

**CONVERSATIONAL IMPLICATURE:
RE-ASSESSING THE GRICEAN FRAMEWORK**

Maria Kasmirli

Submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy in the Department of Philosophy
The University of Sheffield, October 2016

Word count (excluding abstract and references): 67,959

Abstract

Conversational implicature is (roughly) the practice of conveying one thing by saying another. Philosophical and linguistic work on the topic has been dominated by the approach proposed by Paul Grice — the *Gricean framework*, as I call it — according to which implicatures can be calculated from principles of cooperative behaviour. The framework faces numerous objections and counterexamples, however, and this thesis reassesses it in the light of recent work in the area. Chapters 1 and 2 introduce the topic, provide a detailed exposition of the Gricean framework, and highlight a problem concerning the role of speaker intentions in implicature. Chapter 3 sets out some problems for Grice’s approach and argues that we can address them by reinterpreting his framework as a normative one. It proposes some revisions to the framework to make it more compatible with this reading and shows how the tension in Grice’s view of speaker intentions can be resolved. Chapter 4 then argues that, despite its attractions, the revised theory has a serious flaw, being unable to establish norms of implicature that are speaker-independent. The chapter proposes instead an intention-centred account, which abandons the requirement of calculability and allows a direct role for speaker intentions, while still preserving a normative element. Chapter 5 looks at *neo-Gricean* theories, which use Gricean principles to explain a range of supposedly context-independent implicatures. It sets out some problems for neo-Griceanism, comparing it with rival approaches and surveying relevant experimental evidence. The chapter concludes that implicature is more context-sensitive than neo-Griceanism allows and that general principles have at best a limited role in its explanation. Chapter 6 draws some conclusions, arguing that implicature is less rational than Grice supposed and more dependent on context and speaker intention. It also offers some speculations about the social role and ethics of implicature.

This thesis is dedicated with love
to Harry, Mikis, and Matéa.

Contents

Abstract	2
Figures and tables	7
Acknowledgements	8
A note on grammar	10
Chapter 1	11
Implicature: questions and theories	11
1. The case of Mr Bronston	11
2. Issues and questions	14
2.1 Implicature generation	14
2.2 Implicature recovery	15
2.3 Normative issues	17
2.4 Ethical questions	18
3. Theories	19
3.1 Grice's account and some alternatives	19
3.2 The present thesis	20
3.3 Methodological remarks	24
Conclusion	25
Chapter 2	26
The Gricean framework	26
1. Saying and implicating	26
2. Implicature generation	30
3. Calculability	34
4. Particularized and generalized implicatures	38
5. Implicature and speaker meaning	42
6. Applications	47
Conclusion	50
Chapter 3	51
Problems, reinterpretation, and revision	51
1. Problems for Grice's definition	51
1.1 Problems with the cooperative presumption.	52
1.2 Problems with determinacy and calculability.	56

1.3 Problems with mutual knowledge	64
2. A normative reading	66
2.1 Speaker meaning, implicature, and an extended taxonomy	66
2.2 An enriched Gricean framework	72
3. Some modifications	73
3.1 The cooperative presumption revised	73
3.2 Mutual knowledge revised	74
3.3 Unmeant implicatures	78
4. Implicature and speaker meaning again	80
4.1 Normative and psychological conditions for implicature	80
4.2 Speaker implicature and utterance implicature	82
4.3 The role of intention	84
Conclusion	85
Chapter 4	87
Where the Gricean framework fails	87
1. The argument for speaker-dependency	88
1.1 Normativity and speaker dependency	88
1.2 Background knowledge	90
1.3 What is said	97
1.4 Cooperativeness	99
1.5 Context	104
2. Responding to the argument	105
2.1 Resisting speaker dependency	105
2.2 Consequences of speaker-dependency	107
2.3 An intention-centred account of implicature	110
3. Utterer-implicature and audience-implicature	117
4. Implicature recovery	123
Chapter 5	129
Neo-Griceanism and its rivals	129
1. Neo-Griceanism	130
1.1 Utterance-type meaning	130
1.2 The three principles	133
1.3 Applying the principles	137

2. Alternatives to neo-Griceanism	139
2.1 Relevance theory	139
2.2 Convention theory	143
2.3 Weak neo-Griceanism	147
2.4 Back to Levinson	148
3. Assessing the Q-principle	149
3.1 ‘An X’	149
3.2 Scalar implicatures	153
3.3 Reducing scalar GCIs to PCIs	158
3.4 Q-implicature and T-implicature	165
3.5 Scalar implicature or explicature?	168
4. Assessing the I- and M-principles	172
4.1 Stereotypes and defaults	172
4.2 A deeper problem	174
5. Experimental evidence	176
5.1 Reaction-time studies	177
5.2 Developmental studies	181
5.3 Tentative conclusions	187
6. Conclusions	188
Chapter 6	190
Taking stock and looking forward	190
1. Taking stock	190
1.1 The Gricean framework	190
1.2 Normativity	190
1.3 Speaker intentions	191
1.4 An intention-centred account	192
1.5 Generalized implicatures	192
1.6 Implicature recovery	194
2. Looking forward	195
References	198

Figures and tables

Figure 1: The breakdown of what a speaker means, according to Stephen Neale's interpretation of Grice.	44
Figure 2: The relation between speaker meaning, sentence meaning, and conversational implicature, on Saul's reading of Grice.	68
Table 1: Saying and implicating.	81
Figure 3: The relation between sentence meaning and what is said, and between utterance implicature and speaker implicature, on the view proposed here.	83
Table 2: Saying and implicating according to the intention-centred account of conversational implicature.	116

Acknowledgements

This thesis has taken a very long time to finish. Work on it had to be fitted into the gaps left by full-time employment, and had to compete with the demands of several changes of job, three pregnancies, emigration, caring duties to disabled and elderly relatives, and the extended and unrelenting pressures of the Greek financial crisis. The fact that I didn't give up years ago and actually stuck it out to the end is due principally to two people: my primary supervisor, Professor Jennifer Saul, and my partner, Keith Frankish.

Jenny's skill and dedication as a supervisor deserve the highest praise. She has taught me, encouraged me, and perhaps most importantly, believed in me. Under her guidance and with her support, I have developed as a philosopher and achieved things I did not think I could. She has read many drafts of this thesis, always providing thoughtful, focused, and constructive comments and sound strategic advice, helping me to see the wood when I was lost in the trees. I thank her for her dedication and patience, for her support through some very difficult times, and for not giving up on me when it would have been much easier for her to do so.

Keith has also read many drafts of this material and provided immensely valuable feedback, especially on stylistic matters. He has listened to me talk about conversational implicature for many hours (despite, I fear, not finding the subject a particularly exciting one), and his questions and comments have helped me to sharpen my thinking on the topic in innumerable ways, large and small. Even more importantly, he has supported me, believed in me, and made many sacrifices to give me the time needed to complete the project. I couldn't have done it without him.

I also want to record my gratitude to George Botterill, who was my secondary supervisor during this project and with whom I spent many hours in thoughtful and productive discussion about implicature.

I am grateful to many other friends and colleagues for discussions that have fed directly or indirectly into this thesis. These include (in alphabetical order): Chris Daly, Johnathan Dancy, Wayne Davis, André Gallois, Eve Garrard, Andy Hamilton, Nick Lund, David McNaughton, and John Rogers.

Special thanks are due to my parents, Charis and Toula. My mother taught me to stand up for myself and to pursue my dreams, and my father set me an example of what can be achieved with determination and hard work. I also thank my dear brothers Giorgos, Lefteris, and Iannis, and my parents-in-law, Arthur and Eileen, for all their love and support. I am also very grateful to Eileen for her careful proofreading of the final draft.

I have dedicated this thesis to my children, and I thank them for putting up with my preoccupation with it for so long and for keeping my nose to the grindstone during the final push to finish.

Over the years I have given a number of talks drawing on material related to this thesis, including seminars at the universities of Keele and Sheffield and conferences at Lund, Reading, and Sheffield again. I am grateful to the audiences on all those occasions for their comments and questions.

A version of some of the material in Chapter 5, section 5 was originally prepared as part of my contribution to a paper on dual processes and implicature co-authored with Keith Frankish (Frankish and Kasmirli 2010). My thanks to Keith and the editors of the collection in which the paper appeared for inviting me to contribute to the paper. The material in question has been re-written, expanded, and updated for the thesis.

I have kept this list of thanks to the bare minimum and could have extended it much further. Many other friends and colleagues have helped too, in one way or another, and I am grateful to them all.

A note on grammar

In this thesis I use ‘they’, ‘them’, and ‘their’ as gender-neutral pronouns. This practice has a long history in English and, in the words of *The Cambridge Guide to English Usage*, ‘has become unremarkable — an element of common usage’ (Peters 2004, p.538).

Chapter 1

Impicature: questions and theories

1. The case of Mr Bronston

On 10 June 1966, American movie producer Samuel Bronston was being questioned under oath at a bankruptcy hearing. His production company, Samuel Bronston Productions Inc, had failed two years earlier, and lawyers for its creditors wanted to know what overseas assets it held. (Bronston made films in European countries, where costs were lower, and his company held bank accounts in these countries.) In the course of this questioning, the following exchange took place between Bronston and one of the lawyers:

Q. Do you have any bank accounts in Swiss banks, Mr. Bronston?

A. No, sir.

Q. Have you ever?

A. The company had an account there for about six months, in Zurich.

Q. Have you any nominees who have bank accounts in Swiss banks?

A. No, sir.

Q. Have you ever?

A. No, sir.

(Quoted in *Bronston v. United States* 1973)

Bronston's answers were truthful, but the second of them was misleading. Bronston was asked whether *he* had ever had an account in a Swiss bank, but he replied by saying that *his company* had had such an account. The lawyer took this to indicate that Bronston himself had *not* had a Swiss bank account, and moved on. In fact, this was not true. Bronston had had a personal account with a bank in Geneva for nearly five years during the relevant period. He had made large deposits into the account and transferred money from it to his production company. When this was discovered, Bronston was charged with perjury.

At his trial, the prosecution argued that Bronston had deliberately chosen to give information about his company's Swiss bank account in order to give the impression that he himself had not had such an account. The District Court

instructed the jury that perjury consists in ‘wilfully testifying to the truth of a fact which the defendant does not believe to be true’, but added that Bronston could be convicted of perjury if he had given an answer which was ‘not literally false but [which] when considered in the context in which it was given, nevertheless constitute[d] a false statement.’ The Court gave the following example:

[I]f it is material to ascertain how many times a person has entered a store on a given day and that person responds to such a question by saying five times when in fact he knows that he entered the store 50 times that day, that person may be guilty of perjury even though it is technically true that he entered the store five times. (Quoted in *Bronston v. United States* 1973)

After over six hours of deliberation and a request for a repetition of the instructions given to them, the jury found Bronston guilty of perjury.

Bronston appealed, arguing that his answer had been truthful, even if unresponsive. The Court of Appeals ruled against him (although one judge disagreed), stating that:

an answer containing half of the truth which also constitutes a lie by negative implication, when the answer is intentionally given in place of the responsive answer called for by a proper question, is perjury. (*U.S. v. Bronston* 1971)

Bronston appealed to the Supreme Court. The case was heard in November 1972 and Chief Justice Burger gave the Court’s ruling in January 1973. Burger agreed that Bronston had implied that he had no personal Swiss bank account, and that in casual conversation this might be a reasonable interpretation of his utterance. However, he argued that perjury did not extend to the implications of a witness’s words:

the statute does not make it a criminal act for a witness to willfully state any material matter that *implies* any material matter that he does not believe to be true. (*Bronston v. United States* 1973)

Provided they believe that the answers they give are literally true, Burger stated, witnesses should not be held responsible for any further intentions behind their testimony.

A jury should not be permitted to engage in conjecture whether an unresponsive answer, true and complete on its face, was intended to mislead or divert the examiner; the state of mind of the witness is relevant only to the extent that it bears on whether ‘he does not believe (his answer) to be true.’ To hold otherwise would be to inject a new and confusing element into the adversary testimonial system we know. Witnesses would be unsure of the extent of their responsibility for the misunderstandings and inadequacies of examiners, and might well fear having that responsibility tested by a jury under the vague rubric of ‘intent to mislead’ or ‘perjury by implication.’ (ibid).

Burger concluded that it was the questioner’s duty, not the courts’, to challenge unresponsive but literally true answers. The court reversed Bronston’s conviction.¹

The Bronston case illustrates a familiar but puzzling phenomenon: our ability to convey one thing by saying another. By saying that his company had held a Swiss bank account, Bronston was somehow able to convey to his hearers the message that he himself had not had a Swiss bank account. This communicative

¹ The Bronston case established a principle known as the ‘literal truth’ rule (for example, Anon. 1999; Tiersma 1989–1990). It has been argued that this rule, together with similar legal practices, which seem to show a lack of regard for truth, have affected popular attitudes to truth telling, with the result that ‘society may have abandoned morality in favor of legality’ (Castleman 2004).

phenomenon, which philosophers and linguists call *implicature*, is the topic of this thesis.²

2. Issues and questions

Implicature raises many questions and connects with many wider issues, which will recur in different ways throughout this thesis. Here I shall briefly introduce some of the main ones.

2.1 *Implicature generation*

First, there are questions about how implicatures are determined, or, as it is often put, *generated*. By this, I mean how they come to exist, not how they are generated in the mind of the hearer. (I shall treat the question of how hearers detect implicatures separately.) What makes it the case that Mr Bronston's utterance carried an implicature? The statement 'My company had an account there' does not entail 'I did not have an account there'. The first statement could be true and the second false (indeed, that was so in Bronston's case). Nor would the first statement always carry the implicature that the second was true. If Bronston had been asked whether his company had had an account in Switzerland, then no one would have thought that his answer conveyed anything more than its literal meaning. So it seems that Bronston's implicature was due to some feature of the context. But which feature, or features, exactly? Was the implicature determined by Bronston's intentions? Did it depend in any way on how his hearers interpreted his utterance? Or was the implicature generated by non-psychological features of the communicative exchange, and if so, which ones?

There are further questions about implicature generation. Are all implicatures generated in the same way? As we shall see in the next chapter, unlike Bronston's answer, some sentences carry the same implicature in most contexts, unless words are added to cancel the implicature. (These are known as *generalized* implicatures, as opposed to context-dependent *particularized* ones.) For example, in most contexts the sentence 'Some of the students passed the test' carries the implicature

² Strictly speaking, I shall be concerned with what is called *conversational* (as opposed to *conventional*) implicature. I shall explain these terms in Chapter 2.

that not all the students passed the test. Are these implicatures generated in a different way from ones that are more context-specific, or are the same factors involved? Is it always determinate whether or not an utterance carries an implicature, and if so, what is it? If hearers disagree about the existence of a particular implicature, will there always be (at least in principle) some way of settling the dispute?

These questions are, I take, it, broadly speaking, philosophical questions. The implicated meanings of utterances, like their literal ones, depend on us. They are not intrinsic properties of the sounds involved but properties that depend in some way on how we use and react to those sounds — on our communicative practices and conventions, and our expectations, intentions, and beliefs. So in order to explain how they arise, we need to think about our everyday communicative practices and attitudes, and to analyse the conditions under which we ascribe implicatures to utterances, drawing on our intuitions about different cases and making use of thought experiments and counterexamples.

2.2 *Implicature recovery*

The second set of questions are questions about how implicatures are processed or *recovered* — that is, about the processes by which a hearer comes to interpret an utterance as carrying a particular implicature.³ Does implicature recovery involve inference, and if so, what kind of inference is it and what data does it draw on? Are there general principles of implicature recovery or is implicature derivation context-driven? Implicature recovery is a part of *pragmatic* processing, the recovery of contextual aspects of meaning, as opposed to purely *semantic* processing, which is concerned with the recovery of literal, non-contextual meaning. How is implicature recovery related to other aspects of pragmatic processing and to the processing of semantic meaning? Are semantic and pragmatic processing really distinct?

³ I use the term ‘recovery’ (or, alternatively, ‘derivation’) to contrast with ‘generation’, but I shall use it in such a way that there can be recovery without generation. That is, I allow the possibility that a hearer may interpret an utterance as carrying an implicature that it does not in fact carry, according to our preferred theory of implicature generation.

Unlike questions about implicature generation, these are questions about the mental processes involved in utterance interpretation. Theories of implicature recovery are thus broadly psychological ones, though most of them aim to describe the interpretation process at an abstract level, rather than specifying detailed cognitive mechanisms. Theories of this kind have been developed primarily by linguists and cognitive scientists, drawing on linguistic intuitions, evolutionary considerations, and, increasingly, experimental data.⁴

This distinction between implicature generation and implicature recovery is not always drawn, and some theories of implicature tend to run the two sets of questions together. (Jennifer Saul suggests that some ‘relevance’ theorists do this, taking the hearer’s interpretation of an utterance to determine what the utterance implicates; Saul 2002b.) It is not hard to see why this happens. Questions of generation and recovery are closely related, and the answers to one set may be relevant to the other. Implicatures are typically recoverable by competent human hearers, and any theory of implicature generation that would make their recovery impossible or extremely hard for humans can be ruled out. Thus considerations of recoverability constrain theories of implicature generation. Moreover, implicature recovery must be sensitive to whatever factors make it the case that implicatures exist, so a theory of implicature generation sets the target for a theory of implicature recovery. This probably accounts for why questions about generation and recovery are often run together in a ‘theory of implicature’. (I shall expand on these points later, in Chapter 4.) Indeed, implicature generation might be, in a sense, *dependent* on the recovery process. It might be that an utterance carries an implicature just because hearers are typically disposed to interpret it as doing so, and speakers can rely on this. (The ‘neo-Gricean’ theories discussed in Chapter 5 can be interpreted in this way.)

Thus, theories of implicature generation and implicature recovery are not as independent as they seem at first sight. However, they are conceptually distinct

⁴ Of course, theories about implicature generation may also be in a sense psychological. It may be that implicatures exist in virtue of certain psychological states of the speaker or hearer, or both. That is to say, an answer to the philosophical question of what implicatures are may mention psychological states. However, that does not make the question itself a psychological one.

and involve different methods of investigation, and it is important not to confuse them or to judge a theory of one by the standards appropriate to the other.

2.3 Normative issues

We can also ask *normative* questions about implicature. In the case of literal meaning we can make a distinction between what a speaker means and what their words mean. Suppose Mr Bronston had said explicitly, 'I did not have a bank account in Switzerland.' Then if he had later been presented with evidence that he had had such an account, he would have had difficulty defending himself against a perjury charge by saying that he had really meant that he had not had a bank account in *Swaziland*. There are established norms of literal meaning, and witnesses are expected to respect them. Even if a witness accidentally misspeaks, they may still be held responsible for their carelessness.

Are there similar norms for implicature, which would determine what, if anything, Mr Bronston's utterance implicated? We use implicature widely, and it can be used to convey important messages, such as invitations and consent. (Think, for example, of how a question such as, 'Shall we go upstairs?' might, in certain circumstances, be used to convey an invitation to sexual intercourse.) Implicit communication of this kind is open to abuse (as, arguably, in Mr Bronston's case), and may lead to serious misunderstanding and confusion. Having clear norms governing its use would, therefore, be very useful.

Questions about norms of implicature are obviously closely linked to questions about implicature generation. In asking how implicatures are generated we are in effect asking when it is *correct* to attribute implicatures to utterances. When the conditions for a certain utterance to generate a certain implicature are met, then it will be correct to say that the utterance carries that implicature. However, whether this yields socially useful norms will depend on what the generation conditions are. If the condition for an utterance to implicate a proposition *p* is simply that the speaker *intends* it to implicate *p*, then this would not give us *speaker-independent* norms of implicature, like those of literal meaning. Speakers would, potentially, be able to make their utterances implicate anything they liked. (We might call this the *Humpty Dumpty view* of implicature, after Lewis Carroll's Humpty Dumpty,

who claimed he could make his words mean anything he liked.)⁵ Such norms would be of little use in regulating communication, and speakers could always plead that any supposed implicatures of their utterances were unintended and hence non-existent.

On the other hand, if the conditions for implicature generation are independent of, or at least not wholly determined by, the speaker's intentions, then this might (depending on the details) support a substantive normative theory of implicature. Speakers might be held responsible for implicatures generated by their utterances, even if they had not intended them or been aware of them. If the generation conditions are also independent of the hearer's mental states, then an utterance might generate an implicature that neither speaker nor hearer notice (just as a sentence might carry a conventional meaning that neither speaker nor hearer recognize).

2.4 Ethical questions

Another set of questions concerns the ethics of implicature. What responsibility do speakers have for the implicatures their utterances carry? Was the Supreme Court right to reverse Mr Bronston's conviction for perjury? Even if it was as a matter of law, what about speakers in ordinary conversational contexts? Do speakers have a moral responsibility for beliefs their hearers form as a result of implicatures carried by their utterances? Does it matter how obvious the implicatures are? What if a speaker does not notice that their words carry an implicature? (If that is possible; if implicatures depend on the speaker's intentions, it might not be.) Are they still morally responsible for any effects the implicature has on their hearers? What about hearers? If a hearer misses an implicature, have they been negligent? Can it be negligent to trust an implicature (as the lawyer questioning Mr Bronston did)? What if different hearers disagree about what implicature, if any, an utterance carries?

Although these ethical questions can be considered on their own, we cannot deal with them fully until we have good theories of how implicatures are generated and recovered. In order to properly assess speakers' responsibility in this area, we

⁵ Carroll, 2009, p.190 (originally published in 1880).

need to know when and how implicatures come to exist, and what control speakers have over the factors involved. For example, if implicatures depend on the speaker's intentions, then speakers have a greater degree of control over (and thus, arguably, greater responsibility for) what their words implicate than if implicatures depend on conventions or other aspects of the situation that are not under the speaker's control. Similarly, understanding how implicatures are recovered will help us to assess the extent of hearers' duties with regard to the detection of implicatures and to decide when they have been negligent in missing them.

I shall return briefly to questions about the ethics of implicature in the final chapter, but for the most part I shall focus on the preliminary questions about generation and recovery. In this respect the current thesis prepares the ground for further work on the ethics of implicature.

3. Theories

3.1 Grice's account and some alternatives

More than any other person it was the philosopher Paul Grice (1913–1988) who brought implicature to the attention of philosophers and linguists, and Grice's own account of the nature of implicature generation (first presented in a 1967 lecture series and published in 1975) has provided a hugely influential framework for thinking about implicature (Grice 1975). The core idea of the account is that the link between utterances and the implicatures they carry is not arbitrary or contingent, based on the speaker's intentions or general conventions, but a *rational* one, grounded in general principles of cooperative behaviour. Grice argues that an implicature arises when an utterance would be uncooperative if taken literally, violating one or more maxims about how a cooperative speaker should convey information. Since a presumption of cooperation is essential to communication, Grice argues, in such cases the speaker must be understood to be conveying something other than the literal meaning of their utterance, and this is the implicated meaning. On this view, implicatures can be calculated from general communicative principles, although Grice does not claim that hearers must actually go through this calculation process in order to recover them. This approach aims to provide a unified account of both generalized, context-

independent implicatures and particularized, context-specific ones, and it has been the dominant approach to implicature in the philosophical literature.

Grice's work has also inspired psycho-linguistic theories that approach implicature through the recovery process (for example, Levinson 2000). These so-called 'neo-Gricean' theories hold that when hearers interpret utterances they automatically apply certain heuristics, related to the maxims Grice proposes, which transform and enrich the literal meanings of the utterances in various ways, creating a new level of meaning which speakers can exploit and which makes communication more efficient. According to neo-Griceans, implicatures of the generalized kind belong to this level of meaning. This account does not, however, extend to particularized implicatures, and requires us to make a sharp division between generalized and particularized implicatures.

A radically different approach to implicature generation, advocated by Wayne Davis, completely rejects the Gricean view that implicature depends on general principles of communication (Davis 1998). By contrast, Davis argues that particularized implicatures depend on the speaker's intentions, and that generalized implicatures depend on linguistic conventions. On this view, then, the link between an utterance and the implicature it carries may be to a large extent arbitrary, and implicature detection may require specific knowledge of the speaker or relevant linguistic conventions.

The chief alternative to neo-Gricean theories of implicature recovery is *relevance theory* (for example, Carston 2002; Sperber and Wilson 1995). This 'post-Gricean' approach agrees that interpretation involves the application of general communicative principles, but it posits only one of these: that speakers aim to be maximally relevant (in a certain technical sense). Since the literal meaning of a sentence may not be the most relevant one in the context, this often dictates non-literal interpretations, and implicatures are cases of these. This view does not make a sharp distinction between generalized and particularized implicatures, but treats them all as context-dependent, particularized ones.

3.2 *The present thesis*

Grice's approach to implicature (the *Gricean framework*, as I shall call it) is elegant and powerful, and (as we shall see) it promises to establish norms of

implicature of the kind I suggested it would be useful to have. Despite these and many other attractions, however, the framework faces problems. Some important interpretative questions remain unsettled, including questions about the role of speaker intentions in implicature and about the aims of Grice's theory. Moreover, attacks on the Gricean framework have been mounting in recent years. Wayne Davis, in particular, has presented many powerful counterexamples to the view that implicatures can be calculated in the way Grice proposes. And post-Gricean pragmatists have attacked neo-Gricean accounts of implicature recovery, drawing support from a growing body of experimental work on implicature processing. Perhaps, for all its elegance, Grice's approach was too ambitious, and implicature is a messier, more context-dependent, and less rational phenomenon than Grice supposed?

This is, then, a good time at which to reassess the Gricean framework. This thesis attempts such a reassessment. It is structured as follows. Chapter 2 provides a detailed exposition of Grice's theory of implicature. It discusses key distinctions, claims, and applications, and introduces Grice's well-known three-part definition of implicature, according to which what is implicated by an utterance is (roughly) whatever supposition is required to make sense of it as a cooperative contribution to the conversation. In addition, the chapter explores a tension in Grice's views concerning the role of speaker intentions in implicature. I argue that the issue is not resolved in Grice's work and that we should for the moment distinguish two possible versions of Grice's account. Later chapters will return to this issue.

Chapter 3 turns to the detailed assessment of the Gricean framework. Drawing in part on Davis's work, it sets out a number of problems for each of the three clauses of Grice's definition, showing how Gricean theory conflicts with our intuitions about what implicatures various utterances carry. The chapter then goes on to look at a possible response to these problems based on a proposal by Jennifer Saul (Saul 2002a). Saul argues that Grice's notion of implicature is a normative one, parallel to Grice's notion of sentence meaning, and that additional descriptive notions (of utter-implicature and audience-implicature) are needed in order to account for our intuitions about implicature. This reinterpretation, I point out, gives Griceans a line of reply to the problem cases discussed earlier: They can hold that our intuitions in these cases are simply wrong, and that they refer to utter-

implicatures or audience-implicatures rather implicatures proper. This is an attractive option, but the normative reading is not wholly in line with Grice's definition of implicature and some problem cases remain. In response, I go on to propose a revised two-clause version of Grice's definition that is fully in line with the normative reading and which avoids many of the remaining problems. The final section of the chapter then returns to the issue of the role of speaker intentions in implicature. Drawing on Saul's parallel between implicature and sentence meaning, I argue that the tension in Grice's views can be resolved by making a distinction between what a speaker implicates and what their utterance implicates, where the former, but not the latter, depends on the speaker's intentions. The chapter concludes that the reinterpreted and revised version proposed is the most charitable and consistent form of the Gricean framework.

Having identified the most promising version of the Gricean framework, I go on in Chapter 4 to argue that even this version has a serious flaw. As a normative theory, its aim should be to provide a *speaker-independent* notion of implicature. Although implicatures may depend on features of the context of utterance, they should not depend on the intentions, beliefs, and values of the individual speaker. Otherwise, the theory would threaten to collapse into a Humpty Dumpty one. Yet, I shall argue, the Gricean framework does not provide such an account. It holds that implicatures can be calculated from information about utterances and their context, together with general principles of communication. Yet — I shall argue — there is no way to specify the appropriate premises for such calculations without appealing to the speaker's beliefs, intentions, and values. Thus, although a speaker's mental states do not directly determine what is implicated, they indirectly determine it by establishing the background assumptions relative to which implicatures are calculated. The chapter goes on to examine the consequences of this conclusion, arguing that it seriously undermines the Gricean framework and proposing instead an intention-centred account of implicature, which abandons the requirement of calculability and allows a direct role for speaker intentions. I argue that this account need not collapse into a Humpty Dumpty view, since a normative element can be preserved by requiring that an appropriate audience can work out what is being implicated. Moreover, by employing the notion of a meaning being *made available* to an audience, I argue,

we can draw a distinction between what a speaker implicates and what an utterance does, thus allowing for the possibility of unmeant implicatures. The third section of the chapter supplements the case against the revised Gricean framework by examining the notions of utterer-implicature and audience-implicature proposed by Saul and arguing that they cannot play the role required of them. The last section looks briefly at some of the implications of Grice's theory of implicature generation for the process of implicature recovery, arguing that here too the theory has some unattractive consequences.

