an abductive inference (more about which anon), and then uses Disjunctive Syllogism to get to its conclusion. Contrary to first appearances, the logical structure of the Standard Recipe is perfectly simple.

2. See Evans et al. (1993) for a general overview, and Geurts (2003) and Geurts and van der Slik (2005) for reasoning with non-syllogistic quantifiers.

I suspect that authors who, like Wilson, have complained about the "lengthy chains of inference" required by Gricean reasoning never really tried to go beyond the first appearances. If we are presented with arguments that run to half a page or more, just to explain that "some" implies "not all", it may seem absurd to suppose that this is how we actually do it. If this is what prompted Wilson's complaint, she is mistaken. For one thing, the way an argument is displayed on the page need not reflect how it is implemented by the brain. Verbal descriptions are necessarily sequential, for instance, while the brain is a parallel computer. Therefore, we should be careful not to take such descriptions too literally when considering the psychological plausibility of a theory of interpretation (or any other theory, for that matter). Another problem with Wilson's argument is that it is really too impressionistic. Consider visual perception. What could be easier than seeing a blue train, for example? Practically everybody with normal or corrected-to-normal vision can do it. Even very young children do it with ease. But then textbooks on visual perception invariably run into the hundreds of pages. If these textbooks were right, visual perception should be bafflingly complex. Isn't that enough to prove that they are wrong? Evidently not. What it proves is just that we're unaware of the perceptual feats we perform throughout our waking lives.

When it comes to assessing complexity from a psychological point of view, general impressions count for naught. First, because the complexity of mental processes is not always open to introspection. Secondly, because processes that are easy for us may require descriptions that make them appear much harder. Q-implicatures, and conversational implicatures generally, may well be like this.

4.3 Abduction

In a sense, it is true that abduction is hard. It is form of defeasible inference, and non-trivial defeasible logics tend to be computationally intractable. But even if abduction is hard for logicians and computers, it need not be hard for *us*. After all, we're neither computers nor logical proof systems, and it is plain that in practice most instances of abduction are child's play. Here are a few examples. Upon entering the kitchen, you discover that the light doesn't come on. What to do? More to the point: what *not* to do? There are infinitely many things you don't do. For example, you don't call your electricity supplier to ask if they changed the voltage, or consider the possibility

that some fundamental law of physics has been abolished. Instead, you check the light bulb. Second example: You see a woman entering a bank, and ask yourself how to explain her behaviour. Again, there are indefinitely many candidate explanations, but the most likely one is that she intends to conduct some monetary transaction or other. Last example: You're on a bus driving into town, and the driver announces: "The next stop is Gare de Lyon." Why? The most likely explanation by far is that the next stop is Gare de Lyon, and the driver wants to apprise his passengers of the fact. Any other explanation is much less plausible, and most candidate explanations are wildly improbable.

These are only some of the instances of abductive reasoning that we engage in almost continuously, and it is plain that they are not difficult at all. Indeed, most of them are "irrepressible", to use a phrase of Peirce's (1931). True, we routinely rely on quite a lot of aductive reasoning that is hard, like the diagnostic inferences made for us by doctors, mechanics, and metereologists. But that only shows that it is with abduction as it is with deduction: most of our abductive inferences easy, some are more demanding, and some require special training. In all likelihood, conversational implicatures belong to the easier kind. After all, conversational implicatures are the outcome of our trying to make sense of other people's (linguistic) behaviour, and as a species we don't seem to have problems with that sort of thing (§4.5 will belabour this point). If Gricean reasoning was like diagnosing a brain tumour or a blown head gasket, it would be hard, but fortunately it isn't like that at all.

4.4 Incremental processing

This worry is noticeable especially among researchers working on language processing (e.g., Storto and Tanenhaus 2005, Sedivy 2007), and is prompted by the fact that, according to Grice, conversational implicatures are derived on the basis of what is said, which is taken to imply that implicatures cannot be computed before a sentence is completed: "If Grice is right [...] you need to know the literal meaning or sense of a sentence before you can calculate its implicatures in a context [...]." (Levinson 1983: 117) What is worrying about this is that, as we will see in greater detail in §4.6, there is compelling evidence, introspective as well as experimental, that interpretation is a highly incremental process: we construct the meaning of an utterance on the fly, as the words come in, and it would be wholly remarkable if there were a special class of inferences that could not, as a matter of principle, be drawn before the end of the sentence.

Oddly enough, this worry does not seem to extend to entailments, which should be in the same boat as conversational implicatures, since they, too, can only be drawn from complete sentences. For example, in order to infer from (2) that none of the boys are snoring, we need the entire sentence, obviously:

(2) Only the girls are snoring.

However, it will be equally obvious that we can *begin* to compute entailments long before the speaker has finished his sentence. If an utterance starts with "Only the girls …", we can already infer that whatever predicate is going to fill the dots will not apply to any non-girls. Furthermore, in many contexts it will be pretty clear in advance how a sentence is going to end. If (2) serves to answer the question "Who is snoring?", for instance, it is likely (though not certain) that "Only the girls …" will be followed by "… are snoring", and if he is sufficiently confident that this is how it will end, the hearer can safely derive a full-fledged entailment in mid-sentence.

If entailments can be derived incrementally, the same goes for conversational implicatures. Consider again a context in which one speaker has just raised the question who is snoring, and now another proffers the following answer:

(3) Some of the girls are snoring.

Clearly, in the given context, by the time the speaker has completed the subject phrase of (3), her audience can begin to derive the implicature that not all the girls are snoring. Similarly, if B asks "What are the boys doing?", and A starts answering as follows,

(4) Some of them …

B can begin to derive an implicature of the form "Not all the boys …", without having to wait for A to finish her answer.

Hence, there are two factors that dissipate the apparent tension between conversational implicatures and incremental processing. One is that, often enough, the hearer will be able to guess how a sentence is going to end long before the speaker has completed its delivery. Moreover, even if it is not yet clear how the speaker is going to finish, the hearer may often begin to calculate partial implicatures which are then fleshed out as the speaker proceeds.

While it is true that a full-fledged conversational implicature requires a full-fledged sentence meaning, this doesn't mean that hearers can't process implicatures on the fly, just as there is no reason why entailments can't be computed incrementally. However, just because it is possible, it doesn't mean

that this is what hearers generally do. It has been known since Mitchell and Green (1978) and Just and Carpenter (1980) that the process of interpretation slows down at sentence boundaries: the so-called "sentence wrap-up". Though it remains rather unclear what hearers are doing during these decelerations, it is widely agreed that wrap-up tasks claim a relatively large share of processing effort, which indicates that they are important, and that they are devoted to global rather than local aspects of interpretation. Hence, while it is unlikely that lexical lookup will be deferred until the end of the sentence, it is entirely possible that hearers prefer to compute conversational implicatures (or entailments, for that matter) at the wrap-up stage. It remains to be seen whether or not this is so, but given that a significant part of the interpretation process takes place at the sentence boundary, it would make a lot of practical sense if this was the general strategy.

To sum up, the main take-home points of this section are the following. First, contrary to what has sometimes been suggested, from the fact that conversational implicatures presuppose sentence meaning it doesn't follow that they couldn't be computed incrementally. Secondly, it is quite possible, and consistent with the available data, that implicatures are preferably computed at the wrap-up stage; experimental studies will have to decide whether or not there is such a preference. But be this as it may, the rumour that the very concept of conversational implicature is at odds with the evidence on incremental processing is just that: a rumour.

4.5 The intentional stance

One of the characteristic features of Gricean theories of interpretation is the presupposition that interlocutors are constantly trying to guess what is on each other's minds. In Dennett's (1987) phrase, they require that the hearer routinely adopt the "intentional stance" towards the speaker, trying to figure out what he believes, wants, and so on. Until fairly recently, it was widely accepted that taking the intentional stance is inherently difficult, that it is a skill which takes considerable time to mature, and that it may never become routine even for adults. If this were true, it would seem to imply that the Gricean theory, though perhaps acceptable as an idealised picture, shouldn't be taken too literally.[3]

3. See Breheny (2006) for an excellent discussion of the various positions. The notion that we can take the intentional stance towards other beings, be they human or not, is often

It may be that, pre-theoretically, it is not implausible that the intentional stance should be difficult to apply, and this intuitive idea gained support from early experiments with the famous "false belief task" pioneered by Wimmer and Perner (1983), who employed the following scenario. A little Austrian boy named "Maxi" puts a chocolate bar into a blue cupboard, and then exits the scene. In his absence, Maxi's mother replaces the chocolate from the blue cupboard to a green one. Now the critical question is, "Where will Maxi look for the chocolate when he gets back?" Since it is essential to the concept of belief that beliefs can be out of step with reality, an accomplished mind reader should know that Maxi will be heading for the blue cupboard, where he put the chocolate before leaving. By this criterion, Wimmer and Perner's child subjects were not accomplished mind readers: in the two studies they reported, none of the 3-4-year olds and a mere 57% of the 4-6-year olds pointed at the right cupboard, and it was only the 6-9-year-old children who showed no great difficulties with the task (86% correct).

For some time, Wimmer and Perner's findings were widely interpreted as showing that young children lack a mature concept of belief, and on the strength of Baron-Cohen et al.'s (1985) findings it seemed evident that autistic patients tend to suffer from the same handicap. In the meantime, however, the general mood has clearly changed. To be sure, it was never clear-cut that failure on the false-belief task has anything to do with belief (or the concept of belief) per se, since there are other reasons why subjects may fail this task. It may be, for example, that young children master the concept of belief, but lack the linguistic competence to report on what they know, or that the false-belief task is too complex for other reasons. These suspicions were vindicated by several studies reported over the years, culminating in Onishi and Baillargeon's (2005) demonstration that 15-month-old infants can pass a non-verbal version of the false-belief task (cf. also Southgate et al. 2007, Surian et al. 2007).

Over the past few decades, it has become increasingly probable that people *are* constantly reasoning about each other's intentional states, that they are very good at the game, and that children start playing it at a very early age. I will quickly go through a handful studies to illustrate how early children start developing these skills:

---

expressed by saying that we have a "Theory of Mind". I consider this extremely poor terminology, and will therefore avoid it.

- When they are about nine months of age, babies begin following their parents' line of regard, monitoring what they are attending to. For example, if an adult introduces a new word to name an object in front of him, the child will realise that the new word applies to the object the adult is looking at, not the one the child happens to be looking at himself (Baldwin 1991, 1993).
- Furthermore, there is evidence that even at a tender age, attention following is a quite sophisticated art. For example, 12-months-old babies follow the "gaze" of a faceless robot, which is taken to project off the front end, but they will not do so if the robot fails to interact with them in a meaningful way (Johnson et al. 1998).
- It has been shown by Liszkowski et al. (2006, 2008) that 12-month-old infants are able and willing to help an adult who is pretending to have lost a staple remover (say) by pointing at the "missing object".

The last example is particularly interesting. Not only does it show that even before they start speaking, children can assess other people's needs, it also proves that already at a very early age, children are capable of the cooperative behaviour that is essential to Gricean pragmatics.

None of this is to suggest, of course, that infants are accomplished mind readers. It is generally agreed that mind-reading skills take time to mature, although there is considerable disagreement over the course of development and the time it takes. It seems safe to suppose, however, that it will take at least a number of years before a child fully masters the intentional stance. If this much is right, we should expect that pre-schoolers are less adept at deriving conversational implicatures. Since Noveck's (2001) pioneering study, which was followed up by many others, we know that this is indeed the case (Papafragou and Musolino 2003, Guasti et al. 2005, Pouscoulous et al. 2007). Also pertinent in this connection is Pijnacker et al.'s (2009) finding that autistic patients, who typically have difficulties with the intentional stance, may derive scalar implicatures at lower rates than normal controls.

To return to our main theme, what can we say about the objection that, since reasoning about intentions is inherently difficult, it is unlikely that hearers are doing this all the time, as they should if Grice was right? Up to a point, this objection is valid. It is not unlikely that taking the intentional stance is hard, as even our closest primate relatives are not nearly as good at it as we are. But on the other hand, there can be no doubt anymore that our species has evolved to develop sophisticated mind-reading skills in a matter

of years and that we go on honing these skills in our social interactions. In short: instead of undermining it, psychological research on the attribution of intentionality strongly supports the Gricean view.

Before closing this section, I would like to comment on an objection that is sometimes raised against the notion that speakers are constantly engaged in reasoning about each other's intentional states. The objection is that it might well be that, in reality, people don't have a particularly deep understanding of each other's beliefs and desires, nothing like a genuine "Theory of Mind" (cf. note 3); they merely *seem* to be taking the intentional stance. In fact, so the objection goes, people have an arsenal of quick-and-dirty heuristics that allow them to simulate the intentional stance. I wouldn't be too surprised if something like this turned out to be the case. However, I don't really see what the issue is supposed to be about. Even if people are merely pretending to be taking the intentional stance, there can be no doubt that they are exceptionally good at it. Now, if you are really good at pretending to be a good driver, you *are* a good driver. Similarly, if you are really good at pretending to be a Gricean reasoner, you are one. Put otherwise, it doesn't really matter how the intentional stance is implemented, as long is it is agreed that it is implemented some way or other. And this much is not controversial anymore.

4.6 Alternatives

It is one thing to assume that hearers reason about the speaker's beliefs and desires all the time. It is another thing to claim that, on top of that, hearers take into account alternative expressions the speaker might have used but didn't. According to Noveck and Sperber, it is implausible to suppose that hearers constantly keep track of what the speaker might have said, and on the face of it this worry seems reasonable enough. There are so many things a speaker can say at any given point, how could a hearer ever hope to cut down this plethora of alternatives to a set of manageable proportions? Doesn't it appear highly probable that this would make the hearer's job far too onerous, as Noveck and Sperber say? In my view, this is the most serious threat to the Gricean programme for Q-implicatures, and I will discuss it at some length.

As mentioned already in §4.4, there is a considerable body of experimental evidence showing that hearers have strong expectations about what the speaker will be saying next, and that these expectations are recalibrated continuously and *very* quickly. To get an intuitive feel for the remarkable feats

hearers perform non-stop, consider the following passage from Douglas Adams' "Hitchhiker's guide to the Galaxy", from which I have elided an adjective and five nouns:

> At eight o'clock on Thursday morning Arthur didn't feel very [...]. He woke up blearily, got up, wandered blearily round his [...], opened a [...], saw a [...], found his [...], and stomped off to the [...] to wash.

Although it may be hard to guess which exact expressions are missing from this passage, it will come as no surprise that the first one is "good", the second "room", and the third "window". But it *will* be surprising that the fourth missing word is "bulldozer". It should be intuitively clear that these judgments come quite naturally. Experimental data confirm this intuition, and also show that these judgments are extraordinarily fast. It wouldn't be unusual for a hearer to detect the oddity in Adams' orginal passage halfway through the ill-fitting word.

In order to make this a bit more precise, let us have a brief look at processing studies making use of event-related brain potentials (ERPs), an average measure of electroencephalogram (EEG) activity associated with designated events, that is, in the present context, the occurrence of critical words (like "bulldozer" in Adams' text). ERPs provide a direct record of neuronal activity, with virtually no delay, and therefore they can be used to track various aspects of comprehension with high temporal accuracy.

An ERP consists of multiple waveforms superimposed on one another. One of these is the N400, which was first observed by Kutas and Hillyard (1980). It is a negative-going potential with an onset at around 200 milliseconds and a peak at approximately 400 milliseconds after the critical event (whence the N400's name). Although the precise interpretation of the N400 has not been settled yet, it may roughly be characterised as a measure of the *surprise value* of a given event. For example, Kutas and Hillyard found a strong N400 associated with the last word in "He spread his warm bread with socks", which is caused by the fact that "socks" is highly unexpected in this context.

While in Kutas and Hillyard's experiment, the unexpectedness of the critical word was rather glaring, subsequent studies showed that the N400 registers much milder degrees of surprise, as well. For example, when Hagoort and Brown (1994) presented subjects with sentences like (5), they found that "pocket" elicits a stronger N400 than "mouth", which agrees with the intuition that the first word, though perfectly coherent, is less expected in this context:

(5) Jenny put the sweet in her {pocket/mouth} after the lesson.

It will be clear that hearers' expectations are strongly influenced by the discourse context. In general, a sentence like "The peanut was in love", would yield a massive N400 for the last word. But what will happen when this sentence is embedded in a discourse like the following:

> A woman saw a dancing peanut who had a big smile on his face. The peanut was singing about a girl he had just met. And judging from the song, the peanut was totally crazy about her. The woman thought it was really cute to see the peanut singing and dancing like that. The peanut was {*salted*/*in love*}, and by the sound of it, this was definitely mutual. He was seeing a little almond. (Nieuwland and van Berkum 2006: 1106)

As reported by Nieuwland and van Berkum, in this context it is "salted", not "in love", that elicits an N400. Apparently, the expectations tracked by the N400 take into account contextual information, as well as general world knowledge.

One of the most salient components of the context is the speaker, and van Berkum et al. (2008) showed that the N400 takes into account information about the speaker, too. In van Berkum et al.'s study, people listened to sentences whose meanings might be at odds with reasonable assumptions about the speaker's sex, age, or social status, as inferred from the speaker's voice. For example, when "I have a large tattoo on my back" was spoken in an upper-class accent, there was an obvious discrepancy between social elevation and dorsal decoration, which became apparent at the critical word "tattoo". This type of mismatch, too, turned out to yield an N400.

What these ERP studies show, then, is that hearers check incoming utterances against their expectations practically in real time, word by word. For all intents and purposes, it is established immediately whether or not any given word is unexpected at the point at which it occurs. But this is to say that the hearer must be considering alternative expressions all of the time. Moreover, doing so doesn't seem to be onerous, at all. If the processes underlying the N400 were effortful, they should be relatively slow, and they are anything but. The N400 *peaks* 400 milliseconds after the onset of the critical word, which means that, in many cases, unexpected words are recognised as such well before their delivery is complete.

These results refute Noveck and Sperber's claim that it is psychologically implausible to suppose, as the Gricean account of Q-implicature must, that hearers routinely take into account what the speaker could have said but didn't. On the contrary, the ERP literature suggests that this is precisely what hearers do; if anything, the ERP data argue in favour of the Gricean view.

However, although in general terms ERP data are in agreement with the Gricean approach, they don't show that hearers entertain the sort of alternatives that, on the Gricean view, are instrumental in the derivation of Q-implicatures. What we need is experimental evidence qualitatively similar to that reviewed above, which bears on the question whether hearers routinely take into account alternatives that are more or less informative than what the speaker actually said. Such evidence has been provided by Altmann and Steedman (1988), among others, who presented their subjects with sentences like (7a,b) in the context of either (6a) or (6b):

(6) A burglar broke into a bank carrying some dynamite. He planned to blow open a safe. Once inside, he saw that there was ...
   a. a safe with a new lock and a safe with an old lock.
   b. a safe with a new lock and a strongbox with an old lock.

(7) The burglar blew open the safe with ...
   a. the dynamite and made off with the loot.
   b. the new lock and made off with the loot.

In the context of (6a), (7a) took longer to read than (7b), which is not particularly remarkable, since the first sentence contains a definite noun phrase, "the safe", whose intended referent is not resolved by the context. What is surprising is that in the context of (6b), where both sentences are felicitous, and had equally long response latencies, the reading times for (7b) were longer than in the first context. Apparently, participants were slowed down by the fact that, in the second context, the noun phrase "the safe with the new lock" was overly informative, or in other words: they expected a shorter and less informative expression instead.

While Altmann and Steedman's study shows that hearers monitor the level of specificity of expressions used by the speaker, Sedivy et al. (1999) report data suggesting that preferred interpretations are guided, inter alia, by information of this sort. Using an eye movement paradigm, Sedivy et al. presented their subjects with sentences like (8), while showing them displays in which there was more than one object matching the adjective. For example, in the case of (8), the display would contain a tall glass and a tall pitcher:

(8) Pick up the tall glass.

As in other studies, the subjects in Sedivy et al.'s experiment might look at either of the two objects upon hearing the adjective, and only fastened on the target object when hearing the noun. Interestingly, however, Sedivy et al.

showed that a bias for the glass could be induced simply by adding a short glass to the display. That is to say, if the display contained a short glass in addition to the tall glass and the pitcher, hearers were more prone to look at the target glass upon hearing the adjective "tall".[4]

The most natural explanation of these findings is a Gricean one: upon hearing the adjective "tall", hearers ask themselves why the speaker should want to use this word, and find a plausible answer in the display before them: given the presence of a short glass, the adjective might help to select the tall one. This is Gricean reasoning on a micro-scale, and again it is extremely fast, since the preference for looking at the tall glass begins to emerge even before the onset of the noun "glass".

These effects are Q-based in the sense that, apparently, the participants in Sedivy et al.'s experiment expected that the critical adjective wasn't being thrown in for ornamentation, but was meant to serve a definite purpose. However, following up on the discussion in §1.6, I should like to stress that I'm not calling this a "Q-implicature", the reason being not that the kind of reasoning is qualitatively different (which it isn't), but rather that its business is different: whereas Q-implicatures, as the term is commonly understood, are ampliative inferences deriving from what the speaker has said, the Q-based reasoning observed in Sedivy et al.'s experiment is involved in determining the intended referents of definite noun phrases like "the tall glass". Same sort of reasoning, different purpose.

The main lesson to be drawn from the experiments reviewed in the foregoing is that, rather than arguing against the notion that hearers routinely track speakers' alternatives, processing data suggest that this is precisely what hearers do. The experimental evidence shows that hearers have strong expectations about what the speaker is going to say next, and that these expectations are checked and readjusted on a word-by-word basis. Moreover, hearers monitor the level of specificity of expressions used by the speaker, and use it to guide their expectations about what the speaker wants to say.

4.7 Conclusion

In this chapter, I have argued that there are no good reasons for believing that the Gricean theory of implicature is inadequate from a psychological point of view. Very briefly, the main points were that the logic of Gricean

4. For further discussion of these and related studies, see Sedivy (2003, 2007).

reasoning, especially the Standard Recipe, is not particularly complex; that the kind of abductive inferences involved are not particularly hard; that the experimental evidence on incremental processing is consistent with the Gricean approach; that it is plausible that interlocutors are always trying to read each others' minds; and that it is plausible that, while doing so, they are taking into account what the other could have said but didn't.

While the main concern of this chapter was to defuse potential objections against Gricean theories of interpretation, I believe we can say that the experimental evidence does more than that: it offers strong support for a Gricean approach to the psychology of interpretation. In particular, Grice's key insight, that interpretation crucially involves taking the intentional stance, has been dramatically vindicated by several decades' worth of experimental research.

Nonce inferences or defaults?

5

In §1.4, we reviewed Grice's distinction between generalised and particularised implicatures. Grice considers an implicature generalised if it would "normally", that is, "in the absence of special circumstances", be associated with "a certain form of words" (Grice 1975/1989: 37; see p. 18 for a longer quote). Particularised implicatures, on the other hand, are contingent on specific features of the context. Though the distinction between generalised and particularised implicatures may have been of strategic importance to Grice, he didn't see them as separate species: according to him, both generalised and particularised implicatures spring from the same pragmatic sources; it is just that, due to extraneous circumstances, the former happen to arise "normally". Hence, an obvious question to ask is why this should happen to be so: why should implicatures be standardly associated, or at least *seem* to be associated, with some expressions but not with others? This is one of the questions to be addressed in the following.

As Grice observed, from generalised implicatures it is but a treacherously small step to the notion that implicatures may be *conventionally* associated with certain expressions.[1] If the use of "some" normally gives rise to a "not all" inference, it is tempting to suppose that "not all" is part of the lexical content of "some". Of course, one must be careful not to claim that "not all" is part of the lexical *meaning* of "some", for after all it is cancellable. But it does seem attractive to asssume that the word "some" comes with two types of lexical content: a lower-bounding truth-conditional component

1. If we consider the possibility that implicatures may become conventionally associated with certain expressions, the question arises what is the exact nature of the association, a question that has been discussed in a different but related context (i.e., for indirect speech acts) by Searle (1975) and Morgan (1978), among others. This is a rather intricate discussion, but as far as I can tell it is not of central importance to our topic of discourse.

("at least some"), which cannot be cancelled, and an upper-bounding non-truth-conditional component ("not all"), which is cancellable, as the following example reminds us:

(1) Some of the sheep have the flu. In fact, ...
   a. all of them do. [*cancellation of the upper bound*]
   b. *none of them do. [*attempted cancellation of the lower bound*]

Seen this way, scalar inferences are *defaults*: "some" will imply "not all" unless special circumstances indicate otherwise and the inference is cancelled. It bears emphasising that, on this view, scalar inferences are not conversational implicatures; rather, they are conventionally associated with certain words. (I will use "scalar inference" as a neutral term for referring to what I take to be scalar implicatures.)

Whatever the merits of this line of thinking may be, it should be noted that it doesn't apply across the board. It will be evident, for instance, that inferences like the following are not generated in the lexicon:

(2) A: What did you have for breakfast?
   B: A banana.
      $\rightsquigarrow$ $\textsc{Bel}_{\textsc{B}}$(B only had a banana)

This is readily explained as a Q-implicature, but it surely does not spring from the lexical content of "a banana". Hence, the most promising target of the default view are scalar expressions, and more specifically, the strong scalar inferences associated with such expressions:

(3) Some of the sheep have the flu.
   $\rightsquigarrow$ $\neg\textsc{Bel}_{\textsc{S}}$(all the sheep have the flu) [*weak scalar inference*]
   $\rightsquigarrow$ $\textsc{Bel}_{\textsc{S}}(\neg$(all the sheep have the flu)) [*strong scalar inference*]

If the word "some" comes pre-packaged with a "not all" inference, the default interpretation of (3) would imply that, according to the speaker, not all the sheep have the flu. Hence, the focus of this chapter will be on the strong upper-bounding inferences associated with scalar expressions, though other types of Q-implicature will be touched upon, as well.

In the following, I will distinguish two flavours of defaultism, one stronger than the other, and argue that both are wrong, because scalar inferences are context dependent to such a degree that they can't be defaults in either sense. That is to say, I will argue for a contextualist approach to scalar inferences in particular, and Q-implicatures in general. However, even if defaultism is on

the wrong track, it has to be conceded that the defaultist view has considerable intuitive appeal, and one of my objectives in this chapter is to explain what causes this intuition.

5.1 True defaults

The key concept in this chapter is that of a default inference, and before I start arguing that scalar inferences aren't defaults, I want to develop this notion a bit. The concept of default wasn't invented by pragmatic theorists. Rather, they borrowed it from their colleagues in computer science and cognitive psychology. In these disciplines, a default is an assumption which is taken to hold unless there are specific reasons to suppose that it doesn't. Defaults, thus understood, are ubiquitous in everyday life. When you enter a hospital, you expect to see your typical hospital furniture, as well doctors, nurses, and so on, although this expectation may be disappointed in infinitely many ways: it may turn out that the hospital is being renovated, that the building was converted to a school, or (less likely perhaps) is filled to the roof with chocolate ice cream. When you kick a football, you expect it to fly off, although it may turn out (to your chagrin) to be filled with chocolate ice cream. And so on. Such default expectations are mostly unreflected, automatic, and correct in a sufficiently large number of cases. If every second book was a blank, it wouldn't be long before we gave up the default assumption that books contain pages of print.

Default expectations are essential to everyday reasoning, and since interpretation is to a large extent a matter of everyday reasoning, it will not come as a surprise that they show up in many processing studies. Here, I will briefly discuss two experiments by Carreiras et al. (1996), who studied the effects of gender stereotypes on online interpretation. The gender stereotype associated with the noun "carpenter", for example, is that carpenters are male; i.e., carpenters are male by default. Similarly, nurses are stereotypically female.[2] In one of their experiments, Carreiras et al. presented participants with discourses like the following:

2. Atlas and Levinson (1981) classify such inferences as conversational implicatures, but since it is not clear to me how they help to justify the assumption that the speaker is being cooperative, I'm not convinced that this is right.

(4) a. *Not stereotyped for gender:*
The psychologist studied the students' test results.
{He/She} was concerned about the low standard.
b. *Stereotyped for gender:*
The babysitter settled down to watch a video.
Then {he/she} heard the baby crying.

Carreiras et al. assumed, on the basis of data collected in a pre-test, that whereas the noun "psychologist" is not stereotyped for gender, babysitters are female by default. Therefore, it was expected that, in the first case, reading times for the feminine and masculine pronouns would be in the same range, while there should be a delay for "he", as compared to "she", in the second case. In this condition, the noun "babysitter" evokes the default inference that the person in question will be female, and then the pronoun either confirms or overrides this expectation, which should make "she" easier to process than "he". This is exactly what Carreiras et al. found: while the second sentence in a discourse like (4b) took significantly longer to read with "he" than with "she", no reliable difference between reading times was observed in discourses like (4a).

This finding shows that it takes effort to overwrite a default assumption once it has been activated. But what if the noun occurs at a point where the inference is preempted already? This question was addressed in a second study, which was conducted in Spanish, and exploited the fact that the Spanish articles are marked for gender. Here is a sample item:

(5) {El carpintero/La carpintera} tomó las medidas para hacer el armario.
"The carpenter [masc./fem.] took measurements to make the cupboard."

In these cases, the carpenter's gender is established even before the noun is encountered. Nevertheless, Carreiras et al. still found a strong stereotype effect: with "la carpintera" the sentence took considerably longer to read than with "el carpintero".

Both experiments show that it takes time, and hence effort, to cancel defaults. The second study shows, moreover, that even if the context settles the matter before the default could possibly arise, it still takes effort to cancel the default inference.

This is how genuine defaults manifest themselves. They are fast and reflex-like and, as Carreiras et al.'s experiments indicate, they can be quite powerful. If scalar inferences are lexical defaults, it is only natural to suppose that they should behave similarly. Therefore, the question arises: Are scalar inferences anything like this? It will soon become clear that they are not.

5.2 Strong defaultism

I will use the term "strong defaultism" to refer to the position that scalar inferences are like the gender stereotypes studied by Carreiras et al. On this view, a scalar inference is triggered by a word and the triggering process is fast: since "not all" is basically part of the lexical content of "some", it becomes available as soon as "some" is retrieved from the mental lexicon.

One problem with the doctrine of strong defaultism is that there has been, and maybe still is, confusion over who its proponents are. When, around the turn of the millennium, the first experimental studies on scalar implicature started to appear, it was said that "neo-Griceans" are committed to strong defaultism, but such claims are misleading at best: not all neo-Griceans are strong defaultists, and in particular the author of the neo-Gricean *Urtext*, Horn (1972), has never even come near this view.[3] The only neo-Gricean who is committed to strong defaultism is Levinson (2000), though even in his case it is hard to be completely sure, because Levinson's views on conversational implicature are neither univocal nor consistent. However, it is evident that strong defaultism is one of the main threads in his thinking, for he argues at length that what he (misleadingly) calls generalised conversational implicatures enhance the efficiency and speed of language processing, which would make little sense if he thought they are laborious and time consuming. Furthermore, Levinson sees such inferences as

> [...] powerful heuristics that give us preferred interpretations without too much calculation of such matters as speakers' intentions, encyclopedic knowledge of the domain being talked about, or calculations of others' mental processes. (p. 4)

Yet another clue is that Levinson puts scalar inferences on a par with inferences whose default status is indubitable, like the gender stereotypes discussed in the previous section, for example. Hence, there can be little doubt that Levinson is a strong defaultist.

3. Bott and Noveck (2004: 438) say that what they call "a strong default approach" has been defended "to some extent" by Horn (1984), implying that Horn's views are closely related to Levinson's. However, although this may have been so in 1984, Levinson went on to espouse a stronger and blatantly psychologistic version of defaultism, which Horn never subscribed to. (Bott and Noveck are not overly explicit about what they mean by "a strong default approach"; the notion is not clearly separated from what I call "weak defaultism".)

Other authors are equally outspoken about their allegiance to the strong defaultist view:

> Our experimental hypothesis amounts to claiming that the implicated content that is normally associated with uses of the disjunction *or* does not differ much from lexical content. Like lexical content, the exclusive meaning of *or* should be "closely tied" to the utterance of this lexical item and become available as soon as the disjunction is heard. (Storto and Tanenhaus 2005: 433)

This quote is interesting not only because of its admirable candour, but also because it exemplifies the natural affinity between strong defaultism and the notion that scalar inferences are "closely tied" to certain words. (It isn't clear to me why the authors decided to enclose this phrase in quotation marks.)