Chapter 5 turns to implicature recovery and neo-Gricean theories. Such theories hold that hearers derive generalized implicatures by applying interpretative principles similar to those proposed by Grice, but (I shall argue), they can also be seen as offering an account of how generalized implicatures are generated. Thus, if the neo-Gricean approach is sound, then the Gricean framework will be at least partially vindicated. The chapter focuses on Stephen Levinson's influential version of neo-Griceanism (Levinson 2000), comparing and contrasting it with rival approaches, including relevance theory, a cognitive version of convention theory, and a weakened form of neo-Griceanism. Levinson identifies three core interpretative principles from which generalized implicatures can be derived, and the chapter examines each of these in turn. In each case I highlight numerous problem cases, arguing that they indicate that implicature recovery is more context-sensitive than Levinson supposes and that a rival approach may offer a more attractive explanation. This chapter also surveys recent work in experimental pragmatics and shows that its results do not fit well with the predictions of neo-Griceanism. The chapter concludes that the prospects for neo-Griceanism are not bright, although Gricean principles may have a limited role to play in implicature recovery. It does not attempt to adjudicate between alternative non-Gricean theories, however, and suggests that a pluralistic approach to implicature recovery may be called for.

A short final chapter reviews the previous chapters, pulling threads together and drawing some tentative conclusions concerning the various questions raised earlier. The chapter and the thesis concludes with some speculations about the social function of implicature and related ethical issues.

3.3 Methodological remarks

I shall add some brief remarks on methodology. First, I shall assume that propositional attitudes and reasoning involving them can be, and often are, nonconscious. So when I describe a speaker as having certain beliefs or intentions, or a hearer as making certain inferences, I should not be understood to be claiming that the attitudes and processes in question are conscious (though I should not be understood to be claiming that they are *not* conscious either). There are interesting questions about the relative roles of conscious and non-conscious processing in implicature recovery, but for the most part I shall not address them here (for some discussion of the topic within the context of ‘dual-process’ theories of reasoning, see Frankish and Kasmirli 2010).

Second, I shall assume that implicatures are psychologically real for us — that we typically intend them, notice them, and act on them. Thus, as suggested earlier, theories of implicature generation cannot ignore psychological questions about how utterances are interpreted and implicatures recovered. This is not to deny that claims about implicatures may have a normative aspect and that speakers and hearers can make mistakes about what implicature an utterance carries, or even fail to notice an implicature altogether. But I assume that most of us are good at detecting implicatures and that our careful judgements about them are usually sound. Thus, our intuitions about particular cases can provide evidence for our theories of implicature.

Third, and relatedly, in arguing for my position, I shall employ a mixture of philosophical analysis (drawing on our intuitions as data) and psychological theorizing. The former is primarily relevant to questions of implicature generation and the latter to questions of implicature recovery, but since the answers to one set of questions bear on those to the other, the two methodologies overlap. I do not think this mixture of methods is objectionable. It is common nowadays for philosophers of mind and psychology to adopt an eclectic approach, combining conceptual analysis with reflections on experimental results and broad psychological theorizing.

Conclusion

With this introduction, I turn now to exposition of Grice's theory of implicature, which, whatever faults it may or may not have, is a masterly piece of philosophical analysis.

Chapter 2

The Gricean framework

The term ‘implicature’ was coined by Paul Grice, who was one of the first to identify and analyse the phenomenon. Grice proposed a theory of how implicatures are generated, according to which they arise from general principles of rational communication. This account forms the background to all subsequent work on the topic, and in this chapter I shall set it out and discuss a problem concerning its interpretation.

1. Saying and implicating

Grice’s first detailed presentation of his views on implicature was in his 1967 William James lectures, given at Harvard. The ideas appeared in print in his 1975 paper ‘Logic and Conversation’ (Grice 1975), which was later reprinted, together with the rest of the William James lectures, in his 1989 collection *Studies in the Way of Words* (Grice 1989).¹ An earlier 1961 paper ‘The Causal Theory of Perception’ (Grice 1961/1989) also contains some discussion of implicature. Grice expressed his views tentatively, so it is not always easy to pin hard-and-fast commitments on him.² I shall discuss an important interpretative issue later in this chapter (see section 5), but for the most part I shall be concerned with the view of implicature that is proposed in Grice’s writing and that he is commonly taken to endorse, without worrying whether Grice himself would in fact have endorsed it without qualification. I shall refer to this view as *the Gricean framework*.

Grice introduces the notion of implicature by contrasting it with that of *saying*. According to Grice, what a person says by an utterance is ‘closely related to the conventional meaning of the words (the sentence) he has uttered’ (Grice

¹ Where a paper is reprinted in the 1989 collection, I shall cite it 19XX/1989, where 19XX is the original publication date. Where page numbers follow, they refer to the 1989 reprint edition.

² For examples of Grice’s caution in expressing his views about implicature, see Turner 2001. Grice himself humorously notes the suggestion that his remarks employ a new form of speech act, to be represented by an operator called *quessertion*, read as ‘It is perhaps possible that someone might assert that ...’. (Grice 1982, reprinted in Grice 1989, p.297).

1978/1989, p.25), and for a hearer to grasp what is said, he or she will need to know these conventional meanings (and resolve any ambiguities), together with any relevant references (of proper names, indexicals, and so on). For many purposes, what a speaker says by an utterance can be identified with the proposition they express. Strictly speaking, however, in order to say something, in Grice's terms, it is not sufficient to produce an utterance with an appropriate conventional meaning. As Stephen Neale stresses (Neale 1992, p.523), for Grice an utterer counts as *saying* that p only if they *mean* that p — which for Grice involves having a self-referring intention of the sort described in Grice's theory of meaning (Grice 1957, 1968, 1969, all reprinted in Grice 1989). Omitting many complications, a speaker S means that p by utterance x, if S intends to get their hearer H to believe that p (or to believe that S believes that p), and to achieve this in part via the hearer's recognition of this very intention. If an utterer does not have an appropriate intention of this kind (for example, because they are being ironic), then they merely *make as if to say* p, rather than saying p (Grice 1978/1989, p.41, p.53). In short, the speaker must both have meant what they say and have found words with the correct conventional meaning to convey it. The *utterer's meaning* (or *speaker's meaning*) fixed by their communicative intention must coincide with the *sentence meaning* conventionally associated with the words used (Grice 1969/1989, pp.87–8, 1968/1989, pp.120–1). I shall return to the distinction between saying and making as if to say later in this chapter.

On the one hand, then, we have what a speaker literally said (or made as if to say) by an utterance. However, this may not exhaust what is communicated by the utterance. In the exchange discussed in the previous chapter, Mr Bronston said that his company had an account in Zurich, but additionally implied that he himself had not had an account there. Thus, on the other hand, we have what (if anything) the speaker additionally implied by it. Grice notes that various everyday words might be used in this context, including 'imply', 'suggest', 'indicate', and 'mean' (1989, p.86, also 1975/1989, p.24, 1968/1989, p.118). To avoid choosing between these terms, he introduces the semi-technical term 'implicate' and the nouns 'implicature' (the act of implicating) and 'implicatum' (what is implicated).

For Grice, what a speaker says (or makes as if to say) is the vehicle of implicature. From the speaker's point of view, what is said is, in part, the means

to successfully implicating something, and from the hearer's point of view grasping what is said is, in part, the means to recovering the secondary meaning that is implicated. ('In part' since it will not be possible to infer what is implicated *solely* from what is said; other factors too will play a role, such as knowledge of conversational principles, context, background knowledge, and so on.) Speakers implicate one thing by saying another, and hearers recover what is implicated by understanding what is said. Although the implicata of our utterances are often essential to our communicative exchanges, they do not affect the truth conditions of the utterances, which are determined only by what we say. In this respect implicatures differ from the *presuppositions* of an utterance (that is, propositions which must be true in order for the utterance to have a truth value), and from *entailments* of utterances (propositions whose falsity entails the falsity of the utterance). (Of course, implicata have truth values; but their truth values are independent of those of the utterances that generate them.)

In 'Logic and Conversation' Grice introduces two broad categories of implicature, which he calls *conventional* and *conversational*. These are similar in that neither affects the truth-conditions of the utterance that is used to convey them, but in other respects they are very different. Consider the following example, which is Grice's own:

- (1) He is an Englishman; he is, therefore, brave. (Grice 1975/1989, p.25)

This conveys (a) that the person referred to is both an Englishman and brave, and (b) that the person's bravery follows from his being an Englishman. But, Grice claims, what is *said* is simply (a). From a truth-functional perspective, 'therefore' functions simply as a conjunction, and, strictly speaking, an utterance of (1) would not be false if it turned out that the person's bravery was not a consequence of his being an Englishman (Grice 1975/1989, pp.25–6). Hence, (b) is implicated rather than said. However, this implicature is different from ones such as Mr Bronston's, since it is determined by the conventional meaning of the words used. It is part of the meaning of the word 'therefore' that it carries the implicature that the second thing followed from the first. This implicature would be recognized by any competent hearer, no matter what the context, and it cannot be stripped away or

cancelled. This is what Grice calls a *conventional* implicature. Thus, knowledge of the conventional meaning of an uttered sentence (together with knowledge of relevant references) suffices to fix both what is said and what (if anything) is conventionally implicated by the utterance. As Stephen Levinson puts it, commenting on Grice's programme, 'what is *coded* by the linguistic system is the sum of what is *said* (roughly the truth-conditional content) and what is *conventionally implicated*' (Levinson 2000, p.14).

Now contrast (1) with an utterance of:

(2) Some Englishmen are brave.

Taken literally, this says that there exist brave Englishmen, which is compatible with all Englishmen being brave. However, (2) would normally be taken to imply that not *all* Englishmen are brave. Unlike the implicature in (1), however, this implicature is not part of the conventional meaning of the words used, and it could be cancelled — for example, by adding 'In fact, all Englishmen are'.

Or take the following sentences:

(3) You obviously think tenacity pays.

(4) Jones has beautiful handwriting.

Taken out of context, (3) makes a claim about the hearer's attitudes. But if uttered in response to a curious and persistent colleague, it will convey something about the speaker's attitude — namely, that the speaker finds the hearer tiresome and won't cooperate. Similarly, in many contexts (4) would simply express praise for a talent Jones possesses. However, if it were uttered by a professor of philosophy in response to a request for an opinion of a student's academic ability, it would convey the message that Jones is a poor philosopher. (This now famous example was first used by Grice in his 1961, p.130.)

Examples (2) to (4) are cases of what Grice calls *conversational* implicatures. These are nonconventional, pragmatic implicatures, which are not part of the conventional meaning of the words used. Unlike conventional implicatures,

conversational implicatures are at least to some degree context dependent, and they can be cancelled by a subsequent utterance. (I shall say more about cancellability below.) Grice allows that there may be other types of nonconventional implicature, in addition to conversational ones, but says little about them.

My focus in this thesis is on conversational implicatures, which are defeasible and not determined by the conventional meaning of the words used. (I shall, however, consider the suggestion that some of these implicatures are themselves conventional in another sense; see Chapter 5.) When I use the word ‘implicature’ without qualification, it should be understood to refer to conversational implicature.

2. Implicature generation

How is ‘going beyond what is said’ supposed to work in the case of conversational implicature? Grice claims that hearers can arrive at the implicated meaning by a process of inference, guided by the assumption that the speaker is trying to be cooperative. He points out that conversational exchanges are typically cooperative:

Our talk exchanges do not normally consist of a succession of disconnected remarks, and would not be rational if they did. They are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, or at least a mutually accepted direction. (Grice 1975/1989, p.26)

He proposes a ‘rough general principle’ which speakers are expected to observe:

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. (ibid.)

Grice calls this the *Cooperative Principle* (henceforth CP), and he claims that adhering to it involves respecting various maxims, which he assigns to four broad categories: Quantity, Quality, Relation, and Manner (1975/1989, pp.26–7).

Quantity concerns the amount of information provided and includes the submaxims: 'Make your contribution as informative as is required (for the current purposes of the exchange)', and 'Do not make your contribution more informative than is required'. Quality includes the supermaxim: 'Try to make your contribution one that is true' and the submaxims 'Do not say what you believe to be false' and 'Do not say that for which you lack adequate evidence'. Relation comprises the general maxim 'Be relevant', the application of which may be very complex. Finally, Manner includes the supermaxim 'Be perspicuous' and the submaxims 'Avoid obscurity of expression', 'Avoid ambiguity', 'Be brief (avoid unnecessary prolixity)' and 'Be orderly'.

Grice suggests that these conversational maxims are instances of more general maxims which govern other kinds of purposive behaviour, such as helping to fix a car or bake a cake. In such cases, too, the parties involved are expected to be cooperative and to make contributions that are appropriate in quality, quantity, relevance, and manner. Grice also speculates that it is not a contingent fact that we observe the CP and its maxims, but that anyone engaging in communication is rationally required to observe them (Grice 1975/1989, pp.29–30).

Grice allows, of course, that on occasions a speaker may fail to follow these maxims. He mentions four cases (1975/1989, p.30). First, the speaker may covertly *violate* a maxim, usually in order to mislead their hearer. Second, a speaker may explicitly *opt out* of a maxim, for example by indicating that they are unwilling to tell all they know. Third, a speaker may find that two maxims *clash*, forcing them to choose between them. For example, if a speaker has information that is important but of doubtful reliability then they will not be able to simultaneously respect the maxims of Quantity and Quality. Finally, a speaker may openly *flout* a maxim in a way that is obvious to their hearer. Cases of the last type, Grice proposes, are the ones that typically generate conversational implicatures. By openly flouting the CP in what they *say*, yet without ceasing to observe the CP by opting out, a speaker signals to their hearer that they wish to convey some further message that is consistent with the CP. Thus, even though the speaker seems to be flouting the maxims, they are in fact following them at another level. As Grice puts

it, they are *exploiting* the conversational maxims for the purposes of generating a conversational implicature (1975/1989, p.36).³

Take the following example:

(5) Ada: Do you like my new outfit?

Bea: You shouldn't be allowed to buy clothes.

Bea's utterance appears to violate the CP — specifically the maxims of Quantity and Relation (and probably Quality too). Bea must know that more, and more relevant, information, is required, and she ought to be able to provide it, since the question was about her personal opinion. But Ada has no reason to think that Bea has opted out of the conversational exchange; Bea is her friend and knows that choices of outfit are important to her. Ada can reconcile these facts only by supposing that Bea is seeking to convey something else — that Ada's outfit is horrible — which is informative and relevant but which for some reason Bea does not wish to say explicitly. Ada assumes that Bea thought Ada could work this out, and concludes that Bea is implicating that her outfit is horrible.⁴

In the case just described a maxim is actually flouted (exploited), but actual flouting is not necessary in order to generate a conversational implicature, in

³ Grice also allows that implicatures can also be generated by other maxims, such as aesthetic, social, and moral ones. Such implicatures form the class of *nonconversational*, nonconventional implicatures mentioned earlier. However, he holds that the conversational maxims and implicatures are the most important ones as far human communication is concerned (Grice 1975/1989, p.28).

⁴ It might be objected that, taken literally, some of the maxims do not allow for indirect adherence of this kind, at the level of what is implicated rather than what is said. For example, the maxim of Quality tells us not to *say* something that we believe to be untrue, so a speaker cannot follow it by saying something they believe to be false, even if they thereby *implicate* something they believe to be true. Neale responds on Grice's behalf that this probably reflects a looseness of phrasing, and that adherence to the maxim at the level of what is implicated should be allowed to compensate for a violation of it at the level of what is said. Thus 'blatantly violating a maxim at the level of what is said but adhering to it at the level of what is implicated would not necessarily involve a violation of the Cooperative Principle' (Neale 1992, p.526).

Grice's view. In some cases, a conversational implicature is generated in order to *avoid* flouting. To borrow an example from Grice, suppose someone asks me where they can get petrol, and I reply 'There is a garage round the corner' (1975/1989, p.32). Here, what I say *would* flout the maxim of Relation ('Be relevant') if I didn't believe that the garage round the corner has petrol and is currently open. Thus, my hearer must assume that I believe those things in order to preserve the assumption that I am following the CP, and I thereby implicate those propositions. However, what I said didn't actually flout the maxim of Relation, since it is relevant that the station is round the corner. This contrasts with the case where Ada says that Bea should not be allowed to buy clothes, which is neither true nor relevant. In the latter case the implicature serves to *repair* a flouting by supplying a relevant meaning where one was lacking, whereas in the petrol station case it serves to *prevent* a flouting, by supplying additional information which makes the literal meaning relevant. Both cases, however, fit the same broad pattern, in that the implicated meaning must be presupposed in order to maintain the assumption that the speaker is being cooperative, and this is the heart of Grice's account.

More formally, Grice offers the following three-part definition of the conditions that must be satisfied for a conversational implicature to occur:

A man who, by (in, when) saying (or making as if to say) that *p* has implicated that *q*, may be said to have conversationally implicated that *q*, provided that (1) he is to be presumed to be observing the conversational maxims, or at least the Cooperative Principle; (2) the supposition that he is aware that, or thinks that, *q* is required in order to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with this presumption; and (3) the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required. (Grice 1975/1989, p.30–1)

Simplifying slightly, for a speaker to implicate *q* by saying *p*, it must be the case

that (1) the speaker is presumed to be being cooperative;⁵ (2) this presumption requires the supposition that the speaker thinks that q, and (3) the speaker thinks that their hearer can work this out. I shall follow Wayne Davis in referring to these three conditions as, respectively the *cooperative presumption*, *determinacy* (since it is the condition that q specifically is required), and *mutual knowledge* (Davis 1998, p.13). Grice's definition does not explicitly state *who* does the presuming in (1), but the natural interpretation is that it is the hearer.

Note that Grice says that the speaker need only *make as if to say* that p. Though they utter a sentence that conventionally means p, they themselves need not mean that p — the sentence meaning does not need to be backed by a speaker meaning (which for Grice would be constituted by an intention to get their audience to believe that p by recognizing this intention). Thus, for example, when Bea utters the sentence 'You shouldn't be allowed to buy clothes', she does not really mean that Ada should not be allowed to buy clothes, and so (in Grice's terminology) does not *say* it, but merely makes as if to say it. What Bea makes as if to say is merely the means to implicating that Ada's outfit is horrible, which is Bea's real communicative aim. The same will go in many cases where a speaker actually flouts a maxim as a way of generating an implicature, and it will almost always be the case where they flout the maxim of Quality.

3. Calculability

Grice claims that, in the process of working out what a speaker is conversationally implicating (that is, determining what is required to maintain the cooperative presumption), a speaker will draw on the following pieces of information:

- (1) the conventional meaning of the words used, together with the identity of any references that may be involved;
- (2) the Cooperative Principle and its maxims;
- (3) the context, linguistic or otherwise, of the utterance;
- (4) other items of background knowledge; and
- (5) the fact (or supposed fact)

⁵ This is how Grice's first condition is often stated (and even misquoted), though he actually says that the speaker is 'to be presumed' to be being cooperative. I will say more about this issue in the next chapter.

that all relevant items falling under the previous headings are available to both participants and both participants know or assume this to be the case (Grice 1975/1989, p.31).

And he suggests that the process will involve a calculation of the following general kind:

[The speaker] has said that p ; there is no reason to suppose that he is not observing the maxims, or at least the Cooperative Principle; he could not be doing this unless he thought that q ; he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that q is required; he has done nothing to stop me thinking that q ; he intends me to think, or is at least willing to allow me to think, that q ; and so he has implicated that q . (Grice 1975/1989, p.31)

An argument which derives an implicated meaning in this way is sometimes referred to as a *Gricean calculation*.

It is not clear what kind of argument a Gricean calculation is supposed to be. The way Grice sets out the calculation and the fact that he uses the word ‘required’ in clause (2) of the preceding definition of implicature (‘the supposition that he is aware that, or thinks that, q is required in order to make his saying or making as if to say p ... consistent with this presumption’) suggest that the argument is meant to be deductive — that the cooperative presumption and the other items of information mentioned entail that the speaker is implicating that q . But, as Wilson and Sperber note, the third stage in the Gricean calculation (‘he could not be doing this unless he thought that q ’) does not follow from the claims before it, and Grice does not explain how it is derived (Wilson and Sperber 1991, p.378). Moreover, it is doubtful that we could construct a deductive argument for it — at least without giving a complete list of the speaker’s background beliefs. (Given suitably strange background beliefs, a speaker could regard any utterance as cooperative under its literal meaning.) Accordingly, many writers hold that the process of deriving an implicature is not one of deductive, demonstrative inference, but of abductive inference — inference to the best explanation of the data (Bach and Harnish 1979,

pp.92–3; Brown and Yule 1983, p.34; Leech 1983, pp.30–1; Levinson 1983, pp.115–6). On this view the conclusion that the speaker has implicated *q* is not logically required by what they have said (or made as if to say), given the context, the assumption that they are following the CP, and so on; rather, it is the most likely hypothesis given that information. It is not clear whether Grice would accept this, but for interpretative purposes I shall adopt a pluralistic position on which a Gricean calculation can be either deductive or non-deductive (or include both deductive and non-deductive elements), and I shall treat ‘required’ in this context as meaning ‘required or highly probable’. (I shall say more about implicature recovery in Chapter 5.)

Grice adds that the hearer need not *actually* go through a process of inference of this type in order to see that a particular conversational implicature is present. They may just ‘grasp intuitively’ that it is (Grice 1975/1989, p.31). However, he insists that the implicature must be ‘capable of being worked out’ — the intuition must be replaceable by an argument. Otherwise, it will count as a conventional implicature, not a conversational one (1975/1989, p.31; 1978/1989, p.43).⁶ The thought seems to be that if an implicature could not be calculated in the way described, then it could only arise from the conventional meaning of the words used (assuming, that is, that it is not some other kind of nonconventional implicature, derivable from maxims of a different sort).

The claim that the intuition must be replaceable by an argument accords with Grice’s wider views about reasoning, as set out in his posthumously published *Aspects of Reason* (Grice 2001). Here Grice distinguishes a laborious ‘hard’ way of reasoning, in which every step is spelled out, and an easier ‘quick’ way, which leaves gaps. The quick form still counts as reasoning, Grice argues, provided that

⁶ Compare the following passage from a 1981 paper, in which Grice makes the same point:

[T]he final test for the presence of a conversational implicature had to be, as far as I could see, a derivation of it. One has to produce an account of how it could have arisen and why it is there. And I am very much opposed to any kind of sloppy use of this philosophical tool, in which one does not fulfill this condition. (Grice 1981, p.187, quoted in Cummings 2009, p.137)

the agent *intends* that each step could be filled out in such a way as to create a valid argument and has the ability to do this filling out:

we could say (for example) that x reasons (informally) from A to B just in case x thinks that A and intends that, in thinking B, he should be thinking something which would be the conclusion of a formally valid argument the premisses of which are a supplementation of A. ... The possibility of making a good inferential step (there being one to be made), together with such items as a particular inferer's reputation for inferential ability, may determine whether on a particular occasion we suppose a particular transition to be inferential (and so to be a case of reasoning) or not. (Grice 2001, p.16)

As Richard Warner notes in his introduction to the volume, this offers an attractive view of the nature of the reasoning Grice attributes to speakers and hearers in his theory of speaker meaning (reasoning about the speaker intentions, the hearer's recognition of these intentions, and so on). People do not go through this reasoning in the hard way, but they can be regarded as doing so in the quick way, provided they intend their interpretations of each other's utterances to be rational and have the ability to produce reasoning of the relevant kind. The arguments Grice sets out can be thought of as the ones they *would* construct if they were to fill in the steps (Warner in Grice 2001, pp.xxxii-v). (It might be objected that most people do not have the ability to construct these arguments, and that it took a highly trained philosopher of Grice's talents to produce them. Grice might reply that it is sufficient that people would recognize and endorse the arguments when presented with them.) We can take a similar view of the reasoning involved in deriving implicatures. Hearers may not actually go through a Gricean calculation when they interpret an utterance as carrying an implicature, but they intend their interpretation to be a rational one and the calculation sets out the sort of argument they would produce if they were to rationalize their interpretation.

To sum up, according to the Gricean framework, if an utterance U carries a conversational implicature q, then it must be possible to construct an argument (deductive or nondeductive) that derives the claim that the speaker is implicating

q from the conventional meaning of U, the CP, the context, background knowledge, and the fact that all this information is openly available. Following Davis, I shall refer to this claim as the *Calculability Assumption* (Davis 1998, p.14). The assumption serves as a supplement to the definition of conversational implicature, expanding on clause (2). To say that a proposition is required in order to make a speaker's utterance consistent with the cooperative presumption is to say that it is the one that would be uniquely identified by a Gricean calculation.

4. Particularized and generalized implicatures

Grice introduces a further distinction among conversational implicatures, between two sub-categories: *particularized* and *generalized*. Particularized implicatures are one-offs — cases where a person implicates a particular message by saying that p, but there is, as Grice puts it, 'no room for the idea that an implicature of this sort is normally carried by saying that p' (1975/1989, p.37). For example, take:

(6) It's chilly here.

Uttered by a person stepping off an aeroplane into the suffocating heat of a tropical country, this might carry the implicature 'It is extremely hot here'. Given the conditions, the speaker cannot really believe that it is chilly, so their utterance would flout the maxim of Quality if taken literally. To preserve the assumption that the speaker is observing the CP, the hearer must suppose that they are speaking ironically and expressing the thought that it is the very opposite of chilly. However, the reasoning depends on particular facts about the context of the utterance (that it is suffocatingly hot) and in other contexts the same sentence would generate a different implicature, or none at all. For example, if uttered in response to the question 'Do you want to go home?' it might generate the implicature that the speaker does want to go home. The implicatures generated by uttering this sentence are particularized, context-dependent ones.

Generalized conversational implicatures, on the other hand, are not context-dependent in this way; the words used 'would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature' (Grice 1975/1989 p.37). For example, sentences of the form 'Some F are G' will normally

generate an implicature of the form ‘Not all F are G’. If the speaker believed that all F were G, they would have flouted the maxim of Quantity by not saying that they all were. Given the presumption that the speaker is observing the CP, the hearer therefore infers that the speaker does not believe that all F are G. Here the inference does not depend on facts about the particular context of utterance and will go through by default in all contexts.

Levinson helpfully provides a more explicit formulation of the difference between the two types of implicature:

- a. An implicature *i* from utterance *U* is *particularized* iff *U* implicates *i* only in virtue of specific contextual assumptions that would not invariably or even normally obtain
- b. An implicature *i* is *generalized* iff *U* implicates *i* *unless* there are unusual specific contextual assumptions that defeat it.

(Levinson 2000, p.16)

Grice acknowledges that it may not be easy to distinguish generalized conversational implicatures from conventional implicatures (Grice 1975/1989, p.37), but as a noncontroversial example of the former he offers expressions with the form ‘an X’, which, he notes, normally generate the implicature that the X in question ‘does not belong to, or is not otherwise closely connected with, some identifiable person’. For example, an utterance of ‘John is meeting a woman this evening’ would normally carry the implicature that the woman in question was not John’s wife, relative, or close friend. However, this is not a conventional implicature of the phrase ‘an X’, since there are some contexts in which the phrase does not generate the implicature, and ones in which it generates the opposite one. (Grice cites ‘I broke a finger yesterday’, which implies that the finger *does* belong to the speaker.) It is better, Grice argues, to see this as a case of generalized conversational implicature, generated by the mechanisms described above:

When someone, by using the form of expression *an X*, implicates that the X does not belong to or is not otherwise closely connected with some identifiable person, the implicature is present because the speaker has

failed to be specific in a way in which he might have been expected to be specific, with the consequence that it is likely to be assumed that he is not in position to be specific. (Grice 1975/1989, p.38)

That is, by not being more specific about the identity of the X referred to, the speaker has seemingly violated the maxim of Quantity, and the assumption that he or she is being cooperative can be preserved only on the supposition that he was not in a position to be specific — that is, did not think that the X belonged to or was closely connected with some identifiable person. (I shall discuss this example further in Chapter 5.)

Grice identifies some other properties that are distinctive of conversational implicatures as opposed to conventional ones. The two most important of these (in addition to calculability and nonconventionality, discussed above) are *cancellability* and *nondetachability* (Grice 1975/1989, pp.39–40, 1978/1989, pp.43–4). Grice sometimes refers to these as *tests* for the presence of a conversational implicature (1981/1989, pp.270–1), though he says they are more like *prima facie* indications than knock-down tests (1978/1989, p.43). I shall consider them in turn.

First, cancellability. Because conversational implicatures depend on the assumption that the speaker is observing the CP, and because a speaker can explicitly opt out of doing this, it follows that a conversational implicature can be cancelled. This can be done either explicitly, by adding a further statement which indicates that one is opting out, or implicitly, by the context. For example, using ‘somebody’, rather than a more specific expression, normally generates the implicature that the speaker cannot identify the person referred to (following a similar line of reasoning to that for ‘an X’), but this implicature can be cancelled, as the following examples illustrate (the examples are my own):

- (7) *Explicit cancellation.* I heard somebody robbed you yesterday. In fact, it was your brother.