The main tenet of strong defaultism is about processing. The guiding intuition is that scalar inferences aren't really inferences; as Storto and Tanenhaus say, the idea is that scalar inferences are retrieved with the lexical content of a word; that's why they are defaults. In the past five years, a series of experimental studies have been conducted to test this theory; in the remainder of this section, I will discuss what I take to be the most important ones: Bott and Noveck (2004) and Breheny et al. (2006).[4]

Bott and Noveck (2004) One of the landmark papers in the experimental literature on scalar inferences is Bott and Noveck's (2004). In one of the experiments reported in that paper, Bott and Noveck presented participants with sentences like the following, and measured the time participants took to decide that these sentences were true or false.

(6) a. Some elephants are mammals.
b. Some mammals are elephants.
c. Some elephants are insects.

(7) a. All elephants are mammals.
b. All mammals are elephants.
c. All elephants are insects.

The critical item in this experiment was (6a); the others were controls. If a hearer infers from (6a) that not all elephants are mammals, thus obtaining

4. Other experimental studies that are relevant here include Bezuidenhout and Cutting (2002), Noveck and Posada (2003), Feeney et al. (2004), Storto and Tanenhaus (2005), Huang and Snedeker (2006), Chemla (2007), De Neys and Schaeken (2007), and Pijnacker et al. (2009).

an upper-bounded reading, he should reject the statement. Furthermore, if scalar inferences are strong defaults, then rejecting the sentence should take less time than accepting it: according to the default view, a hearer who accepts the sentence first has to cancel the scalar inference in order to obtain an interpretation without an upper bound. On the other hand, if scalar inferences lag behind the computation of truth-conditional meaning, we should see the opposite pattern: rejecting the sentence should take longer than accepting it.

In the control conditions, Bott and Noveck found that in a clear majority of cases (over 90% on average), participants gave the expected responses: (6b) was accepted 89% of the time, (6c) was rejected 93% of the time, and so on. By contrast, the mean acceptance rate for critical sentences like (6a) was a lot closer to chance at 41%. Furthermore, most of the participants in this study (72%, to be exact) failed to be consistent in their responses to the critical sentences. Obviously, these results sit uneasily with *any* form of defaultism, but since our chief concern in this section is with processing, I will not belabour this point now (see §5.3 for discussion).

Turning to Bott and Noveck's main finding, they report that the average "false" response for (6a) was significantly slower than all the other responses in this experiment, including the average "true" response for the same item. In all cases except (7a), the difference was well over half a second. Given that the average response time, which included reading the sentence, was on the order of 2.5 seconds, this is a considerable delay. Hence, the pattern of response latencies observed in this experiment disconfirms the claim that scalar inferences are strong defaults.[5]

A further experiment by Bott and Noveck used the same materials and basically the same design. The main difference with the previous study was that in this version the time available for responding was controlled by the experimenters, and was either long (3 seconds) or short (900 milliseconds). This difference didn't have an effect on the control items, but it did have a significant effect on the critical items: with the short time lag, participants responded affirmatively in 72% of the cases; with the long time lag, this rate went down to 56%. Bott and Noveck interpret this result as confirming that scalar inferences take time to derive, hence as arguing against strong defaultism.

5. Bott and Noveck's findings have been replicated in many studies, including Noveck and Posada (2003), Huang and Snedeker (2006), De Neys and Schaeken (2007), and Pijnacker et al. (2009). The odd one out is Feeney et al.'s (2004) study, which, for reasons that remain unclear, failed to reproduce Bott and Noveck's results.

From a logical point of view, Bott and Noveck's materials contained two different types of sentences, universal and existential, which between them gave rise to three types of construal:

All A's are B's.	[*universal*]
Some A's are B's.	[*existential without upper bound*]
Some A's are B's but not all A's are B's.	[*existential with upper bound*]

Clearly, the upper-bounded construal is more complex than the other two, but it is important to note that this alone does not suffice to explain Bott and Noveck's data. On the strong defaultist view, an upper-bounded interpretation is the default for *all* existential sentences, including (6b) and (6c), but these didn't slow down participants as the upper-bounded interpretations of (6a) did. Though it is tempting to assume that this is because (6b) and (6c) are true (on either interpretation), while the response associated with the upper-bounded interpretation of (6a) is "false", there is no independent evidence that, in this experiment, positive responses were faster than negative responses. In fact, with the universal sentences, "true" responses were markedly slower than "false" responses, although Bott and Noveck don't report whether or not this was statistically significant. But anyway, the mere fact that upper-bounded construals are more complex does not explain the experimental data.

A potential worry about these experiments is the following. Bott and Noveck asked their participants to decide, for any given sentence, if it was "true" or not. But what does it *mean* for a sentence to be true? Many philosophers would say that conversational implicatures have no bearing on the truth value of a sentence. After all, a conversational implicature is calculated on the basis of a sentence with a certain meaning, and the truth value of a sentence is determined by its meaning; any conversational implicatures derived from it are immaterial as far as its truth value is concerned. We might call this the "Fregean" view on truth. If the Fregean view is right, then a sentence like (6a) is true regardless whether a scalar implicature is derived or not.

As a matter of fact, only 41% of Bott and Noveck's participants said that (6a) is true, so it was surely not the case that all of them held the Fregean view on truth. Could it be the case, then, that participants disagreed on the notion? Could it be that, whereas the "true" responses were given by people endorsing a Fregean notion of truth, the "false" responses betrayed a stricter construal, according to which the truth value of a sentence is determined not only by its meaning but also by such conversational implicatures as it gives rise to? If this is on the right track, Bott and Noveck's results tell us very

little about how implicatures are processed. What they show, rather, is that people disagree on the meaning of the word "true".

There are at least three reasons for rejecting this explanation. One is that, as I argued in §1.1, it is doubtful that untutored informants have a Fregean notion of truth. Secondly, it should be borne in mind that, according to the view under discussion, scalar inferences aren't conversational implicatures, but rather aspects of the sentence's conventional content. And thirdly, if this explanation was on the right track, people should become less likely to adopt the Fregean notion of truth when given more time to think, which seems rather improbable to me.

All things considered, the explanation proposed by Bott and Noveck is better than any other I can think of, so I tend to accept that their data argue against strong defaultism. Nevertheless, it is not an open-and-shut case, for it can hardly be doubted that upper-bounded construals are more complex than others, and it may well be that participants' interpretation of the word "true" is less straightforward than Bott and Noveck supposed. Still, I don't see how these considerations could provide us with an alternative story that gets strong defaultism off the hook.

Breheny et al. (2006) There is a another concern about Bott and Noveck's studies: it is that their design is not very natural. Presumably, the question whether some elephants are mammals is not very likely to arise in everyday life. True, experimental conditions are artificial by definition, but the less so the better. In this respect, Breheny et al.'s (2006) experiments 2 and 3 clearly constitute an improvement; I will review the former here. Consider the following sentences:

(8) a. Some of the consultants had a meeting with the director.
    b. Only some of the consultants had a meeting with the director.

(9) a. The director had a meeting with some of the consultants.
    b. The director had a meeting with only some of the consultants.

Suppose that "some" in (8) triggers a scalar inference to the effect that not all of the consultants attended the meeting. Then the content conveyed by (8a) and (8b) is the same; the only difference is that the upper bound on the interpretation of "some" is cancellable in the former case but not in the latter, where it is made explicit by means of "only". The same, mutatis mutandis, for (9a) and (9b). Now let an utterance of (8a) or (8b) be followed by:

(10) The rest did not manage to attend.

In either case, the anaphor "the rest" refers to the consultants that didn't attend the meeting. If scalar inferences are triggered by default, it is reasonable to expect that this anaphoric link is equally facilitated by (8a) and (8b); for, on this view, both sentences imply that there were consultants who didn't attend the meeting. Again, the same should hold, mutatis mutandis, for (9a) and (9b), though in these cases establishing an anaphoric link might be faster than for (8a) and (8b), because here the anaphor is closer to its antecedent. In a reading time study, Breheny et al. found that these predictions are confirmed for the sentences in (8) but not for the ones in (9): while in the context of (8a) and (8b), the reading times for (10) were the same, this sentence took significantly longer to read when it was preceded by (9a) than by (9b). Hence, (9a) falsifies the claim that scalar implicatures are strong defaults.

Still, one might be tempted to say, the default view is confirmed by (8a). However, this cannot be right. If scalar implicatures are defaults, they should always be triggered: that's what it means to be default. So how can the contrast between (8a) and (9a) be accounted for? Breheny et al. propose the following contextualist explanation. Utterances are always interpreted in context, even when they are presented "in isolation", as the sentences in (8) and (9) were. If a sentence appears out of the blue, hearers will try to construct a context in which it might plausibly have been uttered, and in doing so they will take into account such cues as are provided *by the sentence itself*. For instance, when a sentence like

(11) The prince was snoring.

is presented in an experimental setting, participants will conjure up a context in which some prince is part of the common ground between speaker and hearer.

In (11), it is the definite noun phrase that serves as a cue to what kind of context the sentence might be uttered in. In the sentences used by Breheny et al., the relevant cue is part of speech. While in (8a), "some of the consultants" serves as subject, in (9a) it is the object of the sentence, and since subjecthood is correlated with topicality, a suitable context for (8a) would be one in which the consultants are the discourse topic, while (9a) will tend to evoke a context in which the director is the main character. Put otherwise, the consultants and their exploits are more *relevant* in the former context than in the latter, and therefore (8a) is more likely to give rise to the question whether or not all consultants were at the meeting, and thus to the implicature that some of the consultants did not attend.

In this way, a contextualist explanation can be given for the contrast between (8a) and (9a). It may be that the proposed explanation is vaguely paradoxical, in that it asks us to adopt a contextualist stance on sentences occurring in isolation. But the air of paradox is dispelled by the observation that sentences contain context indicators, which provide the hearer (or reader) with clues as to what sort of context the sentence might be uttered in.

There are two important general lessons to be learned from this special case. One is that, if an expression α is regularly consorted by an inference β, it need not be the case that α triggers β by default; only the converse is true. The second is that, contrary to what one might expect, contextualist theories do have the wherewithal for explaining data that, prima facie, suggest that implicatures are defaults. When considered on its own, (8a) could easily be mistaken for proof that "some" implies "not all" by default. However, the contrast with (9a) shows that this is not right, and moreover, we can explain on strictly contextualist principles why there *seems* to be a default inference in the case of (8a).

Epitaph on strong defaultism In §5.1, we saw how strongly processes of interpretation can be affected by true defaults (in this case, defaults created by gender stereotypes). In the meantime, it should have become clear that scalar inferences aren't like this at all. Between them, the experimental data procured by Bott and Noveck, Breheny et al., and others (cf. note 5) render it highly unlikely that scalar inferences are defaults in the true, strong sense of the word. By the same token, they argue against theories claiming that the upper-bounding inferences associated with scalar expressions are hard-wired into the lexicon. For if they were, they should become available as soon as a word is retrieved from the lexicon, and experimental data suggest that they don't.

5.3 Weak defaultism

Motivating weak defaultism Having dealt with strong defaultism, let us now turn to weak defaultism, which is the view that scalar expressions give rise to upper-bounding inferences as a matter of course; it is what will happen "normally". This description is admittedly vague, but then the doctrine itself isn't always easy to nail down, as the following passage illustrates:

> By default interpretation, I simply mean the one that most people would give in circumstances in which the context is unbiased one way or the other. (Chierchia 2004: 51)

Since *every* context is "unbiased one way or the other", I'm not at all sure what this says, but the author's intended meaning is clear enough: ceteris paribus, scalar expressions give rise to upper-bounding inferences.

Unlike its strong cousin, weak defaultism makes no claims about processing. It merely says that upper-bounding inferences are the norm. But why should they be? Oddly enough, this question is rarely addressed. More generally, authors who claim that implicatures are defaults (be they weak or strong) almost never endeavour to explain why (Levinson 2000 is an exception). Now in the case of strong defaultism, such an explanation isn't hard to come by: strong defaults can be motivated by the assumption that they are encoded in the lexicon. I have argued that this is wrong, but at least it is an attempt at explanation. What about weak defaults, then? It is much less obvious what kind of mechanism might account for weak defaults. I will return to this question later, and until then treat weak defaultism as a purely descriptive claim.

What is the evidence for weak defaultism, thus understood? This question, too, has received less attention than it deserves, but one possible argument is suggested by the popular notion that implicatures are cancellable:[6]

(12) a. Betty's grades are good—in fact, they are excellent.
b. Fred broke an arm or a leg, if not both.

If scalar inferences were defaults, we could say (indeed, we would *have* to say) that the second half of (12a) cancels the "not excellent" inference triggered in the first half, and similarly for (12b). However, sentences like (12a,b) don't show that scalar inferences are defaults, because there is no compelling reason for analysing them in terms of cancellation, in the first place, as I argued in §1.5. The "if not" construction, for instance, may be construed as signaling that the speaker is not prepared to commit himself to a stronger claim. Evidence in favour of this view is provided by attested examples like (13a-c), cited from Horn (1989: 241):

(13) a. Overt antifeminism, if not homosexuality, may be the result of such experience in the male.

6. Although I'm not aware of any explicit versions of this argument in the literature, it seems to me that it is a natural corollary of the cancellation concept.

b. In the Netherlands the crowds [for the Pope] were small, the welcome lukewarm if not cold.
c. Most photographers were inarticulate if not subhuman.

Surely, it would be far-fetched to suppose that antifeminism implies heterosexuality *by default*, so (13a) shows that the "if not" construction doesn't demand a default inference for it to cancel; the same, mutatis mutandis, for (13b) and (13c).

Another argument in favour of granting scalar inferences default status is given by Levinson (2000). It is based upon certain patterns of lexicalisation documented in detail by Horn (1989). If we consider the positive duals "all" and "some", we see that the negation of the latter is lexicalised while that of the former is not: there is no single word for expressing "not all", no word that stands to "all" as "no(ne)" stands to "some". The same holds for other dual pairs, like "necessary/possible", "obligatory/permissible", and so on. In all cases, there is a lexical gap in the same place, as the following diagram shows:

some : no(ne)	permissible : forbidden	sometimes : never
all : [not all]	obligatory : [not obligatory]	always : [not always]

Horn's explanation for this pattern, which is found across languages, is that "some", "possible", etc., tend to implicate "not all", "not necessary", etc., thus obviating the need for special lexical entries for the latter group. Levinson endorses Horn's proposal, and maintains that it supports the view that scalar inferences are defaults:

> Anyone who rejects [the default view] will have to find another explanation for this lexicalisation pattern. The [contextualist] cannot explain it—according to any such theory, there are no general tendencies to be found, or if there are they have the status of mere behavioral tendencies, playing no role in the systematic generation of implicatures. (Levinson 2000: 71)

This reasoning is plain wrong (as indeed I pointed out well before Levinson's book was published; Geurts 1998b). Horn's analysis does not require that scalar inferences be defaults. The keystone of Horn's explanation is that, generally speaking, when a speaker should want to use "not all" he can use the shorter form "some" in its stead, because *on such occasions* "some" will tend to implicate "not all". This, together with the assumption that positive concepts have a stronger tendency to be lexicalised than negative ones, suffices

to explain the lexicalisation pattern observed by Horn. What Levinson wants to prove is that "some" implies "not all" by default *simpliciter*; and this simply doesn't follow.

To sum up: thus far, it isn't clear that there is any explanatory work to do for the hypothesis that scalar inferences are weak defaults. However, we haven't yet considered the possibility that there might experimental evidence supporting that hypothesis. We'll go into that soon, but before doing so I would first like to address the question what provides the *intuitive* motivation for weak defaultism.

Defaultist intuitions I suspect that, in the final analysis, the idea that scalar inferences are defaults owes its popularity to a gut feeling: it just seems intuively plausible that this should be the case. When we ask ourselves whether we would take an utterance of (14a) to imply that, according to the speaker, (14b) holds as well, it seems evident that we would:

(14) a. Some of the goats have the flu.
b. Not all of the goats have the flu.

This reasoning seems so compelling that it may be hard to believe that it could be wrong. Still, the way introspection is employed in this argument is arguably misleading (Geurts 1998b, 2006, 2009). Obviously, to ask oneself whether or not (14a) implies (14b) is to suggest already that it *might* be implied, but more importantly, this question raises the issue of whether or not all of the goats have the flu, or in other words, it makes it *relevant* to establish whether this is the case. And even if our intuitions about this case are to be trusted (which is as it may be), they don't necessarily tell us anything about how (14a) is interpreted in situations where the issue is not raised. In short, the introspective evidence used to buttress claims to the effect that scalar inferences are defaults is dubious, at best.

In the last section, I argued that even if it were the case that a sentence like (14a) tends to imply (14b), we shouldn't jump to the conclusion that this inference is a default triggered by "some", because there would be a perfectly good contextualist explanation that doesn't require defaults. On the contextualist view, since "some of the goats" is the grammatical subject of (14a), the hearer will try to set up a context for this sentence in which the goats are the discourse topic, and in such a context it is relevant to ask, apropos of an utterance of (14a), whether or not all the goats might have the flu. Thus far, that question is implicit. The point of the last paragraph is that, when we

consult our intuitions about (14a), the question is asked explicitly. Hence, our intuitions are not about how this sentence would be interpreted in general, but are confined to a limited range of contexts, that is, contexts in which it is relevant to establish if all goats have the flu. And if our intuitions tell us that in such contexts (14a) would imply (14b), that doesn't further the defaultist cause in any way.

This argument is supported by the following experiment, which compared two different methods for probing scalar inferences (Geurts and Pouscoulous 2009). In this experiment, native speakers of Dutch were presented with the same sentence in two different tasks:

(15) Some of the B's are in the box on the left.

In one task, the inference condition, participants were asked whether this implies that not all the B's are in the box on the left. In the verification condition, participants had to decide whether the same sentence correctly describes the following situation:

Following Bott and Noveck (2004) and others (cf. §5.2), we assumed that someone who interprets (15) as implicating that not all the B's are in the box on the left will deny that the sentence gives a correct description of the facts. Note that, whereas in the inference task the scalar inference is associated with a positive response, in the verification task it is associated with a negative response; therefore we had to control for the possibility of a positive response bias, which is why we included a liberal number of foils that were superficially similar to the critical item.

The outcome of the experiment was that scalar inferences were observed almost twice as often in the inference condition (62%) as in the verification condition (34%). Furthermore, since participants' performance on the distractor items in the verification task was nearly perfect (97% correct), it is unlikely that the low rate of "no" answers for the critical item was caused by a positive response bias. These data confirm that the question of whether or not a sentence φ implies a scalar inference ψ is a biasing one. This is not to suggest that we should stop asking such questions altogether. Despite their bias, they remain an invaluable source of information. However, they are useless if we want to gauge how common these inferences are: the intuition that "some" normally implies "not all" counts for naught.

Lack of experimental support In order to test the claim that scalar inferences are weak defaults, quantitative data are needed. In particular, we need experiments that probe scalar inferences in a context that is as neutral as possible. What these experiments should *not* do is ask informants whether or not a sentence φ is likely to imply another sentence ψ, where ψ spells out a scalar inference φ. As it turns out, the literature contains a fair number of studies that meet these requirements. In fact, we have already encountered two of them. In the verification tasks employed by Bott and Noveck (2004) and Geurts and Pouscoulous (2009), people were asked to decide if sentences of the form "Some A's are B's" are true in a situation in which all A's are B's. While in the former study, negative responses were recorded 59% of the time, the corresponding rate in the latter was 34%. Neither result agrees with defaultist expectations. If defaultism was right "some" should imply "not all" in the absence of evidence to the contrary, so unless it can be argued that these experiments provided inadvertent cues that the scalar inferences were to be cancelled, the negative response rates should be a lot closer to 100% than they actually were. As far as I can tell, no such cues were provided, and if they weren't, the reported data argue against weak defaultism.

Let me give one further example, this time of an experiment on disjunction. In an acquisition study, Paris (1973) presented children and adult controls with disjunctive sentences with arbitrary content, like the following:

(16) The bird is in the nest or the shoe is on the foot.

Paris's materials contained sentences with "or" as well as "either/or". Participants were asked to determine whether or not such sentences were true of a pair of pictures. In the critical trials, pictures were shown that made both disjuncts true. If "or" is construed exclusively, i.e. as implying "not both", participants should say that the sentence is false, whereas on an inclusive construal, sans exclusive inference, the sentence should come out true. Paris's main finding was that, overall, inclusive interpretations were preferred for 82% of the sentences with "or" and 76.5% of the sentences with "either/or". (The difference between "or" and "either/or" was significant, though much smaller than one might have expected.) For the adult participants, the rates were 75% and 68.5%, respectively. These data suggest that the normal interpretation of "or" is inclusive, and therefore go against the defaultist view.

The experiments reviewed in the foregoing are representative of what one finds in the literature: when contextual factors are factored out and the experimental paradigm is as neutral as possible, rates of scalar inferences are

source	*scalar term*	%
Paris (1973)	or	25
Chevallier et al. (2008), exp. 1	or	25
Pijnacker et al. (2009)	or	54
Paris (1973)	either/or	32
Evans and Newstead (1980), exp. 2	either/or	33
Braine and Rumain (1981), exp. 3	either/or	41
Noveck (2001), exp. 3	some	59
Bott and Noveck (2004), exp. 3	some	59
Feeney et al. (2004), exp. 1	some	65
Geurts and Pouscoulous (2009), exp. 2	some	34
Noveck (2001), exp. 1	might	65

Table 1: A sample of experimental data on scalar inference. The rightmost column gives the rates at which upper-bounded interpretations were observed.

below 50%, on average, and never higher than 65%. Table 1 gives a sample of relevant experimental studies reported since the early 1970s. Clearly, there is no experimental support for the view that scalar inferences are weak defaults.

5.4 Contextualism

Contextualisation and quasi-defaults I believe we can safely conclude from the foregoing discussion that there is no empirical evidence that words like "some" or "or" come with default scalar inferences attached to them. On the contrary, such evidence as is available militates against this view. Hence, the original intuition underlying defaultism proves to be mistaken. Even if it should turn out that the derivation of quantity inferences isn't entirely dependent on the context, the truth is a good deal closer to the contextualist end of the table.

Still, I believe it may be argued that, when charitably construed, weak defaultism may contain a grain of truth, after all, for two reasons. The first reason has to do with contextualisation; the second, with the possibility of what I will call "default elements" in the derivation of implicatures. Starting with the first point, we have seen already that hearers always try to contextualise a sentence, even when it is presented in the "null context", i.e. without explicit information as to what kind of context the sentence is supposed to be

uttered in. We have seen that the contextualisation process is guided by various factors, including word order and the presence of a question the sentence is supposed to answer. These factors have nothing to do with implicature per se, but since they affect the availability (or salience) of alternatives, they may have an indirect effect on the derivation of implicatures. Due to such factors, it *may* be that a sentence like (17) will by default be interpreted as implying that, according to the speaker, not all the goats have the flu:

(17) Some of the goats have the flu. (= (14a))

But again, *if* there is a default inference in this case (which is as it may be), it is not hard-wired into the lexical meaning of "some" and it doesn't have anything to do with implicature as such. A contextualist view on implicature allows for the possibility that some sentences have a strong tendency to give rise to certain implicatures.

Let me briefly mention three further factors that, like word order and questions, may influence the derivation of Q-implicatures by raising the availability of alternatives:

- *Focus:* It is quite likely that scalar inferences can be facilitated by means of focus. Compare, for example, (18a) to (18b):

  (18) a. The director had a meeting with some of the consultants.
  b. The director had a meeting with SOME of the consultants.

  We encountered (18a) in our discussion of Breheny et al.'s experiment (§5.2, example (9a)), which suggested that this sentence is less likely to give rise to the scalar inference that the director didn't meet all the consultants. However, it seems intuitively plausible that this likelihood will increase when "some" is highlighted, as in in (18b). This intuition is confirmed in Chevallier et al.'s (2008) experimental study on disjunction, which shows that focus on "or" causes higher rates of exclusive interpretations, though it has to be noted that this is a matter of degree: focus doesn't force an exclusive construal, it just makes it more likely. (Chevallier et al.'s study is in line with Paris's findings on the difference between "or" and "either/or", discussed in the last section.)

  Focusing is standardly viewed as a device for evoking alternatives: by using (18b) the speaker draws attention to the fact that he is claiming that the director had a meeting with some rather than α of the consultants, where α would be an expression of the same category as "some". Naturally, one of the more obvious candidates for α-hood will be "all", and if the focus

on "some" is thus interpreted, the likelihood of a "not all" inference is bound to increase.

- *Partitives:* Existential determiners like "some" are special in that for them a scalar construal doesn't always make sense:

  (19) There are some oranges in the fridge.

  Normally speaking, it would be odd to infer from this that, according to the speaker, not all the oranges are in the fridge. Such an interpretation only can be opportune if the said oranges are part of a larger set, or in other words, it requires a partitive construal of "oranges", which in (19) is blocked by the existential "there" construction. There are various ways of promoting a partitive construal, some of which we have encountered already: it will help if the indefinite is the grammatical subject and/or "some" is focused. Another way is to use a partitive construction, as in (20b):

  (20) a. The fridge contains some oranges.
  b. The fridge contains some of the oranges.

  Whereas in (20a) a partitive construal is comparatively unlikely, in (20b) it is enforced by the partitive construction "some of the oranges", and accordingly this sentence is more liable to give rise to a "not all" inference (cf. Milsark 1977, Horn 1997).

- *Relative complexity of alternatives:* For reasons discussed in §3.4, it may be that the likelihood of a Q-implicature is affected by the relative complexity of the alternatives involved in its derivation. To illustrate, consider the following type of case:

  (21) Fred or Barney made a mistake.
  $\rightsquigarrow \neg \textsc{Bel}_S$(Fred made a mistake)
  $\rightsquigarrow \neg \textsc{Bel}_S$(Barney made a mistake)

  According to my intuitions, these inferences are quite robust, and although I have been cautioning that such intuitions should not be taken at face value, I'm fairly confident that this is right. Assuming it is, how could we account for the robustness of such inferences? The answer, I would like to suggest, is that in the case of (21) the relevant alternatives are simpler than the sentence itself. A speaker who utters (21) goes out of his way to deliver a sentence that takes more effort to produce and interpret than either "Fred made a mistake" or "Barney made a mistake", and this is why the ignorance inferences associated with (21) are so robust. (See §6.6 for further discussion of this point.)

Default elements in the derivation of implicatures I have suggested a number of factors that may affect the contextualisation of a sentence, and thus may conspire to foster the illusion that implicatures are defaults. It should be emphasised that none of these suggestions imply a concession to defaultism: everything I have said thus far is consistent with a strictly contextualist stance on Q-implicatures in general, and scalar implicatures in particular. Still, I'm about to suggest that some such concessions may be in order. Not that I will back down from the claim that defaultism is false. As far as I can see, the notion that Q-implicatures of any kind are default inferences has nothing to recommend it. However, I am prepared to concede that there may be default *elements* in the derivation of Q-implicatures:

- *Competence:* The notion of competence plays a conspicuous role in the standard Gricean treatment of Q-implicatures. As discussed in Chapter 2, Gricean theories first derive a weak implicature of the form $\neg\textsc{Bel}_{\textsc{S}}(\varphi)$, which is strengthened to $\textsc{Bel}_{\textsc{S}}(\neg\varphi)$ if speaker S is taken to be competent with respect to φ. In §2.3, I suggested, as a working hypothesis, that hearers adopt the Competence Assumption by default. If this is correct, it doesn't alter the case against defaultism in the least. However, it does imply that, *if* a Q-implicature arises, it will be strengthened by default. This prediction seems entirely reasonable to me; whether or not it is correct is an empirical question.
- *Level of specificity:* In §2.4, we discussed Matsumoto's (1995) observation that, in some cases, Q-implicatures appear to be language specific. One of Matsumoto's examples is that, whereas the English sentence (22) does not normally implicate that the speaker doesn't know whether Michio is older or younger than Takashi, its Japanese equivalent does:

  (22) Michio is Takashi's brother.

  The contrast between the two languages is obviously related to the circumstance that, besides words corresponding to English "brother" and "sister", Japanese also has a system of sibling terms that encode sex as well as seniority. However, this fact alone does not suffice to explain the difference. What is needed, in addition, is the assumption that, by default, speakers of Japanese expect each other to use the more specific sibling terms.

  Although Matsumoto's example is especially vivid, I don't think it is an isolated case. Speaking more generally, it is plausible to suppose that hearers will have default expectations about the minimum level of specificity

called for in ordinary situations. Normally speaking, we expect chairs to be called "chairs" rather than "pieces of furniture" or "things", dogs to be called "dogs" rather than "mammals" or "animals", and so on (cf. Brown 1958, Cruse 1977, Rosch 1978, Hirschberg 1985). Whenever such default expectations are disappointed, Q-implicatures are bound to ensue, but it is important to note that it is not these implicatures that have default status, but rather hearers' expectations about what a thing would normally be called.

5.5 Conclusion

Prima facie, defaultism seems plausible enough. On reflection, it is obvious that this view is applicable, if at all, to a ragbag of cases only, which don't make up anything like a natural class. It is clear that not all Q-implicatures could be defaults. Defaultism is *prima facie* credible when restricted to scalar inferences, but even so it is easy to come up with instances of scalar inferences that are strongly context-dependent, and experimental evidence suggests rather forcefully that the stock-in-trade licensers of scalar implicature, like "some" and "or", do not trigger upper-bounding inferences by default. In short: defaultism is a lost cause.

None of this is to deny that some implicatures may arise more or less systematically, either in general or (more likely) within a certain range of contexts. However, as we have seen in the last section, such observations, if correct, are consistent with the Gricean view that there is nothing about implicatures *per se* that should lead us to expect that they, or some of them, have default status. There is simply no explanatory work to do for the hypothesis that (some) implicatures are defaults.

Intentions, alternatives, and free choice

6

In Chapter 2, we discussed the Standard Recipe for deriving Q-implicatures, which is the best-known way of fleshing out the Gricean notion of quantity implicature. In this chapter, we will see that there is another way, as well, which despite its prima facie similarity to the Standard Recipe, does a much better job of explaining Q-implicatures; I call it the "intention-based" approach. Intention-based treatments of Q-implicature have been advocated by Robert van Rooij and Katrin Schulz (both jointly and separately), and Benjamin Spector,[1] and though it may be less popular than the standard account, I will argue that the intention-based approach is the better choice, mainly because it is more general. While the Standard Recipe works rather well with scalar implicatures, it becomes unwieldy or worse when applied to other types of Q-implicature. By contrast, the intention-based approach offers a unified explanation of all varieties of Q-implicature.

Here's a first taste of the difficulties we run into when we try to extend the Standard Recipe beyond scalar implicatures. Suppose George asks Mildred, "What did you have for lunch?", and Mildred answers:

(1) I had some of the strawberries.

Mildred's answer might yield a scalar implicature to the effect that she didn't have all the strawberries. According to the standard view, this implicature is generated by negating an alternative sentence Mildred could have used in lieu of (1), namely:

(2) I had all the strawberries.

So far so good. But there is a further Q-implicature, as well. In the given scenario, George might well be surprised, apropos of her uttering (1), that

1. See, e.g., van Rooij and Schulz (2004), Schulz and van Rooij (2006), Schulz (2005), and Spector (2006).

Mildred had such a light lunch: in this scenario, (1) would be liable to implicate that Mildred only had strawberries for lunch. Now the question is: What alternative sentence(s) might this implicature be the negation of? The only answer, as far as I can see, is that the implicature negates each of a set of alternatives like the following:

> I had some of the strawberries and a biscuit.
> I had some of the strawberries and a banana.
> I had some of the strawberries and some muesli.
> I had some of the strawberries, a banana, and a biscuit.
> I had some of the strawberries, a banana, and two biscuits.
> I had some of the strawberries, a banana, and three biscuits.
> *etc.*

The chief merit of this solution is that it does its job, but apart from that it is clumsy and not very plausible from a psychological point of view. The number of things Mildred could have had for lunch is huge, and if we consider all possible combinations the number becomes astronomical. It is hard to believe that George, or any other hearer for that matter, would have to entertain all these possibilities merely to infer that Mildred had a light lunch.

The intention-based approach offers an alternative way of deriving Q-implicatures that avoids this problem. The idea is really very simple. Whereas according to the Standard Recipe, George begins by asking why Mildred failed to produce one of a set of stronger alternative statements available to her, the intention-based method has him pose the following question:

> If Mildred believed (knew) that her lunch comprised more than just some strawberries, would she have said so?

If George's answer to this question is "yes", then he will infer that Mildred's lunch contained nothing more than some strawberries. If this is how George reasons, he need not have any particular alternatives in mind. He merely has to decide if a hypothetical belief of Mildred's would have prompted her to deliver a different statement.

In the following, I will develop the intention-based view using so-called "free choice" inferences as my main exhibit. The reason for this is partly that free choice is an interesting and much-debated topic in its own right, and partly that it causes a serious predicament for the Standard Recipe, while the intention-based approach explains it without further ado. Or so I will argue.

6.1 Free choice

My son asks, "What's for dessert?", and I say:

(3) You can have fruit or chocolate cake.

Bright kid that he is, my son swiftly arrives at the conclusion that he can have chocolate cake. But what justifies his inference? The proposition that x has fruit or chocolate cake is entailed by, hence weaker than, the proposition that x has chocolate cake, and in this sense the permission to have fruit or chocolate cake grants *less* than the permission to have chocolate cake. It seems that my son derived his conclusion from a premiss that is logically weaker, and that therefore his reasoning can't be valid. In a nutshell, this is the problem of "free choice permission" (von Wright 1968, Kamp 1973).[2]

It bears emphasising, perhaps, that from a semantical point of view, (3) *is* a weaker statement than either "You can have fruit" or "You can have chocolate cake". (3) is a possibility statement, and possibility is closed under entailment: if φ entails ψ and φ is possible, then ψ is possible, too. Assuming, as always, that a disjunctive proposition is entailed by each of its disjuncts, it follows that "It's possible that φ" entails "It's possible that φ or ψ" (as well as "It's possible that ψ or φ"), but not the other way round. The problem of free choice permission is precisely to explain why it *appears* to be the other way round.

Note that, apart from the free choice inferences it gives rise to, (3) would probably also be construed as implying that the addressee is not allowed to

2. Some authors, including Zimmermann (2000) and Geurts (2005), have assumed that the underlying form of (3) is more like (i), and taken the analysis from there:

(i) You can have fruit or you can have chocolate cake.

There has been some discussion over whether this move is legitimate from a syntactical point of view (Kamp 1979, Jennings 1994, Geurts 2005), but even if it was, it wouldn't suffice to solve the problem. Having expanded (3) into (i), Zimmermann and Geurts have to resort to a modal analysis of disjunction in order to obtain free choice inferences, which is not without its problems. Furthermore, it should be noted that a periphrastic expansion is not always available. For instance, (6) below is obviously not synonymous with:

(ii) Many people had pizza or many people had pasta.

The account of free choice to be presented in this chapter does not apply to cases like (i). I still feel that the disjunction in this type of sentence has a non-Boolean flavour, but have ceased to believe that they are the proper starting point for a general theory of free choice.

have fruit and chocolate cake. However, this exclusivity inference is entirely unproblematic (the relevant alternative is "You can have fruit and chocolate cake") and orthogonal to the problem under discussion. That is to say, free choice is a problem no matter whether or the sentence is read exclusively or inclusively. Therefore, I will ignore this aspect in the following.

Although the paradigm cases of free choice inference involve performative (i.e., permission giving) uses of deontic statements, as in the opening example, it is important to note that they range more widely. To begin with, reportative uses of deontic statements license free choice inferences, as well. If I utter (3), I may be giving my son permission to have fruit or chocolate cake, or alternatively, I may be reporting on what the restaurant's menu has to offer, for example. Free choice inferences arise in the latter case just as they do in the former. Secondly, we observe free choice inferences with non-deontic modals:

(4) George may be British or Canadian.
   $\rightsquigarrow$ He may be British.
   $\rightsquigarrow$ He may be Canadian.

(5) Betty can balance a fishing rod on her nose or chin.
   $\rightsquigarrow$ She can balance a fishing rod on her nose.
   $\rightsquigarrow$ She can balance a fishing rod on her chin.

The most natural construal of (4) is epistemic, and on this reading the statement clearly licenses free choice inferences, as does the ability modal in (5).

Finally, it has been observed that free choice construals don't even require modals (Eckardt 2007, Fox 2007, Klinedinst 2007):

(6) Many people had pizza or pasta.
   $\rightsquigarrow$ Some people had pizza.
   $\rightsquigarrow$ Some people had pasta.

Surely, (6) doesn't entail that both pizzas and pastas were consumed, yet this is what we would normally infer. These observations suggest that the problem of free choice interpretation doesn't have anything to do with modality per se, but is more general.

The free choice inferences observed in the foregoing are of course very similar to the distributive inferences licensed by Boolean "and", and may therefore foster the impression that, on a free choice interpretation, "or" is construed as "and". This impression is false, however:

(7) a. You can have fruit or chocolate cake.
b. You can have fruit and chocolate cake.

(8) a. George may be British or Canadian.
b. George may be British and Canadian.

(9) a. Many people had pizza or pasta.
b. Many people had pizza and pasta.

It is obvious that (9b) is not an adequate way of paraphrasing the free choice interpretation of (9a), and the same goes for (8a,b). The contrast between (7a) and (7b) is perhaps less obvious. These sentences are harder to compare because, even if it is decided that (7a) should have a free choice construal, both sentences admit of more than one reading. However, (7b) is ill-suited as a paraphrase for what I take to be the most likely interpretation of (7a), according to which the addressee can have fruit or chocolate cake but not both. In sum, free choice phenomena cannot be dispatched by supposing that "and" occasionally gets mispronounced as "or".

The problem of free choice inference has prompted a bewildering variety of responses: logical, semantic, pragmatic, and even syntactic.[3] Needless to say, the course I will be pursuing in this chapter is a pragmatic one. There are at least three reasons for supposing that free choice inferences are conversational implicatures. One is that, as a matter of methodological hygiene, we have to assume that free choice inferences are conversational implicatures, unless there is compelling evidence that they aren't. I've said it before (§1.3), but the point is important enough to be reiterated every now and then: if there is a way of deriving an inference—*any* inference—by way of conversational implicature, then it behooves us to do so.

The second reason is that free choice inferences are cancellable:

(10) a. Many people had pizza or pasta—I forget which.
b. George may be British or Canadian, but I'm not allowed to tell you which it is.
c. You can have fruit or ice cream, but I'm not going to tell you which.

Granted, in some cases the speaker's motives have to be rather special, if not downright perverse, for free choice inferences to be cancelled, but as the examples in (10a-c) show, these inferences are non-conventional, and therefore a Gricean approach suggests itself.

3. Recent proposals for dealing with free choice include Zimmermann (2000), Asher and Bonevac (2005), Geurts (2005), Schulz (2005), Chemla (2007), Eckardt (2007), and Fox (2007).

The third reason is that some free choice inferences bear a strong resemblance to run-of-the-mill Q-implicatures:

(11) George is British or Canadian.
  ⇝ He may be British.
  ⇝ He may be Canadian.

(12) George may be British or Canadian.
  ⇝ He may be British.
  ⇝ He may be Canadian.

The only difference between (11) and (12) is that the latter contains a modal expression which the former lacks; otherwise they are the same. In view of this similarity, it is natural to expect that (11) and (12) should receive a unified analysis, and since the inferences in (11) are Q-implicatures, the null hypothesis should be that the same holds for (12).

6.2 Problems with the Standard Recipe

Oddly enough, while the Standard Recipe explains the inferences in (11) without much ado (as we have seen in §2.3), it dramatically fails to account for the parallel inferences in (12). To explain this remarkable discrepancy, let's first recapitulate how the Standard Recipe deals with (11), as uttered by Mildred (say):

i. Instead of saying (11), Mildred could have made a stronger statement, like "George is British" or "George is Canadian".

ii. Had she believed that either of these propositions is true, Mildred would have asserted it rather than (11).

iii. Hence, Mildred doesn't believe that George is British, nor does she believe that he is Canadian: $\neg\textsc{Bel}_{\textsc{M}}$(George is British) and $\neg\textsc{Bel}_{\textsc{M}}$(George is Canadian). Or, equivalently, for all Mildred knows, George may not be British and he may not be Canadian: $\textsc{Poss}_{\textsc{M}}(\neg$(George is British)) and $\textsc{Poss}_{\textsc{M}}(\neg$(George is Canadian)).

iv. From (*iii*) and the fact that Mildred believes (11), it follows that for all Mildred knows, George may be British and he may be Canadian: $\textsc{Poss}_{\textsc{M}}$(George is British) and $\textsc{Poss}_{\textsc{M}}$(George is Canadian).[4]

4. If the speaker believes (*i*) that φ or ψ and (*ii*) that φ may be false and (*iii*) that ψ may be

The key step in the argument is the first one: it starts with the observation that Mildred could have made stronger statements than she actually did, on the basis of which it is inferred that she doesn't believe that those statements are true. Formally, this reasoning transforms alternative sentences φ into inferences of the form "The speaker doesn't believe φ".

The trouble is that when we apply the Standard Recipe to sentences like (12), the result is nothing short of disastrous:

i. Instead of saying (12), Mildred could have made a stronger statement, like "George may be British" or "George may be Canadian".
ii. Had she believed that either of these claims is true, Mildred would have asserted it rather than (12).
iii. Hence, Mildred doesn't believe that George may be British and doesn't believe that he may be Canadian.
iv. From (*iii*) and the fact that Mildred believes (12), it follows that Mildred's beliefs about George's nationality are inconsistent.[5]

It will be clear that something went seriously wrong, but where? After all, we employed the exact same reasoning that served us so well with (11), and the crucial premisses, (*i*) and (*ii*), wouldn't seem to be more doubtful than in the previous example. In fact, the situation is doubly embarrassing: not only does the Standard Recipe fail quite badly, it also seems out of the question that it can deliver the unified analysis invited by the obvious similarity between (11) and (12).

These problems have persuaded some authors that a pragmatic explanation of free choice is not forthcoming, and that a solution will have to be found somewhere else altogether (e.g., Asher and Bonevac 2005, Fox 2007). I see no reason for being so defeatist. In my view, the difficulties caused by free choice inferences aren't endemic to the Gricean programme as such, but rather to the way it is standardly implemented. There is another way, which not only is more faithful to Grice's original intentions, but will also give us a unified analysis of free choice inferences and (other) Q-implicatures.

---

false, then he has to allow that φ may be true (in case ψ turned out to be false) and that ψ may be true (in case φ turned out to be false). More succinctly: if $\textsc{Bel}_{\textsc{S}}(\varphi \vee \psi)$ and $\textsc{Poss}_{\textsc{S}}(\neg\varphi)$ and $\textsc{Poss}_{\textsc{S}}(\neg\psi)$, then it follows that $\textsc{Poss}_{\textsc{S}}(\varphi)$ and $\textsc{Poss}_{\textsc{S}}(\psi)$.

5. So as to heighten the dramatic impact of the problem, I assume that if Mildred doesn't believe that φ is possible, then she believes that $\neg\varphi$. I take it that this assumption is reasonable, but even if it were dropped the problem would be bad enough.

6.3 Intentions first

When one comes to think of it, there is something odd about the way the Standard Recipe implements the Gricean notion of Q-implicature. The Standard Recipe begins by asking why the speaker didn't use this or that *linguistic form* instead of the one he did use. But shouldn't the first question be what the speaker's intentional state is, what he wants, believes, etc.? In the following I will try to show that if we put intentional states first, as any true-blooded Gricean should, then a unified treatment of free choice and quantity implicatures will fall in our laps, with none of the problems that beset the Standard Recipe.

I should note that I am using the term "intentional" in the philosophical sense of the word. Whereas in ordinary parlance this word is related to action, philosophers use it to highlight the directedness of certain mental states: beliefs, hopes, fears, and desires are all directed at possible states of affairs. If George believes that it will be sunny tomorrow, the object of his belief is a proposition, viz. that it will be sunny tomorrow. If Mildred hopes that it will be sunny tomorrow, her intentional state is different from George's, but it is directed at the same proposition. For our purposes, then, intentional states are mental states directed at propositions.

In the following pages, I will be talking a lot about what speakers might believe, want, allow, consider possible, as well as about what they might have said if they were in this or that intentional state. Since it will be vitally important to separate clearly between intentional states and linguistic means for expressing them, I would like to emphasise that the noun "alternative" will be reserved exclusively for referring to *linguistic expressions*, usually sentences, that the speaker could have used instead of one he did use.

Scalar implicatures In order to explain the main ideas underlying the intention-based approach, let's start with a simple case of scalar implicature. George says:

(13) Many of the nurses were drunk.

Assuming that George is competent, Mildred infers from his utterance that $\textsc{Bel}_{\textsc{G}}(\neg(\text{all nurses were drunk}))$. According to the Standard Recipe, Mildred achieves this by asking herself why George didn't make a stronger statement, i.e. "All nurses were drunk." She then conjectures that he didn't do so because

$\neg$BEL$_{\text{G}}$(all nurses were drunk), from which she infers, assuming that George is competent with respect to the proposition that all nurses were drunk, that BEL$_{\text{G}}$($\neg$(all nurses were drunk)). What is crucial about this story is that Mildred's train of reasoning is supposed to start with the observation that George could have produced a stronger statement and then proceeds to the inference that he doesn't believe that claim to be true. In brief, the derivation of Q-implicatures is driven entirely by alternatives.

To see how it could be otherwise, consider three possible ways George's beliefs might relate to the proposition that all the nurses were drunk:

i_1 : BEL$_{\text{G}}$(all nurses were drunk)

i_2 : BEL$_{\text{G}}$($\neg$(all nurses were drunk))

i_3 : neither i_1 nor i_2

Each of i_1-i_3 is a belief state George might be in, and as a matter of logical necessity, George has to be in one and only one of these states. That is, he must either believe that all nurses were drunk or believe that not all nurses were drunk or have no opinion one way or the other. Armed with this three-way partition of belief states, we wind back to the scenario in which George has just uttered (13), and observe that the following alternative statements were available to him:

(14) a. All nurses were drunk.
b. Many but not all nurses were drunk.
c. Many and maybe all nurses were drunk.

Given that George wanted to convey, at a minimum, that many of the nurses were drunk (for that is what he said), he could have uttered (14a) if he was in belief state i_1, (14b) if he was in belief state i_2, and (14c) if he was in belief state i_3. As far as I can see, each of these sentences is amongst the most economical alternatives available to George for conveying explicitly that he was in the corresponding state (while entailing (13)). For example, if i_1 obtained, George could have said, "Every nurse was sloshed", but that wouldn't have been more efficient than (14a). The same, mutatis mutandis, for (14b) and (14c).

Now Mildred reasons as follows. Since (14b) or (14c) are both more complex, hence less economical, than (14a), George would need a better excuse for avoiding the latter than either of the former. If George believed that all nurses were drunk, there would in general be no good reason for him uttering

(13) rather than (14a). Put otherwise, between the states i_1, i_2, and i_3, the first one would have been the most likely to give rise to an utterance different from the one that was in fact produced, and therefore i_1 is the first candidate for elimination. Once i_1 is discarded, Mildred has decided that George's belief state is either i_2 or i_3, or equivalently, that $\neg\textsc{Bel}_{\textsc{G}}$(all nurses were drunk); so at this point she has derived a weak scalar implicature. If, furthermore, Mildred is prepared to suppose that George is competent with respect to the proposition that all nurses were drunk, i_3 drops out, too, and she is left with i_2, which yields the strong implicature that $\textsc{Bel}_{\textsc{G}}(\neg$(all nurses were drunk)). Thus, intention-based reasoning replicates the results obtained previously by means of the Standard Recipe.

It is important to see exactly how the two methods differ from each other. The Standard Recipe begins with a set of one or more alternative sentences. For each of these, the hearer asks himself why the speaker didn't utter it, and the answer is that he didn't because he doesn't believe that the sentence in question is true. The intention-based procedure, on the other hand, starts with a set of possible intentional states, and then proceeds out to eliminate those that are unlikely to obtain for some reason or other. The intentional states that survive this weeding process are the ones that, for all the hearer knows, the speaker might be in. In the example above, there is only one such state, but note that, without the Competence Assumption, Mildred would have been left with two possible belief states for George to be in.

In the intention-based analysis, alternatives are less central than in the Standard Recipe, but they continue to play a very important part. If the hearer feels that, had the speaker been in state i, he would have spoken differently than he did, then he will discard i from his roster of possible intentional states. In brief, alternatives are used for selecting between possible intentional states.

Simple disjunctions Suppose now that we have two atomic propositions: that George is British and that George is Canadian. Suppose, moreover, that we don't know anything about Mildred's beliefs, as yet, so she may consider it possible or not possible that George is British. Since the same holds for the proposition that George is Canadian, we have four possible belief states (assuming that George can have dual citizenship):

$$
\begin{array}{lrcl}
i_1: & \textsc{Poss}_{\textsc{M}}(\text{George is British}) & \wedge & \textsc{Poss}_{\textsc{M}}(\text{George is Canadian}) \\
i_2: & \textsc{Poss}_{\textsc{M}}(\text{George is British}) & \wedge & \neg\textsc{Poss}_{\textsc{M}}(\text{George is Canadian}) \\
i_3: & \neg\textsc{Poss}_{\textsc{M}}(\text{George is British}) & \wedge & \textsc{Poss}_{\textsc{M}}(\text{George is Canadian}) \\
i_4: & \neg\textsc{Poss}_{\textsc{M}}(\text{George is British}) & \wedge & \neg\textsc{Poss}_{\textsc{M}}(\text{George is Canadian})
\end{array}
$$

Each of i_1-i_4 is a belief state Mildred might be in vis-á-vis the proposition that George is British and the proposition that George is Canadian, and Mildred has to be in one and only one of these states: i_1-i_4 are mutually disjoint and jointly exhaustive.

Now Mildred issues the following statement:

(15) George is British or Canadian.

Assuming that Mildred believes what she says (i.e., that she is not violating the Quality maxim), i_4 can be discarded rightaway, for if she was in this state, Mildred would believe that George is neither British nor Canadian. For the remaining states, the following alternatives would have been available for Mildred to utter:

(16) a. George is British.
b. George is Canadian.
c. George may be British, he may be Canadian, and he is at least one of the two.

Given that Mildred wanted to convey, at a minimum, that (15) is true, she could have uttered (16a) if she was in belief state i_2, (16b) if in belief state i_3, and something along the lines of (16c) if she was in belief state i_1. (If the last sentence seems too prolix, bear in mind that it has to entail (15) as well as describe i_1.) It will be clear that expressing i_1 would have been much more of a hassle than expressing either i_2 or i_3, and therefore i_2 and i_3 are the prime candidates for elimination. Once these states are eliminated, only i_1 is left for Mildred to be in, which is to say that, for all she knows, George may be American and he may be Canadian.

Summing up: The intention-based method employs the following ingredients:

- A sentence φ, as uttered by a speaker S.
- A partitioning i_1, …, i_n of S's possible intentional states. Intentional states may be partitioned in many different ways. In the foregoing examples, it was the speaker's beliefs that were at issue, but in other cases different flavours of intentionality will be relevant (e.g., free choice permission is not about S's beliefs; see below).
- Alternatives: sentences S could have used instead of φ, had he wanted to express any of i_1, …, i_n. Since we can assume that S intended to signal his commitment to φ, each of these alternatives should entail φ.

Next hearer H proceeds to weed out those states in i_1, ..., i_n that fail to meet certain conditions, which broadly speaking fall into two categories. On the one hand, there are all manner of things H may already know, or at least assume, about S. For example, if φ is a statement, H will assume that $\textsc{Bel}_S(\varphi)$, and any of states in i_1, ..., i_n that are inconsistent with this assumption can be discarded. H may also be assuming that S is competent with respect to a given proposition, which may cause further casualties amongst i_1, ..., i_n. More generally, anything H assumes about S's beliefs, hopes, desires, etc. may help to narrow down the space of possible intentional states.

On the other hand, H can try to weed out candidate intentional states by considering the alternatives associated with them, using the following alternative heuristics:[6]

Alternative heuristics

i. If S could have expressed that he is in state i by using a sentence that is no more complex than φ, then probably S is not in state i.

ii. For any pair of intentional states, i and j, if it had been easier for S to express that he is in i than that he is in j, then S is more likely to be in j than in i.

Again, it is presupposed that S's utterance of φ is sincere: if φ is a statement, then S believes φ; if φ is a permission sentence, S intends to grant H the permission that φ; and so on. Hence, any alternative must be at least as strong as φ.

Both alternative heuristics are based on the idea that there is an inverse correlation between an alternative's complexity and the likelihood that the speaker is in the corresponding intentional state. If ψ is among the simplest alternatives S could have used for expressing that he is in state i, then the more complex ψ, the more likely it becomes that ψ's complexity is what explains why S failed to utter it, thus making it proportionally less likely that i holds. This idea has a long pedigree in the history of Q-implicatures,[7] but it raises some hairy issues, which will be addressed in §6.5.

6. This is to be understood as a compound nominal, i.e., it is "ALTERNATIVE heuristic", not "alternative HEURISTIC".

7. See, for example, Strawson (1952), Gazdar (1979), Levinson (1983), Horn (1984, 1989), and Blutner (2000).

6.4 Free choice explained

Epistemic free choice One of the signal virtues of the intention-based analysis of Q-implicatures is that it extends to free choice inferences without further ado. To explain how, I will begin by showing how the analysis of the simple disjunction in (15) carries over to its free choice cousin. This time around, Mildred's utterance is slightly weaker than previously:

(17) George may be British or Canadian.

As in the case of (15), we begin by observing that Mildred must be in one of the following belief states:[8]

$$
\begin{array}{lrcl}
i_1: & \textsc{Poss}_{\mathrm{M}}(\text{George is British}) & \wedge & \textsc{Poss}_{\mathrm{M}}(\text{George is Canadian}) \\
i_2: & \textsc{Poss}_{\mathrm{M}}(\text{George is British}) & \wedge & \neg\textsc{Poss}_{\mathrm{M}}(\text{George is Canadian}) \\
i_3: & \neg\textsc{Poss}_{\mathrm{M}}(\text{George is British}) & \wedge & \textsc{Poss}_{\mathrm{M}}(\text{George is Canadian}) \\
i_4: & \neg\textsc{Poss}_{\mathrm{M}}(\text{George is British}) & \wedge & \neg\textsc{Poss}_{\mathrm{M}}(\text{George is Canadian})
\end{array}
$$

Of these states, i_4 can be discounted straightaway, because it is inconsistent with Mildred's believing (17). For the remaining states, the following alternatives would have been available for Mildred to utter:

(18) a. George may be British.
b. George may be Canadian.
c. George may be British and he may be Canadian.

Assuming that Mildred believes what she says, she could have uttered (18a) if she was in belief state i_2, (18b) if in belief state i_3, and something along the lines of (18c) if she was in belief state i_1. Now the second alternative heuristic tells us that, since (18a) and (18b) are considerably simpler than (18c), Mildred's intentional state probably is i_1 rather than i_2 or i_3, which is to say that, as far as Mildred knows, George may be British and George may be Canadian. Free choice explained.

Thus the intention-based theory of Q-implicature explains how (17) gives rise to free choice inferences and, moreover, the explanation it offers is prac-

8. I'm assuming here that the propositions at issue are that George is British and that he is Canadian, not that he *may* be British and that he *may* be Canadian. I believe that this is a natural assumption to make, but even if it simplifies things somewhat, it is not critical, as will become clear presently.

tically the same as the one it gives for simple disjunctions. It is a unified treatment if there ever was one.

Let's try to pinpoint why the intention-based account succeeds where the Standard Recipe failed. Consider the alternative in (18a). In a sense, the two theories agree that this alternative is to be rejected, but they disagree on what exactly this means. According the Standard Recipe, what it means is that the hearer is entitled to infer that Mildred doesn't believe that (18a) is true—which is wrong. On the intention-based account, by contrast, what is rejected is a possible belief state that would have caused Mildred to utter (18a), to wit: that she believes that George is not Canadian but may be British—and this prediction is correct. Herein lies the main difference between the two accounts.

Once we are clear about the distinction between alternatives and intentional states, it is obvious enough where the Standard Recipe went wrong: it conflates alternatives with the objects of possible intentional states. It is this conflation that obstructed a unified treatment of free choice and Q-implicatures.

Free choice permission Now let us return to the opening example of §6.1: my son asks, "What's for dessert?", and I say:

(19) You can have fruit or chocolate cake.

There are various ways of construing this answer, but let's suppose that the modal is being used deontically and performatively: I'm granting my son permission to have fruit or chocolate cake. Thus construed, the deontic case has the same logical structure as the epistemic one. True, (19) speaks of desserts rather than citizenship and its modality is deontic instead of epistemic, but none of this makes any difference for the computation of free choice inferences.

Suppose that, as far as my son is concerned, I am the dessert authority, and until I answered his question he had no idea what would be my position in the matter. Hence, as far as my son was aware, he might or might not be allowed to have fruit, and he might or might not be allowed to have chocolate cake:

$$
\begin{array}{lrcl}
i_1: & \textsc{Perm}_{\mathrm{B}}(\mathrm{F}) & \wedge & \textsc{Perm}_{\mathrm{B}}(\mathrm{C}) \\
i_2: & \textsc{Perm}_{\mathrm{B}}(\mathrm{F}) & \wedge & \neg\textsc{Perm}_{\mathrm{B}}(\mathrm{C}) \\
i_3: & \neg\textsc{Perm}_{\mathrm{B}}(\mathrm{F}) & \wedge & \textsc{Perm}_{\mathrm{B}}(\mathrm{C}) \\
i_4: & \neg\textsc{Perm}_{\mathrm{B}}(\mathrm{F}) & \wedge & \neg\textsc{Perm}_{\mathrm{B}}(\mathrm{C})
\end{array}
$$

where F is the proposition that my son has fruit and C is the proposition that he has chocolate cake. From this point on, the reasoning is the same as in the previous example, so we can go through it fairly quickly. (19) is inconsistent with i_4, and i_2 and i_3 can be discarded, too, because if either of them had obtained, I would have said something else, and shorter, than what I did say, e.g., "You can have fruit" for i_2 and "You can have chocolate cake" for i_3. Therefore, my intentional state can only be i_1, and my son is entitled to conclude that he can have fruit and that he can have chocolate cake (though perhaps not both).

Free choice reports Having shown how (19) can give rise to free choice inferences when construed as a permission-giving speech act, let's consider next how such inferences can arise when the same type of sentence is used reportatively, as in the following scenario. We're having dinner in a restaurant, and my son again wants to know what the dessert options are. I consult the menu, and then utter (19). This time around, I'm not granting my son permission to choose between fruit and chocolate cake; rather, what I'm conveying is my *belief* that the restaurant permits him to make that choice. Hence, in this case, we have not one but two layers of intentionality to consider: my beliefs and the restaurant's permissions.

So, my son will want to consider my beliefs about two propositions: that the restaurant permits him to have fruit and that it permits him to have chocolate cake (of course, his main interest is in the latter proposition, but for argument's sake we will pretend that he has at least an academic interest in the former, as well). If we notate these propositions as $\textsc{Perm}_{\mathrm{R}}(\mathrm{F})$ and $\textsc{Perm}_{\mathrm{R}}(\mathrm{C})$, respectively, my possible belief states can be partitioned as follows:

$$
\begin{aligned}
i_1 &: \quad \textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{F})) \;\wedge\; \textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{C})) \\
i_2 &: \quad \textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{F})) \;\wedge\; \neg\textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{C})) \\
i_3 &: \neg\textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{F})) \;\wedge\; \textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{C})) \\
i_4 &: \neg\textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{F})) \;\wedge\; \neg\textsc{Poss}_{\mathrm{B}}(\textsc{Perm}_{\mathrm{R}}(\mathrm{C}))
\end{aligned}
$$

Note that, save for the fact that the objects of Poss are more complex, this is essentially the same as the partitioning we employed in the case of epistemic free choice, so my son's reasoning will be the same, too, and its outcome will be that, for all my son knows, I must be in belief state i_1: I consider it possible that he can have fruit and I consider it possible that he can have chocolate cake. This is correct as far as it goes, but for my son's purposes it doesn't

go far enough, yet, because what he would like to have is that I believe he can have fruit and (more importantly) that I believe he can have chocolate cake. Fortunately, he can infer this if he is prepared to assume that I am competent with respect to the propositions $\textsc{Perm}_{\textsc{R}}(\text{F})$ and $\textsc{Perm}_{\textsc{R}}(\text{C})$. Since it is quite plausible that the Competence Assumption holds in these cases (I did see the menu, after all), it follows that $\textsc{Bel}_{\textsc{B}}(\textsc{Perm}_{\textsc{R}}(\text{F}))$ and (finally!) $\textsc{Bel}_{\textsc{B}}(\textsc{Perm}_{\textsc{R}}(\text{C}))$.

Pizza or pasta To conclude our tour of the free choice garden, let's see how the intention-based account applies to examples involving ordinary quantifiers, like:

(20) Many people had pizza or pasta.
   $\rightsquigarrow$ Some people had pizza.
   $\rightsquigarrow$ Some people had pasta.

Unlike many other accounts, which are often geared to the modal cases (like Geurts 2005, for example), the present analysis handles this type of example as a garden variety of free choice, which itself comes out as a garden variety of Q-implicature. We start with the following partitioning of speaker S's belief states:

$$
\begin{array}{lrcl}
i_1: & \textsc{Poss}_{\textsc{S}}(\text{some people had pizza}) & \wedge & \textsc{Poss}_{\textsc{S}}(\text{some people had pasta}) \\
i_2: & \textsc{Poss}_{\textsc{S}}(\text{some people had pizza}) & \wedge & \neg\textsc{Poss}_{\textsc{S}}(\text{some people had pasta}) \\
i_3: & \neg\textsc{Poss}_{\textsc{S}}(\text{some people had pizza}) & \wedge & \textsc{Poss}_{\textsc{S}}(\text{some people had pasta}) \\
i_4: & \neg\textsc{Poss}_{\textsc{S}}(\text{some people had pizza}) & \wedge & \neg\textsc{Poss}_{\textsc{S}}(\text{some people had pasta})
\end{array}
$$

Applying the same reasoning as in the preceding examples, we draw the conclusion that, most likely, S is in belief state i_1, which gives us two weak implicatures: that S considers it possible that some people had pizza and that S considers it possible that some people had pasta. If the Competence Assumption holds with respect to these propositions, it follows from this that $\textsc{Bel}_{\textsc{S}}$(some people had pizza) and $\textsc{Bel}_{\textsc{S}}$(some people had pasta). These predictions seem to be correct.

Taking stock Prima facie, the intention-based method for deriving Q-implicatures may not be very different from the standard one. Yet, while the key ingredients—alternatives and intentional states—are the same in both cases, the ways they are deployed aren't. The Standard Recipe assumes a close

coupling between alternatives and intentional states: if the speaker didn't use alternative φ it is always because he doesn't believe that φ. On the intention-based approach this isn't necessarily so. For instance, Mildred's failing to say "George may be British" may cause her audience to reject the possibility that, according to Mildred, George may be British but not Canadian. On this view, the connections between alternatives and intentional states aren't always as simple as assumed by the Standard Recipe.

On the intention-based view, the derivation of Q-implicatures is not driven by alternatives, as it is according the Standard Recipe. However, to say that alternatives are not in the driver's seat anymore is not to deny their importance; on the contrary, they are as important as ever, because candidate intentional states are vetted on the basis of, inter alia, the relative complexity of the alternatives that could have made them explicit. This part of the theory carries a considerable share of its explanatory burden, and it raises some challenging issues, as we are about to see.

6.5 Comparing alternatives

To start explaining the issues just alluded to, let's revisit the analysis of scalar implicatures of §6.3, where we discussed the following example:

(21) Many of the nurses were drunk. (= (13))

Starting out from a three-way partition of the speaker's (i.e., George's) belief states,

i_1 : $\textsc{Bel}_{\text{G}}$(all nurses were drunk)
i_2 : $\textsc{Bel}_{\text{G}}(\neg$(all nurses were drunk))
i_3 : neither i_1 nor i_2

I argued that i_1 is the prime candidate for elimination, because i_1 could have been expressed by a sentence (viz. "All nurses were drunk") that is no more complex than (21) and less complex than any of the sentences that could have been used for expressing i_2 or i_3. If this is right, (a) economy of expression is a key factor in the derivation of Q-implicatures and (b) when assessing an intentional state i, the hearer not only has to compare i's alternative(s) with the sentence uttered by the speaker, but may also have to take into account other possible intentional states and *their* alternatives. Hence, in general, candidate intentional states cannot be evaluated individually; rather, they and their alternatives have to be weighed against each other.