- (8) *Implicit cancellation.* Somebody forgot to turn the headlights off.
(Uttered in a context and with a tone of voice that makes it obvious that the speaker is referring to the hearer.)

Here the second sentence in (7) and the context in (8) make it clear that the speaker's use of 'somebody' was genuinely uncooperative (the speaker could have used a more specific term but didn't, presumably for stylistic reasons), and the usual implicature is cancelled.

The second indicative feature of implicature is nondetachability. Since conversational implicatures are generated by general inferential principles applied to the conventional meanings of the sentences uttered (together with context and background information), utterances with equivalent conventional meanings will generate the same implicature in the same context. Thus in most cases 'it will not be possible to find another way of saying the same thing, which simply lacks the implicature in question' (Grice 1975/1989, p.39). That is, conversational implicatures are typically not *detachable* from the content of the utterance. An exception is where the implicature is generated by flouting the maxim of Manner, in which case *how* a content is expressed will be crucial. Grice gives the following example: 'Miss X produced a series of sounds that corresponded closely to the score of "Home Sweet Home"' (1975/1989, p.37). Because of the roundabout way it is expressed, avoiding use of the word 'sing', this implicates that Miss X's singing was very bad. Saying the same thing in a simpler way would not generate this implicature.

A final feature of conversational implicature mentioned by Grice is indeterminacy. What is implicated by an utterance is, according to Grice, what must be supposed in order to preserve the assumption that the speaker is being cooperative. But in any given case there may be many different suppositions that could play this role. For example, Bea's utterance of 'You shouldn't be allowed to buy clothes' might be taken to implicate that the outfit Ada has chosen is dull, or that it is extravagant, or that it is too young for Ada, or that it is too old for Ada, or that it possesses some other negative feature. More generally, it might implicate that Ada's taste in clothes is poor, or that Ada should listen to Bea's advice, or some other, related claim. Grice accepts this point. In such cases, he claims, the

implicatum will be the disjunction of the various possible suppositions, and if this disjunction is open-ended, the implicatum will be simply indeterminate (1975/1989, p.40).

5. Implicature and speaker meaning

We saw earlier that a speaker S counts as saying that p, by uttering an appropriate sentence, only if they also *mean* that p — that is, only if they have an appropriate communicative intention. Otherwise, they merely make as if to say that p. This naturally raises the question of whether implicatures, too, must be backed by speaker meanings. For a speaker to implicate that q, must they also mean that q, on Grice's view? (And if they go through the motions of implicating q without actually meaning q, do they just *make as if to implicate* q?) Although this is a basic and important question, Grice himself does not address it directly, and says nothing about speaker meaning in his account of implicature in 'Logic and conversation'. In fact, there is a dispute over the correct interpretation of Grice here (Davis 2007; Neale 1992; Saul 2001, 2002a), and a case can be made for both positive and negative answers to the question, as I shall now explain.

There are several reasons for holding that Grice thought that speakers must mean what they implicate. When a person implicates something, it is natural to say that the implicated content is what they really *meant*, in contrast with what they literally said. Indeed, 'mean' is one of the everyday words (along with 'imply', 'suggest', and 'indicate') for which Grice introduces 'implicate' as a technical replacement (1975/1989, p.24, 1989, p.86). It is notable, too, that in Grice's general schema for working out an implicature, quoted earlier, the hearer's train of thought concludes with 'he [the speaker] intends me to think, or is at least willing to allow me to think, that q; and so he has implicated that q' (Grice 1975/1989, p.31) — which suggests that the speaker's intentions (or at least their wishes) are relevant in deciding whether or not an implicature is present (for more discussion of this passage, see Davis 2007, p.1659). Moreover, in 'The Causal Theory of Perception' (Grice 1961) Grice specifically says that an implicature must be backed by a communicative intention. Discussing an imaginary case where he has reported on the abilities of a student, Jones, by saying simply 'Jones has beautiful handwriting and his English is grammatical', Grice comments:

I (the speaker) could certainly be said to have implied that Jones is hopeless (provided that this is what I intended to get across) ... (Grice 1961, p.130)

(It is true that the section in which this passage appears was not included when this paper was reprinted in *Studies in the Way of Words*, but there is no reason to think this was because Grice had changed his mind. In the reprint Grice says that the section was omitted because the material it contained was ‘substantially the same’ as that in ‘Logic and Conversation’ (Grice 1989, p.229)).

Another reason for thinking that conversational implicatures must be meant comes from Grice’s views about the role of the presumption of cooperation in generating them. Hearers posit implicatures in order to preserve the assumption that speakers are being cooperative. But if speakers need not mean what they implicate, how does what they implicate support the presumption that they are being cooperative? How can behaviour that is unmeant and unintended be genuinely cooperative?

Finally, on the basis of a reading of Grice’s other work, Stephen Neale argues that Grice held that what a speaker implicates is a part of what they mean overall (Neale 1992, pp.523–4). Neale points out that in ‘Utterer’s Meaning, Sentence-Meaning and Word-Meaning’, Grice introduces the notion of what a speaker *conventionally means*, which breaks down into what they say and what (if anything) they conventionally implicate, and which is part of what they mean overall (Grice 1968/1989, p.121). Neale suggests that Grice would also have recognized the parallel notion of what a speaker *nonconventionally means*, which breaks down into what (if anything) they conversationally implicate and what (if anything) they nonconversationally nonconventionally implicate. (Nonconversational nonconventional implicatures are ones that are dependent on nonconversational maxims, such as aesthetic or moral ones.) Neale concludes that it is reasonable to think that Grice would have accepted the breakdown of what an utterer *U* means illustrated in Figure 1.

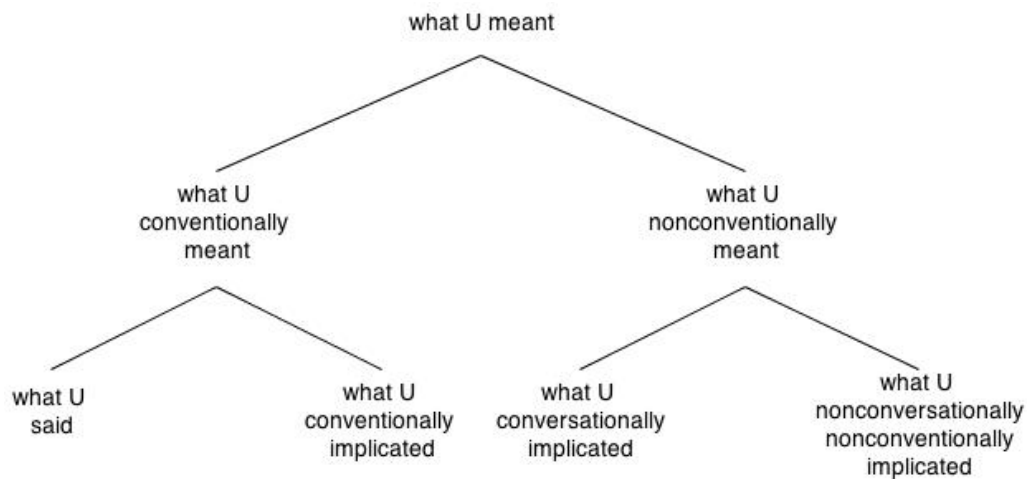


Figure 1: The breakdown of what a speaker means, according to Stephen Neale’s interpretation of Grice. (Adapted from Neale 1992, p.523.)

On Neale’s view, then, Grice’s theory of conversational implicature is part of his wider project of explaining the conventional and nonconventional components of what a speaker means, and the reason there is no mention of speaker meaning in ‘Logic and Conversation’ is simply that Grice’s theory of speaker meaning is assumed as part of the background.

There is a strong case, then, for thinking that Grice held that implicatures must be backed by speaker meanings. However, there are also objections to this view. The main objection is that there is no mention of speaker meaning in Grice’s three-part definition of conversational implicature in ‘Logic and Conversation’ (discussed above), which, as Jennifer Saul stresses, focuses on the attitudes of the *hearer*, not the speaker (Saul 2001, pp.632–3, 2002a, p.241). According to the definition, the hearer must presume that the speaker is observing the conversational maxims (clause 1), and be able to work out that the supposition that the speaker has a certain belief is required in order for his utterance to be consistent with that presumption (clause 2). The only attitude of the speaker that is mentioned is the belief that the hearer can work out that the supposition mentioned in clause 2 is required (Grice 1975/1989, p.30–1). There is no mention of the speaker’s *intentions* at all.

It may be replied that Grice does not offer this definition as a complete account of the conditions necessary for conversational implicature, but only of those conditions necessary for an implicature to count as a *conversational* implicature,

as opposed to a conventional or nonconversational nonconventional one. (The definition begins ‘A man who, by (in, when) saying (or making as if to say) that *p* has implicated that *q*, may be said to have conversationally implicated that *q*, provided that...’; *ibid*). On this interpretation, other conditions will have to be met in order for an implicature to be present at all, and it may be that Grice took these to include the presence of a corresponding speaker meaning. This reading of Grice is adopted by Wayne Davis (Davis 2007, p.1660). It is true, however (as Davis acknowledges), that Grice elsewhere speaks of conversational implicatures being *generated by*, or *present because of*, or *arising from* the conditions mentioned in the definition, which suggests that those conditions are sufficient for the presence of a conversational implicature, rather than merely necessary for an implicature to count as conversational (for example, Grice 1975/1989, p.28, p.38, 1989, p.370). This reading, on which the conditions are sufficient, is often adopted in the subsequent literature (for example, Harnish 1991, p.330; Levinson 1983, p.100, p.103; Sadock 1991, p.366).

There is another reason for doubting that Grice would have accepted that implicatures must be backed by speaker intentions. Suppose that speakers must mean what they implicate — that is, they must have an appropriate communicative intention. But then why couldn’t implicatures be recovered simply by recognizing these intentions, rather than going through Gricean calculations? Indeed, on Grice’s view, for a speaker to mean *q*, they must intend to get their hearer to believe that *q* in part *by recognizing this very intention*. So, it seems, if implicatures must be meant, then a speaker who implicates *q* intends their hearer to come to believe *q* at least in part by recognizing their intention to communicate *q*, and not by going through a Gricean calculation. Thus, the speaker’s conception of how the implicated message is to be recovered seems to be different from the one suggested by Grice. Finally, if a speaker intends their hearer to recover the implicated message by recognizing their intention to communicate it, why is it necessary for them also to believe that the hearer can work out that the implicated message is required to uphold the presumption that they are being cooperative (clause 3)?

It may be replied that Grice claims only that implicatures must be calculable, not that they must actually be calculated. Perhaps they can also be recovered by recognizing the speaker’s intentions straight off, without any actual calculation

(though Grice himself does not mention this possibility). Moreover, the belief that the implicated message will be recovered by detecting the intention to communicate it is not incompatible with the belief that the message can be recovered by a Gricean calculation. The Gricean calculation could be the *means* to recognizing the intention; the speaker might (a) intend to get the hearer to believe that q in part by recognizing this very intention and (b) believe that the hearer can recognize this intention by going through a Gricean calculation. Indeed, in his informal description of the process of working out an implicature Grice makes it clear that he expects the hearer to move from the supposition that the speaker thinks that q to the belief that the speaker intends him (the hearer) to think that q :

he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that q is required; he has done nothing to stop me thinking that q ; he intends me to think, or is at least willing to allow me to think, that q ; and so he has implicated that q . (Grice 1975/1989, p.31)

We might compare the role of sentence meaning in saying. Understanding what someone says involves recognizing the speaker's communicative intentions, but it also involves recognizing the meaning of the sentence the speaker utters, and recognizing the latter is typically the means to recognizing the former.

This is a possible position, but, as it stands, it is still rather puzzling. Why should it be necessary for the speaker to believe that their communicative intention can be recognized by a Gricean calculation if it is not necessary for them to believe that it can be recognized *only* by that means (as Grice must allow, given that he denies that the hearer must actually go through the calculation process)? It is plausible that the speaker must believe that the hearer has *some* means of recognizing their intention, but it is unclear why they must believe that the hearer can do so specifically by a Gricean calculation.

What should we conclude from this discussion? There are cases for both positive and negative answers to the question about Grice's view of the role of speaker meaning in implicature, and I do not think we are justified in attributing to him a settled view on the matter. (Here, perhaps, we should take notice of what

he says about the tentative nature of his proposals.) However, lack of an answer constitutes a serious gap in the Gricean framework, and I shall return to the question in the next chapter, where I shall suggest a way of reconciling this tension in Grice's theory of implicature.

6. Applications

Grice's theory is primarily an account of what we might call *commonsense* implicatures — cases where it is intuitively obvious that an utterance conveys something beyond its literal meaning, as in the examples we have considered. However, the theory also has applications to more technical issues in philosophy of language and linguistics. In particular, it can offer an economical account of the conventional meanings of particular words. By making a distinction between what is said and what is further implicated, it is possible to hold that a word has a single conventional meaning while at the same time explaining how it typically conveys a further meaning. Grice uses 'or' as an example (1978/1989, pp.44–7). Sometimes 'or' is used in a 'weak' way equivalent to logical disjunction. In this sense, to say that p or q is simply to rule out the claim that both p and q are false, and a person could legitimately assert 'p or q' because they knew that p was true or that q was true or that both were true. Thus if one knows that p, then one can assert that p or q for any q at all. However, we typically use 'or' in a stronger sense, to indicate that we have a reason for thinking that p or q *other than* the fact that we think one or both of p and q are true, such as evidence that p and q are the only possible alternatives. In this sense, a person could not legitimately assert 'p or q' just because they knew that p was true or that q was true or that both were, and they could legitimately assert it without knowing any of those things. Faced with these different uses, we could say that 'or' is ambiguous, with two different conventional meanings, but Grice points out that we could instead explain the stronger meaning as a generalized implicature. If a speaker knows that p (or that q, or that p and q), then in most communicative contexts it will be more informative to assert it, rather than asserting that p or q. Thus if a speaker says that p or q, the hearer can uphold the presumption that they are observing the CP (and thus the maxim of Quantity) only by supposing that they do not believe that p (or that q, or both), and thus that they have some reason for asserting the disjunction other than the fact that they

believe one or both of the disjuncts to be true — which corresponds to the stronger sense of ‘or’. Thus in normal circumstances this meaning is generated automatically by conversational implicature.

Grice suggests that we should prefer such explanations to ones that posit further conventional meanings, and he proposes a principle he calls ‘Modified Occam’s Razor’ (and which has subsequently become known as ‘Grice’s Razor’): *Senses are not to be multiplied beyond necessity* (1978/1989, p.47). That is, we should not treat a word as having multiple conventional meanings unless there is no other way of explaining the different ways in which it is used. In particular, if we can explain one of a word’s meanings as due to a generalized conversational implicature arising from the word’s conventional meaning, then we should prefer that explanation to treating the meaning as a second conventional meaning. As François Recanati notes, because the implicature explanation derives from general assumptions and principles that are independently motivated, it is more economical than positing an extra sense, which would be an ad hoc move (Recanati 1989, p.296). Applying Grice’s Razor in a particular case involves showing how the mechanisms of implicature could generate the secondary meaning in question, and also applying the various ‘tests’ for implicature, in particular, non-detachability and cancellability. In Stephen Levinson’s words, this approach

allows one to claim that natural language expressions do tend to have simple, stable and unitary senses (in many cases anyway), but that this stable semantic core often has an unstable, context-specific pragmatic overlay — namely a set of implicatures. (Levinson 1983, p.99)

Grice argues that this approach can be applied to deal with objections to some philosophical theories that use terms in a way that seems to clash with everyday usage. For example (a case that played a role in prompting Grice to develop his theory of implicature), some theories of perception characterize the sort of experience one has when seeing (say) a red object as the experience of *seeming to see* something red, and it may be objected that this conflicts with ordinary usage. One can have the experience in question when one is *sure* that one is seeing something red, but we would not normally say that we *seem* to see something red

unless we had some doubt about whether we really were seeing something red. In response, Grice argues that we can treat the indication of doubt, not as part of the conventional meaning of 'seems to see', but as a generalized conversational implicature, like the strong sense of 'or', which is generated by apparent violation of the maxim of Quantity. If one is sure one is seeing something red, one would normally say that one is seeing something red, not that that one seems to see something red. Thus in saying the latter, one implicates that one is not sure that one is seeing something red (Grice 1961/1989, ch.15). In this way, the objection is removed. Since implicatures do not affect the truth conditions of the utterances to which they attach, the existence of this implicature does not affect the truth of the statement that one seems to see something red, or of the theory to which it belongs.

Grice suggests there are many other contexts in which this approach might be applied to reconcile philosophical theories with the existence of layers or shades of meaning not accounted for by the theory. As examples he mentions claims involving the terms 'see', 'know', 'cause', 'responsible', 'actual' (1961/1989, p.237). He also sketches an application to the word 'true' (1978/1989, p.55–7) and considers (without fully endorsing) the suggestion that the existence claims implicit in utterances such as 'The present king of France is bald' are conversational implicatures, rather than presuppositions (1981/1989, ch.17). Later theorists have followed these hints, applying the Gricean framework (or an extended, revised version of it) to various other problems in philosophy of language (for example, Neale 1990; Recanati 1993; Salmon 1989).

These applications of the Gricean framework differ from the earlier uses of the framework to explain commonsense cases of conversational implicature, such as (5) and (6) above. First, the implicatures in the application cases are all generalized ones and arise from the use of particular words or concepts. Second, as Recanati notes, the applications extend the *scope* of the phenomenon of implicature (Recanati 1989, p.327). In the commonsense cases, it is intuitively obvious that something is implied beyond what is actually said, whereas in the application cases this is not so, and argument is required to establish that an implicature is present. The fact that the Gricean framework can be extended and applied in this way suggests that it is a fruitful research programme, and thus offers further support for

it. However, the applications depend on the soundness of Grice's basic account of how conversational implicatures are generated, in both particularized and generalized cases, and I shall focus primarily on that.

Conclusion

This completes my introduction of the Gricean framework. The framework has been, and continues to be, hugely influential in theorizing about the nature of implicature and (as we shall see in Chapter 5) about how implicatures are recovered. There are attractions to the idea that implicatures can be derived from general principles of communication, as opposed to being, on the one hand, one-off psychological interpretations or, on the other hand, conventions of language or of language use. The Gricean framework suggests that implicature is a rational phenomenon, which can be universally understood, and it thus holds out hope for an ethics of implicature based on general principles.

The framework faces many problems, however. In particular, there are some basic issues with Grice's definition of conversational implicature and the supplementary Calculability Assumption. The next chapter will look at these.

Chapter 3

Problems, reinterpretation, and revision

This chapter will examine some fundamental issues concerning Grice's account of implicature generation. I shall begin by setting out some problems for Grice's definition of conversational implicature and the Calculability Assumption, and arguing that they are serious ones. I shall then consider a proposal by Jennifer Saul, who argues that Grice's notion of conversational implicature is a normative one, and that additional notions are needed in order to capture all the psychological aspects of implicature. I shall argue that this is an attractive reading of Grice, and that the resulting enriched Gricean framework avoids many of the problems discussed. I shall then go on to propose some further modifications to the Gricean definition in order to bring it still more closely in line with the normative reading and to avoid some remaining problems. The aim will be to set out the most plausible version of the Gricean framework, revised as necessary.

1. Problems for Grice's definition

This part of the chapter will consider some problems arising from Grice's account of the way implicatures are generated. Recall Grice's three-part definition of conversational implicature (with Davis's terminology added in brackets):

A man who, by (in, when) saying (or making as if to say) that *p* has implicated that *q*, may be said to have conversationally implicated that *q*, provided that (1) he is to be presumed to be observing the conversational maxims, or at least the Cooperative Principle ['the cooperative presumption']; (2) the supposition that he is aware that, or thinks that, *q* is required in order to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with this presumption ['determinacy']; and (3) the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required ['mutual knowledge']. (Grice 1975/1989, pp.30–1)

This definition is supplemented with the Calculability Assumption, which says that *q* can, in principle, be identified by a Gricean calculation, as described in the previous chapter. I shall treat this assumption as a supplement to the determinacy clause.

As we saw, there is debate about whether these conditions are sufficient for the existence of a conversational implicature (perhaps speaker intentions are needed as well), but it is clear that Grice thinks they are necessary, and this in itself creates problems. (Again, in what follows I shall often drop the qualification ‘conversational’; unless otherwise indicated, ‘implicature’ always means ‘conversational implicature’.)

1.1 Problems with the cooperative presumption.

The first of the three conditions for the presence of an implicature is the cooperative presumption: the speaker is presumed to be being cooperative (observing the CP).

Before going on, I want to mention an interpretative issue. The cooperative presumption is standardly quoted or paraphrased in the way I have just done, as the descriptive claim that the speaker *is* presumed to be being cooperative (following the maxims or at least the CP).¹ However, Grice actually says that the speaker is *to be* presumed to be being cooperative (see the quotation above). This phrasing is not easy to interpret, but it seems to indicate a normative claim: the hearer *ought* to presume the speaker to be being cooperative. The standard version, by contrast, makes a descriptive claim: the hearer *does* presume the speaker to be being cooperative. Whether this interpretation is right and exactly how much importance should be placed on the wording is not clear (for some discussion, see Davis 2007; Green 2002; Saul 2010). Since my aim here is to evaluate the Gricean framework as it is commonly understood, I shall focus primarily on the descriptive version, but I shall also indicate how things might differ if the normative reading were adopted (though, given the uncertainty surrounding the interpretation of

¹ Writers who use this wording in quoting or summarizing Grice’s definition include, among many others, Davis (1998, p.13, 2014), Levinson (2000, p.15, 171), Saul (2002a, p.231), and Soames (2009, p.26, p.299).

Grice's words, these remarks should be taken as tentative). With this preliminary point made, I shall now argue that the cooperative presumption has several counterintuitive consequences.

The first problem is that, assuming it is the hearer who is supposed to do the presuming (and it is unclear who else it could be), then this means that implicatures are dependent for their existence on something hearers do. This is counterintuitive. Usually, we do not think that the meaning of an utterance depends on the particular person hearing it, and this seems to go for non-conventional meaning as much as conventional. In trying to work out whether a speaker is implicating something, our sense as a hearer is that we are trying to ascertain a fact that is independent of us, not one that is dependent on what we ourselves do. This assumes the standard descriptive reading of the condition, of course, but a similar problem threatens to arise on the normative reading of it. For it is no more plausible to think that implicatures depend on what the hearers *ought* to do than on what they actually do. (Note that the normative claim is not that speakers are to *be* cooperative, but that they are *to be presumed* to be cooperative, where the presuming is to be done by (I assume) the hearer.)² Whether or not a hearer ought to presume that a speaker is being cooperative plausibly depends on what the hearer believes about the speaker. If they believe that the speaker is trying to mislead them, then they ought not to presume that they are being cooperative. But this means that whether or not the cooperative presumption holds in any given case, and thus whether or not there is an implicature, depends on facts about the hearer.

A second problem is that the cooperative presumption has the consequence that if the speaker is *not* presumed to be cooperative (or *ought* not to be) — say, because the hearer has reason to think they are lying — then the speaker cannot be implicating something. And this seems wrong. Davis gives the following examples (Davis 1998, p.116):

- (1) Karen: Were you out with Jennifer last night?
George: I was out drinking with the boys.

² Compare 'The prisoner is to be watched closely', which is the passive form of the statement that someone should watch the prisoner closely.

- (2) Alice: Do you like my new dress?
Brett: I like all your dresses.

Karen may suspect, or even know, that George is trying to deceive her, and therefore not presume that he is observing the CP. Nevertheless, she will still interpret his words as carrying an implicature to the effect that he was not out with Jennifer. Similarly, Alice may believe that Brett is merely being polite, and thus not cooperative (in the sense that involves being truthful); yet she will still interpret him as implicating that he likes her new dress. Indeed, a speaker can openly *refuse* to cooperate, yet still implicate something (Sterelny 1982, p.189). Rather than answering Alice's question, Brett might try to change the subject by saying 'Is that the time?'. Though this utterance is clearly not cooperative, it still carries an implicature: namely, that it is late and that Brett must leave (and perhaps that he does not wish to talk about Alice's dress).

This problem with the cooperative presumption can be stated in a more general way. People often use implicature in order to *mislead*. They say something that is strictly true but that implicates a falsehood, hoping to get their hearer to believe the falsehood without having actually said anything untrue. (This, of course, is what Mr Bronston was accused of doing in the case discussed in Chapter 1.) Now in itself this is not incompatible with Grice's definition. The cooperative presumption does not say that speakers *are* cooperative, only that they are *presumed* to be (or ought to be presumed to be). However, there is still a difficulty for Grice. For it is common knowledge that people often implicate falsehoods, and when we detect an implicature we may wonder whether we can trust it. If I ask my son, 'Did you eat all the chocolates?' and get the reply, 'I ate some of them', I might wonder whether I can trust the implication that he did not eat them all. Yet on Grice's view, questions of this sort should not arise. Given the cooperative presumption, if one ceases to regard a speaker as being cooperative (or has good reason to think one ought not to regard them as being cooperative), then one should cease to regard them as implicating anything. In interpreting an utterance, the only options should be that it carries no implicature or that it carries a sincere implicature, which the speaker believes to be true (though it might, of course,

actually be false). But in fact we can, and often do, interpret utterances in a third way, as carrying insincere implicatures.

It might be suggested that we could avoid this problem by weakening the CP. For example, we might say that an utterance is cooperative if it is informative, relevant, and suitably expressed, regardless of whether or not the speaker believes it to be true or has evidence for it. (That is, we might drop the commitment to the maxim of Quality.) The problem with this, however, is that it would threaten to undermine the second condition of Grice's definition, determinacy, which is that there should be a *unique* proposition required to preserve the cooperative presumption. For if all that cooperativeness requires is informativeness, relevance, and appropriate expression, then this condition will be met in few cases, if any. In general, if 'p' is an informative, relevant, well-expressed contribution to a conversation, then 'It is not the case that p' would be an equally informative, relevant, and well-expressed contribution. So the supposition that the speaker believes the latter would preserve the presumption of cooperation just as well as the supposition that they believe the former, and thus neither of them could be singled out as *required* to preserve the presumption.

It is likely that any other attempted weakening of the CP cooperation would suffer from similar problems. As we shall see in the next sub-section, there are problems with the determinacy clause as it stands, and weakening the notion of cooperation would only make these worse. Besides, even if a weaker notion of cooperation were adopted, the cooperative presumption would still fail for cases where the speaker blatantly fails to cooperate by changing the subject.

Another problem case for the cooperative presumption is monologue, in which there is no audience at all and hence no one to make the presumption (or to be under an obligation to make it). Grice's definition implies that in such cases speakers cannot implicate at all, and Grice explicitly endorses this conclusion:

I take it as being obvious that insofar as the presence of implicature rests on the character of one or another kind of conversational enterprise, it will rest on the character of concerted rather than solitary talk production. Genuine monologues are free from speaker's implication. (Grice 1989, p.369).

However, this is counter-intuitive. Suppose Inspector Clouseau is talking to himself, trying to reason out the chauffeur's motives. 'The chauffeur is clearly not the murderer', he says, 'But some of his statements were lies... Perhaps he is covering up for the real murderer.' Intuitively, his second sentence carries the implicature that not all of the chauffeur's statements were lies. Similarly, in his reasoning Clouseau might employ figures of speech which depend on implicature, such as metaphor or irony. He might say to himself, 'Maria Gambrelli cannot be the murderer; she is a saint!' — meaning that she is virtuous, not that she has been canonized. Grice might reply that Clouseau's reasoning does not constitute a *genuine* monologue; perhaps in this case Clouseau is his own hearer, or there is an imagined hearer. But if this is not a genuine monologue, then it is not clear what would count as one.

It might be suggested that Grice could accept the possibility of implicature in monologue if he were to allow that the presumption of cooperation can be made by speakers themselves. However, this would threaten to make the first clause of his definition of implicature redundant. For speakers will always (and probably should always) presume that they are being cooperative *with themselves*.

A final problem with the cooperative presumption is that in making implicature hearer-dependent, it also seems to make it hearer-*relative*. Suppose speaker S has two hearers, A and B. S says that p, aiming to implicate that q and thinking that A and B will realize this. A has no reason to distrust S, and therefore presumes (and ought to presume) that S is being cooperative. B, however, has been told (on seemingly good authority, that S is untrustworthy, and therefore does not presume (and ought not to presume) that S is cooperative. Has S implicated that q? Assuming the other conditions are met, it seems that they have done so as far as A is concerned but not as far as B is concerned. But this would mean that implicature is hearer-relative — something which Grice's definition does not acknowledge or allow for.