To return to example (21): once i_1 is discarded we are left with i_2 and i_3. The alternatives associated with these belief states are about equally complex:

(22) a. Many though not all of the nurses were drunk. (i_2)
b. Many and perhaps all of the nurses were drunk. (i_3)

Therefore, there is no reason for preferring either belief state to the other on the grounds that would have been less economical to express. Nevertheless, I believe that, ceteris paribus, there may be a preference for eliminating i_3, resulting in the strong implicature that, according to George, not all the nurses were drunk. If there is such a preference, it could be due to a tendency to assume by default that the speaker is competent (cf. §2.3, §5.4), but this is as it may be.

The limits of economy Now consider the following:

(23) It must have been a great party. I hear that well before midnight all the doctors, many though not all of the interns, and (21) many of the nurses were drunk.

I think that in this case we would be more likely to infer that George doesn't know whether all the nurses were drunk, which puts him in belief state i_3. If this is so, the reason must be that the discourse in (23) raises the availability of the "many though not all" construction, which is employed in the i_2-alternative in (22a). Therefore, the hearer may reason that if George was in state i_2, then given that he used "many though not all" just an instance ago, he should have uttered (22a) rather than (21). If these observations are on the right track, what they show is that comparing between alternatives isn't just a matter of economy (cf. Katzir 2007). It would have been convenient if we could have made do with a complexity ordering that ranked alternatives by, say, the number of syllables they contain, but the example in (23) suggests that this will not always work.

It is sometimes suggested that economy of expression isn't implicated in the derivation of Q-implicatures, to begin with, but it seems to me that this conclusion is way too strong. While examples like (23) suggest rather forcefully that economy is not the only factor, they don't even begin to show that it is not a factor at all. To the best of my knowledge, the only attempt at arguing for the stronger conclusion has been made by Matsumoto (1995), using the following example (p. 44):

(24) It was warm yesterday, and it is a little bit more than warm today.
$\rightsquigarrow$ It was not "a little bit more than warm" yesterday.

Apparently, the first conjunct of (24) gives rise to a scalar implicature despite the fact that the relevant alternative (i.e., "It was a little bit more than warm yesterday") is clearly more complex. Therefore, according to Matsumoto,

> [...] the relative prolixity of [the stronger alternative] cannot by itself constrain the production of [a quantity implicature]. If a stronger item is regarded as carrying necessary information, that expression is expected to be used even if it is prolix. (*ibid.*)

Matsumoto's second claim is clearly correct: necessary information must be given no matter how much effort it will take. But the first claim is a dubious one. Observe, first, that the preferred interpretation of (24) is forced by juxtaposing "warm" with "a little bit more than warm", which raises the possibility that it isn't due to a Q-implicature in the first place. Rather, the upper-bounding effect may have the same cause as in sentences like (25), which are generally agreed not be amenable to an analysis in terms of implicature (see §7.1 and §8.5 for further discussion):

(25) Actually, it wasn't WARM yesterday: it was a little bit MORE than warm.

Secondly, even if (24) is a case of implicature, it can be explained along the same lines as (23). On this analysis, it is the speaker's own use of the expression "a little bit more than warm" that renders it a viable alternative to "warm". If the speaker goes out of his way to use this form in the second half of his statement, the reason why he did not do so in the first half cannot be that it was too complex, and thus the scalar inference is accounted for.

Availability I would like to replace economy with more general measure for comparing alternatives, which I propose to dub "availability" (cf. Horn 2000, Katzir 2007). Economy is one of the factors that determine availability, but there are further factors, as well, including previous use, frequency of use, familiarity, and general or particular expectations of specificity. In the foregoing, we have seen how previous use can enhance the availability of alternatives. I will give a few more examples. Consider the term "megalopolis", which means something like "oversized metropolis". Given that the word is relatively rare, and much rarer than "metropolis", the use of the latter would not normally implicate that the predicate "megalopolis" doesn't apply. Hence, if someone should claim that there is a metropolis on the back of the moon, his statement would not normally give rise to the implicature that the said city is not a

megalopolis. However, if the word "megalopolis" came to be used more frequently, and thus became more available, such an implicature would be more likely to occur.

A somewhat more pedestrian case is discussed by Horn (2000), who observes the following pattern:

(26) a. I hurt my finger.
$\rightsquigarrow$ I didn't hurt my thumb.
b. I hurt my toe.
$\not\rightsquigarrow$ I didn't hurt my big toe.

Although thumbs and big toes are structural analogues, and thumbs surely count as fingers ("or we'd all have eight fingers instead of ten"; Horn 2000: 308), (26a) is easily read as implicating that the finger in question wasn't the speaker's thumb, while a parallel inference seems rather unlikely in the case of (26b). Evidently, this contrast is related to the fact that English has a monolexical expression for thumbs but not for big toes, and Horn conjectures that, "if the colloquial language replaced its *thumb* with the polymorphous *pollex* (the Latin and scientific English term for both "thumb" and "big toe"), the asymmetry [between (26a) and (26b)] would instantly vanish." (ibid.) It is important to note, however, that the adjective "colloquial" is doing real work in this statement: it is not enough for an alternative word to be in the language; it has to be sufficiently salient, as well: if the word "thumb" was rarely used, then presumably the asymmetry between (26a) and (26b) would vanish, too.

In §2.4, we discussed the interpretation of sibling terms in Japanese (Matsumoto 1995). Recall that, in addition to "kyoodai", which is synonymous with English "brother", Japanese also features words meaning "younger brother" ("otooto") and "older brother" ("ani"), and that the use of "kyoodai" will normally be taken to implicate that the speaker doesn't know which of the hyponyms is applicable. By contrast, the use of "dog" in English does not normally implicate that the speaker doesn't know which of its hyponyms apply in the case at hand, even though colloquial English has words like "poodle" and "terrier", which are by no means rare. In brief, it seems that Japanese "ani" and "otooto" are more readily available than English "poodle" or "terrier", the reason being, presumably, that to speakers of Japanese seniority is more important than are the differences between dog breeds in most of the Western world.

What I'm suggesting, then, is that the availability of "ani" and "otooto" is relatively high because, in Japanese culture, seniority is of general interest,

and therefore relevant by default. But, again, it doesn't follow that availability is the same thing as relevance, because it is crucial that Japanese makes a *lexical* distinction between older and younger brothers. To see that this matters, compare Japanese "kyoodai" to English "cousin". Surely, the dichotomy between men and women is of universal interest (and more important, presumably, than the distinction between older and younger brothers), and yet, if I say that George will be visiting a cousin of his, there will not normally be an implicature to the effect that I don't know whether the cousin is male or female, though there almost certainly would be such an implicature if the distinction between male and female cousins was lexicalised in English, as it is many other languages.

To sum up, I have argued that in the process of weeding out possible intentional states, the hearer may have to take into account the availability of alternatives associated with those states, and that availability is not just a matter of economy, but is also affected by a variety of factors including previous use, frequency of use, familiarity, and general or particular expectations of specificity. How precisely these factors conspire to determine the availability level of an alternative is an issue that will be difficult to resolve, if only because the various factors just mentioned are interdependent to some degree at least, and I will not try to sort out these matters here. However, there is a related point that is at least as important, which is that availability cannot always play a role in the derivation of Q-implicatures, because in many cases there are no alternatives to be compared, in the first place. It is to this issue that we now turn.

6.6 Two flavours of Q-implicature

In the introduction to this chapter we briefly discussed the following example:

(27) George: What did you have for lunch?
Mildred: I had some of the strawberries. (= (1))

We observed that, aside from scalar "not α" inferences (where α is a quantifier that outranks "some"), Mildred's answer might also implicate that all she had for lunch was some strawberries, and that it seems unlikely that the latter inference can be accounted for by the Standard Recipe, because this would involve rejecting an indefinite number of alternatives:

I had some of the strawberries and a biscuit.
I had some of the strawberries and a banana.
I had some of the strawberries and some muesli.
etc.

Apparently, George must be able to infer that Mildred only had strawberries without having a precise notion of what else she might have had for lunch. True, George may have good reasons for believing that, e.g., muesli would have been a highly likely lunch option for Mildred to settle upon, and that she would never even consider consuming boeuf bourguignon in the middle of the day. But normally speaking the set of possible lunches is open-ended, and therefore the same must hold for the set of alternative answers Mildred could have given.

We might say that Mildred's answer to George is associated with two sets of alternatives. First, there is the alternative set involved in the scalar implicatures connected with "some"; this set is more or less clearly defined and fairly small. Secondly, there is the set of alternatives associated with the "strawberries and nothing else" implicature, which is open-ended and potentially huge. Let us say that the first alternative set is "closed", while the second one is "open".

The distinction between open and closed alternative sets correlates with another one. On the one hand, the members of a closed alternative set are, as a rule, no more complex than the sentence uttered by the speaker. On the other hand, the members of an open alternative set are typically, though not perhaps necessarily, more complex than the speaker's sentence. If Mildred had wanted to convey that, besides the strawberries, she had a wedge of cheddar, as well, her answer would have had to be longer, obviously.

Although the correlation between the open/closed distinction and the complexity of alternatives is by no means a logical necessity, there aren't that many exceptions I can think of. The examples in (23) and (24) show that the members of a closed alternative set are sometimes longer than the speaker's sentence, but these are out-of-the-way cases, and I can't think of any good examples of open alternative sets whose members aren't more complex than the speaker's sentence.[9] Be this as it may, it seems quite safe to say that the correlation is reliable enough.

9. If the context imposed clear limits on the what Mildred might have had for lunch, she could have said: "I ate everything." However, in such a case Mildred's set of alternatives would be closed, so this is an apparent exception only.

Most of the discussion in this book has been about Q-implicatures involving closed sets of alternatives: scalar implicatures, ignorance inferences associated with disjunction, and free choice inferences. In these cases, there are always small and determinate alternative sets, and it makes sense to assume that hearers compare alternatives, using the heuristics proposed in §6.3. But clearly, this presupposes that alternative sets are closed; if a set of alternatives is open, and it is indeterminate what the speaker's alternatives are, then the alternative heuristics don't apply. This conclusion is in line with our discussion of (27). For, in this case, George simply cannot derive the relevant Q-implicature by way of rejecting alternatives. Rather, we observed, what he has to do is ask himself:

(28) If Mildred believed (knew) that her lunch comprised more than just some strawberries, would she have said so?

And this question he may have to answer without having *any* idea what Mildred's alternatives were.

Two flavours of Q-implicature If these observations are on the right track, there are two rather different flavours of Q-implicature, which in the remainder of this chapter I will distinguish by means of subscripts. Q_c-implicatures are derived on the basis of closed sets of alternatives; they comprise scalar implicatures, ignorance inferences licensed by "or", and free choice inferences. Q_o-implicatures, by contrast, do not involve closed set of alternatives. One specimen of this class is the "strawberries only" inference prompted by Mildred's answer in (27); further examples will be discussed shortly.

What is the point of distinguishing between Q_c- and Q_o-implicatures? One good reason for introducing this distinction is that these two kinds of Q-implicature have somewhat different etiologies: while alternatives play a key role in the derivation of Q_c-implicatures, Q_o-implicatures are not contingent on the availability of alternatives. Basically a Q_o-implicature is derived on the assumption that if there was more to say, the speaker would have done so, no matter how; unlike Q_c-implicatures, Q_o-implicatures are not constrained by the *form* of what the speaker might have said.

Another reason for distinguishing between Q_c- and Q_o-implicatures is that the former appear to be more robust and determinate than the latter. To illustrate the difference, consider again the free choice inferences prompted by our mainstay example:

(29) You can have fruit or chocolate cake.

In this case, it is perfectly clear cut what the relevant (Q_c-) implicatures are: you can have fruit and you can have chocolate cake. Q_o-implicatures, by contrast, lack this determinacy, as (27) illustrates. Suppose that, in addition to a handful of strawberries, Mildred nibbled, say, two or three peanuts or a diminutive wedge of cheddar. Would that falsify the implicature that she only had some strawberries? Probably not. If we increase the number of peanuts or the size of the wedge, we may reach a point at which we start wondering whether Mildred really had only strawberries for lunch (though it might be argued that the peanuts don't count, because they're not lunch food), but clearly the cutoff point is nowhere as clear as it is in the case of (29).

To see that (27) is not an isolated case, consider the following examples:

(30) Mildred: What did you do when you got home?
George: I fell asleep in front of the television.

(31) George: Who did you see at the reception?
Mildred: Jack, Jill, and Janet.

It is rather unlikely that falling asleep in front of the television was the only thing George did when he got home. He may have taken off his coat and shoes, made himself a cup of tea, and so on. What George's answer in (30) is meant to convey is that falling asleep in front of the television was the only noteworthy thing he did. But clearly, noteworthiness is a matter of degree, and therefore the Q_o-implicature prompted by George's answer is much less determinate than your average Q_c-implicature. Similarly, Mildred's answer in (31) will not normally implicate that she saw no more than three people. George is entitled to suppose that, had she seen anybody else of note, she would have said so, but this is vague in the same way George's own answer in (30) is.

The same indeterminacy can be observed in conditional statements like the following:

(32) If you scratch my back, I'll scratch yours.

Someone who utters (32) will normally convey that, unless the addressee scratches her back, she will not scratch his, and I take it that this is a Q-implicature: if there were other conditions under which the speaker would be willing to scratch the addressee's back, she would have said so, and since she didn't, the addressee can infer that scratching the hearer's back is a condition sine qua non. Strictly speaking, however, this can't be true, because there

are zillions of other situations in which the speaker would scratch the hearer's back:

> If you threaten me with a machete, I'll scratch your back.
> If you give me one million euros in cash, I'll scratch your back.
> If I suddenly feel an irrepressible urge to scratch your back, I will do so.
> *etc.*

More likely than not, each of these alternatives is true. Hence, it is more accurate to say that (32) implicates that, apart from the addressee's scratching the speaker's back, there are no probable circumstances that will cause the reverse to happen—which, again, is vague.

Q_c-implicatures are more determinate than Q_o-implicatures because the latter are driven purely by considerations of relevance, while the former are constrained, in addition, by the availability of closed sets of alternatives that, as a rule, are at least as economical as the sentence uttered by the speaker. Someone who says, "Some of the nurses were drunk", could have made a stronger statement without more personal effort and without inflicting greater processing effort on the hearer, so even if it is doubtful that more information is needed, considerations of economy (or, more generally, availability) are not likely to explain why the speaker decided to use "some" instead of a stronger quantifier. In Q_o-implicatures, by contrast, not only is it unclear what the alternatives are, it also is a practical certainty that whatever additional information the speaker might have wanted to provide would have required more effort from all parties to the discourse. As a consequence, Q_o-implicatures are driven by considerations of relevance alone, and therefore inherently indeterminate.

Conditional perfection The Q-implicature associated with the conditional statement in (32) is an instance of what, in linguistics, is often called "conditional perfection", a term coined by Geis and Zwicky (1971), whose most famous example is:

(33) If you mow the lawn, I'll give you $5.
$\rightsquigarrow$ If you don't mow the lawn, I won't give you $5.

In traditional logic, the same pattern of reasoning is known as "denial of the antecedent", and is considered to be a logical fallacy. It should be observed, though, that even if it is not logically valid, this form of reasoning is pragmatically valid in a great many cases, so writing it off as a fallacy is not particularly helpful.

Although it is widely agreed that conditional perfection (to stick to the linguistic terminology) is an implicature, there is no consensus on what sort of implicature it is. The suggestion that it might be a Q-implicature goes back to Horn (1972) and Noordman (1977) and more recently has been defended at length by van der Auwera (1997a, 1997b), among many others. This is my view, as well, but it is not without its detractors. Indeed, it is opposed by a more recent persona of Horn's (2000), whose criticism focuses on van der Auwera's proposal. I believe that the dispute between van der Auwera and Horn is based on a misunderstanding, which is that both parties in the debate fail to appreciate the Q_c/Q_o-distinction. Let me dwell on this point for a few paragraphs.

Van der Auwera construes the inference in (33) as a Q-implicature deriving from the fact that the speaker could have made any number of stronger claims:

– If you mow the lawn or wash the car, I'll give you $5.
– If you mow the lawn, wash the car, or do the dishes, I'll give you $5.
– If you mow the lawn, wash the car, do the dishes, or prepare dinner, I'll give you $5.
  ...

Each of these alternatives entails (32), but not vice versa, and if each of these gives rise to the implicature that the speaker doesn't believe that it is true, then it follows that unless the addressee mows the lawn, the speaker will not give him $5.

Horn (2000) criticises van der Auwera's proposal on two counts. First, he notes, as we have done in the case of (27), that the alternatives hypothesised by van der Auwera don't seem to figure explicitly in the reasoning underlying (33); quite simply, it is highly implausible to assume that the addressee has to go through all these alternatives just to arrive at the conclusion that he will have to mow the lawn if he wants to earn $5. Secondly, Horn observes that van der Auwera's alternatives are "invoked here as *dei ex machina* to assure a successful [...] account" of conditional perfection (p. 309); that is to say, van der Auwera's story is ad hoc. Having shown to his own satisfaction that van der Auwera's proposal is flawed, Horn decides to drop the idea that conditional perfection is a garden variety of Q-implicature, and proceeds to develop an alternative account in which several maxims collaborate to secure the coveted inference. However, I will not discuss Horn's own proposal here, because in my opinion he is too quick to dismiss the feasibility of a straightforward Q-based explanation.

Despite their disagreement over the true nature of conditional perfection, van der Auwera and Horn share the view that deriving Q-implicatures is a matter of dismissing alternatives. What they disagree about it just whether or not conditional perfection can be explained in these terms. In this chapter I have argued at length that the alternatives-first view is wrong, and that in many instances of Q-implicature alternatives aren't involved at all. If this much is accepted, it is hard to avoid the conclusion that conditional perfection is a perfectly ordinary case of Q-implicature. Hence, I partly agree with Horn as well as van der Auwera. I believe the latter is right in claiming that conditional perfection is a Q-implicature, and that the former is right in criticising the specifics of the latter's proposal. On the other hand, both parties share the tacit assumption that if conditional perfection is a Q-implicature, then it will have to be of the Q_c-type. This is not the case: conditional perfection is a Q_o-implicature.

6.7 Conclusion

In this chapter I have argued that the proper treatment of Q-implicatures should be founded on two key distinctions. The first and most important distinction is that between the speaker's intentional states and the alternative sentences he could have used. It is obvious enough that these are entirely different things, but the standard view on Q-implicatures comes dangerously close to conflating the two, and at the very least it assumes that they are tightly coupled: if the speaker didn't say φ, then it must be because he doesn't believe that φ.

On the intention-based view, hearers do not primarily reason in terms of alternatives but rather in terms of intentional states the speaker might be in. Still, alternatives continue to play an important ancillary role in the theory: they are used to weed out candidate intentional states. However, the availability of alternatives is not a prerequisite for Q-implicatures to be derived. Thus, the intention-based approach prompts the introduction of a second distinction, between Q_c- and Q_o-implicatures, which we couldn't make as long as we were adhering to the view that Q-implicatures must be derived according to the Standard Recipe. This distinction yields the following taxonomy of Q-implicatures:

(34) *Q_c-implicatures*

a. Many of the nurses were drunk. [*scalar implicatures*]
   ⇝ Not all the nurses were drunk.

b. You may have fruit or chocolate cake. [*free choice inferences*]
   ⇝ You may have fruit.
   ⇝ You may have chocolate cake.

(35) *Q_o-implicatures*

a. [What did you do last night?] [*Q/A-implicatures*]
   I fell asleep in front of the television.
   ⇝ That's all I did.

b. If you scratch my back, I'll scratch yours. [*conditional perfection*]
   ⇝ If you don't scratch my back, I won't scratch yours.

The main argument in favour of the intention-based view is that it is more general than the standard, alternatives-first approach. The Standard Recipe works well for scalar implicatures, but doesn't provide a plausible account of Q_o-implicatures, and it fails quite miserably with free choice inferences.

Embedded implicatures: the problems

7

Even before the William James Lectures had started appearing in print, Grice's critics were arguing against his theory of conversational implicature using examples in which, prima facie, conversational implicatures occurred within the scope of conditionals and other operators; which on a Gricean view amounts to a contradiction in terms, since conversational implicatures can only be derived on the basis of a full-blown speech act. The first author to raise this issue was Cohen (1971), who presented cases like the following:

(1) If the old king has died of a heart attack and a republic has been declared, then Tom will be quite content.

Cohen observed that (1) may be understood as:

(2) If the old king has died of a heart attack *and afterwards* a republic has been declared, then Tom will be quite content.

If it is assumed that, within the Gricean scheme of things, the sequential reading of "and" is due to a Manner implicature, this interpretation is problematic, because it would seem to require that the implicature must be derived within the antecedent of the conditional; which is simply impossible (see Carston 2002 for extensive discussion of the pragmatics of conjunction). Hence, Cohen concluded, we have to accept that, in sentences like (1), the truth-conditional meaning of the word "and" has a temporal component.

After some initial skirmishes following in the wake of Cohen's article (e.g., Walker 1975, Wilson 1975), the discussion about "embedded implicatures" petered out, but it became livelier again when Landman (1998), Levinson (2000), Chierchia (2004), and Fox (2007) began calling for drastic departures from the Gricean party line, provoking defensive responses from Sauerland (2004), van Rooij and Schulz (2004), Horn (2006, 2009), Spector (2006), Russell (2006),

and myself (2009), among others. The current stage of the debate concentrates, as do this chapter and the next, on the status of Q-implicatures, and more particularly, though not exclusively, the status of scalar implicatures.

Theorists contending that embedded implicatures cannot be accommodated by the Gricean programme raise the spectre of what I call "conventionalism". In a nutshell, conventionalism is the view that Q-implicatures aren't implicatures at all, but rather soft entailments projected from the lexicon or the grammar. In its simplest form, the conventionalist view is that "some" really *means* "some but not all". Of course, the "not all" part of the meaning of "some" is cancellable, which is what makes the entailment "soft", but the leading idea is that it is included in the lexical content of the word nonetheless. Alternatively, some conventionalists hold that scalar inferences are generated in syntax, by way of a covert operator that may be freely inserted into the parse tree. Either way, the upper-bounding inference associated with "some" is recategorised as a non-pragmatic aspect of meaning whose origin lies in some linguistic convention or other—whence "conventionalism".

If conventionalism is right, the upper bound associated with "some" may go directly into the truth-conditional content of the sentence at the point where the word occurs; and if the word is sitting in the scope of another expression, that is where the upper bound will be inserted, too. Hence, on this view, "embedded Q-implicatures" are in fact lexico-syntactic enrichments that are factored in locally.

The ensuing discussion of embedded implicatures will extend over two chapters and fall into three main parts. First, I will survey the problem area and discuss a variety of cases that, prima facie, qualify as embedded implicatures (§7.1). Next, I turn to conventionalist theories (§7.2) and argue at some length that they are on the wrong track (§7.3). Finally, in the following chapter, I try to show that embedded implicatures can be explained on strictly Gricean principles, after all.

There are two main reasons why this discussion is going to take so long. One is that conventionalism is not one doctrine but many, so it will take a while to sort through the various flavours, and to show that all of them are problematic for some reason or other. The second reason is that, on the Gricean view to be defended in the following chapter, there isn't just one but several problems of embedded implicatures. More concretely, I will argue that there are various ways in which pragmatic inferences may arise that *seem* to be embedded within the scope of, e.g., an existential quantifier, but in reality aren't embedded in any sense, and discussing all of these various ways will

keep us occupied for the best part of the chapter.

Before we get to work, a couple of preliminary remarks are in order. First, although I have just spent an entire chapter advertising an intention-based approach to Q-implicatures, to replace the established alternative-based view, we cannot ignore the latter entirely, for the simple reason that some versions of conventionalism are alternative-based. For the most part, however, the differences between the two approaches don't matter for the following discussion, an important exception being the analysis of existential expressions of §8.4, which hinges on the intention-based view.

One theme that will reappear in the following is defaultism, the reason being that there is a natural affinity between conventionalism and defaultism. However, although at least some varieties of conventionalism squarely run into the problems discussed in Chapter 5, I propose to pretend that discussion never happened, partly because not all species of conventionalism are alike in this respect, and partly because there will be other issues that demand our attention.

Finally, since on a proper Gricean construal, "embedded implicatures" is a contradiction in terms,[1] my use of this term is not to imply that implicatures can be embedded. However, since it has gained some currency as a descriptor of a problem (or rather, a collection of problems), that's how I will use it, too, though I will continue to employ scare quotes as a reminder that the term should not be taken literally.

7.1 The problems

Belief reports A good way of introducing the problems raised by embedded Q-implicatures is by way of belief reports like the following:

(3) Tony believes that Cleo had some of the figs.

Let's say that (3) is uttered by one Julius. Then it may well give rise to the inference that, according to Julius, Tony believes that Cleo didn't have all of the figs. This inference involves two layers of belief, so to speak, and is of the form:

(4) $\textsc{Bel}_{\textsc{j}}(\textsc{Bel}_{\textsc{t}}(\neg(\text{Cleo had all the figs})))$

1. As is Levinson's (2000) notion of "truth-conditionally relevant implicatures", and for much the same reasons.

The problem is that, if we treat this as regular case of scalar implicature, the outcome will be somewhat weaker, namely:

(5) $\textsc{Bel}_{\textsc{J}}(\neg\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$

(Note that (4) entails (5), but not vice versa.) Here is how this implicature comes about. Upon hearing (3), the addressee can reason as follows:

i. Could it be the case that, for all Julius knows, Tony believes that Cleo had all the figs, i.e. $\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$?
ii. Presumably not, for then Julius would have said, "Tony believes that Cleo had all the figs." Hence, $\neg\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$.
iii. It seems reasonable to suppose that Julius is competent with respect to the proposition that $\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs})$. That is, either $\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$ or $\textsc{Bel}_{\textsc{J}}(\neg\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$.
iv. If so, it follows that $\textsc{Bel}_{\textsc{J}}(\neg\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$ (= (5)).

This result is not wrong, and may be quite satisfactory in many cases, but it falls short of the coveted inference in (4).

On reflection, it isn't surprising that, as things stand, this inference is out of reach for the Gricean approach. All scalar implicatures we have encountered so far are either of the form $\neg\textsc{Bel}_{\textsc{S}}(\varphi)$ or $\textsc{Bel}_{\textsc{S}}(\neg\varphi)$, where φ is a proposition that is stronger than the one stated by speaker S. In the case of (3), φ is $\textsc{Bel}_{\textsc{T}}$(Cleo had all the figs), so the strongest implicature we can hope to derive is indeed (5), where the scope of the negation operator is too wide for capturing the inference in (4).

Prima facie, it looks as if, on the problematic reading of (3), a scalar inference is derived within the scope of the verb "believes". Therefore, it seems that, in order to account for this reading, a conventionalist solution is called for: if the "not all" inference is attached to the *word* "some", then it can be factored into the truth-conditional meaning of the embedded sentence, so that (3) simply *means* that Tony believes that Cleo had some but not all of the figs, and the problem doesn't even arise.

One of the potential drawbacks to the conventionalist solution is that it is prone to overgenerate. Whatever the details of the conventionalist analysis of (3) may look like, it will yield analogous predictions for attitude reports in general, and predict that (6a) and (7a) can be read as (6b) and (7b), respectively. Indeed, all conventionalist theories that are currently on the market predict that these readings should be preferred:

(6) a. Tony hopes that Cleo had some of the figs.
b. Tony hopes that Cleo had some but not all of the figs.

(7) a. Tony regrets that Cleo had some of the figs.
b. Tony regrets that Cleo had some but not all of the figs.

This doesn't seem to be right, though: it is doubtful that the (b) sentences are plausible paraphrases of the (a) sentences. The reason why I'm making this observation here is not that it undermines conventionalism, though it does that, too. Rather, what I want to point out is that, interestingly, the problem caused by belief sentences appears to be a *local* one: it arises with "believe" and kindred verbs like "think" and "know", but doesn't seem to extend to attitude verbs in general. If this impression is correct, then it may be a bad idea to deploy general-purpose machinery for generating scalar inferences in embedded positions, as the conventionalist does. That is of course what I will be arguing later, but for now my point is just that "believe" and its kin appear to be special.

Factives I just said that attitude reports with "know" cause the same sort of trouble as do belief reports. The following variation on (3) shows that they do:

(8) Julius knows that Cleo had some of the figs.

Intuitively, it doesn't seem too far-fetched to construe (8) as conveying that Julius knows that Cleo had some but not all of the figs. In this respect the case is no different from (3). But actually (8) is more problematic than (3), for it may license a further inference, as well, namely that, according to the *speaker*, Cleo didn't have all the figs. Where did that come from?

Chierchia (2004) claims that examples like (8) lend support to the conventionalist approach, but in order to explain his idea I will have to digress a little into the topic of presupposition. Factive verbs like "know" and "regret" are presupposition triggers, which is to say that they serve to convey, inter alia, that the speaker is taking it for granted that their complement is true. If I say, "Julius knows that Tony is a debaucher", I'm presupposing that Tony is a debaucher; what I'm asserting is merely that Julius knows it. In this respect, factive verbs are similar to definite descriptions, for example, which likewise trigger presuppositions. For example, if I say,

(9) Tony didn't have breakfast with the queen today.

I presuppose that there is a queen and I assert that Tony didn't have breakfast with her. (9) also exemplifies a hallmark of presuppositions: they tend to be invisible to scope-bearing operators like negation. As a rule, if I utter a sentence φ containing an expression that triggers a presupposition ψ, then it will be as if I had said "ψ and φ", no matter how deeply embedded the presuppositional expression may be. Hence, (9) will not normally be construed as, "It isn't the case that there is a queen that Tony had breakfast with", but rather as, "There is a queen and Tony didn't have breakfast with her." Similarly, despite the fact that the definite description is within the scope of "suspect", "may", and "not", (10a) would normally be construed as conveying (10b):

(10) a. Julius suspects that Tony may not have had breakfast with the queen today.
b. There is a queen and Julius suspects that Tony may not have had breakfast with her today.

This "projection behaviour", as it is generally known, makes sense in view of the fact that presuppositional information is taken for granted. A speaker who refers to "the queen" doesn't state that there is a queen; rather, he signals that he takes it as given that there is a queen, using the definite article to do so. As the foregoing sentences illustrate, what an utterance says is, at least in part, based on what it presupposes, rather than the other way round. Hence, the broad picture, which is widely accepted in the literature, is that presuppositions come before sentence meaning, which in its turn is the basis for calculating conversational implicatures.[2]

It follows from this that there is no way a bona fide implicature could be included in a presupposition. But that is what we *seem* to observe in examples like (8), where it looks as if the inference that Cleo didn't have all the figs is part and parcel of presupposition that Cleo had some of the figs. How could that be? Chierchia's answer is straightforward: the scalar inference is not an implicature. Rather, "not all" is part of the lexical content of "some"; therefore, "Cleo had some of the figs" already entails that she didn't have all the figs even before the sentence combines with the main verb; therefore, the sentence as

2. See, e.g., Geurts (1999a) or Beaver and Geurts (2010) for discussion and references. To say that this picture is widely accepted is to implicate, of course, that there is no full consensus. The most prominent detractor is Grice himself, who held rather idiosyncratic views on presupposition (Grice 1981).

a whole presupposes that Cleo had some but not all of the figs. Thus the conundrum posed by (8) is solved courtesy of a non-Gricean analysis of scalar inferences.

Disjunction The data discussed so far were problematic for the Gricean theory of Q-implicatures because it yielded predictions that were too weak. The following examples are potentially more serious, for in these cases the Gricean approach seems to yield incorrect predictions. To begin with, there is a problem with disjunction (Chierchia 2004, Sauerland 2004):

(11) Cleo either had a mango or some of the figs.

One of the stronger alternatives available to someone who utters (11) is:

(12) Cleo either had a mango or all the figs.

If the Gricean account is right, we should expect that (11) can give rise to the implicature that, according to the speaker, (12) is false, and if (12) is false, then Cleo had neither a mango nor all of the figs, which entails that she didn't have a mango. But that prediction is patently wrong. Prima facie, conventionalist theories make much more sensible predictions in this case: if "some" is construed as "some but not all", (11) comes out meaning that Cleo either had a mango or some but not all of the figs. Not only is this a plausible interpretation of (11), it also fails to predict that (12) is false, which is a good thing.

Indefinites A second type of case in which Gricean reasoning seems to produce incorrect results is the following. Suppose Tony wants to regale us with an anecdote about a particular Dutch sailor, Julius, and he starts off like this:

(13) Last year, a Dutch sailor showed some of the symptoms of delirium.

If the Gricean account is right, we should expect that (13) may give rise to the implicature that, according to Tony, (14) is false:

(14) Last year, a Dutch sailor showed all the symptoms of delirium.

But if (14) is false, then last year no Dutch sailor showed all the symptoms of delirium, and this is plainly wrong: Tony's story is about Julius, and was not intended to have implications for Julius's fellow sailors.

As in the disjunctive example in (11), a non-Gricean approach seems to deliver better predictions in this case. If the scalar inference is folded into the truth-conditional content of the verb phrase, (13) will end up meaning that,

last year, a Dutch sailor showed some but not all of the symptoms of delirium. This interpretation is a sensible one, but unattainable to an implicature analysis—or so it would seem.

Contrastive contexts The problem cases to be discussed in the remainder of this section are very different from the ones mustered so far, in that they involve quite marginal, and therefore strongly marked, uses of scalar expressions. This may be the reason why they haven't played much of a part in the back-and-forth between Gricean and conventionalist theories. Still, I believe these data deserve more attention than they have received, because they provide stronger prima facie support for the conventionalist approach than anything we have seen in the foregoing, though in the end I will argue that it doesn't get these facts right, either. But that argument will have to wait until the next chapter (§8.5).

By way of introduction, let us briefly recap from §3.2 some facts about negation:

(15) a. We like coffee.
b. We don't like coffee.

On the classical Gricean construal of scalar expressions (Horn 1972), the truth-conditional meaning of (15a) doesn't exclude the possibility that we *love* coffee, though it may be implicated that that isn't the case. On the Gricean account, (15b) simply means that we aren't sufficiently fond of coffee to warrant the claim that we like it, which entails that we don't love coffee, either. The diagram in (16) depicts how, according to this analysis, the truth-conditional meanings of "like coffee", "love coffee", and their respective negations relate to each other:

(16)
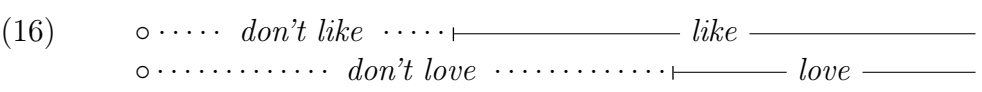

The negation of "like coffee" denotes all individuals outside the denotation of that predicate, which is to say that, since the meaning "like coffee" doesn't impose an upper bound on its denotation, someone who is in the denotation of "don't like coffee" cannot be in the denotation of "love coffee", either. Now, consider how the picture would change if "don't love coffee" was entailed by the truth-conditional content associated with "like coffee", as it might be on a conventionalist analysis:

(17)
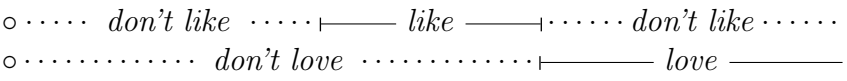

If this were the case, someone who doesn't like coffee might well be a coffee lover, and (15b) would not rule out the possibility that we love coffee, which doesn't seem to be right. Hence, the standard interpretation of negated sentences like (15b) buttresses the Gricean view that the meanings of scalar expressions such as "like" don't impose an upper bound on their denotations.

Unfortunately, there is a fly in the ointment, as the following examples collected by Horn (1989: 382) demonstrate:

(18) a. Around here, we don't LIKE coffee, we LOVE it.
b. I'm not HAPPY he's gone — I'm ELATED.

In (18a), "don't like" does *not* rule out "love", apparently, and likewise in (18b), the negation of "I'm happy" does not rule out that the speaker is more than happy. It is hard to see how we can avoid concluding that, in (18a), the truth-conditional meanings of "like" and "love" exclude each other, which is to say that in this case the picture in (17) gets things right and (16) is wrong; the same, mutatis mutandis, for "happy" and "elated", as these words are employed in (18b).[3]

It is a moral certainty that examples like (18a,b) are not instances of Q-implicature, for the simple reason that these construals are weaker than the standard reading of negation, as illustrated by (15b). (If it isn't immediately clear that this is so, note that, as construed in (16), "don't like" is more informative than it is according to (17): its denotation covers a smaller area in the former case than it does in the latter.) By the same token, the contrast between (15b) and (18a) is grist to the conventionalist mill. It very much looks as if in (18a) the truth-conditional meaning of "like" is restricted to "like but not love", and this is just the sort of thing conventionalist theories are designed to account for. True, it remains to be explained why this type of construal is abnormal, and why in the normal cases, exemplified by (15b), "like" does not have an upper-bounded meaning, but there may be ways of dealing with those problems, as we will see in the next section.

The examples in (18) are special in at least two ways. One is that they impose unusually strict interpretations on the scalar predicates they contain.

3. Alternatively, it might be suggested that the examples in (18) are instances of "metalinguistic negation" (Horn 1985, 1989). For reasons discussed at length (if not ad nauseam) by McCawley (1991), Carston (1996), and myself (1998a), I repudiate this option.

The other is that these interpretations are flagged by contrastive stress. Apart from denials such as (18a) and (18b), there are other constructions that exhibit these features. Take, for instance, comparatives:

(19) a. I'd rather have a WARM bath than a HOT one.
b. It's better to get FEW of the answers RIGHT than to get ALL of them WRONG.

If the lexical meaning of "warm" didn't exclude "hot", then the truth-conditional meaning of (19a) would be that having a bath that is warm-and-possibly-hot is preferred to having a hot bath. This would make little sense, because normally speaking one term of a comparison may not be entailed by another, as the following examples show:

(20) a. ?I'd rather live in France than in Paris.
b. ?Vegetables are healthier than peas.

Therefore, it seems more than likely that, in the context of (19a), "warm" means "warm but not hot". The same, mutatis mutandis, for "few" in (19b).

Disjunctions are similar to comparatives in that one term of "or" is not supposed entail the other (Hurford 1974):

(21) ?Cleo is allergic to peas or vegetables.

Nevertheless, we observe sentences like the following:

(22) a. Is a parallelogram SOMETIMES or ALWAYS a square?
b. MOST or ALL of the appartment must be infested with cockroaches.

Intuitively, the first option considered by (22a) is that a parallelogram may but need not be a square. Similarly, it would appear that, in the context of (22b), "most" has an upper-bounded meaning which excludes "all". Conventionalist theories can account for these construals; the Gricean theory of Q-implicatures cannot.

One last pair of examples (from Horn 2006) before we get to the general point:

(23) a. If you're convicted of a FELONY, you'll spend at least a year in jail. And if you're convicted of MURDER, you'll be executed.
b. If it's WARM, we'll lie out in the sun. But if it's VERY warm, we'll go inside and sit in front of the air-conditioner.

Given that murder *is* a felony, (23a) should make little sense, unless the speaker meant to convey that a murderer has to spend at least one year in jail before

being executed. Since this probably not what the speaker had in mind, the antecedent of the first conditional will have to receive an upper-bounded meaning which leaves murder out of account. Ditto for (23b).

It will be fairly evident, I trust, what the moral of the foregoing observations is going to be: there's something very special about contrastive sentences. If you make sure that it is properly contrasted with a stronger scalemate, "some" can be used to mean "some but not all", "warm" can mean "warm but not hot", etc. And by "mean" I mean the real thing: these are not implicatures, but honest-to-goodness truth-conditional effects, and therefore they cannot be separated from the truth-conditional content of the sentence. In this respect, too, there is a stark contrast between "embedded implicatures" in contrastive sentences and, e.g., belief reports. For instance, if (24a) is understood as implying that Tony believes that Cleo didn't have all the figs, we can paraphrase it as (24b):

(24) a. Tony believes that Cleo had some of the figs. (= (3))
b. Tony believes that Cleo had some of the figs and he believes that she didn't have them all.

By contrast, there is no way you could paraphrase (25) along the same lines, separating "don't love" from the truth-conditional content of "like":

(25) Around here, we don't LIKE coffee, we LOVE it. (= (18a))

Here, the negative statement may be paraphrased as "Either we don't like coffee or we don't love it", but crucially, this is a disjunctive proposition, whose second half is not detachable. Put otherwise, on a contrastive construal, a scalar expression acquires an upper-bounded meaning which is truly local and enters into the truth-conditional content of the sentence.

Time to take stock. In the foregoing, we have reviewed two types of examples that cause embarrassment to the Gricean theory of Q-implicatures. On the one hand, there are the contrastive sentences we discussed last. On the other hand, we have collected a ragbag of non-contrastive cases in which the Gricean theory either fails to make the right predictions, as in (26) and (27), or seems to make predictions that are incorrect, as in (28) and (29):

(26) *Belief reports:*
Tony believes that Cleo had some of the figs.
$\rightsquigarrow$ Tony believes that Cleo didn't have all the figs.

(27) *Factive verbs:*
Tony knows that Cleo had some of the figs.
$\rightsquigarrow$ Cleo didn't have all the figs.

(28) *Disjunctions:*
Cleo either had a mango or some of the figs.
$\not\rightsquigarrow$ Cleo didn't have a mango.

(29) *Indefinites:*
Last year, a Dutch sailor showed some of the symptoms of delirium.
$\not\rightsquigarrow$ Last year, no Dutch sailor showed all the symptoms of delirium.

In the next chapter, I will try to show that, initial impressions notwithstanding, the second batch of observations can be explained in terms of Q-implicature, after all, but this doesn't hold for the contrastive cases. Hence, from a Gricean perspective, there cannot be a unified account of all the data surveyed in this section. One of the main attractions of the conventionalist approach is that, perhaps, it can offer such a treatment.

Before we turn to a conspectus of the varieties of conventionalist theories currently on the market, I would like to enter a caveat. Having gone through a series of problematic facts, you shouldn't come away with the impression that "embedded implicatures" are a rampant phenomenon. They aren't. For one thing, contrastive interpretations of scalar expressions are out of the ordinary by definition: that's why they are "marked". For another, each of the examples in (26)-(29) represents an isolated case. For instance, as pointed out before, the problem caused by belief sentences like (26) does not extend to attitude reports with "hope", "want", and so on. Therefore, it is still fair to say that *on the whole* the Gricean theory of Q-implicatures gets the facts right. All the putative counterexamples discussed in this section are special in one way or other. The question is whether conventionalist theories, which were designed with these special cases in mind, can offer a general account that is anywhere nearly as convincing as the Gricean one. There are no bonus points for guessing what my answer is going to be.

7.2 Varieties of conventionalism

The fundamental intuition underlying all conventionalist theories is that scalar inferences really aren't inferences, to begin with, and a fortiori they aren't pragmatic inferences. Rather, a scalar inference is part of the lexico-grammatical content of a sentence, and not too different from conventional truth-conditional meaning. Whence the label "conventionalist".

Conventionalist theories come in two main families: lexicalist and syntax-based. According to theories of the lexicalist persuasion, "some" implies "not all" by virtue of its lexical content, though of course the implication is cancellable, which regular truth-conditional content is not. Syntax-based theories, on the other hand, maintain that a scalar expression acquires an upper-bounded interpretation when it is within the scope of a covert syntactic operator whose meaning is essentially that of overt "only". When in the scope of this covert operator, "some" will be interpreted, in effect, as "only some", and thus come to exclude "all".

Lexicalist variants of conventionalism have been around since the early 1970s (Cohen 1971), and since they are relatively easy to explain, we will start with them. Syntax-based theories are a recent development (Chierchia 2006, Fox 2007); they are more complex than their lexicalist brethren, but also more versatile in their predictions, as we will see.

Lexicalist conventionalism According to the beginner's version of lexicalism, upper bounds are hard-wired into the lexical entries of scalar expressions like "many", "or", and "like".[4] These upper bounds are cancellable, and therefore they cannot be part of lexical meaning in the conventional sense, but they are part of the lexical content of a word. This simple account readily explains stock-in-trade observations like the following:

(30) Tony had many of the symptoms of delirium.
$\rightsquigarrow$ He didn't have all the symptoms.

If one embraces this version of lexicalism, it is natural to assume that strong implicatures are the norm: if "not all" is pre-packaged with the lexical content of "many", it naturally follows that (30) implies by default that, according to the speaker, Tony didn't have all the symptoms of delirium. We have seen

4. What I call "lexicalism" is formally indistinguishable from the view that, e.g., "some" is ambiguous between a strong and a weak reading, where the former is preferred. The first explicit defence of the idea that many conversational implicatures are better analysed as lexical inferences is Cohen (1971, 1977). Cohen didn't extend this view to scalar implicatures, though he certainly would have done so if prompted. More recently, Levinson (2000) and Chierchia (2004) have proposed theories which are effectively lexicalist, though for these authors the mechanism that associates, e.g., "not all" with "some" is not purely lexical. (Incidentally, since his 2004 paper, Chierchia has moved from a lexicalist to a syntax-based position; see Chierchia 2006, Chierchia et al., to appear). Wainer (2007) presents a formalisation of a lexicalist theory of scalar inferences (among others), using non-monotonic logic.

in Chapter 5 that there are rather compelling arguments against the notion that such inferences have default status, but more importantly for the current discussion, there are problems with embedding, as well:

(31) Tony didn't have many of the symptoms of delirium.

As discussed in the last section, the preferred reading of (31) is that Tony had less-than-many symptoms of delirium. The beginner's version of lexicalism doesn't predict this; what it predicts is that the many-but-not-all reading of (31) should be the default, which is clearly wrong.

There are various ways of getting around this problem, but they all have a common core: they drop the assumption that the preferred reading of a scalar expression is determined by its lexical content alone, since examples like (31) show that the linguistic environment in which such expressions appear has a role to play, as well. Bearing this in mind, we could amend our initial version of lexicalism along the following lines:

(32) *Enhanced lexicalism (strong version)*
A scalar expression α prefers an upper-bounded construal unless this weakens the interpretation of the sentence in which α occurs.

In (31), an upper-bounded reading of "many" weakens the statement made by the speaker, and therefore (32) predicts that, in this case, the preferred construal of "many" is without upper bound. If there is a strong version of enhanced lexicalism, then of course there must be a weak version, too:

(33) *Enhanced lexicalism (weak version)*
A scalar expression α prefers an upper-bounded construal provided this strengthens the interpretation of the sentence in which α occurs.

This is similar to (32) and makes the same predictions in many cases. In particular, both principles suppress upper-bounded readings in downward-entailing contexts like (31). But they are not equivalent: (33) is the weaker of the two, since it predicts preferred upper-bounded readings in fewer instances than (32) does. To illustrate the difference, compare the following sentences:

(34) a. All the senators supported most of the bills.
$\overset{?}{\rightsquigarrow}$ All the senators supported most but not all of the bills.

b. Thirty nine senators supported most of the bills.
$\overset{?}{\rightsquigarrow}$ Thirty nine senators supported most but not all of the bills.

Whereas in (34a) "most" occurs in an upward-entailing environment, its context in (34b) is neither upward nor downward entailing. Consequently, a

stronger interpretation of "most" will result in a stronger proposition in the former case, while in the latter case the resulting proposition will be neither weaker nor stronger: "Thirty nine senators supported most but not all of the bills" may be true while "Thirty nine senators supported most and possibly all of the bills" is false, and vice versa. Given these entailment properties, (32) predicts that an upper-bounded construal of "most" will be preferred in both cases, whereas (33) only predicts this for (34a). In the next section, I will argue that both predictions are false, but I'm happy to concede that either principle constitutes an improvement over the beginner's version of lexicalist conventionalism. Let us ponder a bit why. The main innovation is that both varieties of enhanced lexicalism have relinquished the simplistic idea that scalar inferences are entirely arranged in the lexicon (cf. §3.2). The linguistic environment in which a scalar expression appears plays an essential part, as well, which means that the original lexicalist intuition that, e.g., "not all" is part of the lexical content of "most" will have to be compromised somewhat.

It will have to be compromised more in order to account for the following examples:

(35) a. Cleo didn't have many of the figs.
$\rightsquigarrow$ She did have some of them.
b. You can have some of the dates.
$\rightsquigarrow$ You cannot have them all.

None of the versions of lexicalism presented so far can capture these inferences. The problem with (35a) is that the scalar inference requires that we take into account the weaker scalemates of "many": since "many" occurs in a downward-entailing context, (36) would have been a stronger statement,

(36) Cleo didn't have any of the figs.

and the inference in (35a) results from denying (36). Clearly, any theory which merely assumes that "many" may or may not have an upper-bounded reading cannot explain this.

(35b) causes trouble in a rather different way. To see how, consider (37), which paraphrases the lexicalist interpretation of (35b), with an upper-bounded construal of "some":

(37) You can [have some and not all of the dates].

Since the "not all" rider is inserted in the scope of "can" (it has no other place to go, if it is part of the lexical content of "some"), (37) actually *permits* the addressee not to have all the dates, rather than *forbidding* him to eat the

lot. The heart of the problem is that the preferred reading of (35b) requires wide-scope negation and narrow-scope "all": ¬[you can have all]. A lexicalist analysis obviously cannot provide this.

Syntax-based conventionalism As I hinted already, these problems, but especially the last one, can only be solved by compromising lexicalist principles even further, which brings us to the syntax-based branch of conventionalism. Syntax-based theories stipulate that the grammar is equipped with a covert operator, whose meaning is essentially that of overt "only". The intuitive motivation underlying this hypothesis is that, as observed by Levinson (1983), at least some scalar implicatures can be made explicit by means of overt "only":

(38) a. Julius had some of the symptoms of dengue fever.
b. Julius had only SOME of the symptoms of dengue fever.

When construed as implying that Julius didn't have all the symptoms of dengue fever, (38a) can be paraphrased as (38b), where "only" associates with "some". The intuition informing the syntax-driven programme is that, whenever (38a) licenses this scalar inference, its underlying syntactic form resembles that of (38b) in that (38a), too, contains a form of "only", albeit a covert one, which I will render as "$\mathcal{O}$".

As it turns out, the interpretation of $\mathcal{O}$ is rather tricky (see Fox 2007 for discussion), but (39) gives the gist of it:

(39) *Semantics of* $\mathcal{O}$
$\mathcal{O}(\varphi)$ is true iff φ is true and every alternative not entailed by φ is false.

To explain how this works, let's consider sentence (40a), and assume that the relevant alternatives, besides (40a) itself, are (40b) and (40c):

(40) a. Many of the barbarians were pacified.
b. Some of the barbarians were pacified.
c. All the barbarians were pacified.

Of these alternatives, the only one that is not entailed by (40a) is (40c), and therefore, according to the definition in (39), $\mathcal{O}$(40a) is true iff (40a) is true and (40c) is false. Hence, the meaning of $\mathcal{O}$(40a) is that many but not all of the barbarians were pacified.

Although syntax-based versions of conventionalism are considerably more complex than their lexicalist cousins, they are more powerful, too, simply because, in general, they will generate more readings for any given sentence, owing to the fact that $\mathcal{O}$ is assumed to be a syntactic operator that may be

freely inserted into the parse tree. Hence, there are many different parses for (41), for example, the most interesting of which are displayed in (42).[5]

(41) Cleo didn't have many of the figs. (= (35a))

(42) a. ¬[Cleo had many of the figs]
b. ⫛¬[Cleo had many of the figs]
c. ¬⫛[Cleo had many of the figs]

Besides the ⫛-less parse in (42a), there is an analysis in which ⫛ takes scope over the negation and one in which it is the other way round. (42b) expresses the "many but not all" reading of (41); (42c) represents the marked reading, which contrary to (42a) and (42b) is consistent with Cleo's eating all the figs. Let's see how these interpretations come about.

In (42b), ⫛ combines with ¬[Cleo had many of the figs], so the relevant alternatives might be the following:

(43) a. ¬[Cleo had some of the figs]
b. ¬[Cleo had many of the figs]
c. ¬[Cleo had all the figs]

Since negation reverses the direction of entailment, (43a) is stronger than (43b), which in its turn is stronger than (43c). Therefore, according to the semantics of ⫛ as given in (39), ⫛(43b) (= (42b)) is true iff (43b) is true and (43a) is false, or in other words: Cleo had less than a large proportion of the figs, but she did have some. This is the unmarked interpretation of (41), but enriched with a scalar inference that lexicalist theories cannot account for, as we have seen.

Concerning the syntactic parse in (42c), where ⫛ combines with [Cleo had many of the figs], let us suppose that the relevant alternatives are the following:

(44) a. Cleo had some of the figs.
b. Cleo had many of the figs.
c. Cleo had all the figs.

5. In fact, unless ⫛-insertion is prevented from applying recursively, every sentence will have infinitely many syntactic analyses. For example, in addition to the parses in (42), we also have the following analyses for (41):

⫛⫛¬[Cleo had many of the figs]
⫛¬⫛[Cleo had many of the figs]
⫛⫛¬[Cleo had many of the figs] … and so forth

⫛-recursion tends to be truth-conditionally inert, but it doesn't have to be. Recursive application of ⫛ is a key ingredient in Fox's (2007) proposal for dealing with free choice inferences.

$\mathbb{O}$(44b) is true iff (44b) is true and (44c) is false; that is to say, iff Cleo had many but not all of the figs. Hence, $\neg\mathbb{O}$(44b) (= (42c)) is true iff Cleo had either less than a large proportion of the figs or all of them. This is the marked reading of (41), which would be opportune in the context of (45), for example:

(45) Cleo didn't have MANY of the figs: she had them ALL.

In general, $\mathbb{O}$-theories will generate a great many readings for any given sentence. In the case of (41), free $\mathbb{O}$-insertion will yield at least three distinct interpretations, and there is only one scope-bearing expression to reckon with. Had there been more, the number of readings might have been considerably higher. In order to convey an impression of how quickly this number will escalate, consider (46), where "many" occurs in the syntactic scope of three expressions: "some", "thought", and the negation:

(46) According to some sources, Tony thought that Cleo didn't have many of the figs.

If we allow ourselves to use at most *two* copies of $\mathbb{O}$, this sentence will have no less than *sixteen* different parses, most of which come with distinct meanings. Evidently, the formidable predictive power of syntax-based theories comes with a hefty price tag: most of the readings they predict are not wanted. Hence, this family of theories is burdened with an equally formidable disambiguation problem: how to choose amongst all those readings?

There are various ways of dealing with this issue, but for the moment let's confine our attention to the only one that has been seriously considered in the literature. The idea is to adopt what I will call the "Power Principle":

(47) *Power Principle*
If we have two readings, φ and ψ, and φ is stronger than ψ, then φ is the preferred reading.

(Note, incidentally, that both versions of enhanced lexicalism conform to this principle, too.) In the case of (41), for example, the Power Principle predicts that the reading in (42b) is preferred to the ones in (42a) and (42c). Surely, this prediction is at least partly correct: (42c) is a marked reading by any standard, and should be less preferred than either of the others, which according the Power Principle it is. In view of the discussion of defaults in Chapter 5, it is more doubtful that (42b) should be preferred to (42a); we will return to this issue shortly.

Although the Power Principle takes the sharp edges off the ambiguity problem, it should be noted that it will not solve the problem completely. The

trouble is that, in general, amongst the various readings generated by free $\mathbb{D}$-insertion, there will be some that are incommensurable with respect to strength; or more accurately, there may be pairs of readings that aren't entailed by a common third reading and neither of which entails the other. In such cases, the Power Principle will not resolve the ambiguity, and contextual factors will have to decide which is best.[6]

To conclude the presentation of the syntax-based approach to "embedded implicatures", let us see how it deals with:

(48) You can have some of the dates. (= (35b))

Recall that lexicalist versions of conventionalism fail to explain how (48) can give rise to the inference that the addressee is not allowed to have all the dates, the reason for this failure being that, if "not all" is associated directly with "some", it is bound to remain within the scope of the modal "can", whereas on the reading we need to account for, the negation should take scope over "can", while leaving the universal quantifier in situ. Which is precisely what $\mathbb{D}$-insertion will give us if it is applied to the entire sentence. To see that it will, suppose that, apart from (48) itself, there is only one alternative to consider, namely:

(49) You can have all the dates.

Then $\mathbb{D}$(48) is true iff (48) is true and (49) is false, which is to say that the addressee can have some of the dates but is not allowed to have them all. Problem solved.

The common conventionalist core Conventionalist theories of scalar inference differ from each other along several dimensions, the most important of which is that some are lexicalist, whereas others are syntax-driven. Either way, conventionalist theories generate ambiguities which will have to be resolved somehow; this problem is much worse for syntax-based theories, but it affects their lexicalist cousins, too. The most promising way of solving, or at least mitigating, the ambiguity problem is by adopting some version of the Power Principle, and it turns out that different formulations of this principle may yield rather different results. The same holds true for the definition of the $\mathbb{D}$-operator (Chierchia et al., to appear). In short, if you go shopping for

6. Actually, the situation is somewhat more complex, because the precise definition of the Power Principle is a factor, too. But the main point holds true in any case. See Chierchia et al. (to appear) for discussion.

a conventionalist theory, there are quite a few choices to be made, and the variety of possible theories may seem a bit daunting.

Still, there are commonalities, too. First off, it bears emphasising that conventionalist theories are non-pragmatic to the core, and in this sense are stipulative by nature. Basically, they solve the problem of embedded implicatures by fiat. This is particularly evident in the lexicalist versions of conventionalism, which stipulate an upper-bounded lexical meaning for each scalar expression. But it is no less true of syntax-based versions of the conventionalist doctrine. Hence, the conventionalist approach is at a methodological disadvantage from the start.

The second common thread in conventionalist theories is that, having been designed specifically for dealing with "embedded implicatures", they are committed to the view that upper-bounded interpretations systematically occur in embedded positions. Consequently, by their very nature, conventionalist theories are prone to overgenerate readings, and typically seek to relieve this problem by adopting some version of the Power Principle. As a result, these theories tend to converge on a range of predictions. For example, as we have seen in the foregoing, most lexicalist and syntax-based theories agree that denial uses of negation, as in (50), are marked:

(50) Around here, we don't LIKE coffee, we LOVE it. (= (18a))

That this use of the negation is dispreferred is explained by the fact that, on this reading, the truth-conditional meaning of "We don't like coffee" is quite weak, compared to other possible construals. Hence, the Power Principle may be called upon to explain why denial readings are exceptional.

Since most conventionalist theories adopt a version of the Power Principle, they agree that scalar expressions occurring in upward-entailing environments will prefer upper-bounded construals:

(51) a. All the ladies were wearing pink gloves or shawls.
$\rightsquigarrow$ All the ladies were wearing pink gloves or shawls but not both.
b. Most of these children stutter some of the time.
$\rightsquigarrow$ Most of these children stutter some but not all of the time.

It remains to be seen whether these predictions are correct, but for the moment my point is merely that this is what most conventionalist theories predict. Furthermore, many though not all varieties of conventionalism are committed to the view that the same holds true of non-monotonic environments; that is, contexts that are neither upward nor downward entailing:

(52) a. Twelve ladies were wearing pink gloves or shawls.
⇝ Twelve ladies were wearing pink gloves or shawls but not both.
b. About half of these children stutter some of the time.
⇝ About half of these children stutter some but not all of the time.

We have seen already that the strong version of enhanced lexicalism predicts that these inferences hold by default; the weak version does not make this prediction. As to the syntax-based branch of conventionalism, it depends on the details of the Power Principle whether or not a theory of this family is committed to the inferences in (52). On the definition of the Power Principle given in (47), these inferences don't follow, but on some other formulations of the principle they do (Chierchia et al., to appear).

To sum up, while conventionalist theories generally predict that scalar expressions occurring in upward-entailing environments will have a preference for upper-bounded construals, some but not all of them make this prediction for non-downward-entailing environments in general. In the upcoming section, we will see that both predictions are wrong.

7.3 Against conventionalism

In this section, I will argue that conventionalism is an unpromising doctrine, even when judged on its own merits. (The rider is not gratuitous, because it should always be borne in mind that, in the final analysis, conventionalism is not to be judged on its own merits. Rather, the question to ask is whether conventionalism offers a significant improvement over the Gricean approach to Q-implicatures, for if it doesn't, we should stick with the latter on methodological grounds. Cf. §1.3.) I will begin by showing that there are some phenomena that are beyond the scope of any conventionalist theory, but are readily explained by implicature theories. This is an important point not only because it shows that conventionalism has its limits, but also because it means that conventionalists are compelled to carve up a range of empirical data for which Gricean theories offer a unified explanation.

Having charted the natural limits of conventionalism, I proceed to show that conventionalist theories make false predictions in a wide variety of cases, focusing my attention first on intensional contexts created by attitude verbs and modals, and then on quantificational contexts. The upshot of this discussion will be that, against the most fundamental conventionalist tenet, upper-bounded interpretations do *not* occur systematically in embedded positions.

While this is bad news for conventionalism, it is good news for Gricean theories, because it is precisely what they predict.

The limits of conventionalism For all its predictive potential, there are a number of phenomena that are beyond the purview of any conventionalist approach. By way of introducing the first such phenomenon, consider the following scenario. Tony asks Cleo, "Which places did you see on your trip to Italy?", and she answers:

(53) I saw Naples, Rome, and Ravenna.

Cleo's statement is likely to license the inference that Naples, Rome, and Ravenna are the only places she visited. While lexicalist variants of conventionalism will be hard put to account for this inference, it isn't much of a problem for syntax-based theories. All they have to assume is that the alternatives to (53) are of the form "I saw x", where x is either an Italian city or a list of Italian cities, and then stick a copy of $\mathcal{O}$ in front of the sentence: with these alternatives, $\mathcal{O}$(53) entails that Cleo only saw Naples, Rome, and Ravenna.[7] So far so good. But now consider the following variation on the example in (53). Tony asks the same question, but now Cleo says:

(54) Julius and I first went to Naples and Rome together. Then, while he went to see Milan, I visited Ravenna.

In the given context, this answer can (and probably will) license the same inference as did the previous one, but now the inference derives from a *pair* of sentences, and since the grammar can only insert copies of $\mathcal{O}$ into the syntactic analysis of individual sentences, there is no way this inference could accounted for by standard conventionalist means.

Needless to say (but I'll say it anyway), this is very odd. Intuitively, it shouldn't make a difference for the inference in question whether Cleo chooses to convey her travel information by way of one sentence or two, and a Gricean theory need not distinguish between these cases: the inference is derivable as a regular Q-implicature either way (cf. §1.6). Therefore, whereas conventionalists can claim to have an analysis for the one-sentence case, they will have to leave the multi-sentence cases to pragmatics; the implication being that the two cases are unrelated. Which, as I said, is odd.

7. This is merely to say that an analysis along these lines is technically feasible, not that it is plausible. See the introduction to Chapter 6 and §6.6 for discussion.

And that's not the end of it, as the following stock-in-trade type of example will serve to show:

(55) Some of the figs are poisoned.
$\rightsquigarrow$ Not all the figs are poisoned.

Recall how Gricean theories handle this case: first a weak implicature is derived to the effect that $\neg\textsc{Bel}_S$(all the figs are poisoned), and if the Competence Assumption holds, it follows that $\textsc{Bel}_S(\neg$(all the figs are poisoned)). By contrast, conventionalist theories derive the stronger inference directly, and there is no way they could derive the weaker one. On a conventionalist construal, (55) has two readings, "At least some of the figs are poisoned" and "Some but not all of the figs are poisoned", and neither of these captures the weak implicature routinely predicted by the Gricean account.

This is odd, too. Whereas on a Gricean view weak and strong implicatures are intimately related, for the conventionalist they couldn't be more different. On the conventionalist view, strong implicatures aren't implicatures; rather, they are aspects of sentence meaning. Conventionalism doesn't have a story about weak implicatures, but since they cannot be analysed in conventionalist terms, it would seem to follow that they have little if anything to do with strong implicatures. I take it that the Gricean picture is a lot more plausible than the one painted by the conventionalist.