1.2 Problems with determinacy and calculability.

Grice's second condition for the existence of an implicature with content q is that the supposition that the speaker believes q is required to preserve the cooperative

presumption. This condition is supplemented by the Calculability Assumption, according to which a hearer can work out that this supposition is required from general principles and background knowledge.

The first thing to say here is that this further extends the hearer-dependency of implicature. If the supposition that the speaker thinks that *q* is required in order to make their utterance consistent with the presumption that they are following the CP, then if the hearer has not identified *q*, they cannot consistently continue to presume the speaker to be observing the CP (and surely ought not to either). But according to (1), that presumption is necessary for the existence of the implicature. Thus, it seems, the identification of an implicature by the hearer is necessary for the existence of the implicature itself. This again is counterintuitive. Common sense tells us that a hearer may fail to spot an implicature. A third party, listening to our account of an earlier conversation, may alert us to an implicature we had missed. ('How could you be so silly?', we might say to a friend, 'He was implying that he wanted to ask you out!') Likewise, common sense tells us that a hearer might believe that a speaker is implicating something without being able to work out exactly what. (I shall give an example shortly.) It may be objected that this is an uncharitable reading of Grice's definition, but it follows from the wording as it stands, and it is in line with the hearer-dependency of implicature indicated by condition (1).

Another problem with the determinacy condition is the role it gives to the demands of consistency. The condition says that a speaker *S* implicates *q* in saying *p* only if one must suppose that *S* thinks that *q* in order to make the claim that *S* said *p* consistent with the claim that *S* is observing the CP. The problem is that such a supposition, attributing a specific belief to *S*, will *never* be required; a hearer can always reconcile the two claims simply by supposing that *S* believes (and is trying to communicate) *some*, unidentified, proposition that would be a cooperative contribution to the conversation.³ The demand for consistency would never push us to go beyond this general supposition — with the consequence that no one ever implicates anything specific! Obviously, this is not what Grice intended — though, again, it follows from his wording. To get the intended result,

³ Thanks to Keith Frankish, who pointed this out to me.

the required supposition should be one that does not merely make S's saying that p consistent with their being cooperative, but one that *explains how* it manifests their cooperativeness — what cooperative contribution it makes. That is, (2) should be something like the following:

(2') the supposition that he is aware that, or thinks that, *q* is required to explain how his saying or making as if to say p (or doing so in *those* terms) reflects his presumed observance of the CP.

Having noted this, for convenience I shall continue to use Grice's original wording in what follows; none of the points I make will be affected by this.

Further problems arise from the Calculability Assumption. As we saw in the previous chapter, this states that for an utterance U to carry an implicature, *q*, it must be possible to arrive at the supposition that the speaker believes that *q* by a process of inference from information mutually available to speaker and hearer, including the CP and its maxims, the conventional meanings of the words used and the identity of any references involved, the context of the utterance, and appropriate background knowledge. This assumption has been the subject of much discussion, and a large number of problem cases have been described, in which a Gricean calculation appears to produce the wrong result — either predicting an implicature that intuitively isn't there (a false positive, as Davis calls it; 1998, p.63), or failing to predict one that intuitively is there (a false negative). The latter cases are the more problematic ones for Grice, since Grice does not hold that calculability is *sufficient* for implicature (condition 3, mutual knowledge, must be met too), but only that it is necessary. The hard cases for Grice are ones where a hearer believes that an implicature is present (and may even be able to identify it) but cannot calculate it by Gricean means. Grice must deny that there is any implicature in such cases, but, as we shall see, this is often highly counterintuitive. I shall not summarize this literature here (for a careful presentation of many problem cases, see Davis 1998, Chapters 2–3). Rather, I shall focus on certain core cases that highlight fundamental problems with the Calculability Assumption. (In addition, some problems concerning the calculability of *generalized* implicatures will be considered in Chapter 5.)

One problem lies with the starting point of the calculation process. According to Grice, the process begins when a hearer detects an apparent violation of the CP — when a speaker says something that, taken literally, flouts one or more of the conversational maxims. The problem is that in some cases it may not be clear whether or not a speaker is doing this. Some utterances can be taken either literally or figuratively, and yet be equally cooperative either way. (Davis calls this ‘the rhetorical figure problem’; Davis 1998, pp.65–70.) For example, suppose Danny has taken Candy to the cinema to see Quentin Tarantino’s latest violent action movie. Throughout the movie Candy sits silent and emotionless. When it is over, the following exchange takes place:

- (3) Danny: What did you think of the movie?
 Candy: It was sublime and beautiful.

How should Danny interpret Candy’s reply, given the assumption that she is observing the CP? Intuitively, there are two options: Candy might be speaking literally, or she might be being ironic — saying the opposite of what she believes (flouting the maxim of Quality) in order to implicate that the movie was brutal and ugly. Either message would be an informative and relevant reply to Danny’s question, so the assumption that Candy is observing the CP does not distinguish between the two interpretations. Given Candy’s lack of response during the movie, the context of the conversation does not help to decide between them either. It might be argued that further contextual information or background knowledge could settle the matter. For example, if Danny knows that Candy dislikes violent action movies and was reluctant to see this one, then he might suspect she is being ironic. But even so, Candy might still be speaking literally. She might have been so impressed by Tarantino’s film that she had changed her mind about action movies. And if Danny himself thinks that the film really was sublime and beautiful, then he might regard this as an equally plausible reading of Candy’s remark.

It is true that there might be further items of information that would settle the matter. Suppose that Candy did in fact intend her words to be taken ironically (and believed that Danny could work this out). And suppose Danny, who knows Candy well, senses this, and concludes that she thinks the movie was brutal and ugly.

Still, this does not mean that the implicature is calculable in Grice's sense. For Danny can work out what Candy thinks from the literal meaning of her words together with the fact that she meant them ironically, without appealing to the CP at all. Since the calculation does not depend on the assumption that Candy is observing the CP, it does not show that the supposition that Candy thinks the movie was brutal and ugly is required in order to preserve the assumption that she is observing the CP.⁴ This is not surprising. Since her utterance would be cooperative under both literal and figurative readings, the assumption that she is being cooperative cannot help Danny decide which interpretation to prefer — which means that a Gricean calculation is not possible.

Since Grice holds that calculability is necessary for implicature, he must deny that there is an implicature in this case; the ironic reading is not required, so there is no irony. This clashes with our intuition that there might be an implicature, with Candy's intentions that there should be one, and with Danny's conclusion that there was one. The example here involves irony, but similar examples could be generated for other figures of speech — understatement, overstatement, metaphor, and so on.

A second type of problem case is one where a Gricean calculation can at best show that *something* is being implicated without being able to identify *what*. Consider the following example (based on an actual exchange between the author and a former colleague). Jill enters her workplace and greets a colleague Finn, whom she hasn't seen for several days:

- (4) Jill: How are you feeling?
Finn: I need to find Suleiman again.

Jill knows that Suleiman is someone Finn met on a recent holiday in Turkey. Finn has mentioned him several times and shown her a photo of him. So she knows the

⁴ Davis makes the same point: if facts about the figures of speech speakers are using are treated as part of the context of their utterances, then implicatures can be calculated directly from this context and what is said, without establishing that they are required to maintain the cooperative presumption (Davis 1998, p.70).

reference of ‘Suleiman’. But she thinks it is unlikely that Finn literally means that he needs to find this person again, and, even if he does, the information is not relevant to her question. Now Jill knows that Finn has a liking for non-literal uses of speech, so she suspects that he is implicating something. And she can confirm this by Gricean reflections. Finn’s utterance appears uncooperative, violating the maxims of Relation and Quantity. However, Jill thinks it is unlikely that Finn is being uncooperative or that he has produced an uncooperative utterance by mistake, so she infers that Finn means to convey some relevant information, and that he thought she could and would work out what this was. However, Jill cannot do this. She can infer that the relevant information must concern Finn’s state of physical or mental well-being, and perhaps further that he is lacking something that Suleiman could provide. But she cannot move beyond this, since she has no idea *what* Suleiman could provide. She cannot identify the relevant attribute of Suleiman.⁵

It may be objected that Jill lacks some crucial background knowledge. If she knew more about Suleiman and Finn’s relationship with him, then she could make the calculation. But even if Jill did have more background knowledge, she still might not be able to make the calculation. Suppose she knows that Suleiman is an amusing conversationalist, a skilled masseur, a good cook, and that he makes herbal teas to treat headaches. Then, perhaps, she can narrow down what Finn is implicating. It may be that Finn is bored and wants to be diverted, or that his back is hurting, or that he is hungry, or that he has a headache. But, even so, Jill cannot work out *which*; she doesn’t know which of Suleiman’s attributes is the relevant one.

⁵ In the actual conversation on which this example is based, my colleague said ‘I need to find myself a Suleiman’, and in previous presentations of this material I used this form of words. I have changed the wording in the present version in order to avoid the objection that in the phrase ‘a Suleiman’, ‘Suleiman’ serves, in the speaker’s idiolect, as a common noun (meaning, say, ‘an interesting conversational partner’) and thus that if the hearer knew its literal meaning they would be able to work out the implicature. The objection does not apply to the present version, where ‘Suleiman’ is a singular term. (Thanks to André Gallois for drawing my attention to this point.)

Of course, there is information that would allow Jill to work out what Finn is implicating. She could simply ask Finn why Suleiman would help. If he mentions Suleiman's conversational talents, then Jill will conclude that Finn is implicating that he is bored. But this would still not establish the sort of rational connection between literal meaning and implicature that Grice requires. Even if Jill could now work out that Finn was expressing the belief that he is bored, she still wouldn't be in a position to establish that he *must* have been doing this (or even that it is probable that he was), in order to preserve the assumption that he was observing the CP. For, given the range of Suleiman's talents, Finn's utterance would have been equally cooperative if it had expressed the belief that his back hurt, or the belief that he was hungry, or the belief that he had a headache.

Again, then, it seems that Grice must deny that there is an implicature here. If a Gricean calculation cannot identify a unique belief that must be attributed to Finn to make sense of his utterance, then the determinacy condition fails to hold and there is no implicature, despite the fact that Gricean considerations would lead a hearer to believe that he is implicating *something*.

It might be objected that Grice allows that implicatures can be indeterminate. He writes:

Since, to calculate a conversational implicature is to calculate what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed, and since there may be various possible specific explanations, a list of which may be open, the conversational implicatum in such cases will be a disjunction of such specific explanations; and if the list of these is open, the implicatum will have just the same kind of indeterminacy that many actual implicata do in fact seem to possess. (Grice 1975/1989, pp.39–40)

Thus, since there are various equally plausible explanations for Finn's utterance (that Finn is bored, that his back is hurting, that he is hungry, that he has a headache), perhaps his utterance should be interpreted as the disjunction of these claims: that Finn is bored or has a sore back or is hungry or has a headache.

There are problems with this suggestion, however, which illustrate a general difficulty for the determinacy condition. It is true, as Grice notes, that many implicata are indeterminate; many metaphors, for example, are effective precisely because they express an open-ended range of related thoughts (Martinich 1991). However, as Davis points out, a metaphor is not equivalent to a disjunction of the thoughts it expresses. We would not regard a metaphor as appropriate if only one of the many thoughts it expressed were true, even though that would be enough to make the disjunction true. ‘My love is like a red rose’ would not be an apt metaphor if its subject possessed only one of the many properties which a human might share with a rose — say, being sweet-smelling — while being utterly unlike a rose in every other respect (Davis 1998, pp.71–2). Moreover, it is very unlikely that a speaker who uses a metaphor believes that their hearer will construct a disjunction of all the ideas their words express and attribute a belief in it to them. Yet if they do not, then the determinacy clause will not hold and there will be no implicature. For example, in Finn’s case it is very unlikely that he is seeking to convey a complex disjunctive proposition to the effect that he has at least one of a series of needs, and thinks Jill can work this out. And even if he were, it would be hard to reconcile the supposition that he is doing so with the presumption that he is being cooperative. Since Finn must know which of the disjuncts is true (all the claims are about his own feelings), it would be uncooperative of Finn not to indicate which of them it is (violating the maxim of Quantity). And if more than one of the disjuncts is true, then it would be uncooperative not to indicate that their conjunction was true (again violating the maxim of Quantity).

Another problem with Grice’s approach is that it makes some implicatures dependent on the hearer’s state of knowledge. Suppose Finn wants to implicate that he is bored. He thinks that all Jill knows about Suleiman is that he is a good conversationalist, and so believes she will work out what he means. And if that were all Jill knew about Suleiman, then the implicature would succeed. Jill presumes that Finn is observing the CP, and, given what she knows about Suleiman, the only way to make sense of his utterance is to suppose that he is indicating that he is bored. And Finn believes that she can work this out. So the three conditions are met, and Finn successfully implicates that he is bored. Now consider another case. Everything is the same except that Jill knows more about

Suleiman: that he is a skilled masseur, a good cook, and so on. Now there are other equally plausible suppositions available to Jill, and the disjunctive reading is required. But since Finn does not know that Jill has this additional knowledge, he still believes she will suppose that he is indicating that he is bored, and he doesn't even consider the disjunctive reading. So the determinacy condition does not hold, and nothing is implicated at all. This is counterintuitive, and it leaves us with a puzzle as to what to say about cases where a speaker has multiple hearers. What if Finn's remark had been addressed to a group of people, each of whom had a different level of knowledge about Suleiman? For Finn to successfully implicate that he was bored, would it be sufficient that one of his hearers was required to attribute that belief to him (given their particular background knowledge) or must all of them be required to do so (given their different levels of background knowledge)? Neither option seems plausible. It seems too weak to allow that one hearer is sufficient (especially as that hearer might be the one with the least background knowledge), but too strong to require that all are necessary, since then the addition of one new poorly informed or over-informed hearer could undermine an implicature that everyone else agreed existed.

1.3 Problems with mutual knowledge

The third of Grice's conditions for the existence of an implicature with content q is that the speaker should think (and expect the hearer to think that they think) that the hearer can work out, or grasp intuitively, that the supposition that the speaker believes q is required to make their utterance consistent with the presumption that they are observing the CP. (For simplicity I shall omit the parenthesis in what follows; nothing will turn on it.) Again, there are problems with this condition.

One problem is that the condition requires speakers to possess a sophisticated understanding of the role that (according to Grice) cooperation plays in the generation of implicatures. Yet, it seems, young children can implicate things without having this understanding. When my five-year old son says he ate some of the chocolates or asks me if I can find his shoes, he implicates things, though I doubt if he believes I can work out what these things are from the presumption that he is observing the CP (despite his having overheard me talk about this at great length!). Likewise, a person can doubt that hearers can derive implicatures from

the CP without thereby becoming unable to implicate. Wayne Davis notes that his rejection of Gricean theory has not reduced his power to implicate things (Davis 1998, p.121).

Moreover, even if speakers do sometimes form the beliefs Grice mentions, it seems possible to implicate without them. I implicate things to my five-year old (for example, in saying that he can have *some* of the sweets), without believing that he can work out that the supposition that I believe that he may not have all the sweets is required in order to preserve the assumption that I am being cooperative. And though I do believe that he can grasp my meaning intuitively, I do not believe that he can do this in the ‘quick’ way Grice describes, which involves having the ability to construct, or at least recognize, the full Gricean calculation (see the discussion in Chapter 2, section 3).

There are many other cases where a speaker implicates something without believing that their hearer can calculate or intuitively grasp the implicatum. A speaker might say something with the intention that it should carry an implicature that they do not expect their hearer to recognize. Talking to a particularly annoying acquaintance, I might say something that carries a subtle and insulting implicature just for my own satisfaction and with the intention that it should go over my acquaintance’s head. Or a speaker might say something with the *hope* that their hearer will get the implicature but without being confident that they will or even believing that they won’t. (Think of a relative talking to a coma patient in the hope of triggering a response.) Or a person might say something without realizing that it carries an implicature until it is pointed out to them later — an embarrassing situation with which most of us are familiar. Or, finally, as noted earlier, one might use implicatures in a monologue, where there is no hearer involved at all.

It might be objected that the beliefs Grice mentions need not be conscious ones, and that in the cases mentioned the speakers have *nonconscious* beliefs of the required kind, which in some cases conflict with their conscious beliefs (for the view that we have separate conscious and nonconscious systems of belief, whose contents may conflict, see, for example, Frankish 2004). However, without independent evidence for the existence of implicit beliefs with the contents in question, this looks like an ad hoc move designed to save the theory. Moreover, it

is hard to see how the suggestion could be extended to cases of unintended or solitary implicature.

2. A normative reading

Many of the problem cases considered above are ones in which Grice's theory of conversational implicature seems to conflict with our intuitions as to whether an implicature is present or what its content is. But it may be that the conflict is only apparent. Perhaps the theory and the intuitions concern different things. In particular, it may be that our intuitions concern what speakers *intend* to implicate by their utterances or what hearers *take* utterances to implicate, whereas Grice's theory concerns what utterances actually *do* implicate, in a normative sense. If so, then many of the supposed problems for Grice will disappear. This approach has been proposed by Saul (Saul 2002a), and I shall look at it in this section of the chapter.

2.1 Speaker meaning, implicature, and an extended taxonomy

Saul proposes that Grice's notion of conversational implicature is intended to capture a normative aspect of language use, parallel to his notion of sentence meaning. Grice holds that for a speaker to succeed in saying *p* by uttering sentence *S*, it is not enough for *the speaker* to mean that *p* (to have the right communicative intentions); it must also be the case that the *sentence S* means that *p* — which, for Grice, is, roughly, to say that people typically use it to mean *p*. The sentence's meaning must match the speaker's meaning. Similarly, Saul proposes, for a speaker to implicate that *q* it is not enough for them to mean *q*; they must also produce an utterance that implicates *q*, where this is not determined by the speaker's intentions. The idea is that as speakers we do not have complete control over what we implicate, any more than we have complete control over what we say. We cannot implicate whatever we like by a given utterance, any more than we can say whatever we like by it. (I cannot say that I am dyslexic by uttering the words 'I am dialectic', even if that is what I mean to say, and even if I think that that is what the sentence means.) To implicate that *q* by uttering *S*, an objective, normative condition must be met as well as a subjective, psychological one. However, this normative condition cannot be that *S* is typically used to implicate

q (as the parallel with sentence meaning would suggest), since many implicatures are context dependent. So, Saul suggests, Grice identifies it instead with the meaning that the hearer is required to attribute to the speaker in order to preserve the assumption that the speaker is being cooperative. This may vary from context to context, but it is not controlled by the speaker's intentions, and so gives the required objective condition (Saul 2002a, p.241).

On this reading, implicature is not a form of speaker meaning, and we must reject the view (which Saul suggests is close to being an orthodoxy) that for Grice what a speaker *means* divides into what they say and what they implicate. (This is the view by defended Stephen Neale, discussed in Chapter 2.) Saul argues that this is an implausible reading of Grice, given that his definition of speaker meaning is framed wholly in terms of the speaker's intentions, whereas his definition of implicature focuses on what *the hearer* presumes and supposes. It follows that a hearer's attitudes cannot constrain what a speaker means but can constrain what they implicate, which suggests that the two can diverge. Saul argues that one can mean things that one does not implicate (and does not say either), and she suggests that it may be possible to implicate things one does not mean.

Saul illustrates the former case (meanings that are not implicated) with variants of an example used by Grice, in which a philosophy professor is writing a reference letter for a student who is applying for an academic job (Grice 1975/1989, p.33). In the first variant, the professor is writing a reference for Fred, who is a poor philosopher and a thief. She doesn't wish to say this explicitly, however, so she devotes her reference letter to praising Fred's genuine typing skills, intending to implicate that Fred is a poor philosopher. However, her audience does not interpret her letter in the way she expects. Fred is in fact applying for a typing job, and the employer takes the letter literally. In the second variant, the professor writes a similar letter for Cedric. Again she focuses on irrelevant matters, such as Cedric's typing skills, intending to implicate that Cedric is a poor philosopher. Again, she fails, however. For the appointing committee have been told that she disapproves of writing reference letters, and they therefore assume that she is simply being uncooperative, and do not search for an implicated meaning. In both cases, then, the speaker means something that is not successfully implicated (and not said either) (Saul 2002a, p.230, pp.234–5).

As a possible case of unmeant implicature, Saul uses another variant of the same example. This time, although the student, Roland, is a poor philosopher, the professor likes him and wants him to get the job. So she writes a long and detailed letter containing lots of information about Roland’s academic career but no judgements on his ability. Since the letter lacks important information, it can be read as cooperative only on the supposition that the professor thinks that Roland is a poor philosopher, and the professor believes that the audience can work this out. So the letter implicates that Roland is a poor philosopher (in the example it is the letter that implicates rather than any particular sentence in it). However, the professor does not *intend* the audience to form this belief; in fact, she hopes that they will read the letter superficially and form a positive impression of Roland. So she does not intend the audience to form the belief that Roland is a poor student and so does not *mean* that, though her letter implicates it (Saul 2002a, p.237–8). If this is right, then speakers can mean things that they do not implicate (or say either) and implicate things that they do not mean. This is illustrated in Figure 2.

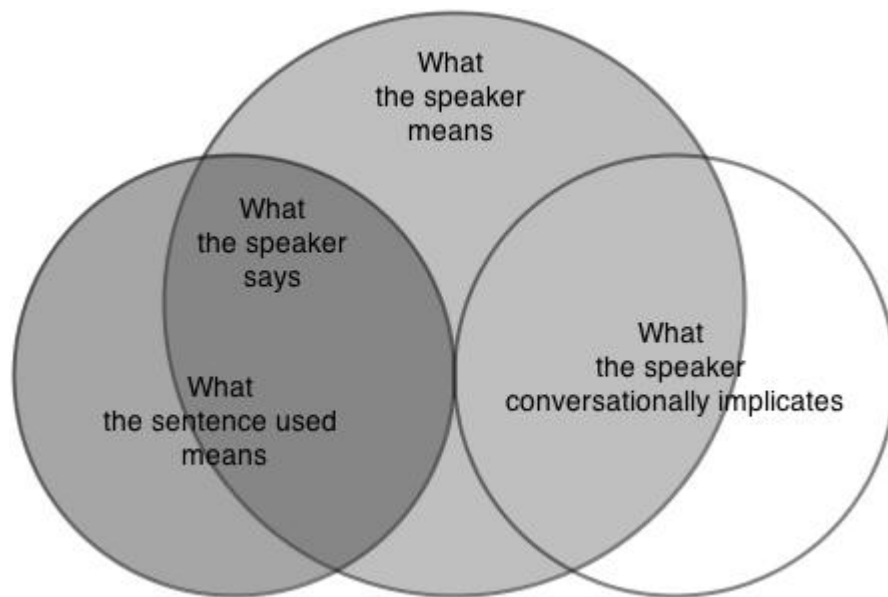


Figure 2: The relation between speaker meaning, sentence meaning, and conversational implicature, on Saul’s reading of Grice. A given proposition may fall in any of the circles or their defined overlaps. (The Roland case would fall in the unshaded area on the far right.)

Saul notes that if there are aspects of meaning that are neither said nor implicated, then Grice's taxonomy is incomplete. To rectify this, Saul introduces the notion of *utterer-implicature*. The definition of utterer-implicature is the same as that of conversational implicature, except that it is not necessary for conditions (1) (the cooperative presumption) and (2) (determinacy) to hold, but only for the speaker to *think* they hold.⁶ That is (1) and (2) are replaced by:

(1*) *The speaker thinks that* he is presumed to be following the conversational maxims, or at least the Cooperative Principle.

(2*) *The speaker thinks that* the supposition that he [the speaker] is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with this presumption.

(From Saul 2002a, p.235)

In the Fred and Cedric cases, the professor utterer-implicates that their student is a poor philosopher, without conversationally implicating it.

Saul also proposes a corresponding notion of *audience-implicature*, which replaces clauses (2) and (3) in the definition of conversational implicature with the following:

(2A) *The audience believes that* the supposition that he [the speaker] is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with this presumption.

(3A) *The audience takes the speaker to think that* it is within the audience's competence to work out that the supposition mentioned in (2) is required.⁷

(From Saul 2002a, p.242)

⁶ Saul adopts the usual descriptive phrasing of the cooperative presumption, rather than Grice's own normative one, and I shall follow her in this in what follows.

⁷ Saul says 'the supposition mentioned in (2)' but to make the definition more self-contained we could change this to 'the supposition mentioned in (2A)'.

As an example of an audience-implicature, Saul uses another reference-letter case. This time the professor is writing a letter for a student, Felix, who she thinks is applying for a job as a typist. Although Felix is an excellent philosopher, the professor says nothing about his philosophical abilities and writes about his typing skills and punctuality. However, Felix is in fact applying for a philosophy post, and the audience interpret the letter as implicating that Felix is a poor philosopher. This is not a genuine conversational implicature, since, although conditions (1) and (arguably) (2) are met, condition (3) is not.⁸ However, it is an audience-implicature, since (2A) and (3A) hold (Saul 2002a, p.242).

As Saul notes, an utterer-implicature is (roughly) what the speaker is trying to implicate, and an audience-implicature is what the audience takes the speaker to be implicating (Saul 2002a, p.243). Thus, if a claim *q* is both utterer-implicated and audience-implicated, then it will have been successfully communicated. This might suggest that a conversational implicature is simply a combination of the two: *q* is conversationally implicated if it is both utterer-implicated and audience-implicated. Saul rejects this suggestion, however, arguing that something can be conversationally implicated without being audience-implicated. She illustrates this with two final examples, in which a professor writes reference letters for two students, Trigby and Wesley. In both cases, the professor seeks to implicate that the student is a poor philosopher by writing about irrelevant matters, such as their rock-climbing skills or wide knowledge of illegal drugs. The letters can be read as cooperative only on the assumption that the professor thinks that the students are poor philosophers and believes the audience will realize this, so the conditions for conversational implicature are met. However, the audiences fail to interpret the letters as intended. In Trigby's case, they see that the supposition that Trigby is a

⁸ (2) is met if it is assumed that the writer understood who they were addressing and why. Although this assumption may be false (as in the Felix case), it seems reasonable for hearers to make it when deciding how to interpret an utterance. If it is not made, then a non-literal reading of an utterance will be required only if there is no conceivable misunderstanding on the speaker's part which would make the utterance cooperative on a literal reading — with the consequence that non-literal readings are required far less often than we think.

poor philosopher is needed, but don't realize that the professor intended them to work this out. (They think she was trying to trick them into forming a positive opinion of Trigby.) Hence (3A) is not met. In Wesley's case, the audience read the letter quickly, notice some positive words, and think that the professor is recommending Wesley. Here (2A) is not met. Thus in these cases the claim that the student is poor is conversationally implicated but not audience-implicated (Saul 2002a, p.244).

Saul uses the Wesley case to illustrate what the role of conversational implicature actually is, on the interpretation she proposes. Suppose the audience (the hiring committee) complain that the professor misled them about Wesley. It would not be sufficient, Saul argues, for the professor (who in the examples is Saul herself) to reply that she *utterer-implicated* that Wesley was a poor student — that she *believed* that the supposition that she thought he was a poor student was required in order to make sense of her letter. For if her belief was not justified, then she could still be blamed for the miscommunication.

Saying that I utterer-implicated that Wesley is a poor philosopher is not much of a defense: I could have utterer-implicated that Wesley was Elvis if I was crazy enough to suppose that attributing this belief to me was required to make sense of my utterance, and that the audience could work this out. (Saul 2002a, p.244)

But (Saul continues) it *would* be a good defence to claim that she *conversationally implicated* that Wesley was a poor student:

What I can do, however, is maintain that I conversationally implicated it: It *was* required in order to understand me as cooperative, and my audience *was* capable of working it out. (ibid.)

By conversationally implicating something, Saul argues, one has *made it available* to one's audience, and thereby fulfilled one's communicative responsibilities in the matter — whether or not one's audience actually grasps it. Grice's notion of

conversational implicature, Saul concludes, is designed precisely to play this normative role.

2.2 *An enriched Gricean framework*

By treating (conversational) implicature as a normative notion, and by introducing the additional concepts of utterer-implicature and audience implicature, Saul enriches the Gricean framework and enables it to avoid some of the problems discussed earlier. In particular, the enriched framework provides at least partial solutions to the problems concerning determinacy and calculability. If we understand implicature as normative, then it becomes easier to accept that there is no implicature present in cases like those of Candy and Finn, where calculability fails. We can agree that the utterer has not done enough to make the information available. Our intuitions that there is an implicature in place in these cases (or that there is a more determinate one than we can calculate) can be vindicated by the existence of appropriate utterer-implicatures. Thus, Candy utterer-implicated that Tarantino's movie was ugly and brutal, and Finn utterer-implicated that he was bored, since they believed that those interpretations were required to make sense of their utterances as cooperative and they thought their hearers could work that out. However, they were wrong about this, since in neither case was the intended meaning required in the Gricean sense. Candy's utterance did not require a non-literal reading at all, and although Finn's did require such a reading, he did not do enough to narrow down the possible non-literal meanings to the one he had in mind. Thus, neither conversationally implicated the things they utterer-implicated.

However, Saul's account doesn't resolve all of the problems for Grice. First, it does not address the problems that arise from condition (1) (the cooperative presumption). As noted earlier, it is plausible to think that an utterance can carry an implicature even if the hearer does not (or should not) presume the speaker to be being cooperative (for example, where the hearer thinks the speaker is trying to mislead them, or where the speaker is changing the subject or engaging in monologue). We might respond by saying that in such cases there is only an utterer-implicature, not a conversational one. That is, though (1) does not hold, (1*) does: the speaker *thinks* that he or she is presumed to be being cooperative. However, this is not adequate. For in many of the cases discussed earlier, even

(1*) will not hold: typically, speakers who openly change the subject or engage in monologue will not even *think* that their hearers presume them to be being cooperative.

Second, the enriched framework still faces the problems created by condition (3), mutual knowledge, which says (simplifying somewhat) that a speaker thinks that the hearer can work out the implicated message in the Gricean way. As we saw, intuitively it is possible for an utterance to carry an implicature even where this condition is not met (for example, where the speaker is a child, or where an implicature is intended to go over a hearer's head, or where the speaker is engaged in monologue or talking to a coma patient). And the enriched framework still cannot explain this. There cannot be an utterer-implicature in these cases any more than there can be a conversational one, since the definition of utterer-implicature includes condition (3) unchanged. Nor is it plausible to think that there is an audience-implicature. For an audience-implicature exists only if the audience thinks that condition (3) holds (that is, if they take the speaker to think that they can work out the implicated content in Gricean fashion) and in the problem cases discussed (where the speaker is a child, or the hearer uncomprehending or comatose or non-existent) this is no more likely than it is that condition (3) will actually hold.

3. Some modifications

Despite its attractions, there are also some difficulties for Saul's reading of Grice, and in this section I shall highlight these and propose some modifications to the definition of conversational implicature in order to bring it more in line with the normative view proposed by Saul. Although this means departing from the letter of Grice's account, the modified version will remain broadly faithful to his approach, and, as we shall see, the modifications will also remove some of the remaining problems for the Gricean framework mentioned in the first section.

3.1 The cooperative presumption revised

The first difficulty concerns condition (1), which runs against the spirit of the normative reading. If implicating something involves *making it available* to one's hearer, then it should not depend on the *hearer's* doing something. A content might

be made available to a hearer even if they did nothing to pick it up — not even presuming that the speaker was being cooperative. From a normative perspective, it doesn't matter whether anyone actually makes the presumption that the utterer is being cooperative.⁹ What matters is what is *rationally required* to make their utterance consistent with that presumption, and this is captured by (2) on its own, with slight rephrasing as follows:

(I) (= 1 & 2 revised) The supposition that the speaker is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with the presumption that he is observing the CP.

This suggests that if implicating involves making a content available, then we should revise its definition, dropping (1) altogether and replacing (2) with (I) above. This would be more consistent with the idea that the notion of implicature is a normative one. Moreover, it has the additional advantage that it removes all the problems for the Gricean framework arising from clause (1). On the modified view, a speaker can conversationally implicate something (make it available) no matter what attitude their audience takes towards them and even if there is no audience at all. (In the latter case, the content will be made available to merely *potential* hearers.)

3.2 Mutual knowledge revised

The second difficulty concerns condition (3), which also does not fit in well with the normative approach. If implicating something involves making it available according to an objective standard (fulfilling one's communicative responsibilities with regard to it), then there ought to be more to it than the speaker merely *believing* that certain conditions are met. Thus, when discussing clauses (2) and (2A), Saul points out that if a speaker is accused of not properly communicating a

⁹ A similar point holds on the 'to be presumed' reading of the cooperative presumption. Whether or not a speaker makes a piece of information available should not depend on the *hearer's* obligations.

piece of information, p, it would not be enough for them to reply that they had *believed* that attributing the belief that p to them was required in order to make sense of their utterance (Saul 2002a, p.244). If attributing that belief to them was not *in fact* required, then they had not made p available. Similarly, I suggest, it would not be enough for a speaker to say that they *believed* their audience could work out that attribution of the belief that p was required: their belief about their audience's abilities must be a reasonable one. To see this, consider another reference letter example. Again, the professor wants to implicate that one of her students, Dido, is a poor philosopher. This time, however, she tries a different tactic. She writes a glowing letter full of relevant details, but closes the letter with the sentence 'In closing, it is of the utmost importance to stress that I am a biscuit'. She believes that the only way to make sense of this surreal comment is as conveying the belief that all she has just written about Dido is unreliable, and she believes that her audience are able to work this out. It is not implausible to think she is right about the first point (how else could one make sense of the comment?), so the conditions for implicature are met (the message that Dido is a poor philosopher is required to make sense of the speaker's utterance and the speaker believes the audience can work this out). As it turns out, however, the professor has overestimated her audience's inferential abilities. Her tactic is too subtle for the hiring committee, who are simply baffled by the closing sentence. So although the conditions for implicature are met, the professor has not made the relevant information available, since she had a mistaken view about her audience's ability to work out what is required to make sense of her utterance. (This resembles Saul's Wesley case, where the audience read the professor's letter quickly and failed to work out that the intended meaning was required in order to make sense of it. However, in that case the professor was right to think that the audience could work out the intended meaning (they could have worked it out if they had read the letter more carefully), and responsibility for the failure of communication lay with the audience. In the Dido case, by contrast, the professor is wrong to think that the audience can work out that the intended meaning is required, and responsibility for the failure of communication lies with her. Thus, in the Wesley case the professor made her meaning available, whereas in the Dido case she did not.)

Since (3) is the cause of the problem here, this suggests that we should revise it so that it becomes a claim about what the audience *can* do, rather than about what the speaker *believes* they can do:

(3') It is within the audience's competence to work out, or grasp intuitively, that the supposition mentioned in (I) is required.

Together with the previous modification, this change has the effect of making the notions of utterer-implicature, audience-implicature, and conversational implicature more clearly parallel to each other: the first concerns what the utterer thinks is the case, the second concerns what the audience thinks is the case, and the third concerns what is in fact the case. Thus, an utterance U carries the utterance implicature q if *the utterer thinks* that the supposition that they think q is required to make sense of U and that the audience can work this out. U carries the audience-implicature that q if *the audience thinks* that the supposition that the utterer thinks that q is required to make sense of U and that the utterer thinks they can work this out. And U carries the conversational implicature q if the supposition that the utterer thinks that q *is in fact* required to make sense of U and the audience *can in fact* work this out. Or, more concisely, q is utterer-implicated if the utterer thinks U makes q available; q is audience-implicated if the audience thinks U makes q available, and q is conversationally implicated if U does make q available.

There are problems with this suggestion, however. If (3) was not demanding enough on speakers, the revised version, (3'), seems *too* demanding. For it would mean that an implicature could fail for reasons outside the speaker's control. A speaker might produce an utterance which is designed to implicate that q and which most hearers would interpret as implicating that q, yet fail to implicate that q because their actual hearer is unable to work out that q is implicated — say, because they are confused, ill, or suffering from some mental disability. A normative standard which requires speakers to take account of the specific abilities of individual hearers seems too strict. After all, on the normative reading, implicature is supposed to have a similar role to that of sentence meaning, and that is not relativized to individual hearers. Moreover, many of the original problems for condition (3) would still remain on the revised version. For example, it would

still be impossible to implicate something when talking to a coma patient or in monologue, or to implicate something that went over the hearer's head.

One option here would be to drop condition (3) altogether, and say that an utterance implicates a proposition *q* if the supposition that the speaker thinks that *q* is required to make sense of the utterance, regardless of whether the speaker thinks the audience can work this out or whether the audience can in fact work it out. This would of course solve all the problems arising from condition (3), and it seems to get at the core of Grice's account, as reflected in his briefer presentations of it.¹⁰

However, dropping (3) would again give us a rather loose normative standard, as the unrevised version of (3) did. There could be cases where a certain supposition is required to make sense of an utterance, but where it is beyond the scope of a typical audience to work this out, and in such cases the speaker would not have done enough to convey the supposition. The Dido reference letter serves again as an example. Perhaps the best option, then, would be to revise (3) further, so that it becomes a claim about what a typical, or normal, audience would be capable of:

¹⁰ For example:

[W]hat is implicated is what it is required that one assume a speaker to think in order to preserve the assumption that he is observing the Cooperative Principle (and perhaps some conversational maxims as well), if not at the level of what is said, at least at the level of what is implicated. (Grice 1989, p.86)

and

Implicatures are thought of as arising in the following way; an implicatum ... is the content of that psychological state or attitude which needs to be attributed to a speaker in order to secure one or another of the following results; (a) that a violation on his part of a conversational maxim is in the circumstances justifiable, at least in his eyes, or (b) that what appears to be a violation by him of a conversational maxim is only a seeming, not a real violation ... (Grice 1989, p.370)

(II) (= 3 revised) It would be within the competence of a typical audience to work out, or grasp intuitively, that the supposition mentioned in (I) is required.

(I call the revised claim (II) to match (I), which is our revised version of (1) and (2)). Here ‘typical’ means something like ‘alert, averagely informed, competent adult speaker of the relevant language’. (This definition will do for the present, but in the next chapter I shall argue that we must define the typical hearer more narrowly, to take account of the background assumptions and other information on which the hearer is expected to draw.) This also has the advantage of avoiding the other problems for (3). A speaker can implicate things even if their hearer cannot calculate the implicatum or even if there is no hearer at all, provided a typical hearer could make the calculation. Thus, one can implicate something in monologue or while talking to a coma patient.

3.3 *Unmeant implicatures*

Another consequence of this revision of clause (3) is that there will be clear cases of unmeant implicatures. Although Saul suggests that there may be unmeant implicatures, she does not commit herself and says that the issues are ‘incredibly tricky’ (2002a, p.247 n.28). This is true so long as we stick with Grice’s definition, with the original clause (3). On the one hand, if conversational implicature is the pragmatic parallel of sentence meaning, then it should be possible to produce utterances with implicatures one does not intend, just as it is possible to utter sentences with meanings one does not intend. On the other hand, it is not easy to think of cases where a speaker believes that their audience can work out that the supposition that they (the speaker) believe that q is required to make sense of their utterance *and yet* does not intend to express the belief that q. This is not to say that implicature *conceptually* requires intention. As we saw in the previous chapter, Grice’s views on that are not clear. The point is that it is not easy to find *psychologically plausible* cases where a person fulfils the conditions for

implicating that *q* without also intending to communicate *p*. (Saul offers the Roland case as a possible example, but does so only tentatively.)¹¹

The revised version of Grice's definition, with (II) instead of (3), avoids this problem. Since (II) does not make reference to what the speaker believes, there is no obstacle to a speaker's producing an utterance that carries an unmeant implicature. If a certain reading is required to make sense of a given utterance as cooperative, and if a typical audience would work out this meaning, then it is implicated, regardless of whether or not the speaker intended it.

This is, I suggest, an intuitively desirable result. The following example may help to illustrate this. A professor, Donald, phones a colleague, Rita, to ask her opinion of one of her former students, Omar, who has just applied to Donald's institution. As it happens, Rita thinks that Omar is an excellent philosopher. However, she also knows that he is an exceptionally kind, caring, and inspiring human being, who spends all his spare time on voluntary work and charitable activities, and on the spur of the moment she is so eager to say what a wonderful person Omar is that she forgets to say anything at all about his philosophical abilities. Donald interprets this failure to mention Omar's philosophical work as implicating that Omar is a poor philosopher, and decides to reject his application. Here I think we would intuitively say that Rita (or perhaps her utterance — see below) had implicated that Omar is a poor philosopher, and if she were later to reflect on what she had said, she might well come to that conclusion herself and feel guilty for misrepresenting Omar, albeit unintentionally. But Grice's definition does not support this intuitive conclusion, since clause (3) would not be met. Rita would not count as having implicated that Omar was a poor student, since she did not believe that Donald was able to work out that that belief must be attributed to her to preserve the supposition that she was being cooperative. (If she had believed that, she would have spoken very differently.) However, the revised definition, with (II) instead of (3), does support the intuitive conclusion, since it does not

¹¹ To recap: The professor writes a reference letter that clearly implicates that Roland is a poor philosopher, yet she does not intend the hiring committee to form the belief that he is a poor philosopher, and hopes they will not notice the implicature (Saul 2002a).

mention Rita's beliefs at all. Since (plausibly) attributing to Rita the belief that Omar was a poor philosopher was required to uphold the presumption that she was observing the CP and following the maxims, and since it was within the power of a typical hearer to work that out, Rita (or her utterance) implicated that message.

Although implicatures like Rita's are unintended, the speaker may still be held responsible for them, since they neglected to consider the way in which their utterance would be understood. The recognition of the possibility of unmeant implicatures may have consequences for questions about the ethics of implicature.

4. Implicature and speaker meaning again

I now want to return to the role of speaker meaning in implicature. As we saw in the previous chapter, Grice does not clearly commit himself as to whether implicatures must also be meant, and we left the question open. However, Saul's normative reading, revised as proposed, suggests a way to resolve this issue, as I shall now explain.

4.1 Normative and psychological conditions for implicature

Saul proposes that Grice's notion of implicature is intended to play a normative role, similar to that of sentence meaning. Now, for Grice, uttering a sentence which means that *p* (sentence meaning) is only one of two necessary and sufficient conditions for saying that *p*. The other condition is that the speaker must mean that *p* (speaker or utterer's meaning) — where this is a matter of the speaker being in a certain psychological state (having certain communicative intentions, as previously discussed). If the notion of implicature plays the same role as that of sentence meaning, then this suggests that for a speaker to implicate that *q* it is not enough for them to produce an utterance that implicates that *q*; in addition, the speaker must meet some psychological condition, parallel to having an appropriate speaker meaning (see Table 1).

But what is this psychological condition? One option would be to say that it is the existence of an utterer-implicature (in Saul's sense) with content *q*. Utterer-implicature is defined in terms of speaker's beliefs, so this is a psychological condition. However, it does not seem to be the right psychological condition. Roughly speaking, to say that a speaker *S* utterer-implicates *q* is to say that *S* thinks

that their utterance implicates q — that is, S thinks that the normative condition holds. But in the case of *saying* the two conditions are not related in this way. To say that a speaker S means that p is not to say that S believes that the sentence they utter means that p . Speaker meaning is a matter of having certain communicative intentions, and it is more basic than sentence meaning and independent of it.

A better view, I propose, is that the psychological condition is simply that S means that q in Grice’s standard sense — that is (roughly), S intends to get their hearer H to believe that q (or to believe that S believes that q) and to achieve this in part by getting H to recognize this very intention. As we saw in the previous chapter, there are several reasons to think that Grice held that implicatures need to be backed by speaker meanings, and these reasons all support the present proposal. Thus, I propose that the conditions for the performance of speech acts with, respectively, conventional and nonconventional meaning are as in Table 1. (For simplicity, I have assumed that conventional meaning is limited to what is said, and nonconventional meaning to what is conversationally implicated; the table could easily be extended to accommodate conventional implicature and nonconversational nonconventional implicature.)

	Act	Psychological condition	Normative Condition
Conventional meaning	S says that p	S means that p (speaker meaning)	S uses a sentence that means p (sentence meaning)
Nonconventional meaning	S implicates that q	S means that q	S produces an utterance that meets the Gricean conditions for implicating q

Table 1: Saying and implicating. A psychological condition and a normative condition are individually necessary and jointly sufficient for each act.

Of course, as we also saw in the previous chapter, there are also reasons to think that Grice held that implicatures do *not* need to be backed by speaker meanings. However, from our present position we can explain this, as I shall now show.

4.2 *Speaker implicature and utterance implicature*

If what has been said in the previous section is correct, then a speaker might meet the normative condition for implicating that *q* without meeting the psychological one. In the case of conventional meaning, Grice holds that a speaker says that *p* only if both conditions are met, so by parallel, we should say that a speaker implicates that *q* only if both relevant conditions are met. In the case where only the normative condition is met, we might say that the *utterance*, but not the *speaker*, implicates that *q*. (The parallel with conventional meaning would suggest that we should say it is the *sentence* that implicates. But although this might be appropriate for generalized implicatures, where the same sentence generates the same implicature in most contexts, it does not allow for particularized implicatures, which are context-dependent. To accommodate both kinds, we need to consider sentences as uttered in particular contexts — that is, utterances.) Thus, we need to distinguish between *an utterance* implicating something and *a speaker* implicating something. An utterance *U* implicates that *q* when the normative condition holds (that is, on the proposed revised account, when (I) and (II) hold), whether or not the speaker means that *q*. A speaker *S* implicates *q* when both the normative and psychological conditions hold (that is, when (I) and (II) hold, and, in addition, *S* means that *q*).¹² I shall refer to these as *utterance implicature* and *speaker implicature* respectively. (The latter should not be confused with *utterer-implicature* as defined by Saul. A person utterer-implicates that *q* in Saul's sense if they think that their utterance meets the Gricean conditions for implicating *q*. A person speaker implicates that *q* in my sense if (a) their utterance meets the Gricean conditions for implicating that *q* (or my revised versions, (I) and (II)), and (b) they

¹² As a reminder, conditions (I) and (II) are:

(I) The supposition that the speaker is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with the presumption that he is observing the CP.

(II) It would be within the competence of a typical audience to work out, or grasp intuitively, that the supposition mentioned in (I) is required.

4.3 *The role of intention*

This distinction between utterance implicature and speaker implicature enables us to resolve the apparent conflict in Grice's views about the role of speaker meaning in implicature. As we saw in the previous chapter, a case can be made for thinking both that Grice did and did not hold that implicatures must be backed by speaker meanings. Now, on the view proposed here, an *utterance* can implicate something without the speaker meaning it, as in Rita's case, but a *speaker* cannot implicate something without meaning it (though that is not all there is to implicating it; meaning that *q* is necessary but not sufficient for implicating that *q*). Thus, the apparent tension in Grice's views can be resolved. The two contrasting positions on the role of speaker meaning can be thought of as corresponding to two different questions: what is required for an *utterance* to implicate something and what is required for a *speaker* to implicate something. Since the requirements for a speaker to implicate that *q* differ from those for saying that *q* only in that they involve producing an *utterance* which implicates that *q*, it is natural that Grice's discussion of implicature focuses almost exclusively on the conditions for utterance implicature and makes little or no mention of intentions. However, when we set Grice's theory of implicature in the wider context of his theory of meaning, we need to bring in the further psychological conditions for a person to communicate something, and thus to focus on the broader notion of speaker implicature. Thus, the distinction between utterance and speaker implicature makes sense of what is otherwise a mysterious conflict in Grice's views.

I am not claiming that Grice would have endorsed this explanation. For one thing, he does not distinguish between speaker implicature and utterance implicature, and indeed it is difficult to make that distinction while working with the unrevised Gricean definition, which does not easily accommodate unmeant implicatures. Moreover, Grice's definition of implicature is framed in terms of the conditions for a *person* (indeed, a man) to implicate something ('A man who, by (in, when) saying (or making as if to say) that *p* has implicated that *q* [etc.].' (1975/1989, p.30). This is not, however, strictly incompatible with the proposal made here. Grice's definition might give only *necessary* conditions for speaker implicature, while at the same time giving *sufficient* conditions for utterance implicature. At any rate, I suggest that the interpretation I have proposed makes

the best sense of Grice's comments on the topic, which, as noted in the previous chapter, are often tentative and exploratory.

A final comment on the role of intention in implicature. The view just outlined is similar to the hybrid one suggested in the previous chapter, which tried to link Grice's accounts of implicature and speaker meaning by proposing that when a speaker S implicates that q, S must both (a) intend to get their audience to believe that q in part by recognizing this very intention and (b) believe that the audience can do this by going through a Gricean calculation. (These conditions were not supposed to be sufficient for implicature; the other Gricean conditions for implicature were also required.) I noted, however, that the suggestion created a puzzle as to why it should be necessary for condition (b) to hold, given that it is not necessary for S to believe that their intention can be recognized *only* by means of a Gricean calculation. The proposal set out in the present chapter is similar to this, but also significantly different. For if the revisions to the Gricean definition in section 3 above are accepted, it is no longer necessary for (b) to hold. Condition (b) followed from clause (3) of Grice's definition, but we have revised this so that it no longer requires the speaker to believe that the implicated message can be recovered by a Gricean calculation. It is sufficient that the message *can* be recovered (by a typical audience) by a Gricean calculation (condition (II)), and the reason for including this condition is simply that (together with condition (I)) it provides the normative element highlighted by Saul.

Conclusion

In this chapter I have considered some fundamental problems for Grice's definition of (conversational) implicature and set out a response which involves adopting a normative reading of Grice and introducing new psychological notions of utter-implicature and audience-implicature. I showed how this enriched Gricean framework avoided many, but not all, of the problems raised. In the spirit of the normative reading, I then went on to propose further modifications and extensions, replacing Grice's three-part definition with a shorter two-part one, and introducing a distinction between speaker implicature and utterance implicature. I argued that these revisions removed many of the remaining problems for Grice and also resolved a persistent problem concerning the role of intention in implicature. I

suggest that the resulting account is the most charitable and consistent formulation of the Gricean framework, and that it affords the best line of reply to the problems raised in section 1 of this chapter. Finally, construed in this way as a normative theory, the framework promises to provide us with norms of implicature, of the sort discussed in Chapter 1, which may help us to avoid misunderstandings and resolve disputes. So far, the Gricean framework is standing up well, then. However, as we shall see in the next chapter, a serious problem remains.

Chapter 4

Where the Gricean framework fails

The previous chapter argued that, if interpreted as a normative theory and enriched and revised in certain ways, the Gricean framework can avoid many of the objections raised against it. However, some problems remain, and in this chapter I shall set them out and argue that they seriously undermine the framework.

The first section of the chapter deals with the issue of speaker-dependency. On the normative reading of Grice we have been considering, what is conversationally implicated should be independent of the intentions and attitudes of the particular speaker involved. I shall argue, however, that conversational implicature, as defined by Grice, is not speaker-independent in this way. Although speaker intentions do not directly determine what is implicated, they indirectly determine it via their role in fixing relevant background knowledge, cooperative standards, and other inputs to the interpretative process. The next section of the chapter looks at the implications of this argument for the Gricean framework. It argues that they are serious and that the appropriate response is to adopt a different approach to implicature, which drops the requirement of calculability and gives a greater role to speaker intentions. This intention-centred account, I argue, can still retain a normative element, at least of a weak kind. Section 3 of the chapter looks at the concepts of utterer-implicature and audience-implicature, introduced by Saul as part of her revision of Grice's account. I noted that these notions might be employed to prop up the Gricean framework, but I shall argue that they cannot do the work required of them. The final section of the chapter looks at the implications of Grice's theory of implicature generation for the process of implicature recovery, arguing that here, too, it faces problems.

Although I proposed some (sympathetic) modifications to the Gricean definition of implicature in the previous chapter, the problems I shall highlight in this chapter will not depend on these modifications being accepted and would remain even if they were rejected. Where it is important, I shall indicate how the problems apply to both the original and revised versions.

1. The argument for speaker-dependency

1.1 Normativity and speaker dependency

As we saw in the previous chapter, it is plausible to interpret Grice's notion of conversational implicature as a normative one. On this view, speakers cannot make their utterances implicate whatever they want, any more than they can make them carry whatever conventional meaning they want. There is a standard of correctness for implicature that is independent of the speaker. What is implicated by an utterance is *what is required* to make sense of it as cooperative, where this is determined by general communicative principles, not by what the speaker intends to implicate, and it may differ from what the speaker and hearer think is required — notions for which Saul introduces the terms *utterer-implicature* and *audience-implicature*, respectively.

This is supposed to apply to particularized, context-dependent, implicatures as much as to generalized ones — if not to a greater degree. In the case of generalized implicatures, a normative, speaker-independent notion could be defined by generalizing across speakers. We could say that 'Some F are G' conversationally implicates 'Not all F are G' if speakers typically use it to implicate that (as they do). But this tactic cannot be used with particularized implicatures, since these depend heavily on context, and, Saul argues, Grice included audience-related criteria in his definition of implicature precisely in order to provide a speaker-independent element in such cases.¹ What is implicated by an utterance is what the audience is rationally required to believe in order to make sense of it as cooperative

¹ Saul writes:

Despite his focus on speaker intentions, [Grice] wanted what is said not to be entirely subject of the whims of individual speakers. Instead, he defined 'saying' in terms of both speaker meaning and sentence meaning, and defined sentence meaning by generalising across speakers. ... Grice's inclusion of the audience in his definition of 'conversational implicature' serves a similar purpose. ... With conversational implicature, generalising across speakers would be inappropriate given the importance of context. Instead, he looked to the other participant in the conversation — the audience. (Saul 2002a, p.241)

in the context — where again this may not be the same as what the speaker or hearer actually *think* is required. Thus, if a sentence carries a certain implicature in a certain context X (where X includes the language used as well as other factors) when uttered by speaker A in the presence of speaker B, then it should carry the same implicature if uttered in X by speaker C in the presence of speaker D, and so on. What is conversationally implicated should not be dependent on beliefs and intentions specific to the individual speaker or hearer. This is not to say that *none* of the speaker's or hearer's beliefs are relevant. In its original form at least, Grice's conditions for the presence of a conversational implicature include that the speaker believes that the hearer can work out that the implicated message is required and that the hearer believes that the speaker is being cooperative. In the previous chapter I argued that, in the spirit of Saul's normative reading, we should remove or revise these conditions (Chapter 3, section 3). But even if we do not, they are general conditions for the existence of an implicature; the *content* of the implicature does not depend on further beliefs and intentions specific to the speaker and hearer in question.

In this part of the chapter I shall argue that conversational implicature is not in fact independent of the speaker's beliefs and intentions in this way, and thus that the concept of conversational implicature cannot play the proposed normative role. The overall argument is as follows. According to Grice, what is required to make sense of an utterance as cooperative can be inferred from the following items of information:

- (1) the conventional meaning of the words used, together with the identity of any references that may be involved;
- (2) the Cooperative Principle and its maxims;
- (3) the context, linguistic or otherwise, of the utterance;
- (4) other items of background knowledge; and
- (5) the fact (or supposed fact) that all relevant items falling under the previous headings are available to

both participants and both participants know or assume this to be the case
(Grice 1975/1989, p.31).²

That is, what an utterance implicates is what is required to make sense of it as cooperative, given premises concerning (1) conventional meaning and references, (2) cooperativeness, (3) context, (4) background knowledge, and (5) mutual knowledge. Davis refers to these premises as the ‘background constraints’ relative to which an implicated meaning is required (1998, p.63). Assuming that only one conclusion can be derived from the premises, this would give a speaker-independent standard of what is required and hence implicated. However, this assumes that the relevant premises grouped under (1) to (5) can themselves be identified without reference to the speaker’s beliefs, desires, and intentions, and this, I shall argue, is not so. I shall begin with item (4), background knowledge, which presents the biggest difficulty for Grice, and then look more briefly at items (1), (2), and (3).

The points that follow apply especially to particularized implicatures, which are more dependent on background information and contextual detail, though they apply in principle to all implicatures. (Some problems specific to generalized implicatures will be discussed in the next chapter.)

1.2 Background knowledge

I take it that the background knowledge employed in a Gricean calculation includes any information that is not specific to the utterance or its context but that is still necessary for interpreting the utterance. For instance, take the familiar example in which a philosophy professor devotes a reference letter to praising their student’s handwriting skills. This implicates that the student is a poor philosopher, but this conclusion cannot be derived simply from the premise that the professor is being cooperative, together with the conventional meaning of the words and the context. We need to add the background knowledge that handwriting is irrelevant

² I assume that Grice says ‘supposed fact’ under (5) because he requires only that the speaker should *believe* that the hearer can work out the implicature on the basis of the information listed, not that they can *actually* work it out. This was a point on which my revised version differed.

to philosophical ability. If good handwriting were in fact a highly reliable sign of philosophical ability, then the utterance would implicate that the student was a good philosopher, not a bad one. The same holds in many other cases. In general, inferences about contingent matters depend on a large number of background assumptions that are not made explicit in our reasoning. But since different sets of background assumptions will generate different inferences, this raises an important question: what is the relevant set of background assumptions for the generation of conversational implicatures? More precisely, since we are adopting a normative perspective, what is the *correct* set of background assumptions to use in working out what, if anything, an utterance conversationally implicates? The implicated content is the one that is *required* to make sense of the utterance as cooperative, but what are the background assumptions *relative to which* it is required? Given different background assumptions, different implicata will be required. (I speak of *assumptions* rather than *knowledge*, in order to avoid begging the question of whether these attitudes must be true.)

One suggestion is that the appropriate background is simply the truth — the set of all relevant true propositions. This would fit with the idea that what is conversationally implicated is speaker-independent. There is a problem with this suggestion, however. Consider the following exchange:

- (1) Al: Do you think Cally will realize she's been tricked?
Bea: She's not Einstein.