If weak implicatures cannot be explained by conventionalist means, how are they to be accounted for? Chierchia et al. (to appear) suggest that they are to be taken care of by Gricean pragmatics. This suggestion strikes me as profoundly odd, if only because this impromptu recognition of Gricean pragmatics would entail massive redundancy. For, if the conventionalist is willing to leave weak implicatures to pragmatics, they will be forced to concede, by their own theoretical lights, that whenever a sentence φ gives rise to an inference of the form $\textsc{Bel}_S(\neg\psi)$, where ψ is stronger than φ, this inference can always be explained in two ways: as an entailment or as a Q-implicature.

This corollary is curious enough to be belaboured a bit. Consider sentence (55) again, and suppose it is uttered in a context in which it licenses the inference that $\textsc{Bel}_S(\neg$(all the figs are poisoned)). Conventionalist theories propose to explain this in lexico-grammatical terms: on such occasions, they maintain, (55) *means* that some but not all of the figs are poisoned, and it follows, assuming S is sincere, that $\textsc{Bel}_S(\neg$(all the figs are poisoned)). However, if we adopt Chierchia et al.'s suggestion, we may as well bypass the conventionalist machinery, and allow Gricean pragmatics to derive $\neg\textsc{Bel}_S$(all

the figs are poisoned), which will be strengthened to $\textsc{Bel}_S(\neg$(all the figs are poisoned)) provided the Competence Assumption holds. In short, Chierchia et al.'s suggestion pulls the rug from under their own theory.

The long and short of it is that conventionalism has its limits: it cannot aspire to be a wholesale replacement of the Gricean theory of Q-implicature. Hence, we should ask ourselves which of the frameworks can provide us with a more coherent treatment of the facts: one that offers an integrated analysis of weak and strong implicatures and a unified account of one-sentence and multi-sentence implicatures, or one that does neither?

False predictions: intensional contexts The main problem with the conventionalist approach is that, no matter what version one fancies, it is bound to generate predictions that are dubious if not downright false. Here are a few examples:

(56) a. I hope that some of my relatives will remember my birthday.
$\overset{?}{\rightsquigarrow}$ I hope that not all of them will remember it.
b. You have to read some of these papers.
$\overset{?}{\rightsquigarrow}$ You're not allowed to read all of them.
c. I want you to give me some of your biscuits.
$\overset{?}{\rightsquigarrow}$ I want you not to give me all of your biscuits.

Conventionalist theories are designed to predict that these inferences hold. Take (56b), for instance. If we adopt the syntax-based approach, this sentence will have the following parses in addition to an $\mathbb{O}$-less analysis:

(57) a. $\mathbb{O}$[you have to read some of these papers]
b. you have to $\mathbb{O}$[read some of these papers]

On the wide-scope analysis in (57a), (56b) rules out the possibility that the addressee has to read all the papers, which is weaker than the reading associated with (57b), according to which the addressee is obliged to read some but not all of the papers, and is therefore not allowed to read all of them. Since the second reading is stronger than the first one, and both are stronger than the $\mathbb{O}$-less reading, syntax-based theories predict that (57b) is the preferred analysis, and hence that the inference in (56b) should hold by default. Lexicalist theories won't be able to generate a reading corresponding to the analysis in (57a), but otherwise they will make the same predictions. In the case of (56a) and (56c), the story is along the same lines. In short, for all the sentences in (56), conventionalist theories predict that the preferred reading should be one

in which "some" is interpreted, in effect, as "some but not all", and therefore the mooted inferences should hold by default.

According to my intuitions, these predictions are plain false, and my intuitions are confirmed by experimental data reported by Geurts and Pouscoulous (2009), who found that their subjects endorsed inferences like (56b) and (56c) less than 20% of the time, on average. Contrary to what conventionalist theories predict, upper-bounded interpretations are strongly dispreferred in these cases. This conclusion is reinforced by a fact we established in Chapter 5 (§5.3), namely, that the way you probe speakers' intuitions about Q-implicatures may have an effect on the outcome. If you ask someone (yourself, for example) whether any sentence containing "some" implies the negation of the corresponding sentence with "all", then this very procedure will induce at least some bias towards a positive response. Taking this fact into account, 20% is likely to be an overestimation, which only makes things worse for conventionalism.

It would seem, though, that there is an obvious way out for conventionalists, since they might argue that, in these particular examples, the default readings predicted by his theory are cancelled, simply because they are implausible. Unfortunately for the conventionalist, this line of defence will not wash. The cancellation argument presupposes that preferred scalar inferences will quietly withdraw whenever they happen to be implausible, and this presupposition is doubtful, as it is contradicted by the behaviour of bona fide scalar inferences. Let me give a couple of examples to establish this point.

(58) Some of the liberal parliamentarians voted against the bill.

In some democracies, parliamentary fractions tend to vote *en bloc*. In the Netherlands, for instance, it would be quite unlikely that some but not all liberal parliamentarians voted against a bill. Nevertheless, as far as I can tell, this wouldn't diminish the likelihood that (58) is interpreted as implying that the liberal fraction was divided.

(59) In order to prevent the rinderpest from spreading through the herd, some of the cows were vaccinated.

Since rinderpest is a highly contagious disease, it would be decidedly odd if only some of the cows were vaccinated, yet that is how we would understand (59).

(60) Cleo threw all her marbles in the swimming pool. Some of them sank to the bottom.

No doubt, it would be *very* odd if some of Cleo's marbles failed to sink, yet according to my intuitions that is precisely what (60) conveys.

What these examples show is that *genuine* scalar inferences are not so easy to cancel, at all (cf. Geurts 1999a). From a Gricean perspective, this is to be expected: scalar implicatures arise because the speaker goes out of his way to make a statement that is weaker than what he could have said with equal effort, so it stands to reason that it should require special circumstances to suppress a scalar implicature. But in particular, lack of plausibility will generally be insufficient for doing so.

Summing up, the cancellation argument is based on the supposition that the preference for scalar inferences is so weak that it will be overridden whenever the otherwise preferred reading is even mildly improbable. The foregoing observations show that this supposition cannot be justified on independent grounds. Therefore, the fact remains that the example sentences in (56) falsify conventionalist predictions.

More false predictions: quantificational contexts According to my intuitions, none of the following inferences would go through under normal circumstances:

(61) a. All the farmers in this region own goats or sheep.
$\not\rightsquigarrow$ None of them own both.
b. At least 300 of the farmers in this region own goats or sheep.
$\not\rightsquigarrow$ At least 300 of them don't own both.
c. Only six farmers in this region own goats or sheep.
$\not\rightsquigarrow$ Only six farmers in this region own goats or sheep but not both.

However, conventionalist theories of scalar inference unanimously predict that (61a) and (61b) should hold by default, and at least some of them predict this for (61c), as well. In (61a,b), "or" occurs in an upward-entailing environment, and therefore all conventionalist theories agree that an exclusive construal should be preferred. In (61c), "or" occurs in an environment that is neither upward nor downward entailing, and therefore an exclusive reading is predicted by, e.g., the stronger version of enhanced lexicalism and by some syntax-based theories.

These predictions were tested by Geurts and Pouscoulous (2009) in a verification experiment, using sentences like the following:

(62) a. All the squares are connected with some of the circles.
b. There is more than one square that is connected with some of the circles.

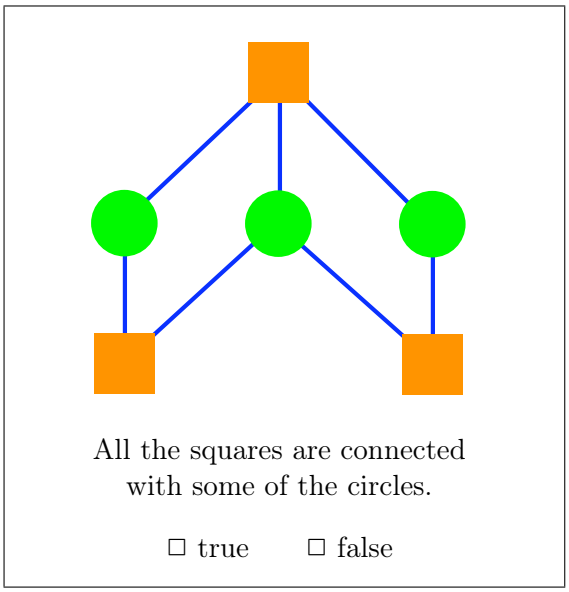

Figure 1: Experimental item used in the verification study by Geurts and Pouscoulous (2009).

c. There are exactly two squares that are connected with some of the circles.

Note that these sentences are parallel to the ones in (61), and therefore all conventionalist theories agree that, in (62a) and (62b), there is a preference for reading "some" as "some but not all", while at least some brands of conventionalism make the same prediction for (62c). In Geurts and Pouscoulous's experiment, each of these sentences was paired with a picture in which the sentence's classical and conventionalist construals yielded conflicting truth values. For example, when interpreted according to the conventionalist canon, the sentence in Figure 1, i.e. (62a), fails to match the depicted situation, because one of the squares is connected to all the circles, but the sentence is true on its classical reading.

Geurts and Pouscoulous's findings were as univocal as they get: not a single one of their participants' responses agreed with conventionalist predictions. For example, participants invariably gave a positive response in the case of (62a) (Figure 1), and the same held, mutatis mutandis, for the other conditions, regardless whether they were upward entailing or merely non-downward-entailing. Hence, these results equally disconfirm weak and strong versions of conventionalism alike.

Thus far, we have been concerned with what we might call "mainstream conventionalism": a spectrum of theories which agree that scalar expressions preferably receive an upper-bounded construal in upward-entailing environments and perhaps in non-downward-entailing environments generally. It seems that these predictions are wrong. What to do? One option would be to retreat into a weaker version of conventionalism that refuses to deal with the ambiguities it begets, and doesn't make any predictions about which construal is the preferred one. Leave it to … pragmatics.

According to this minimal version of conventionalism, each of the sentences in (62) has several readings on any given occasion, but it is for pragmatics to decide which is the right one. To the best of my knowledge, nobody has come forward to advocate such a minimalist take on conventionalism, as yet, and I'm tempted to say that this is just as well, since this view is too weak to be taken seriously. Nonetheless, it instructive to consider how minimal conventionalism might be tested, precisely because it is such a weak position. If this view could be disproved it would be a crippling blow indeed.

Obviously, in its most austere form, minimal conventionalism fails to make any predictions at all: the claim that, e.g., (63a) may or may not be read as (63b) cannot be falsified.

(63) a. All the squares are connected with some of the circles.
b. All the squares are connected with some but not all of the circles.

In order for the theory to have at least a modicum of empirical bite, auxiliary assumptions will have to be made. One hypothesis that naturally comes to mind is that native speakers are able to *detect* the readings that are distinctive of the conventionalist approach. That is to say, even if speakers don't prefer to hear (63a) as synonymous with (63b), it is reasonable to expect them to realise that the latter sentence expresses a possible construal of the former. If this auxiliary hypothesis is adopted, minimal conventionalism predicts that, in a situation that falsifies (63b), native speakers should judge that (63a) is ambiguous between a true reading and a false one. This prediction was tested in the last experiment of Geurts and Pouscoulous (2009)

For this experiment, Geurts and Pouscoulous used the same materials as in the verification study discussed above, but added a third response option: participants were invited to choose between "true", "false", and "could be either". According to minimal conventionalism, participants should have a preference for the last option in the upward-entailing conditions, and perhaps in non-downward-entailing conditions generally. In order to assess whether partici-

pants can cope with this type of task, in the first place, Geurts and Pouscoulous included a number of control sentences which were genuinely ambiguous, for instance:

(64) a. The circles and the squares are connected with each other.
   - The circles are connected with the squares.
   - The circles are connected with each other and the squares are connected with each other.

   b. There are green circles and squares.
   - There are green circles and there are squares.
   - There are green circles and there are green squares.

A sample control item is shown in Figure 2.

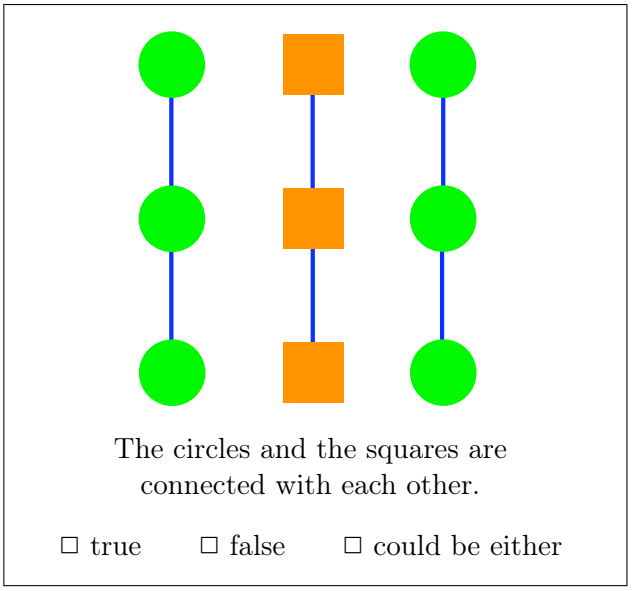

Figure 2: Control item used in Geurts and Pouscoulous's (2009) ambiguity experiment.

The main finding of this study was that, whereas ambiguous controls were recognised as such in 70% of the cases, the corresponding rate for the non-downward-entailing sentences was 6% while for the upward-entailing sentences it was 0%. On a further 5% of the non-downward-entailing trials, participants gave a true/false response that was consistent with a conventionalist construal. In raw numbers, only 9 out of 88 responses were consistent with minimal conventionalism. Moreover, all but one of these "conventionalist" responses

were associated with "exactly two": between them, upward-entailing "all" and "more than one" prompted a conventionalist response on *one trial* only.

In short, while they were rather good at picking out undisputedly ambiguous sentences, participants overwhelmingly failed to produce responses consistent with a conventionalist analysis of scalar expressions, either by classifying target sentences as ambiguous or by assigning them a truth value in line with a conventionalist construal. In my opinion, this is *very* strong evidence against the conventionalist approach to scalar inferences.

7.4 Conclusion

I have argued that the conventionalist approach to scalar inference has two major empirical issues to contend with. On the one hand, we have seen that in some cases conventionalist theories fail to make the right predictions. On the other hand, there are quite a few cases in which conventionalist predictions are just false. As far as I can see, the first issue is incontrovertible: there is no way a genuinely conventionalist theory can hope to explain weak implicatures or multi-sentence implicatures. However, with the second issue we aren't quite done yet. We have seen that conventionalist theories make erroneous predictions about scalar expressions that occur in the scope of various quantifiers and intensional verbs ("hope", "must", "want", and so on). However, our discussion of these problems proceeded on the tacit assumption that we were concerned with relatively *standard* interpretations only. To be sure, this restriction was entirely reasonable, since the theories under discussion make predictions about the preferred interpretation of scalar expressions, and the whole point was to establish that these predictions are wrong.

What complicates the story is that in all the cases in which we don't normally obtain what conventionalist theories predict to be the default interpretation, this interpretation can be *forced*, e.g., by employing contrastive stress. For example, it is possible to construe (65) as meaning that all the squares are connected with some but not all of the circles:

(65) All the squares are connected with SOME of the circles. (= (62a))

Similarly, (66) might be used to convey that the speaker wouldn't like all his relatives to remember his birthday:

(66) I hope that SOME of my relatives will remember my birthday. (= (56a))

These observations suggest that, contrary to what we have supposed so far,

conventionalism might be better cast as a theory of *marked* interpretation. This suggestion makes sense for other reasons as well. First, as argued in §7.1, contrastive construals of scalar expressions have to be analysed as contributing to the truth-conditional meaning of a sentence, and one of the fundamental tenets of conventionalism is precisely that upper-bounding interpretations have truth-conditional effects. Secondly, whereas lexicalist versions of conventionalism are most naturally viewed as default theories, syntax-based accounts would seem to be geared more towards contrastive interpretations. After all, covert ◐ is modeled on overt "only", which is contrastive if anything is.

However, if conventionalism is to be recast as a theory of marked interpretation, what are we to make of *un*marked specimens of "embedded implicature", like (67), for example?

(67) Tony believes that Cleo had some of the figs. (= (3))

It doesn't require contrastive stress to interpret (67) as implying that, for all the speaker knows, Tony believes that Cleo didn't have all the figs, and therefore it would seem that cases like this are outside the purview of a theory that requires upper-bounded construals to be marked. Which brings us back to a point made at the beginning of the last section: that conventionalism cannot aspire to offer a wholesale replacement for the Gricean theory of Q-implicatures. This point now applies with a vengeance: if conventionalism is rebranded as a theory of marked construals (and this is the most promising strategy for salvaging the theory, as far as I can see), it will *ipso facto* be of marginal importance, at the best of times.

To conclude, I have argued that, if it is touted as a general framework for scalar inference, conventionalism does not acquit itself particularly well. However, it is still possible that conventionalism may save its skin by scaling down its ambitions, and reinventing itself as a theory of marked interpretation. We will return to this possibility in §8.5.

Embedded implicatures: a Gricean approach

8

In §7.1, we discussed five lines of evidence that, prima facie at least, are at variance with the Gricean theory of Q-implicature, and more in agreement with a conventionalist approach. Since then, we have seen that conventionalism is not without its problems, either, so the obvious thing to do is to have a closer look at the evidence and re-assess whether a flight into conventionalism is really called for. In this chapter, I argue that it isn't, and to explain what the general form of the argument is going to be, let's remind ourselves what the data were. First, there are cases in which the Gricean theory seems unable to capture inferences that are intuitively real:

(1) *Belief reports:*
Tony believes that Cleo had some of the figs.
$\rightsquigarrow$ Tony believes that Cleo didn't have all the figs.

(2) *Factive verbs:*
Tony knows that Cleo had some of the figs.
$\rightsquigarrow$ Cleo didn't have all the figs.

Secondly, there are cases in which the Gricean theory seems to make false predictions:

(3) *Disjunction:*
Cleo either had a mango or some of the figs.
$\not\rightsquigarrow$ Cleo didn't have a mango.

(4) *Indefinites:*
Last year, a Dutch sailor showed some symptoms of delirium.
$\not\rightsquigarrow$ Last year, there were no Dutch sailors who showed all the symptoms of delirium.

(5) *Contrastive contexts:*
Around here, we don't LIKE coffee, we LOVE it.
$\not\rightsquigarrow$ The extent to which we appreciate coffee is not enough to warrant the predicate "like".

In the following, I will argue that these data fall into two classes, which it is vitally important to keep distinct. The line separating the two runs between examples (4) and (5); that is to say, while contrastive sentences are in a class of their own, the others belong together. What unites the latter class is that, first impressions notwithstanding, they are consistent with the Gricean theory of Q-implicatures, after all: for each of the examples in (1)-(4), there is a perfectly adequate explanation in terms of Q-implicature, and I will spend the better part of this chapter trying to prove so, case by case.

Which leaves us with contrastive sentences. I have already conceded that examples like (5) are beyond the scope of the theory of Q-implicatures (§7.1), and granted this is so, it may seem natural to relegate them to some version of conventionalism. However, for reasons to be explained at the end of this chapter (§8.5), this is not the course I recommend. Rather, I will argue for an approach that yields the same readings as conventionalist theories are designed to deliver, but does so in a truly *pragmatic* way. I will even argue that, though the kind of pragmatic reasoning that is needed operates at a different level from the one at which conversational implicatures are derived, it does have a natural place in the grand Gricean scheme of things.

8.1 Disjunction

Actually, sentences like (6) pose not one but two challenges to the theory of Q-implicatures:

(6) Cleo either had a mango or some of the figs. (= (3))

One is to explain how (say) Julius's utterance of (6) could give rise to the implicature that Cleo didn't eat all the figs. The second is to avoid predicting that, according to Julius, Cleo didn't have a mango. Here is how these challenges can be met. When Julius uttered (6), he had at least the following alternatives at his disposal:

(7) a. Cleo had a mango.
b. Cleo had some of the figs.

c. Cleo had all the figs.
d. Cleo either had a mango or all the figs.

Each of these sentences entails (6) but not vice versa, and none of them is more complex. Hence, for each of these sentences, it would be reasonable to surmise that Julius doesn't believe it is true:

(8) a. $\neg\textsc{Bel}_\textsc{j}$(7a)
b. $\neg\textsc{Bel}_\textsc{j}$(7b)
c. $\neg\textsc{Bel}_\textsc{j}$(7c)
d. $\neg\textsc{Bel}_\textsc{j}$(7d)

Now suppose that the Competence Assumption holds for (7c): $\textsc{Bel}_\textsc{j}$(7c) $\vee$ $\textsc{Bel}_\textsc{j}(\neg$(7c)). In conjunction with (8c), this entails that $\textsc{Bel}_\textsc{j}(\neg$(7c)), i.e. Julius believes that Cleo didn't have all the figs. Thus the first problem is solved simply by assuming that (7c) is an alternative to (6), which seems natural enough.

The implicature in (8c) may be strengthened on the basis of the Competence Assumption. What about the remaining implicatures in (8)? Shouldn't it be possible that they are strengthened, too? Consider (8a). If Julius were competent with respect to (7a), i.e. $\textsc{Bel}_\textsc{j}$(7a) $\vee$ $\textsc{Bel}_\textsc{j}(\neg$(7a)), then it would follow that $\textsc{Bel}_\textsc{j}(\neg$(7a)): Julius believes that Cleo didn't have a mango—which would be patently wrong. However, and this is crux of the matter, there are good reasons for doubting that the Competence Assumption holds in this case.

There are two ways of arguing this point. One is to adopt a line of reasoning proposed by Sauerland (2004), who argues that the Competence Assumption cannot hold for (7a) because Julius's own statement conveys that he is *agnostic* about the truth (7a). That this is so follows from the fact that Julius believes that (6) plus the fact that he doesn't believe that the second disjunct is true, which is what (8b) says. More generally: given that $\textsc{Bel}_\textsc{s}(\varphi \vee \psi)$ and $\neg\textsc{Bel}_\textsc{s}(\psi)$ (i.e. $\textsc{Poss}_\textsc{s}(\neg\psi)$), it must be the case that $\textsc{Poss}_\textsc{s}(\varphi)$, or equivalently: $\neg\textsc{Bel}_\textsc{s}(\neg\varphi)$. Therefore, on the assumption that Julius's statement is truthful, (8b) entails $\neg\textsc{Bel}_\textsc{j}(\neg$(7a)), and since we already have $\neg\textsc{Bel}_\textsc{j}$(7a) (= (8a)), Julius's utterance implies that $\neg\textsc{Bel}_\textsc{j}(\neg$(7a)) $\wedge$ $\neg\textsc{Bel}_\textsc{j}$(7a), which is the negation of the Competence Assumption for (7a) (see §2.3 for a more leisurely presentation of this argument).

There is an another and more direct way of showing that the Competence Assumption cannot hold for (7a), and it goes as follows. If the Competence Assumption held in this case, then either Julius believed that (7a) is true or

he believed that (7a) is false. However, given that Julius believes that (6), he should have made a different statement in either case: if Julius took (7a) to be true, the second disjunct of (6) would have been redundant, and if he took (7a) to be false, he should have said "Cleo had some of the figs." Hence, if Julius is trying to be cooperative, he cannot be competent with respect to (7a), and it does not follow from $\neg\textsc{Bel}_{\textsc{j}}$(7a) that $\textsc{Bel}_{\textsc{j}}(\neg$(7a)).

Given that an utterance of (6) implies that Julius is not competent with respect to the first disjunct, it will be clear that the same holds for the second disjunct, (7b), and therefore we don't predict that (8b) will be strengthened, either. What about (7d), which caused the main problem we set out to solve? Here the reasoning is essentially the same as for the two disjuncts of (6). The simplest version of the argument goes like this. If Julius is competent with respect to (7d), then either he believes that (7d) is true or he believes that (7d) is false. In the former case, he should have uttered (7d) instead of (6); in the latter case, which implies that Cleo had neither a mango nor all the figs, he should have said, e.g., "Cleo had some of the figs."

On the face of it, it would seem that, if (6) is construed as "Cleo either had a mango or some but not all of the figs", "some" is interpreted *in situ* as "some but not all"; that is, it looks as if an upper-bounded reading of "some" is achieved within the scope of "or". On a Gricean analysis, this is a mirage: (6) may *implicate* that, for all the speaker knows, Cleo didn't have all the figs, and if it does the speaker's utterance may be paraphrased as "Cleo either had a mango or some but not all of the figs", but the *literal meaning* of the second disjunct of (6) is the classical one, and not upper-bounded in any sense.

8.2 Belief reports

What I just said about disjunctions holds for belief reports, as well:

(9) Tony believes that Cleo had some of the figs. (= (1))

Though it may seem that in (9) "some" can be construed as "some but not all", this interpretation, too, may be explained in terms of Q-implicature, without resorting to the assumption that the embedded sentence, "Cleo had some of the figs", has an upper-bounded meaning. One possible line of explanation has been proposed by van Rooij and Schulz (2004) and Russell (2006). As observed in §7.1, the strong scalar implicature that standard Gricean reasoning will deliver for (9a), when uttered by Julius, is that:

(10) $\textsc{Bel}_{\textsc{J}}(\neg\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$

Now, depending on the context, it may be plausible to assume that, for all Julius knows, *Tony* knows whether or not Cleo had all the figs:

(11) $\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs})) \vee$
$\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\neg(\text{Cleo had all the figs})))$

That is to say, it may be the case that the Competence Assumption holds not only for the speaker, but for the subject of the belief report, as well. Between them, (10) and (11) entail:

(12) $\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\neg(\text{Cleo had all the figs})))$

Thus the seemingly local scalar inference is accounted for on strictly Gricean principles.[1]

There is an alternative way of obtaining the same outcome, which is due to Spector (2006). To explain, consider the following variation on (9):

(13) a. Tony said that Cleo had some of the figs.
b. Cleo had some of the figs.

Intuitively, (13a) may suggest that, when he uttered (13b), Tony *believed* that Cleo didn't have all the figs. Why? Obviously, because it may be assumed that Tony's utterance of (13b) implicated that she didn't. Armed with this observation, we return to the belief report in (9), and ask ourselves what sort of evidence Julius may have had for making this claim. In principle, Julius may have been relying on all kinds of sources, but there is one in particular that readily comes to mind: he may have known (or at least believed) that (13a) is true. But we have just seen how this proposition might prompt the inference that, according to Tony, Cleo didn't have all the figs, so *if* an occurrence of (9) suggests (13a), it may also suggest (12).

Formally speaking, Spector's analysis is less economical than the one proposed by Russell and van Rooij and Schulz. For the latter to work, we merely have to assume that the Competence Assumption holds for the subject of the

1. As pointed out to me by Larry Horn, the reasoning from $\textsc{Bel}_{\textsc{J}}(\neg\textsc{Bel}_{\textsc{T}}(\text{Cleo had all the figs}))$ to $\textsc{Bel}_{\textsc{J}}(\textsc{Bel}_{\textsc{T}}(\neg(\text{Cleo had all the figs})))$ is reminiscent of the "neg-raising" interpretation of sentences like:

(i) Tony doesn't believe that Cleo had all the figs.
$\rightsquigarrow$ Tony believes that Cleo didn't have al the figs.

It doesn't follow from this, however, that neg-raising can explain why the implicature is strengthened, though it may be that implicature strengthening and neg-raising have a common source, i.e. the assumption that the subject of "believe" is competent (cf. Bartsch 1973).

belief report. This may not be a plausible assumption in general, but in the case at hand it doesn't seem far-fetched (which is good enough, since the inference we're trying to capture does not hold by default). Spector's story is less straightforward: in order to explain how (9) comes to imply (12), we have to assume not only that Tony said that Cleo had some of the figs, but also that his utterance licensed the implicature that she didn't have all the figs. Still, it doesn't seem unlikely to me that, in some cases at least, hearers consider what may be the evidence on which the speaker bases his statement, and if the statement happens to be a belief report, Spector's explanation strikes me as eminently plausible. Moreover, we don't have to choose between the two explanations on offer: since they don't exclude each other, it is possible that hearers rely on both lines of reasoning.

It is important to note that, whichever of the two explanations is adopted, it does not generalise beyond epistemic attitude reports, i.e. statements about people's beliefs, knowledge, and so on. As discussed at length in the last chapter (especially §7.3), conventionalist theories predict that, in principle, scalar inferences can occur in just about any embedded position, and this prediction appears to be false: "embedded implicatures" are the exception rather than the rule. For example, unless forced by contrastive stress on "some", the following should not hold:

(14) Tony hopes that Cleo had some of the figs.
$\not\rightsquigarrow$ Tony hopes that Cleo didn't have all the figs.

In this light, it is a considerable advantage that both Spector's and the Russell/van Rooij/Schulz proposal are restricted to epistemic reports. To see that they are thus restricted, consider first what would have to be assumed for making the latter applicable in the case of (14). Supposing, as seems reasonable, that Julius's utterance of this sentence gives rise to the implicature that

(15) $\textsc{Bel}_{\textsc{j}}(\neg\textsc{Hope}_{\textsc{t}}(\text{Cleo had all the figs}))$

it should be the case that

(16) $\textsc{Bel}_{\textsc{j}}(\textsc{Hope}_{\textsc{t}}(\text{Cleo had all the figs})) \vee$
$\textsc{Bel}_{\textsc{j}}(\textsc{Hope}_{\textsc{t}}(\neg(\text{Cleo had all the figs})))$

in order for (15) to license the inference that

(17) $\textsc{Bel}_{\textsc{j}}(\textsc{Hope}_{\textsc{t}}(\neg(\text{Cleo had all the figs})))$

However, it is evident that, in general, (16) will not be a particularly plausible assumption to make; or at the very least, it will be considerably less plausible

than its counterpart in (11).[2]

Concerning Spector's proposal, it will not be hard to see that it doesn't extend to (14), either. For, if we ask ourselves what evidence Julius might have had for making this claim, the answer is not going to be that Tony has *said*, "Cleo had some of the figs"; so the analogy with the belief report in (9) breaks down forthwith.

To sum up, as in the case of disjunction, there are ways of explaining how a sentence like (9) may come to convey that, according to the speaker, Tony believes that Cleo didn't have all the figs, and neither explanation requires the assumption that anything like a scalar inference is derived within the scope of "believe". Furthermore, neither explanation carries over to other embedding constructions, thus safeguarding the Gricean approach against the proliferation of unwanted readings conventionalist theories are plagued by.

8.3 Factives and other presupposition inducers

One of the issues raised by Chierchia (2004) in his critique of Gricean pragmatics concerns factive verbs:

(18) a. Tony knows that Cleo had some of the figs. (= (2))
b. Cleo had some but not all of the figs.
c. Cleo had (at least) some of the figs.

An utterance of (18a) may convey that, according to the speaker, (18b) holds as well, and since the sentence presupposes that its embedded clause is true, it is natural to suppose that the upper-bounding inference is somehow tied up with the presupposition. That is what Chierchia assumes, too: on his account, the "not all" inference attaches to the word "some", and thus is part of the meaning of the complement clause of "know" even before the presupposition is triggered. Consequently, the upper bound on "some" projects out along with the proposition expressed by the factive complement, and (18a) presupposes (18b) rather than (18c).

While Chierchia's analysis delivers passable predictions for sentences with "know" (though see Russell 2006), it doesn't work for other factive verbs:

2. In other words, in a context in which (16) holds, (14) should imply (17). As far as I can see, this possibility is marginal, at best, but in this case my intuitions are rather labile. More generally, the claim that there is a clear-cut difference between "think/believe/know", on the one hand, and the remaining attitude verbs, on the other, may well be too strong, but even if it is, my main point is not affected.

(19) Tony regrets that Cleo had some of the figs.

On its most natural reading, (19) means that Tony regrets that (18c), and conveys, in addition, that (18b). On Chierchia's analysis, however, the presupposition has to be the same as the object of the attitude verb: Chierchia predicts that (19) must either presuppose (18b) and assert that Tony regrets (18b) or presuppose (18c) and assert that Tony regrets (18c). Either way, the presupposition is the same as the object of Tony's regret, and therefore the most likely reading is not accounted for.

The trouble with Chierchia's analysis is that his conventionalist convictions lead him to assume that the scalar inferences we observe in (18a) and (19) are part of the presupposed content from the start. The obvious alternative is to hypothesise that these inferences *derive from* what is presupposed. Reverting to the classical view that presuppositions come before implicatures and that both sentences presuppose (18c), we apply the standard Gricean reasoning to this presupposition, and thus arrive at the implicature that, for all the speaker knows, (18b) holds, as well. In order to explain that (18a) may convey, furthermore, that *Tony* believes (18b), we can appeal to either of the analyses discussed in the last section, and since neither analysis carries over from epistemic "know" to emotive "regret", we can account for the asymmetry between (18a) and (19) that Chierchia's proposal is unable to explain.

As we saw in §7.1, the hallmark of presuppositions is that they tend to escape from practically any kind of embedding. This being so, the hypothesis I propose is that, whenever a sentence φ presupposes that ψ, and ψ may implicate that χ, then φ may carry the same implicature. The following examples confirm this hypothesis:

(20) a. Does Tony know that Cleo had some of the figs?
b. Please make sure that Tony knows that Cleo had some of the figs.

As is usual with a bona fide presupposition, both sentences suggest rather strongly that Cleo had some of the figs, and they also convey (or, at the very least, are easily construed as conveying) that she didn't have all the figs. This pattern extends to other presupposition triggers, like it-clefts, for example:

(21) a. It was Julius who stole some of Cleo's shoes.
b. It wasn't Julius who stole some of Cleo's shoes.
c. Wasn't it Julius who stole some of Cleo's shoes?

Ceteris paribus, each of the sentences in (21) presupposes that somebody stole some of Cleo's shoes (in the first case this is an entailment, as well), and accord-

ing to my intuitions they will also tend to imply that the person in question didn't steal all of Cleo's shoes. These are admittedly concocted examples, but real data are no different. In (22), I present three sentences I found on the internet that involve, respectively, an it-cleft construction, a factive verb, and a definite description:

(22) a. It wasn't me who got some of the font sizes wrong on the Vacancy Pages, though!
   ⇝ Somebody got some but not all of the font sizes wrong.
   b. I didn't realize that some of the early church fathers and even the great reformers [...] believed in the perpetual virginity of Mary.
   ⇝ Not all the early church fathers believed in the perpetual virginity of Mary.
   c. After all, this was the latest film from Quentin Tarantino, the man who brought some of the greatest films of the last decade.
   ⇝ Tarantino didn't bring all the greatest films of the last decade.

I would like to suggest, then, that presuppositions can give rise to implicatures. More precisely, if a sentence comes with a presupposition φ, implicatures may be derived from φ just as if it was asserted. To show that this is actually a plausible idea, let's have a closer look at (22a), assuming that it is uttered in a context in which it is not yet part of the common ground that somebody got the font sizes wrong. Then the sentence brings two new messages: (*i*) that somebody got the font sizes wrong and (*ii*) that that somebody wasn't the speaker. (*i*) is presupposed and (*ii*) is asserted, which means that they don't have the same discourse status, but since in both cases the speaker brings new information to the hearer's attention, the hearer is entitled to ask himself in both cases how this information is supposed to contribute to the goals of the discourse, whether it is relevant, sufficiently specific, and so on. Briefly, in both cases the hearer is entitled to apply the kind of reasoning that may yield conversational implicatures.

My proposal requires that presuppositional devices may be used for presenting new information, and since prima facie this requirement is at odds with the widespread view that presuppositional devices are supposed to carry given information, I will discuss it a bit further. To begin with, it should be noted that on some accounts, the notion that presuppositions may carry new information is entirely straightforward. For instance, one view which recently has gained some following is that, in general, presuppositions are just parts of information that are less central to the speaker's concern than what he wants

to assert, question, and so on (e.g., Geurts 1999b, Abbott 2000, Simons 2006). Clearly, if one takes this line, presuppositions may carry new information, and thus give rise to implicatures.

Things are not as simple if one adheres to what may still be the received view on presupposition, which originated in Stalnaker's work of the early 1970s (Stalnaker 1973, 1974). For, on this account, to presuppose something is to present it as given, as part of the common ground between oneself and one's audience: a presupposition is "an item of presumed common knowledge" (Stalnaker 1973: 450). However, according to Stalnaker and his followers, just as one can present old stories as if they were new, a speaker can present new information *as if* it were already part of the common ground. In such an event, the speaker dispenses new information by pretending that the hearer already has it. For example, if I come late to a meeting, I may apologise by saying:

(23) I'm sorry I'm late: my bicycle had a flat tire.

In saying this, I presuppose that I have a bicyle. It may be that no one in the audience knew this beforehand, but they will let me get away with my presupposition because, in my country, it is perfectly normal for people to own (and use) bicycles, and because they appreciate that in order to avoid the presupposition I should have resorted to a cumbersome formulation along the lines of:

(24) I'm sorry I'm late: I own a bicycle and it had a flat tire.

Hence, new information can, and often will, be conveyed by way of presupposition:

> 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 openly assert a proposition that one wants to communicate. (Stalnaker 1974: 202)

Now, my main point is simply this. Even on an account that explains presupposition in terms of givenness, it is agreed that, *in practice*, presupposed information is often new, and that it may be common knowledge between speaker and hearer that this is so. More briefly, presuppositions may serve as speech acts for conveying new information. In this sense, they is no different

from assertions (which is not to deny that presuppositions and assertions are different in other respects), and there is no reason why conversational implicatures couldn't be derived from them.

Of course, if presuppositions can give rise to implicatures, we will have to reconsider the standard Gricean line on the relationship between implicatures and "what is said". According to Grice, conversational implicatures derive from what is said, and this is usually taken to imply that they are based on the truth-conditional content of a statement. However, we know already that this isn't quite right, e.g., because non-declarative and multi-sentence speech acts can have implicatures, as well (cf. §1.3 and §7.3). Although the overwhelming majority of examples discussed in this book and elsewhere in the literature involve individual declarative statements, it is just not true that a conversational implicature has to be based on a single asserted proposition.

Still, even if there is no consensus on the matter, there is a widely held view that, in the logical order of events, presuppositions are *prior to* what is said. Isn't this a problem? Not as far as I can see. The reason why a statement can give rise to conversational implicatures is that it is a linguistic act that serves to convey the information that such-and-such is the case. And as we have seen, it is universally agreed that, in some cases at least, presuppositional devices are pressed into service for doing just that. Presupposition and assertion may be different in other respects, but in this respect they need not be, and that's the only thing that matters as far as implicatures are concerned.

To recapitulate, a speaker may choose to present new information by means of a single statement, a series of statements, or by way of presupposition. In all these cases, implicatures may ensue, but what makes the last option particularly interesting is that presuppositions that are triggered in embedded positions typically behave as if they weren't; they "project out" of the scope of negation, modals, and so on. Therefore, though it may seem as if implicatures derived from such presuppositions exhibit the same projection behaviour, in reality, that's not what happens: it is the presuppositions that project out, not the implicatures they give rise to.

Hence, the problem with factive verbs observed by Chierchia (2004) is just an instance of a much broader phenomenon. I have argued that the best way of accounting for the larger pattern is by adopting a discourse perspective: what is crucial is that presuppositions may serve to bring new information to the discourse, and this has nothing to with syntax.

8.4 Indefinites

I'll start this section with a seemingly innocuous detail that, under magnification, will turn out to be deeply significant. One of Grice's most famous examples of conversational implicature is the following (cf. §1.1):

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

Grice glosses this exchange as follows:

> B would be infringing the maxim "Be relevant" unless he thinks, or thinks it possible, that *the garage* is open, and has petrol to sell …
>
> (Grice 1975/1989: 32; emphasis added)

Grice's gloss suggests that A's reasoning should be unpacked thus:

i. B has said, "There is a garage round the corner", so apparently he believes that there is a garage round the corner.
ii. Could it be that B has reason to believe that *the garage in question* might be closed?
iii. Presumably not, because then he would have said so.
iv. Hence, I'm going to assume that, for all B knows, *the garage* is not closed.

What's interesting about this is that, in order to reason sensibly about B's beliefs, A has to *accept* that there is a garage round the corner, or more accurately, that B believes there is such a garage. A must be accepting this, because the question he poses, (*ii*), and the inference he draws, (*v*), are about that very garage. The same holds for other examples discussed by Grice:

> Anyone who uses a sentence of the form *X is meeting a woman this evening* would normally implicate that *the person to be met* was someone other than X's wife, mother, sister, or perhaps even close platonic friend.
>
> (Grice 1975/1989: 37; emphasis added)

In this case, too, the hearer has to accept that, according to the speaker, X is meeting a woman this evening, so that he can ask himself whether or not that woman might be X's wife, mother, sister, or friend.

To see why this is important, consider what an alternative-based analysis would yield in Grice's second scenario:

i. S has said: "X is meeting a woman this evening."
ii. S could have made a stronger statement, e.g., by saying: "X is meeting his mother this evening." Why didn't he do so?
iii. Presumably, it is because S doesn't believe that the stronger statement is true. Hence, S doesn't believe that X is meeting his mother this evening.

This conclusion is weird: though in all likelihood, S's statement is not about X's mother, what we should like to infer is that, as Grice says, "the person to be met" isn't X's mother, but without precluding the possibility that X is going to see his mother, perhaps *en route* to the meeting S has in mind.

What went wrong here? The derivation of the fake implicature started out from the wrong question, namely:

(26) Why didn't S say: "X is meeting his mother this evening"?

The question that should have been asked is, in Grice's own words: "Why didn't S say that the person to be met is X's mother?" The problem with the alternative-based approach is that it cannot ask this question, simply because it would be patently improbable to suppose that "The person to be met is X's mother" is a suitable alternative to "X is meeting a woman this evening".

Fortunately, the intention-based stance advertised in Chapter 6 allows us to ask the right sort of question:

i. S has said: "X is meeting a woman this evening."
ii. Could it be that S believes the person to be met is X's mother?
iii. Presumably not, because then S would have said: "X is meeting his mother this evening."
iv. Hence, S doesn't believe that the person to be met is X's mother.

Apparently, as van der Sandt (1992) was the first to observe, there has to be a kind of *anaphoric link* between the speaker's statement and the hearer's question. When the speaker says, "X is meeting a woman this evening", he asserts the proposition that X is meeting a woman this evening, but he also introduces a *discourse referent* that subsequent utterances can refer back to and that, crucially, other interlocutors can ask questions about, be it overtly or tacitly.

The notion of discourse referent was introduced in a seminal paper by Karttunen (1971), which inaugurated a truly discourse-oriented approach to interpretation that, in the meantime, has become more or less standard, but hasn't had much of an impact on theories of implicature. The main tenet of

this section is that, when we look closer even at the stock-in-trade examples of conversational implicature, it is evident that such an approach is called for. In the following, I will develop this point by showing how an intention-based treatment of Q-implicature might be implemented in Discourse Representation Theory (DRT), which was the first theory built on Karttunen's ideas, and remains one of the most influential frameworks of discourse interpretation to date (Kamp 1981, Kamp and Reyle 1993, Geurts 1999a).[3]

DRT is primarily concerned with the interpretation of discourses, not just sentences. As a discourse unfolds, hearers incrementally construct a model of what the speaker is saying and implying. Such models consist of two main parts: a set of discourse referents, which represent the objects under discussion, and a set of conditions which encode the information that has accumulated on these discourse referents. For example, when uttered by speaker S, the sentence in (27) might give rise to the discourse model in (28):

(27) Cleo is reading a book by Theo.

(28) $\textsc{Bel}_S$:

x
x is a book by Theo
Cleo is reading x

This representation may be given a truth-conditional meaning exactly as one should expect. That is, (28) is true iff S believes that there is an x such that x is a book by Theo and x is being read by Cleo. The discourse referent x is introduced into the model by the indefinite noun phrase "a book by Theo". Other expressions may serve to pick up such discourse referents. To illustrate, suppose S continues his discourse by saying:

(29) It is about logic.

Since this statement is made in the context represented by (28), the pronoun may be construed as picking up x, and if it is, the effect of (29) is that the discourse model is extended:

(30) $\textsc{Bel}_S$:

x
x is a book by Theo
Cleo is reading x
x is about logic

3. I should note that the following is an extremely sketchy exposition of a hybrid between Gricean pragmatics and DRT. There are all sorts of issues, technical as well as conceptual, that remain to be sorted out, and that I have to leave for future research.

One of the attractions of this approach to interpretation is that it has no need for propositions; no need, that is, for sentence-sized truth-value carriers. Consider sentence (29) again. It is obvious enough that, taken on its own, it is neither true nor false. It may be true or false in a given context, such as the one represented by (28), but even then we first have to decide that the pronoun is to be interpreted by the discourse referent x, and once that decision has been made, whatever (29) means has become inextricably tied up with the information in (28). The primary unit of meaning, on this view, is the discourse. In the first instance, it is discourses, not sentences, that have truth conditions.[4]

Let us now return to implicatures, and the problem posed by:

(31) Last year, a Dutch sailor showed some symptoms of delirium. (= (4))

This is interpreted along the following lines:

(32) $\textsc{Bel}_S$:

x
x is a Dutch sailor
x showed some symptoms of delirium

As in example (28), the singular indefinite causes the introduction of a new discourse referent, x (the plural indefinite will do likewise, of course, but we will ignore that here). On the harmless though admittedly unrealistic assumption that the previous context is empty, (32) represents what the speaker said by uttering (31), and it is on the basis of this interpretation that the hearer will ask himself whether or not S might believe that x showed all symptoms associated with delirium, which can be represented like this:

(33)

$\textsc{Bel}_S$:	x
	x is a Dutch sailor
	x showed some symptoms of delirium
?$\textsc{Bel}_S$:	x showed all symptoms of delirium

The key point of this model is that the hearer's question about S's beliefs is about the same discourse referent that was introduced by S's statement. In

4. This is not to deny that individual utterances have truth-conditional import. Utterances serve to increment the discourse, and can be assigned meanings in these terms. For example, we could define the meaning of (29), as used on this particular occasion, as the difference between the truth conditions associated with (28) and (30).

this sense, there is an anaphoric dependence between S's statement and the hearer's tacit question. From this point on, the reasoning is of the familiar sort. If the hearer decides to infer that S doesn't believe that x showed all symptoms of delirium, his inference can be represented as follows:

(34)

$\textsc{Bel}_S$:	x x is a Dutch sailor x showed some symptoms of delirium
$\neg\textsc{Bel}_S$:	x showed all symptoms of delirium

And if the hearer accepts that S is competent with respect to the proposition that x showed all symptoms of delirium, there will be a strong implicature, as well:

(35)

$\textsc{Bel}_S$:	x x is a Dutch sailor x showed some symptoms of delirium ¬(x showed all symptoms of delirium)

According to the model in (35), S believes that there was a Dutch sailor who showed some though not all the symptoms of delirium. Again, it turns out that what looks like an upper-bounded construal of "some" achieved in an embedded position is in reality the result of honest-to-goodness pragmatic reasoning, though it is essential that we adopt a discourse perspective on this process. That is to say, we have to assume that, when making inferences about the speaker's intentional state, hearers reason in terms of discourse referents. Which is precisely what Grice did in the passages cited at the beginning of this section.

I have tried to show how we can handle Q-implicatures that seem to arise within the scope of indefinite noun phrases, without straying from the basic principles of Gricean pragmatics. The proposed solution to this problem opened up a much larger issue, for I have argued that the heart of the matter is that a theory of conversational implicature should be discourse-oriented rather than sentence-oriented. When formulated in such general terms, this is a truism: it is plain that discourse is central to all Gricean thinking about interpretation. Nevertheless, in the meantime we have come across several discourse-related factors that have received little or no attention in the literature, but are critically important to the Gricean theory of conversational implicatures. First, as we observed already in §1.6, in some cases it is a stretch

of discourse, rather than a single utterance, that licenses an implicature. Secondly, I have argued in the last section that implicatures may derive not only from what an utterance says, in Grice's sense of the word, but also from what it presupposes. Thirdly, as argued in the current section, even when implicatures are derived on the basis of what is said, they generally go beyond the propositional content of an utterance: the implicit question that sets off the Gricean train of reasoning is very much like a question the hearer might have asked explicitly in order to clarify the speaker's intentions, and such explicit questions will not just be about what the speaker has said. Rather, such questions will typically accept what the speaker has said, at least in the sense that when the speaker's utterance introduced new discourse referents, the hearer's question can be about those.

My general point, then, is that Gricean pragmatics has a lot to gain from joining forces with a theory of discourse interpretation like DRT, though it bears emphasising that the foregoing discussion has scarcely done more than sketch the bare outlines of such a joint venture. Working out a full-fledged theory along these lines is not going to be easy, and will inevitably turn up a fair number of issues, conceptual as well as technical. But the general direction is clear enough.

8.5 Contrastive construals and lexical pragmatics

Thus far I have argued, for an assortment of allegedly "embedded" implicatures, that they can be accounted for in a principled way as Q-implicatures, or at any rate, as essentially involving Q-implicatures. True, we had to assume that additional factors were involved in each case, but such auxiliary assumptions as had to be made could always be motivated on independent grounds. Thus we are left with a single class of problem cases, a selection of which I repeat here for ease of reference:

(36) a. Around here, we don't LIKE coffee, we LOVE it.
b. I'd rather have a WARM bath than a HOT one.
c. Is a parallelogram SOMETIMES or ALWAYS a square?
d. If it's WARM, we'll lie out in the sun. But if it's VERY warm, we'll go inside and sit in front of the air-conditioner.

For reasons discussed in §7.1, there is no escaping the fact that, in cases like these, the truth-conditional meaning of "like", "warm", etc. has an upper

bound. That is to say, in the context of (36a), "like" *entails* "not love", and the same, mutatis mutandis, for the other examples in (36). But even if we have to concede that marked uses of scalar expressions cannot always be explained in terms of Q-implicature, that doesn't mean they can't be accommodated within the general Gricean framework. In this section we will see how that may be done.

Polysemy and the pragmatics of word meaning One of the first intimations of Grice's nascent theory of pragmatics occurs towards the end of his 1957 paper:

> Again, in cases where there is doubt, say, about which of two or more things an utterer intends to convey, we tend to refer to the context (linguistic or otherwise) of the utterance and ask which of the alternatives would be relevant to other things he is saying or doing, or which intention in a particular situation would fit in with some purpose he obviously has (e.g., a man who calls for a "pump" at a fire would not want a bicycle pump). (Grice 1957/1989: 222)

As pointed out by Neale (1992), this passage is interesting not only because it contains the seeds of the views Grice would go on to develop more fully in his William James lectures, one decade later, but also because his chosen example is *not* a conversational implicature but rather a case of lexical ambiguity. Although one of the established senses of "pump" is "bicycle pump", it is evident that someone who calls for a "pump" at a fire is unlikely to have need for a bicycle pump, because under the circumstances that wouldn't be relevant to his purposes. Hence, in this case, Gricean reasoning is employed not to draw inferences from what the speaker has said, but rather to determine what it is he wants to say (cf. §1.6).

Actually, though, it is slightly inaccurate to say that the word "pump" is ambiguous: it is *polysemous*. A word is standardly considered to be polysemous if it has several senses that are conceptually related to one another. If you consult a good dictionary for the meaning of a word, it is practically inevitable that the search will yield more than one sense. In the vast majority of cases, these senses will be related: polysemy is the rule rather than the exception. To illustrate, here are some of the senses of "pump", as given by the Oxford English Dictionary:

> *i.* A mechanical device for raising water, commonly consisting of a cylinder or tube in which a piston or plunger is moved up and down, and usually incorporating a valve to prevent back-flow; (more generally) any mechanical or electromechanical device for raising or moving fluid or compressing or evacuating gas.

ii. *intr.* To work a pump (in early use a ship's pump); to raise or move water or other fluid by means of a pump.

iii. *trans.* To inflate (a pneumatic tyre, etc.) by means of a pump. Freq. with *up*.

iv. *trans.* orig. *U.S.* To enlarge or strengthen (a muscle, part of the body, etc.) through weight training or steroid use. Chiefly with *up*.

v. *trans.* To fire (shots) at a target; to shoot (bullets, etc.) into a person or thing, now esp. from a machine gun or automatic firearm. Also *intr.*

All told, the OED lists 40-odd related senses for "pump", and this bulk is by no means exceptional: there are multiple senses for the overwhelming majority of English words, and it is not unusual for a word to have ten senses or more. (The longest entry in the latest print edition, AD 1989, is for the verb "set", which uses 60,000 words to describe some 430 senses.) These observations belie the naive view that, as a rule, a word has just a single meaning. Polysemy is clearly the norm, and this fact has an important consequence: it entails that hearers *routinely* have to engage in pragmatic reasoning merely in order to select the right meaning for practically any word used by the speaker.

As illustrated by the excerpt from the OED, your typical lexical entry is a cluster of senses, which are usually derived from one primary sense, either directly or indirectly. In the case of "pump", the oldest sense is (*i*), and it is fairly obvious how the others derive from it. I will also be clear, I trust, that the web of senses associated with a word is essentially pragmatic in nature. The fact that dictionaries try to list, for each word in the language, all the various meanings the word can be used to convey, should not delude us into thinking that all these senses are purely conventional. To appreciate this point, consider the following examples (from Nunberg 1979):

(37) a. He can hit the ball two *football fields*.
b. He hit a home run two *games* ago.
c. *IBM* went up three points last week.

In each of these examples, the lexical meaning of the italicised expression must be shifted so that it will mesh with the rest of the sentence: "football field" means "length of a football field", "game" provides a measure of time, and "IBM" refers to IBM stocks.[5] Obviously, such meaning shifts are based on conceptual relatedness: if α is amongst the established senses of a given word, then in principle the same word can be used to express *any* related

5. For discussion and many more examples, see Nunberg (1978, 1979), Sag (1981), Bosch (1984), Bartsch (1987), Carston (2002), Recanati (2003), Wilson and Carston (2007).

concept β that is rendered salient by α, in a given context, where "context" is understood sufficiently broadly to encompass the world knowledge shared between interlocutors. And if a particular meaning shift occurs often enough, the word in question will have acquired a new established sense: polysemy is rooted in the context-dependence of word meaning.

In general, lexical interpretation tends to be viewed as a rather passive affair, the basic idea being that, in order to determine the meaning of a word, all the hearer has to do is look it up in his mental dictionary. The foregoing observations paint a rather different picture. Interpreting a word is very much an active process, in which the hearer continually has to select between large numbers of possible meanings and has to construct new meanings on the fly. The dynamics of this process is Gricean: the logic that underwrites the selection and construction of word meanings is similar to that of conversational implicatures. To drive home this point, consider how a hearer might deal with the compound nominal "football fields" in (37a):

i. The only use of "football field" I've encountered so far is to refer to (well) football fields.
ii. This is probably not what the speaker has in mind, because if it were, (37a) would be gibberish. (cf. Quality)
iii. Still, I suppose that that sense is relevant to the intended interpretation of "football field". (cf. Relevance)
iv. (37a) would make sense if "football field" denoted a unit which may be used for measuring the distance a ball is hit.
v. If this unit is to be related to the standard meaning of "football field", it might be either the length or width of a football field.
vi. Since a ball is usually hit along the length of a football field, it seems more likely that the speaker has the first unit in mind, so that's what I'm going to assume.

I'm not claiming that hearers actually go through an argument like this every time they encounter a word. Not only is this an stylised example, but it is evident that, in many cases, lexical interpretation is a routinised process. But I do believe it is plausible to assume that the underlying *logic* of lexical interpretation is like this.[6]

6. An interesting corollary of the Gricean take on polysemy is that it gives us no reason to suppose that polysemy should be confined to individual words. If the Gricean view is correct, polysemy might as well occur in complex expressions. There is plenty of evidence that this

In §1.6, I argued that there are two stages at which Gricean pragmatics is involved in utterance interpretation. First, Gricean reasoning serves to derive conversational implicatures from speakers' utterances; it is this use that has received by far the most attention in the literature and in this book. Secondly, but first in the natural order of things, Gricean pragmatics helps to establish the truth-conditional meaning of an utterance, or "what is said", to use Grice's preferred phrase. If the foregoing considerations are on the right track, this "pre-propositional" use is rife: interpretation processes are guided by the Cooperative Principle from the lexical level up.[7]

Narrowing We have seen that one word meaning may spawn another whenever the two are saliently related. One particular way in which two concepts may be related is that one is more specific than the other. On my first visit to the United States, I was somewhat taken aback when a waitress who had just served me returned to ask, "Are you okay?" My first thought was that I must be looking sick or lonely (or both), but then I realised that she wanted to know whether I had everything I needed from her. I had just found out that the waitress was using "okay" in a sense that was somewhat more specific than I had been accustomed to so far. Examples like this abound; the following are due to Wilson and Carston (2007):

(38) a. I'm not *drinking* tonight.
b. Buying a house is easy if you've got *money*.
c. Either you become a *human being* or you leave the group.

Though it is not entirely unthinkable that someone who utters (38a) intends not to drink at all, there is a more natural construal on which the object of

is indeed the case, but as the point isn't central to our purposes, I'll make do with one example:

(i) There was freshly caught fish in the salad.

In (i), it is presumably not just the noun "fish" that acquires a mass interpretation, because the adjectival phrase "freshly caught" requires that its head have a count meaning. Therefore, it must the phrase "freshly caught fish", rather than any of its parts, that is construed in the culinary sense. Apparently, polysemy is not restricted to the lexicon (and it follows, incidentally, that the title of this section is not quite accurate).

7. As discussed in §1.6, Grice himself considered that what is said must be fairly close to conventional meaning. I don't believe this is right, but nor do I believe that this assumption is inevitable. On the contrary, it seems to me that the general Gricean framework naturally leads one to suppose that word meaning may be every bit as context dependent as other aspects of interpretation.

the verb "drink" is constrained to alcoholic beverages. Note that, on this interpretation, the alcohol constraint must be "visible" to the negation operator, or else the sentence will have the wrong meaning. Similarly, in (38b), "money" means something like "considerable means", and in (38c), "human being" is restricted to members of the species that are "normal" in some sense or other, and in these cases, too, the restriction must be part of the truth conditions of the words in question.

In the sentences in (38), sense narrowing happens quite smoothly, because world knowledge alone suffices to bias the hearer towards a specific interpretation. In other cases, narrowing is helped by other factors, as well:

(39) Bankers are bankers.

If both occurrences of "bankers" had the same meaning, (39) would be a tautology, which it is not, or at least, not necessarily. One way of bringing an utterance of (39) in line with the Quantity maxim is by supposing that the second occurrence of the noun is restricted in its application to individuals exhibiting greed, lack of moral fibre, or what have you. Similarly, (40) becomes informative when the second occurrence of "drinks" is construed as entailing binge drinking:

(40) When he DRINKS, he DRINKS.

In this case, the hearer is steered towards a non-standard interpretation by the contrastive intonation, as well, which we also observe in:

(41) Cleo didn't OPEN her handbag: she made an INCISION on the side.

Since Cleo opened her handbag in surgical fashion, her action arguably doesn't count as opening in the context of (41).

Whereas in (39) and (40), narrowing offers a way of avoiding a tautological interpretation, the following sentences would be contradictory if the juxtaposed terms both had their standard meanings. Hence, one of them will have to be construed more narrowly:

(42) a. Julius isn't RICH: he's RICH. (courtesy Richard Breheny)
b. The old count didn't DIE: he PASSED AWAY.
c. They didn't have SEXUAL INTERCOURSE: they FUCKED. (Horn 1989)
d. They didn't FUCK: they had SEXUAL INTERCOURSE. (ibid.)

While normally speaking, "pass away" would be truth-conditionally equivalent to "die", in (42b) it might connote, e.g., an especially stylish manner of expiring, and the meaning of "die" would then be narrowed down to exclude passing

away, thus understood. (42c) and (42d) are similar, if coarser.[8]

It will be clear by now, I trust, where my argument is heading: the examples in (43) (= (36)) can be accommodated in a general Gricean framework simply as special instances of narrowing:

(43) a. Around here, we don't LIKE coffee, we LOVE it.
b. I'd rather have a WARM bath than a HOT one.
c. Is a parallelogram SOMETIMES or ALWAYS a square?
d. If it's WARM, we'll lie out in the sun. But if it's VERY warm, we'll go inside and sit in front of the air-conditioner.

When construed with the ordinary sense of "like", (43a) would be a contradiction, but if its meaning is narrowed down to "like but not love", the sentence makes perfect sense. In other words, (43a) is no different from the examples in (42), and something rather similar goes for (43b), where the standard meaning of "warm" would yield something at least close to a contradiction (as in "I'd rather live in France than in Paris"), which is averted by an upper-bounded reading of the adjective. The remaining examples in (43) can be analysed along the same lines.

What I'm proposing, then, is that the upper-bounded interpretations observed in (43) are special cases of pragmatic narrowing, which itself is a special case of the context-dependence of lexical meaning. Conversational implicatures don't play a role in these examples, but this is not to say that they lie outside the purview of Gricean pragmatics. Rather, they are merely instances of quite general phenomenon, which has a natural place within the Gricean scheme of things.

Last call for conventionalism At this juncture, I would like to return to a suggestion we left dangling at the end of the last chapter (§7.4), where we considered the possibility of recasting conventionalism as a theory of marked interpretations of the sort observed in (43). Now that I have sketched the outlines of a Gricean account of this phenomenon, I can explain why conventionalism is not a viable alternative.

For starters, unlike the Gricean approach, conventionalist theories don't

8. There are alternative ways of interpreting examples (42b-d). These sentences might be used merely to object against the use of an expression on the grounds that it is impolite, overly bland, or uncouth. Thus interpreted, these are not instances of narrowing, though there is a still a shift away from the standard meaning, which might be analysed along the lines of Carston (1996) and Geurts (1998a).

have a unified treatment of narrowing to offer. To explain, let's have another look at Wilson and Carston's examples in (38), repeated here for convenience:

(44) a. I'm not *drinking* tonight.
b. Buying a house is easy if you've got *money*.
c. Either you become a *human being* or you leave the group.

Since it is plausible to suppose that "drink alcohol" is one of the established meanings of the verb "drink", a lexicalist analysis might be defensible for (44a). But consider how a syntax-based theory would have to explain the preferred reading of this sentence:[9]

(45) I'm not [Ɔdrinking] tonight.

In order to make this work, the truth-conditional meaning of "Ɔdrinking" will have to be restricted to alcoholic beverages, and given that Ɔ performs its job by eliminating alternatives, the alternative we need is "drink non-alcoholic beverages", so that we will get:

(46) I'm not [drinking and drinking non-alcoholic beverages] tonight.

This is blatantly stipulative. Basically, the procedure is to postulate the alternatives that are needed in order to obtain the right meaning; it is entirely unclear how these alternatives might be justified on independent grounds.

The other two examples in (44) are problematic for any version of conventionalism, though (44c) is perhaps the clearest case. I take it that "ordinary human being" is not amongst the established lexical senses of the noun "human being". True, it may be that speakers normally have ordinary human beings in mind when they use this expression, but that doesn't mean that it is an established sense. If that were the case, "ordinary N" would an established sense of *any* noun N (and the same, mutatis mutandis, for other lexical categories), and presumably it would be the most prominent one, too. But then the preferred reading of (47a) would be (47b), which isn't plausible at all.

(47) a. If human beings didn't have bones, they couldn't drive cars.
b. If ordinary human beings didn't have bones, they couldn't drive cars.

Hence, it doesn't seem very promising to try and argue that "ordinary human being" is one of the conventional senses of "human being". As for syntax-based versions of conventionalism, they don't do any better in this case than they did before, and for the same reason: technically speaking, an Ɔ-based analysis

9. To simplify the exposition a bit, I assume here that Ɔ can deal with non-propositional arguments, too.

can be made to work, but only at the expense of stipulating alternatives that are wildly ad hoc.