We would naturally interpret Bea as implicating that Cally is stupid and will not realize that she has been tricked (and let us assume that this is the interpretation that Bea intends). Einstein was a genius, and in saying that Cally is not like him Bea implicates, by understatement, that Cally is stupid. But now suppose that, unknown to everyone, the real Einstein was in fact of very low intelligence, and that all the mathematical and scientific work for which he is known was actually produced by someone else and was falsely presented as Einstein's. Then if Bea's utterance is to be interpreted in the light of what is in fact true, a very different interpretation would be required. If the relevant background assumption is that Einstein was of very low intelligence, then Bea's utterance requires us to suppose

that she thinks Cally is actually very smart and *will* realize that she has been tricked. Here is another example:

(2) Raj: Do you think the king will win the battle?

Sal: Is the earth flat?

Here Sal answers by asking a rhetorical question, implying that the answer to Raj's question is the same as — and as obvious as — the answer to it. But of course *what* answer is implicated will depend on what we think is the answer to the rhetorical question. If Raj and Sal are living in a prescientific society where everyone thinks that the earth is flat, then Raj will naturally interpret Sal as implicating that it is obvious that the king will win. (Again, let us assume that this is what Sal intended.) However, if the relevant background is the truth, then this interpretation would be wrong, and Sal ought to be interpreted as implicating that it is obvious that the king will not win.

Now in these cases the speaker would not actually count as implicating the unintended reading (that Cally is smart, that the king will lose), since clause (3) of Grice's definition ((II) of the revised version) would not be met.³ The speaker would not believe that the audience can work out that that reading was required, nor (on the revised version) would a typical hearer (contemporary to the speaker) be able to work out that it was required. So in these cases Grice would have to say that nothing is conversationally implicated, although the intended meaning is both utterer-implicated and audience-implicated. This is implausible, however. It would be inappropriate to base a pragmatic interpretation of an utterance on information unknown, not only to the speaker and the hearer, but to anyone in their society.

³ As a reminder, Grice's clause (3) and my revised version (II) run as follows:

(3) The speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required.

(II) It would be within the competence of a typical audience to work out that the supposition mentioned in (I) [= Grice's (2)] is required.

Such a normative standard would be far too strict. Indeed, if the normatively correct interpretation is the one that is *made available* to the hearer, then it should be the other reading (the intended one) that is the normative one. Bea's utterance makes available the information that Cally will not realize that she has been tricked, and Sal's utterance makes available the information that the king will win — and these utterances would make these pieces of information available to any other typical hearer in Bea's or Sal's societies.

(Note that the previous paragraph assumes that being able to work out that a certain reading is required involves being able to identify the appropriate background assumptions to use as well as being able to make the appropriate calculation from them. However, if working out the implicature simply involves being able to make the calculation itself, then Al and Raj (or any other typical hearer of Bea or Sal) could work out that the truth-based, *unintended* meaning is required, since they *would* be able to make the calculation *if* they were provided with the appropriate background information. In this case, clause (3)/(II) *would* be met and the utterances would carry the unintended implicatures. However, this would not make it any more plausible to think that truth is the relevant background for interpretation. If anything, it is even less plausible to claim that these utterances do in fact carry the unintended implicatures than to claim that they carry no implicature at all.)

These examples suggest that the appropriate background assumptions to use for interpretation are not those that are true but those that are *believed* to be true in one's community. This would mean that in the case of Bea and Sal the speakers' intended interpretations would be the normatively correct ones, since they are the ones that would be derived from the background beliefs current in their communities. However, there are further problem cases. For on many matters, different and conflicting background beliefs exist *within* a community, and these differences will affect the interpretation of utterances. Here is an example (suppose that Don and Ellie are both American citizens):

- (3) Don: Do you think Senator Bloggs took the bribe?
Ellie: Well, he's a Republican.

Here, how one interprets Ellie's utterance depends on what background assumptions one makes about Republicans. If one assumes that Republicans are corrupt, then Ellie's remark would seem to implicate that Bloggs did take the bribe. If one assumes that Republicans are honest, then it would seem to implicate that Bloggs did not take the bribe. (I say *seem* to implicate, since, again, we have to allow for the effect of clause (3)/(II) of the definition. Even if we agree on what meaning is required to make sense of the utterance, that meaning will not actually be implicated unless the speaker believes that the hearer can work out that it is required, or (on the revised version) unless a typical hearer could work out that it is required.) Given that both views are widespread in Don and Ellie's community (assuming that to be the community of American citizens), which is the correct view to use in interpreting Ellie's utterances?⁴

One option would be to say that since there is a diversity of relevant background views in Ellie's community, there is no unique proposition required to make sense of the utterance as cooperative, and hence no conversational implicature is present at all. Instead, there are just an utterer-implicature and an audience-implicature, corresponding to what the speaker and hearer think is required, and reflecting their personal background beliefs. But again, this seems too strict. Suppose that Ellie is well known to her friends (including Don) for her strong dislike and distrust of Republicans. Then Don will naturally interpret her as implicating that the Senator took the bribe. This interpretation would be particularly natural if Don shares Ellie's views about Republicans, but even if he does not, it would be the obvious one to adopt. If Don knows Ellie's views about Republicans, then he will assume a negative view of Republicans for the purposes of interpreting her comment, even if he is himself a Republican supporter. Moreover, this seems to be the *correct* interpretation, given Ellie's beliefs and Don's knowledge of them. Ellie has successfully made her opinion available to Don, and there has been no confusion or lack of attention on either side. Her

⁴ Keith Frankish has suggested another example to me, as follows. Writing a reference for a former student, a philosophy professor includes the comment 'Her impact on professional philosophy may be similar to that of Ludwig Wittgenstein.' Depending on what we think of Wittgenstein, this might be taken as either the highest praise or an accusation of charlatanism.

opinion is not only utterer-implicated and audience-implicated, but conversationally implicated. There will be many cases like this, where different people within a community have different background beliefs or make different value judgements, each of which would generate a different implicature.

Note that Grice's condition (3) may also be met in this case. We may suppose that Ellie believes, correctly, that Don knows her background beliefs and can work out that the anti-Republican reading is required. The revised version of this condition, (II), might also be met. This says that a typical hearer could work out that the anti-Republican reading is required. Whether or not this is the case depends on what we mean by 'typical'. But since the aim of the revised condition was to capture our intuitions about implicature, the fact that we have an intuition that there is a conversational implicature in Ellie's case indicates that we should read 'typical' in a way that is compatible with this. For example, we might say that a typical hearer is one that has (among other things) the sort of familiarity with the speaker's attitudes that the speaker expects them to have — in other words, one that is a member of the community to whom Ellie might address this remark. Thus, we might say that a typical hearer for an utterance is one that is (a) alert, averagely informed, linguistically competent, and (b) a member of the community to whom the speaker might address the utterance (with the particular communicative intentions they have on this occasion). In short, a typical hearer is a competent potential addressee. In some cases, where the interpretation of an utterance depends on specific assumptions shared by few people, the pool of potential addressees might be very small — perhaps including only the person actually addressed. In such cases, an implicature might serve the function of a private code.

The Ellie example suggests, then, that the appropriate background beliefs for derivation of a conversational implicature are those of the speaker. It might be objected that this is covered by the suggestion made earlier in response to the Al and Raj examples — namely, that the appropriate background beliefs are those that dominate in the speaker's community. In this case (the objector may say), the relevant community is that of people who think that Republicans are corrupt. However, this does not remove the speaker-dependency, since the relevant community has to be identified by reference to the attitudes of the speaker. The reason we pick out the people who believe Republicans are corrupt as the relevant

community is that that belief is the relevant background assumption made by the speaker. This is speaker-independency in name only, since it holds only between people who share the same (relevant) beliefs as the speaker.

I think this is nearly right, but I want to make a modification, which gives priority to the speaker's *intentions* rather than their beliefs. Suppose Don is the father of Ellie's boyfriend, and that Ellie is meeting him for the first time. She knows that Don is a committed Republican, and, although she herself deeply distrusts Republicans, she is anxious to make a good impression and not to offend Don. So she conceals her real views, goes along with Don's comments on political matters, and intends her answer to his question about Bloggs to be understood in the light on Don's beliefs, not her own, and thus to implicate that the Senator did not take the bribe. Intuitively, this would seem the correct way to interpret Ellie's utterance. If she purposely conceals her own background beliefs about Republicans, then she has not *made available* her belief that Bloggs took the bribe. (This would remain true even if Don realizes that Ellie is concealing her real views. We know that what a person makes available to us may not be what they really believe.) This suggests, then, that the relevant background assumptions to use in calculating an implicature are those that the speaker *intends* to be used — where these will often, but not always, be ones the speaker actually holds.

If all this is right, then conversational implicature will not be speaker-independent, since in order to work out an implicature we shall need to know what background assumptions the speaker intends us to draw on. It may be objected that this is not a problem for Grice, even on a normative reading. For the speaker's intentions regarding the appropriate background assumptions to use could be treated as part of the *context* for their utterance. Then speaker-independency will be preserved, since all hearers who know the context will, by definition, make the same background assumptions. However, like the previous objection (that the relevant background assumptions are the ones that predominate in the speaker's community), this preserves speaker-independency only in name. On this view, implicatures are speaker-independent relative to a context, but contexts themselves are not speaker-independent, since they are defined by reference to the speaker's intentions. There will be agreement on what is implicated only between people

who share the same (relevant) intentions as the speaker, and bundling these intentions into the context does not change this.⁵

Finally, note that the background knowledge cases discussed above are similar to the Suleiman case discussed in the previous chapter, in which different interpretations are required depending on what background information the hearer has about Suleiman (that he is a good conversationalist, that he is a skilled masseur, and so on). The main difference is that in that case it was a matter of choosing from a range of different but compatible background beliefs (or choosing their disjunction), whereas in the cases discussed here the choice was between incompatible background beliefs. In the Suleiman case, I argued that we would have to say that nothing was implicated, since there were no grounds for picking out one interpretation over another, and the disjunction of all the possible interpretations was ruled out by clause (3). (The speaker did not believe that their hearer could work out that the disjunction was required.) However, given the discussion above, another option offers itself: namely, that the correct background information to apply is that which the speaker intends their hearer to apply. Assuming that the speaker believes their audience can detect these intentions (or, on the revised version, that a typical hearer could detect them) this would mean that there is an implicature in such cases after all.

I shall look at some consequences of speaker-dependency later in this section, and discuss a general objection to it. First, however, I want to look at other inputs to the Gricean calculation process and argue that these are speaker-dependent too.

1.3 What is said

The first item of information Grice lists as entering into the process of calculating what an utterance implicates is the conventional meaning of the words used and any references. This, fixes ‘what is said’ by the utterance, which Grice identifies with what is directly communicated — the proposition that is judged true or false

⁵ Note that if we were to regard information about speaker intentions with regard to background assumptions as part of the context of the utterance, then it becomes uncontroversial that condition (II) is met in cases like Ellie’s. A typical hearer could work out such implicatures, since knowledge of the context *including the speaker’s intentions* is taken as a given in the working-out process.

and that is the starting point for the calculation of anything that may be indirectly communicated by implicature.

There are several issues here. First, as Davis points out, it is not enough to know the conventional meanings of the words used and their referents; as interpreters we need to know their ‘applied’ meaning — their meaning on this particular occasion of use (Davis 1998, p.64). If a word has more than one conventional meaning, we need to know which one is relevant on this occasion. (Grice acknowledges that disambiguation is required in order to fix what is said; Grice 1975/1989, p.25.) However, on Grice’s own view, applied meaning is fixed by *both* conventional (sentence) meaning and speaker meaning (Grice 1968, 1969, both reprinted in 1989). There is, therefore, no completely speaker-independent way of fixing what is said.

Davis extends this thought further, suggesting that even disambiguated conventional meanings are not decisive for interpretation. A speaker may misspeak (for example, saying ‘coroner’ for ‘corner’), or use words in an unconventional way. Provided the hearer understands what the speaker means by the words on this occasion, Davis argues, the utterances in question could still generate conversational implicatures. In such cases, it would be speaker meaning that forms the premise for implicature calculation, not conventional (sentence) meaning (Davis 1998, pp.64–5).

A further problem for Grice arises from his identification of what is said with what is directly communicated. Conventional meaning (even when disambiguated) and referents determine only a minimal form of explicit content which does not correspond to our intuitive understanding of what is directly communicated. Here is an example (borrowed with simplification from Carston and Hall 2012)

(4) Max: How was the party? Did it go well?

Amy: There wasn’t enough drink and everyone left early.

Here what Amy *says* (in Grice’s sense) is simply that there was not enough drinkable liquid and everyone left early — which would be false if there had been plenty of lemonade and if someone somewhere did not leave early. But this is not

what Max will take Amy to be directly communicating. Rather, he will understand her to be saying something like:

- (5) There wasn't enough alcoholic drink to satisfy the people at the party and so everyone who came to the party left it early.

This is the proposition Max will judge to be true or false, and it is from this that he will work out the implicated answer to his question — namely that the party did not go well. Cases like this are very common (for more examples, see Carston and Hall 2012).

The process of filling out the explicit content of an utterance in this way is called 'explicature', a term coined by Sperber and Wilson (1995), and relevance theorists argue that it is a pragmatic process, involving the application of general communicative principles like those involved in the derivation of implicatures (see, for example, Carston 2004a; Hall 2008; Recanati 2002; Wilson and Sperber 2002). This view is controversial, and other theorists argue that explicature is a semantic process involving the filling in of hidden indexicals, demonstratives, and variables present in the logical form of the utterance (for example, Stanley 2000). This is a large and complicated debate, which is not directly relevant to my main topic, and I shall not discuss it here. (The relevance theory approach to implicature recovery and the relation between explicature and implicature will be discussed more in Chapter 5.) Rather, I simply want to note that this gives us yet another reason for doubting that implicature is speaker-independent. If contextualized pragmatic processes are involved in explicature, then they may draw on information about the speaker. And if so, then even the basic step of establishing the explicit content that forms the starting point for implicature calculation will not be speaker-independent.

1.4 Cooperativeness

Another premise in a Gricean calculation is that the speaker is being cooperative, following the CP and its maxims. Moreover, there must be a further judgement to the effect that what the speaker says is *not* cooperative — that the utterance cannot be taken at face value, consistently with the premise that the speaker is being

cooperative. But what counts as being cooperative? A normative account of implicature will need a standard of cooperativeness to guide our judgements on this matter. Grice spells out cooperativeness in terms of adherence to the maxims of Quantity, Quality, Relation, and Manner, but again these need interpretation, and I shall argue that we cannot determine whether an utterance meets them without taking into account the speaker's attitudes. I shall consider each maxim in turn.

The first is the maxim of Quantity, which says that speakers should be as informative as is required and no more informative. But what amount of information is required in any given situation? There is no clear speaker-independent standard, and the same sentence might be regarded as either sufficiently or insufficiently informative, depending on who uttered it. Consider a reference letter case. Suppose a hiring committee receives a reference letter for Rebecca, an applicant for a philosophy lectureship, which consists simply of the following sentence:

(6) Rebecca is a good philosopher and I recommend her to you.

Does this violate the maxim of Quantity? It is plausible to think that the answer depends on who the committee think wrote it. If they think it was the famously uncommunicative Professor A, who rarely agrees to write reference letters at all and never writes more than a few words, then they will probably see it as providing as much information as they can expect and take it at face value. On the other hand, if they think it was written by Professor B, who usually writes detailed reference letters running to several pages, then they will regard it as uncooperatively short and probably read it as implicating that Rebecca is not a very good philosopher and is not really being recommended. This indicates that cooperativeness with regard to information quantity is speaker-relative, and that hearers may need to consider the speaker's attitudes and intentions in order to judge whether it is being violated.

The next maxim is the maxim of Quality, which says that speakers should not say things they believe to be false or for which they lack adequate evidence. The speaker-dependency here is obvious. Since the maxim makes reference to what the

speaker believes and what evidence they have, there is no speaker-independent way of determining whether an utterance violates it. If the speaker believes that *p* is false or if they have inadequate reasons for believing *p*, then they are violating the maxim of Quality if they say that *p*, even if *p* is true and objectively well supported by evidence. (I assume that what matters is the evidence the speaker actually *possesses*. Even if there is good evidence that *p*, the speaker violates the maxim of Quality if they are not aware of the evidence.)

There is also room for speaker-dependency in the notion of adequate evidence. Different speakers may have different conceptions of what counts as adequate evidence (different evidential standards). Some people are quicker to jump to conclusions than others. It is not clear that there are any independent norms available in this area, and even if there were, it is the speaker's own norms that matter in determining whether they are respecting the maxim of Quality. If I assert that aliens exist on the basis of evidence that is in fact flimsy but that I regard as adequate, then it is plausible to think that I have not violated the maxim and should not be taken to be implicating something. But if a speaker with the same evidence but much stricter evidential standards says the same thing, then they would be violating the maxim and might be appropriately interpreted as being ironic.

Next consider the maxim of Relation, which says that contributions should be relevant. Now, as we saw in section 1.2 above, what counts as relevant to a conversation varies with background assumptions. A comment about a student's handwriting might be relevant to a philosophical reference if handwriting is taken to be a good indicator of philosophical ability. Given the right background assumptions, any comment could be relevant to any exchange. So determining whether an utterance violates the maxim of Relation involves determining what background assumptions to apply, and, as we saw, this involves considering the speaker's intentions.

Moreover, even given fixed background assumptions, there is still room for judgements of relevance to vary depending on who the speaker is. For example, consider this exchange:

- (7) Freda: Why should we vote for Smith?
Jack: She has a strong policy on protecting frogs.

Whether Jack's reply is relevant depends on Jack's attitude to frogs. If Jack is an animal-lover obsessed with protecting amphibians, then his reply is a relevant response to Freda's question under its literal meaning. Jack is being straightforwardly cooperative, and there is no reason to think he is implicating anything. But if Jack cares nothing for wildlife and regards frog protection as a trivial issue, then his reply is not relevant under its literal meaning and implicates that there is no good reason to vote for Smith. Thus, whether or not there is an implicature in this case depends on the speaker's attitudes.

In other cases, the speaker's attitudes may determine an utterance's *degree* of relevance or irrelevance, and thus *what* it implicates. Consider:

- (8) Ed: Do you like this shirt?
Joy: It has lovely buttons.

Here Joy's reply is not directly relevant to Ed's question, and its lack of relevance signals that Joy is implicating something. However, *what* Joy is implicating depends on how relevant remarks about buttons are, which in turn depends on Joy's attitude to buttons. If she prizes buttons and judges items of clothing by the quality of their buttons, then she is praising a relevant feature of the shirt and implicating that she likes it. However, if she regards buttons as trivial, then she is praising an irrelevant feature and implicating that she dislikes it.

Note that the different interpretations in these cases depend, not on different background assumptions about animal welfare and clothing accessories, but on different values or preferences. Joy need not think that buttons are objectively important, but just have strong personal feelings about them. There will be many cases like this, where an utterance's relevance to a conversation depends on what the speaker values. In the absence of objective standards of value on the matters in question, there will be no speaker-independent way of interpreting these utterances.

Finally, similar issues arise with the maxim of Manner, which includes submaxims requiring speakers to avoid obscurity and ambiguity and to be brief and orderly. The problem again is that there are no clear speaker-independent

standards on these matters. For example, what counts as obscurity? It will depend on context: wording that would be obscure in a teacher's comment on a third-grade pupil's work might not be obscure in a supervisor's comments on a PhD thesis. Moreover, there is room for different styles within each type of context. For example, what is the correct type of language to use in a reference letter? Suppose a reference letter for Rebecca, a philosopher, concludes with the following statement:

- (9) I can sum up my view of Becks by saying that she is one cool philosophy dude.

Does this violate the maxim of Manner by using obscure, slang terms and phrases ('Becks', 'one' (for 'a'), 'cool', 'dude')? Whether the hiring committee reading the letter think so plausibly depends on who they think wrote it. If they think it was Professor C, who likes to present himself as youthful and hip and usually writes in an informal style reflecting his (rather dated) ideas of youth culture, then they will probably take the comment at face value, as expressing high praise of Rebecca. If they think it was written by the conservative Professor D, who usually writes in a very formal style, then they will probably read it as deliberately flouting the obscurity submaxim and implicating that the professor does not take Rebecca seriously as a philosopher. (We can also imagine the reverse case in which a formal and old-fashioned reference would implicate a negative evaluation if written by Professor C but not if written by Professor D.) Thus, as in the case of Quantity, cooperativeness with regard to Manner is speaker-relative, and judgements on the matter will depend on facts about the speaker's attitudes and intentions.

Similar points could be made with respect to the submaxims relating to orderliness and brevity. Many different styles of speaking and writing are acceptable (compare the styles of different novelists), and what counts as a violation of brevity and orderliness in one writer or speaker might be perfectly normal in another.

1.5 Context

Another premise in the Gricean calculation is a description of context of the utterance. Again, I shall argue that this cannot be given in a speaker-independent way. I want to focus specifically on the *linguistic* context — the topic and purpose of the wider communicative exchange to which the utterance being interpreted belongs. Identifying this is crucial for establishing whether the Cooperative Principle is being followed at the level of what is said or what is implicated. For example, the sentence ‘Fred has excellent handwriting’ implicates that Fred is a poor philosopher only if uttered in the context of a request for an assessment of Fred’s philosophical abilities, and then only if it has a prominent position in the assessment (added at the end of a reference letter, following a long series of highly positive comments about Fred’s philosophical abilities, it would not generate the implicature).

But how do we determine the context for an utterance? We cannot simply read it off from the words uttered. Consider the following exchange between Sally and Sarah, academics who have just listened to a talk by a colleague Phil:

(10) Sally: What did you think of Phil’s talk?

Sarah: It was OK. By the way, did I tell you about the new air freshener I bought?

Is Sarah’s second sentence part of her response to Sally’s question or is she changing the subject? If it is part of her response, then it must be interpreted as carrying an implicature, to the effect that Phil’s talk ‘stank’. But if it is a new topic, then it can be taken literally. It is hard to see how to decide without appealing to facts about Sarah and her intentions. We need to know whether she meant to change the subject with her second sentence, and to decide this we need to know more about what Sarah thinks of Phil and his philosophical views, how she reacted during his talk, whether she has a particular interest in air-freshening products, and so on. And this means that linguistic context cannot be determined in a speaker-independent way. Again, it might be suggested that we could treat all this information about Sarah as included within the context for the utterance, but,

again, this would preserve speaker-independence in name only. The nature of what was implicated would still vary from speaker to speaker.

2. Responding to the argument

I have reviewed all the premises for a Gricean calculation — the background constraints relative to which the implicated meaning is required — and argued that none of them can be established in a completely speaker-independent way. I will consider the consequences of this in a moment, but first I want to deal with a possible objection.

2.1 Resisting speaker dependency

Griceans might try to resist the argument for speaker-dependency by claiming that in the cases considered in the previous section the speakers' attempts at implicature simply fail. Because their utterances do not meet speaker-independent standards of interpretation, the speakers do not identify unambiguous implicata and thus either fail to implicate anything or implicate open-ended disjunctions. We are simply wrong to think that there are determinate implicatures in these cases. Our intuition that something specific is being implicated may correspond to an utterer-implicature or an audience-implicature, but there is no specific conversational implicature. This is counter-intuitive, but if the Gricean theory is a normative one, then we should expect that it will correct some of our everyday judgements.

I do not think this reply is satisfactory. First, it would mean that Saul's notions of utterer-implicature and audience-implicature will have a lot of theoretical work to do, and, as I shall argue in section 3 below, it is doubtful that they are up to the job. Second it means that Grice's theory becomes a radically revisionary one. Grice's aim was, I take it, to analyse the everyday phenomenon we have in mind when we talk of a speaker conveying something indirectly — implying, suggesting, indicating, or meaning one thing by saying another (Grice 1989, p.86). But on the view just proposed the Gricean notion of conversational implicature is very strict, and there will be many cases where the Gricean account of what, if anything, a speaker is implicating is very different from what we would naturally take them to be indirectly implying (suggesting, indicating, meaning). Where we detect a clear indirect meaning, Gricean theory often tells us we should see none

at all, or only a much weaker, disjunctive one. (In many cases, these disjunctive meanings would be *very* weak. For example, given different background assumptions, Ellie's reply in (3) could be interpreted as indicating either that Senator Bloggs did take the bribe or that he did not, so if the implicature is the disjunction of these alternative interpretations, then it is completely uninformative.) Thus, Gricean theory requires us to revise or abandon a wide range of everyday interpretations and replace them with ones that are far less rich and informative. It is hard to see what value such a normative theory has.

Finally, and most importantly, the objection misdescribes the problem. It is not that in the cases we considered in section 1, the utterances *failed to meet* speaker-independent standards. Rather, the cases showed that there *are no such standards*. There are no speaker-independent norms for establishing the correct premises to use in a Gricean calculation. In the case of background assumptions, we saw that truth cannot serve as the standard, and community-wide belief will not do either, since there is no speaker-independent way of identifying the relevant community. Similarly, in the case of cooperativeness and the associated maxims, there are no objective standards governing the amount of information required, the truth or probability of the content, the relevance to the topic, and the manner of expression. As we have seen, the same words may be judged cooperative when uttered by one speaker and uncooperative when uttered by another, and there is no clear speaker-independent standard to which we could appeal to correct these judgements. Similar points hold for linguistic context and explicit (directly communicated) content. These cannot simply be read off from the words used, and fixing them involves reference to the speaker's attitudes.

Griceans might respond that even if there are no speaker-independent standards for determining background constraints, it does not follow that all implicatures are speaker-dependent to any significant degree. First, *generalized* implicatures are not sensitive to context at all and so should not be speaker-dependent. Second, even in the case of particularized implicatures, it will often be obvious what the relevant background constraints are from the non-psychological context, without considering the attitudes of the particular speaker.

I do not think these points are very strong. I will consider generalized implicatures in the next chapter, where I will argue that these are more context-

sensitive than Griceans suppose. But note that even if speaker dependency holds only for particularized implicatures, this still involves giving up the idea that Gricean theory provides global speaker-independent norms of implicature that cover both particularized and generalized varieties. As for the second point, it is not clear that it is correct. Although it may frequently be obvious what the appropriate background constraints are, this may be because it is obvious what the relevant psychological attitudes of the speaker are (perhaps because they are widely shared in the speaker's community), rather than because these attitudes are not relevant. At any rate, the examples used to argue for speaker-dependency in section 1 were familiar, everyday ones, and many more could have easily been offered.

2.2 Consequences of speaker-dependency

Suppose we accept, then, that implicatures are not speaker-independent. Although the notion of *what is required* to make sense of an utterance as cooperative appears to offer a speaker-independent standard, it turns out that the background constraints relative to which an interpretation is required cannot themselves be identified without reference to the attitudes of the speaker. What are the consequences of this?

First, it weakens the normative role of the Gricean framework. As we have seen, the framework is best understood as aiming to establish conditions for implicature that are independent of the speaker's attitudes and intentions. If the arguments in section 1 above are sound, then it fails to do this. Although the framework still establishes conditions for the presence of an implicature, these conditions are not independent of the speaker. Facts about the speaker and their mental states are an essential part of the background for a Gricean calculation. The notion of what is required to make sense of an utterance as cooperative still provides a standard for implicature, but it is a speaker-dependent one. Quite different interpretations may be required to make sense of two utterances of the same sentence in the same context made by different speakers. This is a much weaker normative standard.

Second, if, as I argued, a speaker's intentions partially fix the background assumptions with respect to which utterances are to be interpreted, then speakers

can at least partially *control* what is required to make sense of their utterances as cooperative.⁶ Consider an example from Davis. Carl says ‘I am sick’ and Diane replies ‘A flying saucer is nearby’ (Davis 1998, p.74). Although there is no single interpretation required to make sense of this bizarre comment, Davis claims that Diane may nonetheless be implicating something specific. For she may mean to convey a specific message — say that Carl can get help from the doctors on the flying saucer — and on Davis’s view for a speaker to implicate something is simply for them to mean to convey it by saying something else (where meaning to convey something is a matter of having appropriate communicative intentions) (Davis 1998, p.4–5, p.114, p.122, p.130, 2007, p.1661).⁷ Saul objects that it is counter-intuitive to think that speakers possess such a degree of authority over what they implicate, and uses this intuition to support the normative reading of Grice we have been considering (Saul 2001, p.633; 2002, pp.240–1). However, if a speaker’s intentions determine which background assumptions to use in interpreting their utterances, then it might seem that Diane’s utterance could carry the implicature she intends, even on Gricean principles. Suppose Diane intends her utterance to be interpreted in the light of the following background assumptions: (a) Flying saucers are real and often visit Earth, (b) Flying saucers carry alien doctors, (c) Alien doctors are able and willing to cure human sickness. Then, given

⁶ Not all the background constraints that are speaker *dependent* are under speaker *control*, of course, since speakers may not have control over the relevant facts about themselves. For example, a speaker violates the maxim of Quality in saying that p only if they believe that p is false or have (what they regard as) inadequate evidence for p. These are facts about the speaker, but not ones over which the speaker has any direct control. But as we saw earlier, some background constraints, in particular background assumptions, are plausibly fixed by the speaker’s *intentions*, which are under their control.

⁷ This is how Davis defines what he calls *speaker implicature*, which he identifies with Grice’s particularized implicature (Davis, 1998, p.21). Davis also argues for the existence of what he calls *sentence implicatures* (corresponding to Grice’s generalized implicatures), which depend on conventions of use within a language community. I shall discuss Davis’s views about sentence implicature in the next chapter.

this background, the supposition that Diane believes that Carl can get help on the spaceship *is* plausibly required to make sense of her utterance as cooperative.

This is too fast, however. Given Diane's intentions with regard to background assumptions, the supposition that she is expressing the belief that Carl can get help on the spaceship is required to make sense of her utterance as cooperative, so clause (2) of Grice's definition (= clause (I) of the revised definition) is met. But there is still clause (3) (= revised clause (II)) to consider. On the original version this says that the speaker believes that the audience can work out that the supposition in question is required.⁸ But, as I argued in the previous chapter (section 3.2), this does not provide a strong normative constraint. Diane may believe, wrongly, that Carl shares her beliefs about flying saucers and alien doctors, and thus that he can work out that she is implicating that he can get help on the saucer. However, I proposed a revised version of this clause (II), to the effect that a typical audience can work out that the supposition in question is required — where a typical hearer is a competent potential addressee for the utterance in question (an alert, averagely informed, linguistically competent, member of the community to whom the speaker might address the utterance). Since speakers cannot control whether or not a typical audience can work out what

⁸ As a reminder, Grice's conditions (2) and (3) and my conditions (I) and (II) are as follows:

(2) The supposition that [the speaker] is aware that, or thinks that, *q* is required in order to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with this presumption.

(3) The speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required.

(I) The supposition that the speaker is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with the presumption that he is observing the CP.

(II) It would be within the competence of a typical audience to work out, or grasp intuitively, that the supposition mentioned in (I) is required.

supposition is required to make sense of their utterances as cooperative, this limits speakers' control over what they can implicate. A speaker may have intentions that dictate that a certain reading of their utterance is required to make sense of it as cooperative, but the speaker does not *implicate* that reading unless a typical audience can work out that it is required. Thus, depending on the details of the situation, Diane may or may not succeed in implicating what she intends. If she thinks, mistakenly, that everyone shares her eccentric views about flying saucers and alien doctors, then she will fail to implicate it. However, if she and Carl are members of a UFO cult who all share similar beliefs about flying saucers, and she takes this for granted in making her remark, then a typical hearer (in the specified sense) would be able to work out her intended meaning and she would succeed in implicating it. Thus, on this view there remains a significant constraint on speakers' power to implicate.

2.3 An intention-centred account of implicature

Is the normative element provided by clause (II) sufficient to save the Gricean framework? I don't think so — at least not in its traditional form. In fact, I think the moral of the discussion is that a theory of implicature can and should give a greater role to speaker intentions.

On the surface, the Gricean framework gives no role to speaker intentions in fixing what utterances implicate (though, as argued in the previous chapter, it may allow a role for intentions in determining what *speakers* implicate). But if the arguments above are correct, then there is a hidden indirect role for speaker intentions in fixing what utterances implicate, since implicatures are sensitive to facts about speakers' intentions with regard to the background constraints on interpretation (background assumptions, cooperative standards, and so on). But, once this is admitted, then the question arises of why speaker intentions should not play a greater role. Why not allow that speaker intentions can directly fix what is implicated, as Davis holds? This would immediately resolve many of the problem cases discussed in Chapter 3, where considerations of calculability could not determine whether or not a speaker was implicating something (as with Candy and the Tarantino movie) or provide a non-disjunctive account of what they were

implicating (as with Finn's Suleiman remark). For the speaker may have clear communicative intentions that settle the matter.

As noted, Saul rejects Davis's view on the grounds that it gives speakers too much power over the nonconventional content of their utterances — power they do not have over their utterances' conventional content. (Compare the Humpty Dumpty theory of implicature mentioned in Chapter 1.) As Saul puts it, 'What Grice's theory gives us and Davis' does not is the idea that what is implicated is not wholly up to the speaker' (Saul 2001, p.633). But, as we have seen, the notion of what is required (on Gricean principles) to make sense of an utterance as cooperative does not in itself provide a speaker-independent standard, since a speaker's intentions play a crucial role in the calculation, by establishing the relevant background assumptions, cooperative standards, and so on. The normative element comes in only with the requirement that a typical hearer can *work out* what is required. But a similar constraint could also be incorporated into an intention-based account like Davis's. We might say that a speaker S who says that p implicates that q if (a) S means to convey that q by saying p, and (b) a typical audience could work out that S means to convey that q. (Again, a typical audience is a competent potential addressee, where this might be someone with considerable background knowledge of S's attitudes.)⁹

⁹ It is arguable that a constraint similar to Grice's condition (3) is actually implicit in an intention-based account of implicature — at least given a Gricean view of speaker meaning. According to Grice, for S to mean that q is (in essence) for S to intend to get their hearer to believe that q by recognizing S's intention to get them to do so. But one cannot seriously intend to do something unless one believes one has at least a chance of success. We cannot seriously intend to do things we think are impossible. So if S intends to get their hearer H to recognize their intention to communicate q, then S must at the very least think it is possible for H to recognize it, and thus that it must be possible for H to see the connection between what they say and what they are trying to implicate. Davis himself stresses this point:

On my view, S means or implies I by uttering Σ only if S utters Σ with the intention of providing an indication that he believes I ... Because intention implies expectation, S must have some expectation that uttering Σ will provide

Thus we can avoid the Humpty Dumpty theory. Even if speakers' intentions can directly fix what they implicate, speakers cannot implicate whatever they like. They must ensure that they make their communicative intentions clear. Borrowing a term from Saul, we might say that they must *make* their intended meaning *available* to their audience. They might do this in various ways. They *could* rely on Gricean mechanisms. If the meaning they intend to convey is the one that is obviously required (given the relevant background constraints) to make sense of their utterance as cooperative, then they have made it available. But this is not the only means they might use. They might indicate their intended meaning by tone of voice, expression, gestures, and other non-verbal cues. They might rely on their hearer's knowledge of their beliefs, intentions, preferences, dispositions, and conversational habits. (Remember that a typical audience in our sense may be one that is well acquainted with the speaker.) They might know that they have a rapport with the hearer, which enables him or her to pick up their communicative intentions intuitively from numerous subtle cues. They might, at the extreme, simply *tell* the hearer what they mean. Or they might rely on a combination of these and other means.¹⁰

On this view, then, the core normative constraint on speakers is that they make their intended meaning available *by some means*. The Gricean framework is just one particular account of how a content might be made available, and calculability is no longer a necessary condition for implicature.¹¹ It is true that this move makes

an indication that he believes I. Unless S is psychotic, S will have such an expectation only if S perceives some connection in the context of utterance between I and the proposition E literally expressed by Σ . (Davis, 1998, p.186)

¹⁰ They might also, perhaps, rely on the hearer's knowledge of specific interpretative principles or implicature conventions, which support generalized implicatures. I will discuss this in the next chapter.

¹¹ Compare Davis:

[A] speaker can make her implicature available if she creates a context in which there is enough evidence available to her conversational partners to give them a reasonable chance of figuring out what she has implied on the basis of what she

our account of implicature more vague, but, arguably, this reflects the open-ended nature of the phenomenon itself. Implicature-based communication is complex and subtle and cannot be captured in a precise definition, even one as clever as Grice's. (I shall suggest in Chapter 6 that the difficulty we have in providing a precise theory of implicature tells us something about the social function of implicature.)

I think this is on the right lines, but I want to make an important modification. Above, I said that a speaker *S* implicates that *q* by saying that *p* if (a) *S* means to convey that *q* by saying that *p*, and (b) a typical hearer could work out that *S* means to convey that *q*. However, this formulation states conditions for the *speaker* to implicate and does not allow for the possibility (discussed in Chapter 3, section 3.3) that an *utterance* may implicate something that the speaker does not mean. (I offered the example of Rita, who, in her eagerness to say what a wonderful human being Omar was, forgot to mention his philosophical abilities and thereby implicated that he was a poor philosopher.) If we are to allow for unmeant implicatures, we need to make a distinction between what a speaker implicates and what their utterance implicates, and the latter must thus be independent of (or at least not wholly determined by) the former. As we saw in the previous chapter, with the modifications I proposed, the Gricean framework can allow for this: We can say that an utterance *U* implicates that *q* if the supposition that the speaker thinks that *q* is required to make sense of *U* as cooperative, and that a speaker *S* implicates that *q* if *S* means that *q* and produces an utterance that implicates that *q*. If we want to make a similar distinction on our intention-centred account, then we cannot say that what an utterance implicates is fixed by what the speaker means to convey (even if what they themselves implicate is fixed by that).

said. Knowledge that speakers generally observe the Cooperative Principle and the maxims may be part of this evidence. But that knowledge may be overridden ... Moreover, the evidential base may be completely different. (Davis, 2007, 1668–9)

For more on the range of factors that may be involved in the recognition of an implicature, see Davis 1998, pp.127–31.

I think we can do this, by appealing again to the notion of *making a meaning available*. Let us say that an utterance U *makes q available* (where q is not the literal meaning of U) if a typical audience (in our sense) would identify q as the (or a)¹² intended meaning of U, inferring this from the literal meaning of U, together with any or all of the indications mentioned earlier (Gricean considerations, non-verbal cues, knowledge of the speaker's attitudes and habits, personal rapport, and so on).¹³ Then we can say that an utterance implicates q if it makes q available in this way. In short, what an utterance implicates is what a typical hearer (in our sense) would *judge* the speaker to mean to convey by saying what they do, where this may differ from what the speaker actually meant to convey by saying it. The implicated meaning is the non-literal meaning the speaker *appears* to intend.¹⁴

¹² I say 'or a' since the hearer may also regard the utterance's literal meaning as an intended meaning of the utterance.

¹³ This definition of making available is restricted to non-literal meanings. We might also speak of utterances making literal meanings available: A literal meaning is made available if a typical audience would interpret it as the speaker's intended meaning. If we do this, then it will of course be crucial to distinguish the indirect making available involved in implicature from the direct making available involved in literal communication. It might seem that we can do this by saying that implicature-carrying utterances make their implicated meanings available by making their literal meanings available. However, if 'made available' means 'would be interpreted as intended', then this will not do, since the literal meaning of an implicature-carrying utterance would often *not* be interpreted as intended, as in cases of irony. It would be better to say that an implicature-carrying utterance makes its implicated meaning available by *having* another meaning, which itself may or may not be made available. In the text, 'made available' should always be understood in this way, to mean made available *via another meaning*.

¹⁴ In a fuller treatment, this might need some refinement. For example, we might want to allow for cases where a hearer judges that an utterance implicates something while at the same time realizing that the speaker does not really intend to convey it. (In the Rita example, Donald might realize that Rita's admiration for Omar's general goodness has led her to misjudge the content of her reference.) Many such cases will be ones where the speaker is not paying full attention to what they are saying, so one option would be to say that what an utterance implicates is what a typical

On this view, then, the speaker's communicative intentions are central to what an utterance implicates, since evidence about them will affect how a typical hearer would interpret the utterance. (On the Gricean view, by contrast, such evidence is irrelevant.) However, if that evidence is misleading (if the speaker has given misleading indications of their meaning), then the meaning their utterance makes available, and thus implicates, may differ from what the speaker meant to convey. (In Saul's Roland case, for example, the professor's letter implicated that Roland was a poor student since it gave clear indications that that was what the writer meant to convey, even though she did not in fact want its readers to form that belief.) Similarly, if the speaker fails to give adequate indication of their intended meaning, then their utterance may fail to make any (non-literal) content available and thus fail to implicate anything at all. Thus, we can say that an *utterance* implicates that *q* if it makes *q* available, in the sense just described, and that a *speaker* implicates that *q* if they mean that *q* *and* produce an utterance that implicates that *q*.

To sum up, then, on our intention-centred account (conversational) implicature can be defined as follows:

In saying that *p*, a speaker *S* implicates that *q* if (a) *S* means to convey that *q* by saying that *p*, and (b) *S*'s utterance makes *q* available (in the sense defined earlier, relative to a typical audience).¹⁵

hearer would judge the speaker to intend to convey if they thought the speaker was being fully attentive.

¹⁵ Davis suggests a similar account. Responding to Saul's objections, he stresses that his notion of (speaker) implicature is a purely descriptive one, and that we may also need a normative notion, *proper implicature*. On Davis's view, a person implicates that *q* if they mean or imply *q* by saying something else, but they properly implicate that *q* only if they also fulfil their communicative responsibilities by making their meaning available to their audience (Davis 2007, pp.1662–3).

This preserves the parallel between saying and implicating, both of which involve a psychological condition and a normative condition, as set out in the table below (a revised version of Table 1, updated to reflect our intention-centred definition).

	Act	Psychological condition	Normative Condition
Conventional meaning	S says that p	S means that p (speaker meaning)	S uses a sentence that means p (sentence meaning)
Nonconventional meaning	S implicates that q	S means that q (speaker implicature)	S produces an utterance that makes q available (utterance implicature)

Table 2: Saying and implicating according to the intention-centred account of conversational implicature.

I want to emphasize that, although this account retains a normative component and allows for unmeant implicatures, it does not establish *speaker-independent* norms of implicature. As I explained, in interpreting an utterance a typical hearer may draw on information about the particular speaker (their beliefs, conversational habits, and so on), and such information may therefore determine what the utterance makes available, and so implicates (which may, however, be different from what the speaker actually meant). Thus, the same sentence may generate different implicatures in the same context when uttered by different speakers. The account is, we might say, only *weakly* normative. Given the problems raised earlier in this chapter, I doubt that it is possible to provide speaker-independent norms of implicature, at least for particularized implicatures.

This is only a sketch of the intention-centred account, and further refinements and additions might be needed in order to develop it fully. But I hope I have said enough to show that it is possible to develop an intention-centred account of implicature which drops the requirement for calculability in Grice's sense but does not ignore normative concerns.

3. Utterer-implicature and audience-implicature

Before moving on, I want to return to an issue that I set aside earlier. As part of her case for a normative reading of Grice, Saul introduced the notions of *utterer-implicature* and *audience-implicature*. These are psychological states, and the purpose of positing them was, as it were, to fill in the gaps left by Grice's account. Many apparent cases of implicature turn out not to meet the strict Gricean conditions for implicature (given the argument in section 1, there may be a very large number of such cases). But — the suggestion was — rather than see this as evidence against Grice, we could see it as evidence for the existence of distinct psychological phenomena related to implicature. In cases of apparent implicature, Griceans may say, there is not an implicature, but the speaker or hearer *thinks* there is one. In this way, Griceans might use the notions of utterer-implicature and audience-implicature to respond to the challenge of speaker-dependency. I said earlier that I did not think this tactic would work, and I shall now explain why.

To recap, here is how Saul defines the two notions:

Utterer-implicature:

(1*) The speaker thinks that he is presumed to be following the conversational maxims, or at least the Cooperative Principle.

(2*) The speaker thinks that the supposition that he is aware that, or thinks that, q , is required to make his saying or making as if to say p (or doing so in *those* terms) consistent with this presumption.

(3) The speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2*) is required.

Audience-implicature:

(1) The speaker is presumed to be following the conversational maxims, or at least the Cooperative Principle;

(2A) The audience believes that the supposition that [the speaker] is aware that, or thinks that, q , is required to make his saying or making as if to say p (or doing so in *those* terms) consistent with this presumption.

(3A) The audience takes the speaker to think that it is within the audience's competence to work out that the supposition mentioned in (2A) is required.

(Adapted with minor revisions from Saul 2002a, p.235, 237, 242)

More concisely, there is an utterer-implicature when the speaker/utterer believes that each of the three conditions for a Gricean conversational implicature is met, and there is an audience-implicature when the hearer/audience thinks each of the three conditions is met.

Now, as theorists we have appealed to these notions mainly in cases where there is not in fact a conversational implicature. However, it cannot be supposed that utterer-implicatures and audience-implicatures exist only in cases where conversational implicature fails. It would be crazy to think that people believe that the Gricean conditions are met only when they are not met! Indeed, it is plausible to think that (on the Grice-Saul view) there will typically be an utterer-implicature whenever there is successful implicature generation, and that there will typically be an audience-implicature whenever there is successful implicature recovery. Cases where a person tries and fails to implicate something are subjectively just like cases where they try and succeed (assuming they do not realize they have failed), and cases where a hearer mistakenly thinks an implicature is present are subjectively just like cases where they correctly think one is (again, assuming they do not realize their mistake). So the same range of psychological states will, typically, be present in both cases. Thus, if utterer-implicatures are typically present in cases of failed implicature generation, then they will typically be present in cases of successful implicature generation too, and if audience-implicatures are typically present in cases of failed implicature recovery, then they will typically be present in cases of successful implicature recovery too.

Of course, even on the Grice-Saul view there will not *always* be an accompanying utterer-implicature and audience-implicature whenever there is a conversational implicature. As we saw in Chapter 3, a speaker may generate an implicature without believing that their audience will detect it (the Roland case), or (at least on my revised definition) without even realizing that they have done so

(the Rita case), so there can be conversational implicature without utterer-implicature. And there can be conversational implicatures that the audience does not recognize as having been intended (the Trigby case) or that they simply do not detect at all (the Wesley case), so there can be conversational implicature without audience-implicature.¹⁶ These are, however, atypical cases, where the generation of the implicature is partially or wholly unintentional, or the recovery of the implicature partially or wholly unsuccessful. But when a speaker successfully implicates something in the normal way, there will be an utterer-implicature present, and when a hearer successfully detects an implicature there will be an audience-implicature too. Or, at least, that is a consequence of the Grice-Saul view. There are several problems with this position, however.

First, it re-introduces many of the problems for the Gricean framework that we considered in the previous chapter. The problems arose because the original definition imposed psychological conditions for implicature. Condition (1) ('the cooperative presumption') required that the hearer should think the speaker is being cooperative, and condition (3) ('mutual knowledge') required that the speaker should believe that the hearer can work out that the implicated proposition is required by Gricean principles. And these conditions, we saw, are implausible. Contra (1), speakers can implicate things when thought to be uncooperative (as when they change the subject or try to mislead) or when there is no hearer at all. And contra (3), people unfamiliar with Gricean theory can implicate things even though they do not believe that the implicated meanings are required on Gricean principles, and people can implicate things they do not believe their audience can grasp (for example, when talking to a coma patient). The normative view, under the revisions suggested, resolved these problems by *depsychologizing* the notion of implicature. The new conditions, (I) and (II),¹⁷ did not mention the beliefs of the

¹⁶ For the Roland, Trigby and Wesley cases, see Chapter 3, section 2.1; for the Rita case, see Chapter 3, section 3.3.

¹⁷ Again, as a reminder, these are conditions are:

actual speaker or hearer, but stipulated only that the implicated meaning should be rationally required to make the speaker's utterance consistent with the cooperative presumption (regardless of whether anyone had actually made that presumption), and that a typical hearer could work out that it was required (regardless of whether the actual hearer did so, or whether the speaker thought they could do so). This removed the problems at a stroke.

But the notion of utterer-implicature and audience-implicature are couched in psychological terms, and thus reintroduce those problems, or related versions of them. For utterer-implicature, (1*) is no more plausible than (1).¹⁸ If speakers can successfully implicate things when their hearer does not presume they are being cooperative or when they have no hearer, then it is hard to see why they cannot successfully implicate things when they *doubt* or *disbelieve* that their hearer presumes them to be being cooperative, or *know* they have no hearer. And, since the definition of utterer-implicature includes Grice's original (3), it inherits all the problems arising from that clause.

Similarly, audience-implicature retains Grice's (1), and inherits all of the problems arising from it. And (2A) and (3A) raise an acute version of the problem of non-Griceans. These clauses require, not only that anyone who successfully recovers an implicature should believe that the implicated proposition is required by Gricean principles, but also that they should have complex higher-order beliefs — beliefs about the speaker's beliefs about their (the hearer's) mental abilities. And this requires, not only a knowledge of, and belief in, Gricean theory, but also

(I) The supposition that the speaker is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with the presumption that he is observing the CP.

(II) It would be within the competence of a typical audience to work out, or grasp intuitively, that the supposition mentioned in (I) is required.

¹⁸ Saul's (1*) (which is Grice's (1) prefixed with the 'The speaker thinks that') runs as follows:

(1*) The speaker thinks that he is presumed to be following the conversational maxims, or at least the Cooperative Principle. (Saul 2002a, p.235)

a high level of conceptual sophistication. It seems highly unlikely that laypeople, non-Griceans, and children would form the beliefs mentioned; yet they are all capable of successfully recovering implicatures. When I tell my five-year old son that he may have some of the sweets, he understands that he may not have all of them. But, even as a proud mother, I find it highly improbable to suppose that he believes that I believe it is within his power to work out that the supposition that he may not eat all the sweets is required to make my saying that he may eat some of them consistent with the presumption that I am following the principles of cooperative communication.

We could reduce the problems here somewhat by remodelling the definitions of utterance-implicature and audience-implicature around the revised definition of conversational implicature proposed in the previous chapter, which omits the troublesome (and unnecessary) condition (1). Thus, the revised definition of utterer-implicature would be as follows:

(I*) The speaker believes that the supposition that the speaker is aware that, or thinks that, *q*, is required to make his saying or making as if to say *p* (or doing so in *those* terms) consistent with the presumption that he is observing the CP.

(II*) The speaker believes that it would be within the competence of a typical audience to work out, or grasp intuitively, that the supposition mentioned in (I*) is required.

The definition for audience-implicature would be derived by substituting ‘The audience’ for ‘The speaker’ in the first clause, and inserting ‘The audience believes that’ at the beginning of the second, giving (IA) and (IIA).

But although this removes the problems arising from clause (1), all of these stemming from (3) remain, transferred to clauses (II*) and (IIA). If successful implicature generation and implicature recovery presuppose the existence of utterer-implicatures and audience-implicatures, then those who do not or cannot understand the Gricean theory, or who reject it, or who have trouble forming higher-order beliefs, cannot generate or recover implicatures.

A second problem with the notions of utterer-implicature and audience-implicature is that there do not appear to be any independent reasons to posit these states in order to account for the psychology of implicature — and it is, in fact, rather implausible to do so. As I argued in the previous chapter, the parallel with *saying* suggests that the psychological condition for implicature should be simply that the speaker means that *q* (that is, has the right communicative intentions). On Grice's view, in order to say that *p*, a speaker must use a sentence with the right conventional meaning (the normative condition), and they must mean that *p* (the psychological condition). They need not also *believe* that the sentence they use has the right conventional meaning, and still less that the conditions spelled out in the Gricean analysis of conventional meaning hold. Given this, it is hard to see why a similar condition should be imposed on implicature. By contrast, as we saw in Chapter 2, there are reasons for thinking that Grice held that a speaker does need to mean that *q* in order for them to implicate it (though not for their *utterance* to implicate it).

Similarly, when a hearer recovers an implicature with content *q*, there seems no reason to think that they must believe that the supposition that the speaker is implicating *q* is required on Gricean principles. Grice allows that implicatures may be intuitively grasped — that is, the hearer may realize that they must suppose that the speaker thinks that *q* (and hence is implicating *q*) without explicitly reflecting on the grounds for this supposition. This seems right. We do not consciously engage in Gricean reflections when interpreting implicatures, and indeed we may grasp an implicated meaning without even realizing that it *is* an implicature. This is compatible with the Gricean account giving a correct account of the normative conditions for implicature. We can be sensitive to norms without having explicit beliefs about them and without explicitly calculating when and how they apply. A driver can follow the rules of the road without thinking about them and without even being able to state them clearly. Of course, it could be that in such cases the rules are known and applied at a nonconscious level, and the same might be true of Gricean principles. As we shall see in section 4 below, Grice may in fact be committed to this view. But, as we shall also see, there are problems for that view.

Given this, it is doubtful that implicature generation typically involves the existence of utterer-implicatures and that implicature recovery typically involves

the existence of audience-implicatures, as the Grice-Saul view supposes. And if they do not, then we cannot rely on an appeal to utterer-implicatures and audience-implicatures to account for our intuitions in cases where implicature generation or recovery is not successful. If utterer-implicatures and audience-implicatures do not routinely accompany conversational implicatures, then we cannot appeal to them to account for our intuitions in cases where the conditions for conversational implicature are not met but we still feel there is something implicature-like occurring.

Note that this is not to deny that we may need notions *similar to* those of utterer-implicature and audience-implicature. In fact, I think we do, but they should be weaker and less closely tied to Gricean theory. At a first attempt, we might say that an utterance U carries a *weak utterer-implicature* with content q if the speaker intends U to indirectly convey (imply, suggest, indicate, mean) that q. And we might say that U carries a *weak audience-implicature* if the hearer believes that U indirectly conveys that q or that the speaker is indirectly conveying q via U. These notions are much less theoretically loaded than the original ones, and it is not implausible to claim that they typically accompany attempts at implicature generation and recovery, both successful and unsuccessful. So they may be useful in characterizing the psychology of implicature. However, since they are not linked to a Gricean approach (they are compatible with any account of the precise conditions required for implicature), they offer no specific support for the Gricean framework.

4. Implicature recovery

We have been concerned so far with implicature *generation* — the analysis of the conditions necessary for an implicature to exist, which is the heart of the Gricean framework. I have argued that Grice's theory is best interpreted as a normative one, but that even on this interpretation it still faces serious problems, and that an alternative, intention-centred account is preferable. I turn now to look at the Gricean framework's implications for implicature recovery — the psychological process by which hearers detect implicatures and recover their content. Again, I shall argue that there are problems for the Gricean approach.

As noted in Chapter 1, a theory of implicature generation will have implications for a theory of implicature recovery. We can recover conversational implicatures with some reliability, and we must, therefore, have some means of detecting the existence of the conditions that generate them. Broadly speaking, if an implicature exists in virtue of a certain state of affairs, S, then recovering the implicature will involve detecting S — either by directly detecting S itself or by detecting some state of affairs that reliably co-varies with S. For example, if implicatures are determined by the speaker's communicative intentions, then detecting an implicature must involve detecting those intentions. Assuming we can in fact recover implicatures (not infallibly, of course, but with some reliability), this means that a theory of implicature generation cannot completely ignore psychological questions about implicature recovery, and we can rule out theories on which we would not be able to reliably recover implicatures. Conversely, theories of how implicatures are recovered imply something about how implicatures are generated. If hearers typically detect implicatures by paying attention to a certain property of the communicative situation, then that property must be at least reliably connected with the property that generates the implicature. More fundamentally, the need for implicatures to be recoverable places some general constraints on a theory of implicature generation.

This is not to deny that a theory of implicature generation may play a normative role and may sometimes correct our judgements about what an utterance implicates. But the theory should not come totally apart from our everyday judgements. If the aim is to systematize the principles implicit in our best judgements about implicature, then it should be compatible with the nature of those judgements. If a normative theory says that implicature is determined by feature X, but our psychological theory tells us that our best judgements about implicature actually track feature Y, then the appropriate response, I suggest, would be to revise the normative theory.

Now on Grice's view, for an utterance U to implicate q there must exist a certain rational relation between certain premises (concerning U, its context, general conversational principles, and background assumptions), and the supposition that the speaker believes that q, to the effect that the former entail the latter or make it probable. So recovering the implicature must involve detecting

that this relation holds. Now Grice sketches an inferential process by which the existence of this relation could be established ('He has said that p; there is no reason to suppose that he is not observing the maxims ...'), However, he adds that hearers need not actually go through this process and may simply 'grasp intuitively' (1975/1989, p.31) that the relation holds. This is important for the plausibility of Grice's position, since as hearers we do not typically, if ever, go through Gricean calculations, at least consciously. However, the reference to intuition does not offer an alternative explanation of how implicatures could be recovered. As Daniel Dennett remarks, 'Intuition, after all, is not a particular method of deduction or induction; to speak of intuition is to deny that one knows how one arrived at the answer' (Dennett 1986, p.152). Nor will it do to appeal to the competence view, discussed in Chapter 2, on which a quick, gappy inferential process can be regarded as valid, provided the reasoner intends it to be valid and has the ability to produce a full version filling in all the gaps (Chapter 2, section 3). This may be sufficient for the reference to intuition in clause (3) of Grice's definition of implicature, where it is the responsibilities of the speaker that are at issue ('the speaker thinks ... that it is within the competence of the hearer to work out, or intuitively grasp, that the supposition mentioned in (2) is required'). But it will not do here, where the question is not whether an intuitive inference could be replaced by an explicit version, but *how* an intuitive process actually got to the answer. Implicature recovery is not magic, so there must be some reliable mechanism at work, at a nonconscious level if not at a conscious one. But what could it be? There seem to be only two options for Grice. First, the mechanism could involve a *nonconscious* Gricean calculation. Second, it could involve a shortcut — the detection of some state of affairs that usually occurs when and only when the implicature-generating relation holds.