Next, let's have a closer look at Breheny's example:

(48) Julius isn't RICH: he's RICH. (= (42a))

I take it that the intended interpretation of this sentence is something like: "Julius is not just rich, but filthy rich". Referring to the scales in (49), if the standard, non-upper-bounded meaning of "rich" is RICH_0, then the first occurrence of "rich" in (48) would be interpreted as RICH_1 and the second occurrence as RICH_2:

(49)

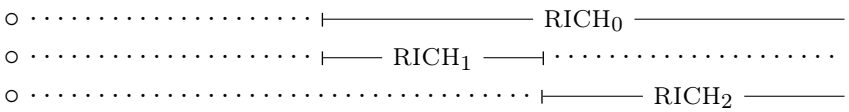

Hence, *both* occurrences of "rich" in (48) must have undergone narrowing, but in different ways: whereas RICH_1 is derived by imposing an upper bound on RICH_0, RICH_2 is obtained by shifting RICH_0's lower bound upwards, in the direction of greater wealth. Now, syntax-based versions of conventionalism are custom-made to handle the first kind of the narrowing; the second kind is simply out of their reach. As for lexicalist theories, they would have to stipulate that RICH_1 and RICH_2 are amongst the lexical meanings of "rich", which would be ad hoc. In brief: lexicalist and syntax-based theories equally fail to provide a principled and unified account of narrowing.

8.6 Conclusion

As always, the important thing is to make the right distinctions. In this chapter, I have argued that, regarding "embedded implicatures", the key distinction to be made is that between contrastive and non-contrastive environments. In a contrastive environment, upper-bounded construals may occur in just about any embedded position. Such construals, which do have truth-conditional import, I propose to analyse in terms of pragmatic narrowing, but they are the oddball cases, since they require contrastive stress or at least a contrast made salient by the context. Apart from contrast-induced narrowings, embedded implicatures are far and few between, they never have truth-conditional import, and they can always be explained in terms of Q-implicature.

This, then, is how embedded implicatures can be accommodated in a Gricean framework. The alternative is to embrace conventionalism, which seeks to capture the same phenomena in non-pragmatic terms. However, conventionalist theories have three main drawbacks. First, they leave explanatory gaps. In particular, weak implicatures, multi-sentence implicatures, and implicatures associated with presuppositions cannot be accounted for within a conventionalist framework. As a consequence, conventionalist theories will have to leave these phenomena to pragmatics, and thus are compelled to carve up the empirical facts in way that goes against nature. Secondly, conventionalist theories are faced with the problem of overgeneration: they produce too many readings. This problem is more serious for syntax-based versions of the theory, but lexicalist versions have it, too. Conventionalists try to solve this problem by stipulating, e.g., that stronger readings are preferred to weaker ones, but apart from the fact that it is stipulative and merely mitigates the problem, this strategy yields false predictions in many cases. The fact of the matter is that conventionalist theories are designed to predict that, in principle, upper-bounded construals are freely available in all embedded positions, and this doesn't seem to be correct. Thirdly, conventionalist theories are inevitably ad hoc: their predictions derive from stipulations about the lexicon and/or the grammar. Hence, the conventionalist approach is at a methodological disadvantage from the start, and since it can hardly be said to outperform its Gricean competitor, the choice between the two is a no-brainer.

Embedded implicatures have been a formidable bone of contention in the literature. In particular, they raise the question of whether a considerable tract of data that, for some decades, was routinely allocated to pragmatics shouldn't be relegated to the lexico-grammatical system instead. I have argued that they shouldn't, but in doing so have proposed several amendments to the Gricean framework, some of which were substantial. Most importantly, I believe that Grice's views on "what is said" need to be rethought in a number of respects. First, in my opinion, there is no escaping the fact that the propositional content of an utterance is determined in large part by pragmatic processes; Grice's hope (1975/1989: 25) that what somebody says is "closely related to the conventional meaning of the words" he uses is not realistic, as far as I can see. Secondly, we have to reconsider the relation between conversational implicatures and what is said, and concede that implicatures are derived not only from the propositional content of a sentence, but also from its presuppositions and even from longer stretches of discourse. Thirdly (and this is related to the second point), we have to get away from the traditional, sentence-centred approach

to interpretation. To be sure, Grice was very much aware that discourses, not sentences, are the primary objects of analysis in pragmatics, but still I believe that this insight should be carried much further than he did.

Afterword

Looking back on the preceding chapters, we can discern a number of major themes: some critical, others more constructive. The critical themes all relate to Grice's distinction between conventional and non-conventional aspects of linguistic and paralinguistic content. The seminal idea underlying the Gricean programme for pragmatics is that conventional content is to be strictly separated from conversational implicature, and one of the recurrent trends in the literature is for interpretive phenomena that admit of a perfectly decent pragmatic explanation to be nonetheless relegated to the domain of linguistic convention. This trend manifests itself variously, but I suspect it is always rooted in the same basic intuition, namely, that certain aspects of interpretation just don't *feel* like bona fide Gricean inferences of the "petrol station" variety.

To use the hackneyed example one last time, there appears to be an widely spread gut feeling that "some" has an upper-bounded meaning which is too robust to be left to Gricean pragmatics. I have tried to dispel this intuition in sundry ways: by clarifying the notion of cancellability; by arguing that Q-implicatures are not defaults; by trying to show that the robustness of some Q-implicatures admits of a pragmatic explanation and that so-called "embedded" Q-implicatures can be accommodated by the Gricean theory. I have even tried to explain what underlies the conventionalist intuition. But the long and short of it is that there is simply no *need* for a conventionalist analysis of (some) Q-implicatures, because the Gricean account will do just fine, and better than any of the non-pragmatic alternatives we have mustered.

Turning to the more constructive strands in the narrative, I have outlined a theory which is broad enough to encompass not only scalar implicatures, which have been the staple of the implicature industry since the 1970s, but

also free choice inferences, for example. This widening of the theory's scope is not only important in its own right; it also brings to light limitations of the standard Gricean approach to Q-implicatures, in which alternatives play too central a role. Not that I would want to suggest that alternatives should be dismissed altogether. On the contrary, alternatives impose vital constraints on the derivation of some varieties of Q-implicature. But: (*i*) other varieties of Q-implicature aren't constrained by the availability of alternatives at all, and (*ii*) if and when alternatives are involved, Gricean reasoning shouldn't begin with the question why the speaker failed to produced this or that alternative statement, as the Standard Recipe has it; rather, it should always begin by considering the speaker's actions and intentions.

Which brings me to the last leitmotif that I would like to bring out: the inferential basis of Q-implicatures, in particular, and Gricean reasoning, in general. The standard picture is very simple, if not downright simplistic: The speaker utters a sentence, which expresses a proposition ("what is said", in Gricean parlance), and there are alternatives to consider: what the speaker could have said but chose not to. On the standard view, these are the main ingredients for deriving implicatures. I have argued that, in point of fact, things are quite a bit more complex than this. First, as mentioned in the last paragraph, Gricean reasoning is about the speaker's intentional state. Not that there can be any doubt that this was Grice's intention all long, but oddly enough, this fundamental point often seems to have been lost in the pragmatic literature. Secondly, while that literature has concentrated its attention on conversational implicatures licensed by utterances of sentences, there is really nothing special about sentences, since Gricean reasoning equally applies above as well as below the sentence level. Thirdly, I have argued that we should take more seriously the undisputable fact that implicature is a discourse phenomenon, which means, for example, that the hearer takes into account discourse referents introduced by the speaker and that presuppositions, too, may give rise to implicatures. In short, the standard dichotomy between conventional content, on the one hand, and conversational inferences, on the other, is really much too simple.

Suppose, if only for argument's sake, that the theory of Q-implicatures I have outlined is on the right track. What, then, remains to be done? A lot, actually, but let me single out three issues. To begin with, we need more experimental data. To date, the bulk of the experimental studies on Q-implicatures have confined their attention to two words: "some" and "or". If it is true, as I have argued, that introspective evidence on implicatures isn't very reliable,

it will be clear that much needs to be done on this score. Secondly, although I hope to have shown that it is fruitful to develop a discourse-based theory of implicature, there is no such theory, as yet; what I have provided is a thumbnail sketch, at best. In my opinion, here lies one of the principal theoretical challenges for the Gricean approach to Q-implicatures and conversational implicature at large. Thirdly, even if, by its very nature, a Gricean theory of pragmatics is about speakers' beliefs and desires, it is not a processing theory. It studies reasoning at a rather lofty level of abstraction, where memory limitations, the fallibility of human reasoning, and hangovers don't play a role. I have argued that Gricean pragmatics provides the best general framework for developing viable processing theories of pragmatic reasoning, but for the time being such theories are pies in the sky.

Notation and abbreviations

Entailment and truth-conditional equivalence

- "$\varphi \Rightarrow \psi$" means that φ entails ψ, which is to say that φ implies ψ by virtue of its truth-conditional meaning. Entailment is primarily a relation between sentences, but it is defined for smaller expressions, as well. For example, "bright red" $\Rightarrow$ "red", "and" $\Rightarrow$ "or", and so on.
- "φ iff ψ" means that φ is true if ψ is true, and vice versa; or in other words: φ and ψ are truth-conditionally equivalent.

Implication

- "$\varphi \rightsquigarrow \psi$" is used, somewhat loosely, to indicate that φ may be construed as implying ψ. In many cases, my argument will be that ψ is conversationally implicated by φ, but occasionally the nature of the implication is left open.

Hence:

- "$\varphi \not\rightsquigarrow \psi$" means that φ doesn't imply ψ, except perhaps under very special circumstances.
- "$\varphi \overset{?}{\rightsquigarrow} \psi$" means that it is moot whether or not $\varphi \rightsquigarrow \psi$.

Propositional connectives

- "$\neg\varphi$" means that φ is not the case.
- "$\varphi \wedge \psi$" means that φ and ψ are both the case.
- "$\varphi \vee \psi$" means that at least one of φ and ψ is the case.

The propositional connectives are used with their classical meanings, and it is especially important to note that disjunction is inclusive: if φ and ψ are both true, then "$\varphi \vee \psi$" is true, as well.

Epistemic notions

- "$\textsc{Bel}_S(\varphi)$" means that S believes that φ.
- "$\textsc{Poss}_S(\varphi)$" means that S considers it possible that φ.

For our purposes, BEL and POSS may be viewed as very basic modal operators, which are interdefinable:

$\textsc{Bel}_S(\varphi)$ iff $\neg\textsc{Poss}_S(\neg\varphi)$

I.e., believing φ is tantamount to rejecting the possibility that φ is false. Conversely:

$\textsc{Poss}_S(\varphi)$ iff $\neg\textsc{Bel}_S(\neg\varphi)$

I.e., to consider it possible that φ is tantamount to not believing that φ is false.

Note, finally, that "$\neg\textsc{Bel}_S(\varphi)$" is not equivalent to "$\textsc{Bel}_S(\neg\varphi)$": if S doesn't have an opinion as to whether φ is true, then "$\neg\textsc{Bel}_S(\varphi)$" is true whereas "$\textsc{Bel}_S(\neg\varphi)$" is false. "$\textsc{Bel}_S(\neg\varphi)$" entails "$\neg\textsc{Bel}_S(\varphi)$", but not vice versa.

Bibliography

Abbott, B. (2000). Presuppositions as nonassertions. *Journal of pragmatics* 32: 1419–1437.

Altmann, G. and M. Steedman (1988). Interaction with context during human sentence processing. *Cognition* 30: 191–238.

Asher, N. and D. Bonevac (2005). Free choice permission is strong permission. *Synthese* 145: 303–323.

Atlas, J. and S. C. Levinson (1981). *It*-clefts, informativeness and logical form: radical pragmatics (revised standard version). In P. Cole (Ed.), *Radical pragmatics*, pp. 1–61. New York: Academic Press.

Bach, K. and R. M. Harnish (1979). *Linguistic communication and speech acts*. Cambridge, Massachusetts: MIT Press.

Baldwin, D. A. (1991). Infants' contribution to the achievement of joint reference. *Child development* 62: 875–890.

Baldwin, D. A. (1993). Infants' ability to consult the speaker for clues to word reference. *Journal of child language* 20: 395–418.

Baron-Cohen, S., A. Leslie, and U. Frith (1985). Does the autistic child have a "theory of mind"? *Cognition* 21: 37–46.

Bartsch, R. (1973). "Negative transportation" gibt es nicht. *Linguistische Berichte* 27: 1–7.

Bartsch, R. (1987). Context-dependent interpretations of lexical items. In J. Groenendijk, D. de Jongh, and M. Stokhof (Eds.), *Foundations of pragmatics and lexical semantics*, pp. 1–26. Dordrecht: Foris.

Beaver, D. and B. Geurts (2010). Presupposition. *Stanford encyclopedia of philosophy*.

Bezuidenhout, A. and J. Cutting (2002). Literal meaning, minimal propositions, and pragmatic processing. *Journal of pragmatics* 34: 433–456.

Blutner, R. (2000). Some aspects of optimality in natural language interpretation. *Journal of semantics* (17): 189–216.

Bosch, P. (1984). *Lexical learning, context dependence, and metaphor*. Indiana University Linguistics Club.

Bott, L. and I. A. Noveck (2004). Some utterances are underinformative: the onset and time course of scalar inferences. *Journal of memory and language* 51: 437–457.

Braine, M. D., D. O'Brien, I. Noveck, M. Samuels, R. Lea, S. Fisch, and Y. Yang (1995). Predicting intermediate and multiple conclusions in propositional logic inference problems: further evidence for a mental logic. *Journal of experimental psychology: general* 124: 263–292.

Braine, M. D. and B. Rumain (1981). Development of comprehension of "or": evidence for a sequence of competencies. *Journal of experimental child psychology* 31: 46–70.

Breheny, R. (2006). Communication and folk psychology. *Mind and language* 21: 74–107.

Breheny, R., N. Katsos, and J. Williams (2006). Are generalised scalar implicatures generated by default? An on-line investigation into the role of context in generating pragmatic inferences. *Cognition* 100: 434–463.

Brown, R. (1958). How shall a thing be called? *Psychological review* 65: 14–21.

Carreiras, M., A. Garnham, J. Oakhill, and K. Cain (1996). The use of stereotypical gender information in constructing a mental model: evidence from English and Spanish. *The quarterly journal of experimental psychology* 49A: 639–663.

Carston, R. (1996). Metalinguistic negation and echoic use. *Journal of pragmatics* 25: 309–330.

Carston, R. (2002). *Thoughts and utterances: the pragmatics of explicit communication*. Oxford: Blackwell.

Chemla, E. (2007). *Présuppositions et implicatures scalaires: études formelles et expérimentales*. Doctoral dissertation, EHESS, Paris.

Chevallier, C., I. A. Noveck, T. Nazir, L. Bott, V. Lanzetti, and D. Sperber (2008). Making disjunctions exclusive. *The quarterly journal of experimental psychology* 61: 1741–1760.

Chien, A. (2008). Scalar implicature and contrastive explanation. *Synthese* 161: 47–66.

Chierchia, G. (2004). Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. In A. Belletti (Ed.), *Structures and beyond*, pp. 39–103. Oxford University Press.

Chierchia, G. (2006). Broaden your views: implicatures of domain widening and the "logicality" of language. *Linguistic inquiry* 37: 535–590.

Chierchia, G., D. Fox, and B. Spector (to appear). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In C. Maienborn, K. von Heusinger, and P. Portner (Eds.), *Handbook of semantics*.

Mouton de Gruyter.
Cohen, L. J. (1971). The logical particles of natural language. In Y. Bar-Hillel (Ed.), *Pragmatics of natural language*, pp. 50–68. Dordrecht: Reidel.
Cohen, L. J. (1977). Can the conversationalist hypothesis be defended? *Philosophical studies* 31: 81–90.
Cruse, D. (1977). The pragmatics of lexical specificity. *Journal of linguistics* 13: 153–164.
De Neys, W. and W. Schaeken (2007). When people are more logical under cognitive load: dual task impact on scalar implicature. *Experimental psychology* 54: 128–133.
Dennett, D. C. (1987). *The intentional stance*. Cambridge, Massachusetts: MIT Press.
Dummett, M. (1973). *Frege: Philosophy of language*. London: Duckworth.
Eckardt, R. (2007). Licensing "or". In U. Sauerland and P. Stateva (Eds.), *Presupposition and implicature in compositional semantics*, pp. 34–70. Houndmills, Basingstoke: Palgrave Macmillan.
Evans, J. S. and S. E. Newstead (1980). A study of disjunctive reasoning. *Psychological research* 41: 373–388.
Evans, J. S., S. E. Newstead, and R. M. Byrne (1993). *Human reasoning: the psychology of deduction*. Hove, East Sussex: Lawrence Erlbaum.
Fauconnier, G. (1978). Implication reversal in a natural language. In F. Guenthner and S. Schmidt (Eds.), *Formal semantics and pragmatics of natural language*, pp. 289–301. Dordrecht: Reidel.
Feeney, A., S. Scrafton, A. Duckworth, and S. J. Handley (2004). The story of "some": everyday pragmatic inference by children and adults. *Canadian journal of experimental psychology* 58: 121–132.
Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland and P. Stateva (Eds.), *Presupposition and implicature in compositional semantics*, pp. 71–120. Houndmills, Basingstoke: Palgrave Macmillan.
Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100: 25–50.
Frege, G. (1918). Der Gedanke. *Beiträge zur Philosophie des deutschen Idealismus* 1: 8–77.
Gazdar, G. (1979). *Pragmatics: implicature, presupposition, and logical form*. New York: Academic Press.
Geis, M. L. and A. M. Zwicky (1971). On invited inferences. *Linguistic inquiry* 2: 561–566.
Geurts, B. (1998a). The mechanisms of denial. *Language* 74: 274–307.
Geurts, B. (1998b). Scalars. In P. Ludewig and B. Geurts (Eds.), *Lexikalische Semantik aus kognitiver Sicht*, pp. 85–117. Tübingen: Gunter Narr.

Geurts, B. (1999a). *Presuppositions and pronouns*. Oxford: Elsevier.

Geurts, B. (1999b). Specifics. In B. Geurts, M. Krifka, and R. van der Sandt (Eds.), *Focus and presupposition in multi-speaker discourse*, pp. 99–129. Utrecht: ESSLLI 99.

Geurts, B. (2003). Reasoning with quantifiers. *Cognition* 86: 223–251.

Geurts, B. (2005). Entertaining alternatives: disjunctions as modals. *Natural language semantics* 13: 383–410.

Geurts, B. (2006). Implicature as a discourse phenomenon. In E. Puig-Waldmüller (Ed.), *Proceedings of Sinn und Bedeutung* 11, pp. 261–275. Barcelona: Universitat Pompeu Fabra.

Geurts, B. (2009). Scalar implicature and local pragmatics. *Mind and language* 24: 51–79.

Geurts, B. and N. Pouscoulous (2009). Embedded implicatures?!? *Semantics and pragmatics* 2(4): 1–34.

Geurts, B. and F. van der Slik (2005). Monotonicity and processing load. *Journal of semantics* 22: 97–117.

Green, G. (1996). *Pragmatics and natural language understanding*. Mahwah: Lawrence Erlbaum.

Grice, H. P. (1957). Meaning. *Philosophical review* 66: 377–388. Reprinted in and cited from Grice (1989: 213-223).

Grice, H. P. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and semantics 3: Speech acts*, pp. 41–58. New York: Academic Press. Reprinted in and cited from Grice (1989: 22-40).

Grice, H. P. (1978). Further notes on logic and conversation. In P. Cole (Ed.), *Syntax and semantics 9: Pragmatics*, pp. 113–128. New York: Academic Press. Reprinted in and cited from Grice (1989: 41-57).

Grice, H. P. (1981). Presupposition and conversational implicature. In P. Cole (Ed.), *Radical pragmatics*, pp. 183–198. New York: Academic Press. Reprinted in and cited from Grice (1989: 269-282).

Grice, H. P. (1989). *Studies in the way of words*. Cambridge, Massachusetts: Harvard University Press.

Guasti, M. T., G. Chierchia, S. Crain, F. Foppolo, and A. Gualmini (2005). Why children and adults sometimes (but not always) compute implicatures. *Language and cognitive processes* 20: 667–696.

Hagoort, P. and C. M. Brown (1994). Brain responses to lexical ambiguity resolution and parsing. In C. Clifton, L. Frazier, and K. Rayner (Eds.), *Perspectives on sentence processing*, pp. 45–81. Hillsdale, New Jersey: Lawrence Erlbaum.

Harman, G. H. (1965). The inference to the best explanation. *The philosophical review* 74: 88–95.

Hirschberg, J. (1985). *A theory of scalar implicature*. Ph.D. thesis, University of

Pennsylvania.

Horn, L. R. (1972). *On the semantic properties of the logical operators in English*. Ph.D. thesis, University of California at Los Angeles.

Horn, L. R. (1984). Towards a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, form, and use in context: linguistic applications*, pp. 11–42. Washington: Georgetown University Press.

Horn, L. R. (1985). Metalinguistic negation and pragmatic ambiguity. *Language* 61: 121–174.

Horn, L. R. (1989). *A natural history of negation*. Chicago University Press.

Horn, L. R. (1997). All John's children are as bald as the king of France: existential import and the geometry of opposition. *CLS* 33: 155–179.

Horn, L. R. (2000). From *if* to *iff*: conditional perfection as pragmatic strengthening. *Journal of pragmatics* 32: 289–326.

Horn, L. R. (2006). The border wars: a neo-Gricean perspective. In K. von Heusinger and K. Turner (Eds.), *Where semantics meets pragmatics*, pp. 21–48. Amsterdam: Elsevier.

Horn, L. R. (2007). Toward a Fregean pragmatics: "Voraussetzung", "Nebengedanke", "Andeutung". In I. Kecskes and L. R. Horn (Eds.), *Explorations in pragmatics: linguistic, cognitive and intercultural aspects*, pp. 39–69. Berlin: Mouton de Gruyter.

Horn, L. R. (2009). WJ-40: Implicature, truth, and meaning. *International review of pragmatics* 1: 3–34.

Huang, Y. T. and J. Snedeker (2006). On-line interpretation of scalar quantifiers: insight into the semantic-pragmatics interface. In *Proceedings of the 28th annual conference of the Cognitive Science Society (CogSci 2006)*, pp. 351–356.

Hurford, J. (1974). Exclusive or inclusive disjunction. *Foundations of language* 11: 409–411.

Jennings, R. (1994). *The genealogy of disjunction*. Oxford University Press.

Johnson, S., V. Slaughter, and S. Carey (1998). Whose gaze will infants follow? the elicitation of gaze-following in 12-month-olds. *Developmental science* 1: 233–238.

Just, M. A. and P. A. Carpenter (1980). A theory of reading: from eye fixations to comprehension. *Psychological review* 87: 329–354.

Kamp, H. (1973). Free choice permission. *Proceedings of the Aristotelian Society* 74: 57–74.

Kamp, H. (1979). Semantics versus pragmatics. In F. Guenthner and S. Schmidt (Eds.), *Formal semantics and pragmatics of natural language*, pp. 255–287. Dordrecht: Reidel.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal methods in the study of language*, pp. 277–322. Amsterdam: Mathematical Centre Tracts 135.

Kamp, H. and U. Reyle (1993). *From discourse to logic*. Dordrecht: Kluwer.

Karttunen, L. (1971). Discourse referents. Indiana University Linguistics Club. Reprinted in: McCawley (1976: 363-385).

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and philosophy* 30: 669–690.

Klinedinst, N. (2007). Plurals, possibilities, and conjunctive disjunction. *UCL working papers in linguistics* 19: 261–284.

Kutas, M. and S. Hillyard (1980). Reading senseless sentences: brain potentials reflect semantic anomaly. *Science* 207: 203–205.

Landman, F. (1998). Plurals and maximalization. In S. Rothstein (Ed.), *Events and grammar*, pp. 237–271. Dordrecht: Kluwer.

Levinson, S. C. (1983). *Pragmatics*. Cambridge University Press.

Levinson, S. C. (2000). *Presumptive meanings*. Cambridge, Massachusetts: MIT Press.

Lipton, P. (2004). *Inference to the best explanation*. London: Routledge.

Liszkowski, U., M. Carpenter, T. Striano, and M. Tomasello (2006). Twelve-and 18-month-olds point to provide information. *Journal of cognition and development* 7: 173–187.

Liszkowski, U., M. Carpenter, and M. Tomasello (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition* 108: 732–739.

Matsumoto, Y. (1995). The conversational condition on Horn scales. *Linguistics and philosophy* 18: 21–60.

McCawley, J. D. (Ed.) (1976). *Syntax and semantics 7: Notes from the linguistic underground*. New York: Academic Press.

McCawley, J. D. (1991). Contrastive negation and metalinguistic negation. In *Papers from the 27th regional meeting of the Chicago Linguistics Society. Part 2: the parasession on negation*, pp. 189–206.

Mill, J. S. (1865). *An examination of Sir William Hamilton's philosophy and of the principal philosophical questions discussed in his writings* (2nd ed.). London: Longmans, Green and Co.

Milsark, G. (1977). Toward an explanation of certain peculiarities of the existential construction in English. *Linguistic analysis* 3: 1–29.

Mitchell, D. and D. Green (1978). The effects of context and content on immediate processing in reading. *The quarterly journal of experimental psychology* 30: 609–636.

Morgan, J. (1978). Two types of convention in indirect speech acts. In P. Cole (Ed.), *Syntax and Semantics 9: Pragmatics*, pp. 261–280. New York: Academic Press.

Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and philosophy* 15: 509–559.

Nieuwland, M. S. and J. J. A. van Berkum (2006). When peanuts fall in love: N400 evidence for the power of discourse. *Journal of cognitive neuroscience* 18: 1098–1111.

Noordman, L. G. (1977). *Inferring from language*. Doctoral dissertation, University of Groningen.

Noveck, I. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition* 78: 165–188.

Noveck, I. and D. Sperber (2007). The why and how of experimental pragmatics: the case of "scalar inferences". In N. Burton-Roberts (Ed.), *Pragmatics*, pp. 184–212. Basingstoke: Palgrave Macmillan.

Noveck, I. A. and A. Posada (2003). Characterising the time course of an implicature. *Brain and language* 85: 203–210.

Nunberg, G. D. (1978). *The pragmatics of reference*. Indiana University Linguistics Club.

Nunberg, G. D. (1979). The non-uniqueness of semantic solutions: polysemy. *Linguistics and philosophy*: 143–184.

Onishi, K. H. and R. Baillargeon (2005). Do 15-month-old infants understand false beliefs? *Science* 308: 255–258.

Papafragou, A. and J. Musolino (2003). Scalar implicatures: experiments at the syntax-semantics interface. *Cognition* 86: 253–282.

Paris, S. G. (1973). Comprehension of language connectives and propositional logical relationships. *Journal of experimental child psychology* 16: 278–291.

Peirce, C. S. (1931). *Collected papers of Charles Sanders Peirce, Volume 1* (edited by Charles Hartshorne and Paul Weiss). Cambridge, Massachusetts: Harvard University Press.

Pijnacker, J., P. Hagoort, J. Buitelaar, J.-P. Teunisse, and B. Geurts (2009). Pragmatic inferences in high-functioning adults with autism and Asperger syndrome. *Journal of autism and developmental disorders* 39: 607–618.

Pouscoulous, N. (2006). *Processing scalar inferences*. Doctoral dissertation, EHESS, Paris.

Pouscoulous, N., I. Noveck, G. Politzer, and A. Bastide (2007). A developmental investigation of processing costs in implicature production. *Language acquisition* 14: 347–376.

Recanati, F. (1989). The pragmatics of what is said. *Mind and language* 4: 295–329.

Recanati, F. (2003). Embedded implicatures. *Philosophical perspectives* 17: 299–332.

Recanati, F. (2004). *Literal meaning*. Cambridge University Press.

Rosch, E. (1978). Principles of categorization. In E. Rosch and B. Lloyd (Eds.), *Cognition and categorization*, pp. 27–48. Lawrence Erlbaum.

Russell, B. (2006). Against grammatical computation of scalar implicatures. *Jour-*

nal of semantics 23: 361–382.

Sag, I. (1981). Formal semantics and extralinguistic content. In P. Cole (Ed.), *Radical pragmatics*, pp. 273–294. New York: Academic Press.

Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and philosophy* 27: 367–391.

Schulz, K. (2005). A pragmatic solution for the paradox of free choice permission. *Synthese* 147: 343–377.

Schulz, K. and R. van Rooij (2006). Pragmatic meaning and non-monotonic reasoning: the case of exhaustive interpretation. *Linguistics and philosophy* 29: 205–250.

Searle, J. R. (1969). *Speech acts: an essay in the philosophy of language*. Cambridge University Press.

Searle, J. R. (1975). Indirect speech acts. In P. Cole and J. Morgan (Eds.), *Syntax and semantics 3: Speech acts*, pp. 59–82. New York: Academic Press.

Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: evidence for effects of informativity expectations. *Journal of psycholinguistic research* 32: 3–23.

Sedivy, J. C. (2007). Implicature during real time conversation: a view from language processing research. *Philosophy compass* 2/3: 275–496.

Sedivy, J. C., M. K. Tanenhaus, C. Chambers, and G. N. Carlson (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition* 71: 109–147.

Simons, M. (2006). Presupposition without common ground. Unpublished manuscript, Carnegie Mellon University.

Soames, S. (1982). How presuppositions are inherited: a solution to the projection problem. *Linguistic inquiry* 13: 483–545.

Southgate, V., A. Senju, and G. Csibra (2007). Action anticipation through attribution of false belief by 2-year-olds. *Psychological science* 18: 587–592.

Spector, B. (2006). *Aspects de la pragmatique des opérateurs logiques*. Doctoral dissertation, University of Paris VII.

Sperber, D. and D. Wilson (1995). *Relevance* (2nd edition). Oxford: Blackwell.

Sperber, D. and D. Wilson (2002). Pragmatics, modularity and mind-reading. *Mind and language* 17: 3–23.

Stalnaker, R. C. (1973). Presuppositions. *Journal of philosophical logic* 2: 447–457.

Stalnaker, R. C. (1974). Pragmatic presuppositions. In M. K. Munitz and P. Unger (Eds.), *Semantics and philosophy*, pp. 197–214. New York University Press.

Storto, G. and M. K. Tanenhaus (2005). Are scalar implicatures computed online? In E. Maier, C. Bary, and J. Huitink (Eds.), *Proceedings of Sinn und Bedeutung 9*, pp. 431–445.

Strawson, P. F. (1952). *Introduction to logical theory*. London: Methuen.

Surian, L., S. Caldi, and D. Sperber (2007). Attribution of beliefs by 13-month-old infants. *Psychological science* 18: 580–586.

Thagard, P. and C. Shelley (1997). Abductive reasoning: logic, visual thinking, and coherence. In M. L. Dalla Chiara, K. Doets, D. Mundici, and J. van Benthem (Eds.), *Logic and scientific methods*, pp. 413–427. Dordrecht: Kluwer.

van Berkum, J., D. van den Brink, C. Tesink, M. Kos, and P. Hagoort (2008). The neural integration of speaker and message. *Journal of cognitive neuroscience* 20: 580–591.

van der Auwera, J. (1997a). Conditional perfection. In A. Athanasiadou and R. Dirven (Eds.), *On conditionals again*, pp. 169–190. Amsterdam: Benjamins.

van der Auwera, J. (1997b). Pragmatics in the last quarter century: the case of conditional perfection. *Journal of pragmatics* 27: 261–274.

van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of semantics* 9: 333–377.

van der Wouden, T. (1997). *Negative contexts: collocation, polarity and multiple negation*. London: Routledge.

van Rooij, R. and K. Schulz (2004). Exhaustive interpretation of complex sentences. *Journal of logic, language and information* 13: 491–519.

von Wright, G. H. (1968). *An essay in deontic logic and the general theory of action*. Amsterdam: North-Holland.

Wainer, J. (2007). Modeling generalized implicatures using non-monotonic logics. *Journal of logic, language and information* 16: 195–216.

Walker, R. C. S. (1975). Conversational implicatures. In S. Blackburn (Ed.), *Meaning, reference and necessity*, pp. 133–181. Cambridge University Press.

Wilson, D. (1975). *Presupposition and non-truth-conditional semantics*. London: Academic Press.

Wilson, D. (2000). Metarepresentation in linguistic communication. In D. Sperber (Ed.), *Metarepresentations: a multidisciplinary perspective*, pp. 411–448. Oxford University Press.

Wilson, D. and R. Carston (2007). A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics*, pp. 230–259. London: Palgrave Macmillan.

Wilson, D. and D. Sperber (2004). Relevance theory. In L. R. Horn and G. Ward (Eds.), *Handbook of pragmatics*, pp. 607–632. Oxford: Blackwell.

Wimmer, H. and J. Perner (1983). Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13: 41–68.

Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural language semantics* 8: 255–290.