Take the second option first. This might work for generalized implicatures, which are not context dependent. If the same sentence generates the same implicature in most contexts, then we could reliably recover the implicature simply by detecting the use of the sentence. But this method would not work for particularized implicatures, which are heavily context-dependent and one-off. There are no repeatable associations here to pick up on and it is hard to see what other feature might offer a shortcut in such cases. It would need to be a feature

which is as context-sensitive as the implicature-determining relation itself. The only option seems to be speaker intentions. If there is a reliable link between what speakers intend to implicate and what their utterances actually do implicate on Gricean principles, and if hearers can reliably detect speaker intentions, then this would work. However, it is doubtful that there is a reliable — or reliable *enough* — link between speaker intentions and Gricean conversational implicatures. As we have seen, there will be many cases where what speakers intend to implicate differs from what their utterances do implicate, on Gricean principles. If it is speaker intentions that guide our everyday judgements in such cases, then it looks as if Gricean theory is not describing the everyday phenomenon of implicature at all, and that an intention-centred theory would be preferable.

If this is right, then Griceans should take the first option and hold that implicature recovery involves nonconscious Gricean calculations — at least in cases of particularized implicature. It is generally accepted by psychologists that complex nonconscious mental processes support everyday behaviour, so this is not in principle an implausible claim. However, there are specific problems in positing nonconscious Gricean calculations.

First, we run up against the problem of laypeople, non-Griceans, and children again. How do those who have not mastered, or have rejected, the Gricean framework recover implicatures? The only option seems to be to suppose that such people are in fact nonconscious Griceans — that although they have no explicit, conscious knowledge of, or acceptance of, the Gricean framework, the nonconscious mental processes that enable them to recover implicatures nevertheless employ Gricean concepts and principles. (It might even be suggested that these concepts and principles are innate.) The idea that nonconscious processing employs concepts and principles to which the person has no conscious access, and which may even conflict with their conscious beliefs, is not uncommon in cognitive science. Moreover, as we shall see in the next chapter, some linguists propose that the recovery of generalized implicatures involves the nonconscious application of simplified Gricean principles. However (as we shall see), the current experimental evidence does not favour this view, and to maintain that all implicatures are recovered by means of nonconscious Gricean calculations would be to make a strong and risky empirical commitment.

Second, there are worries about the feasibility of Gricean calculations, at least for particularized implicatures. As we saw in Chapter 2 (section 3), it is unlikely that Gricean calculations can be formulated as deductive inferences. Grice's own sketch of an implicature calculation is not deductively valid, and many additional premises would have to be added to make it so. It is more plausible to see the calculation as an abductive inference, or an inference to the best explanation — a process central to everyday and scientific reasoning. However, the sort of abductive inference involved in a Gricean calculation is an unusual one. The task is not to find the best explanation of an event — the speaker's making the utterance. If it were, then the best explanation would surely be that the speaker intended to convey something, and in hypothesizing about what this was, the hearer would naturally be guided by evidence about the speaker's beliefs, desires, and other mental states. But the task for the Gricean interpreter is different. What they must find is not an explanation of the speaker's utterance, but a way of reconciling two claims: that the speaker said what they did and that the speaker was being a cooperative communicator, and the data they are supposed to draw on excludes information about the speaker and their mental states. Given that there are, in principle, a limitless number of ways of making any two claims compatible, this could be a very demanding task, and a lot of background assumptions would have to be made in order to home in on a specific, non-disjunctive claim. It is hard to see why hearers should follow this route, especially when the more straightforward option of theorizing directly about the speaker's intentions is available.

Note finally that a closely related problem arises for speakers. If speakers are to non-accidentally succeed in implicating, then they will need to be sensitive to the implicature-generating relations between literal and implicated contents, so that they can choose suitable literal contents to convey the meanings they wish to implicate. And again we can ask how they achieve this sensitivity. In the case of generalized implicatures, they might rely on learned associations between sentences and implicatures, but in particularized cases, this option is not available. Since speakers cannot rely on hearers simply detecting their communicative intentions (which are, after all, irrelevant to implicature generation within the Gricean framework), it seems they will have to calculate the implicature-

generating relation, at least nonconsciously. That is, if they want to implicate q they will have to make a sort of reverse Gricean calculation, working out what they should say in order to make it the case that the supposition that they believe q is required in order to reconcile the claim that they said what they did with the claim that they are being a cooperative communicator. Again, this is a strong, and implausible, psychological hypothesis.

Conclusion

The previous chapter argued that many problems for the Gricean framework can be removed by following Saul in adopting a normative reading of Grice and by revising and extending the framework in sympathetic ways. This chapter has argued, however, that serious problems remain, especially as regards particularized implicatures. The revised Gricean framework does not achieve its aim of establishing speaker-independent norms for implicature, and it has some implausible implications for the psychology of implicature. One moral that emerged was that theories of implicature, whether concerned with generation or recovery, should give a greater role to speaker intentions, and I proposed an account of implicature that gave speaker intentions such a role while still retaining a normative element.

However, we have not finished with Gricean ideas yet. Some linguists have drawn on Grice's work in developing theories of implicature recovery, especially for generalized implicatures. (As we shall see, these accounts also have implications for implicature generation.) We shall consider this *neo-Gricean* approach in the next chapter. This will also give us chance to take a deeper look at the Gricean approach to generalized implicature.

Chapter 5

Neo-Griceanism and its rivals

The previous chapters assessed the Gricean framework as a theory of implicature generation — of what makes it the case that certain utterances carry implicatures. I argued that the framework is best understood as aiming to provide a normative, speaker-independent notion of conversational implicature, parallel to that of conventional meaning. And I went on to argue that it fails in that aim: Gricean principles cannot be applied without appealing at some level to speaker intentions. I also argued that theories of implicature generation cannot be separated entirely from psychological theories of how implicatures are recovered and that it is hard to see how Grice's account could be integrated with such a theory.

But we should not write off the Gricean framework yet. Although it may not succeed in its original aim, it may still explain some important aspects of conversational implicature. In fact, many linguists have seen Grice's ideas as providing the basis for accounts of implicature recovery. These *neo-Gricean* theorists argue that in interpreting utterances we automatically apply certain general principles or heuristics, which are versions of the Gricean maxims. Where applicable, these principles yield non-literal meanings that become the default, or preferred, interpretations of utterances of the type in question. The same principles, it is assumed, guide speakers in their choice of utterance, ensuring that communication is usually successful (for example, Levinson 2000, p.24). Thus, this view treats the Gricean maxims, not as norms of conversation, but as inferential principles that guide conversational behaviour (Levinson 2000, p.35).

Though neo-Griceanism is primarily a theory of implicature recovery (of how hearers derive implicatures), it can also be seen as offering a theory of implicature generation. Neo-Griceans can say that an utterance possesses a generalized implicature *q* if the interpretive principles the theory posits would yield *q* as a preferred interpretation of it. Since neo-Griceans hold that we do typically apply these principles in making and interpreting utterances, this amounts to saying that an utterance implicates *q* if hearers would typically interpret it as possessing it and speakers typically would expect it to. Understood in this way, the theory would licence normative claims to the effect that a *particular* speaker or hearer has

misunderstood what an utterance implicates. Since the neo-Gricean principles are derived from Grice's maxims, and make many of the same predictions, this approach promises to rescue at least part of the original Gricean project. (Neo-Griceans sometimes speak of the principles they posit 'generating' implicatures — meaning that they yield them as interpretations in the minds of hearers. Although I treat neo-Griceanism primarily as a theory of implicature recovery, I shall occasionally follow this usage in the chapter, especially since, as just explained, the principles can also be thought of as generating implicatures in the constitutive sense.)

This chapter will assess neo-Griceanism, focusing on one prominent version of the approach. The first section outlines the theory, and the second section briefly introduces some rival theories with which I shall contrast it. The third and fourth sections look at the neo-Gricean principles, asking whether they accord with and explain our intuitions about what implicatures utterances possess, and the final section reviews some relevant experimental work.

The literature in this area is often technical and deeply involved with wider issues in theoretical linguistics. It is impossible to do justice to it in a chapter, but I shall focus on some key points and test cases.

1. Neo-Griceanism

Neo-Gricean theories propose simple, formalized versions of the Gricean maxims and rules for their application, with the aim of explaining and predicting patterns of implicature. Key figures in the field are Gerald Gazdar (1979), Laurence Horn (1984, 1989, 2004), and Stephen Levinson (1983, 2000). I shall focus on Levinson's presentation in his 2000 book *Presumptive Meanings*, which synthesizes earlier work in the tradition and is a comprehensive and influential presentation of the neo-Gricean approach (Levinson 2000).

1.1 Utterance-type meaning

Levinson adopts a broadly Gricean approach to communication, distinguishing aspects of meaning that are coded (including what is said and what is conventionally implicated) and aspects that are inferred from conversational

principles, including, but not limited to, conversational implicatures (2000, p.14).¹ Levinson follows Grice in accepting that there are two types of conversational implicature — particularized and generalised. (Levinson uses the abbreviations PCI and GCI.) PCIs hold *because of* specific contextual assumptions that do not hold in all, or even many, cases; whereas GCIs hold universally *unless* there are specific contextual assumptions that cancel them. Take, for example:

- (1) Some of the guests got food-poisoning.

Unless the implicature is cancelled (say, by adding ‘In fact all of them did’) this would always carry the implicature that not all the guests got food poisoning, which is a GCI. However, if uttered in response to the question ‘How was the wedding?’ it would also carry the implicature ‘The wedding went badly’, or something similar, which would be a PCI. Levinson suggests that PCIs are the result of applying the maxim of Relevance, where this involves attending to the particular speaker’s goals and plans (2000, p.17, p.380 n.4).

Levinson’s primary interest is in GCIs, which he regards as the central class of ‘presumptive meanings’ — default, or preferred, interpretations, which are ‘carried by the structure of utterances, given the structure of the language, and not by virtue of the particular contexts of utterance’ (2000, p.1). These presumptive meanings, he claims, form a distinct level of meaning, *utterance-type meaning*, which is distinct from both the linguistically coded meaning that a sentence carries in every context (sentence-meaning), and the pragmatically enriched meaning that a sentence carries when uttered by a particular speaker in a particular context (speaker-meaning, or utterance-token meaning). Utterance-type meaning is like

¹ Although he makes this distinction, Levinson regards all communication as, fundamentally, an inferential process, in which even coded aspects of meaning are clues to interpretation. He writes:

From a Gricean perspective, communication involves the inferential recovery of speakers’ intentions: it is the recognition by the addressee of the speaker’s intention to get the addressee to think such-and-such that essentially constitutes communication. (Levinson 2000, p.29)

sentence-meaning in being context-independent and deeply connected with the structure of language, but unlike it in being cancellable. Utterance-type meaning is like speaker-meaning in resulting from pragmatic enrichment of sentence-meaning, but unlike it in resulting from the application of general principles rather than theorizing about the speaker's intentions (2000, p.22). (Levinson notes that utterance-type meaning is not composed only of GCIs, but includes a variety of other pragmatic phenomena, including presuppositions, conventional implicatures (in Grice's sense), and conventions of use (2000, p.23). Levinson's aim is to defend the existence of a level of utterance-type meaning, in opposition to *reductionists*, who would reduce it either to sentence meaning or (as relevance theorists do) to speaker meaning (2000, p.25).

Levinson suggests that our capacity for GCIs is an evolutionary adaption, which developed in order to compensate for an inefficiency in human communication (2000, pp.27–9). He points out that pre-articulation and comprehension processes in the human brain run three to four times faster than the process of phonological articulation (2000, p.28). Our brains can prepare and process utterances much faster than our vocal systems can articulate them, creating a bottleneck in the human communication system. Evolution has eased this bottleneck, Levinson argues, by designing our comprehension systems to apply certain general pragmatic principles, or heuristics, which are defined over formal features of utterances and are applied by default whenever certain expressions are encountered.² These principles yield GCIs — default interpretations that are derived without the need to theorize about the speaker's intentions — and they speed up communication by creating an extra layer of utterance-type meaning,

² Levinson writes:

Now, the solution to the encoding bottleneck, I suggest, is just this: let not only the content but also the metalinguistic properties of the utterance (e.g., its form) carry the message. Or, find a way to piggyback meaning on top of the meaning ... by utilizing the form, the structure, and the pattern of choices within the utterance to signal the extra information beyond the meanings of its constituents. (Levinson 2000, p.6)

which enriches the content of utterances in ways that we all understand and expect. This requires, of course, that GCIs are recovered very swiftly, without complex theorizing. Indeed, Levinson suggests that some of the principles are applied on a word-by-word basis and that a quantifier such ‘some’ will trigger its default interpretation of ‘not all’ even before the predicate it governs has been processed (2000, p.5, p.259)³.

Although the inferencing here is default, Levinson stresses that it is *defeasible*; it goes through automatically unless contrary information is available, in which case the inference is cancelled. The process therefore cannot be a deductive one, since deductive inference is not defeasible, and it must involve some form of non-monotonic reasoning. Levinson reviews various types of defeasible inference, including induction, abduction, practical reasoning, and default logics, and argues that the last offers the most promising model for implicature (2000, pp.45–6).

1.2 The three principles

Levinson proposes that GCIs can be accounted for by appeal to three principles, which he calls the Q-principle, the I-principle, and the M-Principle — the first and second derived from Grice’s maxim of Quantity and the third from his maxim of Manner. The maxim of Quality plays ‘only a background role’ in the production of GCIs, and the maxim of Relation or Relevance plays none (it ‘has pertinence only to the immediate, ever variable, conversational goals: it generates PCIs, not GCIs’) (Levinson 2000, p.74).

The Q-principle can be summarized as ‘What isn’t said, isn’t.’ The idea is that speakers make the most informative statement they can, given what they know, and that hearers assume that they do this. This resembles Grice’s first submaxim of Quantity (‘Make your contribution as informative as is required’). On its own the Q-principle is too vague to be applied automatically, and Levinson explains

³ Levinson writes ‘the phrase *some of the boys* can invoke the default assumption “not all of the boys” even before the predicate has been heard’ (Levinson 2000, p.259). This is an important claim for him, since he holds that it is just these sorts of rapid default inferences that enable hearers to transform semantic fragments into full-blown propositional representations (Levinson 2000, pp.256–9).

that it can be applied only where there is a salient set of contrasting expressions of different informational strength from which the speaker is assumed to have chosen. In choosing a weaker element from the set, the speaker implicates that a corresponding statement substituting a stronger element is false. So, for example, if a person utters the sentence, ‘Some of the students failed’, the contrast set <*all, some*> is salient, and the Q-principle produces the default reading ‘It is not the case that all the students failed.’

The most important class of Q-implicatures are *scalar implicatures*, which depend on an entailment scale (also known as a Horn scale), in which the stronger elements entail the weaker ones. As examples, Levinson gives the following (with stronger items to the left):

quantifiers <*all, most, many, some*>,
connectives <*and, or*>,
modals <*necessarily, possibly*>, <*must, should, may*>,
adverbs <*always, often, sometimes*>,
degree adjectives <*hot, warm*>
verbs <*know, believe*>, <*love, like*>
(Levinson 2000, p.79).

In making an utterance using an expression to the right of one of these scales, one Q-implicates that a corresponding utterance substituting an expression to the left either is false, or might be for all one knows.⁴ Thus ‘Some of the students passed’ Q-implicates that not all the students passed’; ‘I told Jack or Annie’ Q-implicates

⁴ There is debate about the strength of the epistemic commitment involved in Q-implicatures. Does the speaker implicate that they know that the stronger claim is false, or that they believe it is, or simply that it may be false for all they know? The issues are complex and I shall not address them here. (For discussion and references, see Levinson 2000, pp.77–9.) In any case, as Atlas notes, the speaker’s attitude is independent of the *content* implicated. An utterance of ‘Some F are G’ implicates that not all F are G, and invites the hearers to believe that not all F are G, regardless of what attitude the speaker is understood to take towards that claim (Atlas 1993, discussed in Levinson 2000, p.78, p.387, n.10).

that the speaker didn't tell both Jack and Annie; 'You may smoke' Q-implicates that it is not the case that the hearer is obliged to smoke, and so on.

The other main type of Q-implicature is *clausal implicature*, which can arise when a sentence contains an embedded clause. By choosing an expression that does not entail the truth of the embedded clause instead of one that does, the speaker implicates that they do not know whether or not the embedded clause is true. For example, if I say 'John believes there is life on Mars' (rather than 'John *knows* there is life on Mars') I Q-implicate that I myself do not know whether or not there is life on Mars (Levinson's example; 2000, p.76).⁵

The second principle is the I-principle, which Levinson summarizes as 'What is simply described is stereotypically exemplified'. The idea here is that typical ('unmarked') expressions implicate that the thing described is itself typical, prompting hearers to fill in the details according to the appropriate stereotype. Levinson notes that this is related to Grice's second sub-maxim of Quantity: 'Do not make your contribution more informative than is required'. Speakers need not spell out details that hearers will fill in automatically ('one need not say what can be taken for granted') (2000, p.37).

The I-principle is a powerful one, which underpins a variety of linguistic phenomena, including *generality narrowing*, where a general expression is interpreted in a more specific sense ('secretary' is understood as 'female secretary', 'road' as 'hard-surfaced road', 'John's book' as 'the book John read/wrote/borrowed'); *conjunction buttressing*, where a conjunction is interpreted as indicating temporal or causal sequence ('John turned the switch and the motor started' implicates that the switch turning preceded, or caused, the starting, or was done with the intention of causing it); and *conditional perfection*, in which a conditional is read as a biconditional ('If' in 'If you mow the lawn, I'll give you \$5' is interpreted as 'if and only if') (Levinson 2000, pp.37–8; the examples are Levinson's).

The I-principle allows us to enrich the content of an informationally minimal utterance by drawing on background knowledge. It might seem that the reliance

⁵ In this example there is also a separate *scalar* implicature to the effect that John does not know that there is life on Mars.

on background knowledge here undermines the status of I-implicatures as generalized and default. I assume Levinson would reply that the knowledge in question is of stereotypes that are both immediately accessible (hence default) and context independent (hence generalized). Thus applying the I-principle does not involve drawing on knowledge of the specific context of utterance or speculating about the speaker's intentions. (This may not be an adequate reply, however; I will return to this issue in section 4 below.)

Levinson summarizes the third principle, the M-Principle as 'What's said in an abnormal way isn't normal' (2000, p.38). This is the reverse of the I-principle: The use of an untypical, or 'marked', expression indicates that the thing referred to is itself atypical in some way. Levinson notes that the principle is related to Grice's maxim of Manner ('Be perspicuous') and in particular to its first and third⁶ submaxims: 'avoid obscurity' and 'avoid prolixity'. In flouting these submaxims by using unusual or long-winded expressions, speakers indicate that there is something unusual about the thing described. For example, 'Bill caused the car to stop' (Levinson's example) M-implicates that Bill stopped the car indirectly rather than by simply pressing the footbrake, and 'Jack talked and talked' M-implicates that Jack talked at unusual length. As with Q-implicatures, M-implicatures depend on an implied contrast, this time with unmarked expression that could have been used instead ('stopped the car', 'talked').⁷

Levinson notes that there can be conflicts between the three principles. For example, the Q-principle and the I-principle pull in opposite directions: the former tells us that if a speaker doesn't say something then it should be ruled out; the latter that if a speaker doesn't say something then it can be taken for granted. Similarly, the I-principle tells us to adopt standard interpretations, the M-principle to look for non-standard ones. Levinson argues that these potential conflicts are resolved by assigning priorities to the different types of implicature: Q-implicatures and M-implicatures take priority over I-implicatures; Q-implicatures take priority over M-

⁶ Levinson actually calls it the fourth, but this seems to be a slip (2000, p.38)

⁷ As Levinson notes, the M-principle and the Q-principle both involve *negative* inferences: in both cases the hearer infers the implicated message from the fact that the speaker has *avoided* using some other expression, informatively stronger in one case, less marked in the other (2000, p.40).

implicatures, and clausal Q-implicatures take priority over scalar Q-implicatures.⁸ He suggests that all applicable principles are applied automatically, and any inconsistent results subsequently filtered out in accordance with the rules of priority (2000, pp.161–2). Levinson also allows that a Q-implicature can be implicitly cancelled if it conflicts with an entailment of what the speaker says, or is inconsistent with shared background assumptions, or is obviously irrelevant to the speaker’s conversational goals (in the last case considerations of relevance, in Grice’s sense, will play a role). (Levinson 2000, p.49–52).⁹

1.3 Applying the principles

Levinson gathers a huge amount of data to support the existence of GCIs, highlighting the ‘regularity, recurrence, and systematicity’ of pragmatic inferences of the kind he describes (2000, p.22). One important piece of evidence comes from facts about lexicalization (Levinson 2000, pp.64–71). English lacks words for the contradictories of certain logical concepts (the concepts are not *lexicalized*). For example, we have a word for *all*, but none for its contradictory *not all*. This, Levinson, argues, is because that concept is carried by ‘some’ (the contrary of the contradictory of ‘all’) in virtue of a Q-implicature. The same goes for several other concepts that stand in similar logical relations; for example:

⁸ Levinson suggests that the priority of Q- and M-implicatures is due to the fact that they involve a deliberate choice of words (a weaker term or a marked expression) rather than reliance on stereotypical interpretation, and that the priority of Q-implicatures over M-implicatures reflects the greater importance of informational content over nuances of expression (2000, p.161).

⁹ Other neo-Gricean theorists propose closely related taxonomies; for a useful table comparing them, see Levinson 2000, p.41. In particular, Horn reduces the principles to two: the Q-principle and the R-principle (for example, Horn 2004). The former combines Grice’s first submaxim of Quality (be as informative as required) and two submaxims of Manner (avoid obscurity and avoid ambiguity), and it does the combined work of Levinson’s Q- and M- principles. The R-principle combines Grice’s second submaxim of Quantity (do not be more informative than required), maxim of Relation (be relevant), and third and fourth submaxims of Manner (be brief and be orderly), and it corresponds to Levinson’s I-principle.

Not require is unlexicalized since Q-implicated by ‘permit’.

May not is unlexicalized since Q-implicated by ‘may’.

Not always is unlexicalized since Q-implicated by ‘sometimes’.

Not necessary is unlexicalized since Q-implicated by ‘possible’.

Not both is unlexicalized since Q-implicated by ‘or’.

Levinson notes that these patterns arise in other languages too, suggesting that they are due to the operation of a general interpretative principle (2000, p.69).

Levinson reviews the principles and their application in great detail, showing how they can explain a wide variety of linguistic features and intuitions and raising and responding to numerous objections. He also argues that GCIs are deeply involved in processes of disambiguation, indexical resolution, reference identification, ellipsis unpacking, generality narrowing, and co-reference (anaphora), which are necessary to establish the truth-conditional content of an utterance, and have been traditionally thought of as part of semantic processing.¹⁰ As he notes, this creates a problem for Grice’s view that implicatures are determined in part by what is said (the truth-conditional content expressed), since what is said may itself be determined by implicatures. (Levinson calls this ‘Grice’s circle’; 2000, p.186). Levinson himself avoids this problem by arguing that the GCI principles can be applied to utterance fragments (words or phrases), before a complete propositional content has been determined. This view does, however, present a challenge to the traditional conception of the relation between semantics and pragmatics, on which semantic processing yields a fully-fledged propositional content, which is then enriched or supplemented by pragmatic processes. The upshot, Levinson suggests, is that there are two rounds of pragmatic processing —

¹⁰ Levinson also highlights the importance of *intrusive constructions*, such as comparatives and conditionals, where the truth conditions of a sentence depend on an implicature generated by a part of it (Levinson 2000, pp.198–217). An example is ‘Driving home and drinking three beers is better than drinking three beers and driving home’, where the proposition expressed by the whole sentence is determined by I-implicatures of temporal sequence generated by its two component phrases.

a presemantic round, which establishes a truth-conditional content and a postsemantic round, which may produce a further implicature (Levinson 2000, pp.187–8).

2. Alternatives to neo-Griceanism

In assessing neo-Griceanism, it will be helpful to compare it with rival theories of implicature recovery, and I will briefly introduce three of these in this section. As we shall see, there are reasons for thinking that each has some advantages over neo-Griceanism, and it may be that a theory of implicature recovery can draw on elements from all of them.

2.1 *Relevance theory*

In the linguistic literature, the neo-Gricean approach to implicature recovery is usually contrasted with that of *relevance theory* (for example, Carston 2002; Sperber and Wilson 1995; Wilson and Sperber 2004). The theory (really a cluster of closely related theories) is complex and has developed over time.¹¹ Here I shall give a simplified outline, emphasizing the contrast with the neo-Griceanism.¹²

According to relevance theorists, a hearer infers a speaker's meaning from the linguistically coded meaning of their words and contextual information, searching for the interpretation that is the most *relevant* one, in a certain technical sense. The relevance of an utterance is a measure of its *positive cognitive effects* (in particular its *contextual implications* — conclusions one can draw from it in the context), set against the *effort* it takes to process it. Relevance theorists hold that human cognition is automatically geared to maximize the relevance of the inputs it receives, aiming for maximum effects for minimum effort.

¹¹ As Wayne Davis puts it,

Exposition [of relevance theory] is difficult because formulation of the theory varies significantly from presentation to presentation. And many interlocking technical terms require considerable clarification.' (Davis 1998, p. 99)

¹² The following outline draws in particular on Wilson and Sperber 2004 and Carston 2004a, 2004b.

Since speakers want hearers to attend to what they say, utterances carry a presumption of *optimal relevance* — that is, that they are both (a) sufficiently relevant to be worth the hearer’s effort to process them and (b) the most relevant the speaker is able and willing to provide. This presumption (the ‘communicative principle of relevance’) gives hearers specific expectations of relevance and guides how they interpret utterances. A hearer seeks to infer the speaker’s meaning from their words and the context of the utterance, forming and testing hypotheses until their expectations of relevance are satisfied. Note that since clause (b) refers to the speaker’s abilities and preferences, relevance theory, unlike the Gricean framework, treats interpretations as sensitive to information about the particular speaker, and it can allow for the possibility that speakers are uncooperative (see, for example, Carston 1998).

According to relevance theory, the interpretation process starts with the linguistically coded content of the utterance (roughly, the context-independent meaning of the sentence uttered), which will typically be underspecified and fail to express a propositional content. Interpreting the utterance then involves two tasks. First, there is a process of what relevance theorists call *explicature*, which involves enriching the linguistically coded content to produce an explicit propositional content, corresponding to what the speaker literally meant. Relevance theorists argue that this process involves not only resolving ambiguities and identifying references, but also a substantial process of pragmatic enrichment, including filling in missing conceptual elements (for example, expanding ‘It’s raining’ to ‘It’s raining in Sheffield’) and narrowing down (or broadening) the meaning of expressions to express more specific ad hoc concepts (for example, narrowing down the meaning of ‘happy’ to express some contextually salient level or type of happiness).¹³ Second, there is a search for additional implicated

¹³ From a relevance theory perspective, talk of the ‘literal’ meaning of an utterance is thus ambiguous. It might refer either to its linguistically coded content as opposed to an enriched explicature of that content, or it might refer to the explicature of the utterance as opposed to a distinct implicature of it. Moreover, neither the linguistically coded content of an utterance nor the explicature of it correspond to what is said by it, in Grice’s sense. What is said is richer than the linguistically coded content. (Grice allows that determining what is said requires resolving

meanings (implicatures) distinct from the explicit meaning. Crucially, relevance theorists hold that these processes follow a path of least effort, starting with the simplest, most accessible interpretation and progressing to more complex ones only if current expectations of relevance have not been met. (Although the most accessible interpretation of an expression will usually be what we would regard as the literal one, it may not always be so. Sometimes the linguistic context may strongly prime for a pragmatically enriched meaning, and sometimes an enriched meaning may be much easier to process, as with some metaphors; Noveck and Sperber 2012, p.371.)¹⁴

ambiguities and identifying references; Grice 1975/1989, p.25). Yet what is said is weaker than the explicature, since it does not depend on pragmatic enrichment. We might say that what is said by an utterance is the *minimal proposition* it expresses — the minimal filling in of its linguistically coded content needed to generate a propositional content (Recanati 1993). One problem for this view, however, is that such a minimal proposition will often be quite different from what the speaker means, and will often be trivially true or trivially false. (Consider, for example, ‘Everyone screamed’ and ‘It’s snowing’, which, without further specification of the relevant domain, will always be respectively false and true.) Since Grice holds that what is said must be meant by the speaker (1968/1989, p.88), this is an implausible consequence. For more discussion of this tricky topic, see Carston 2004b (from where the examples just given are taken) and for a useful table comparing different theorists’ use of ‘what is said’, ‘explicature’, ‘implicature’, and related terms, see Levinson 2000, p.195. Saul defends the Gricean notion of what is said, arguing that it is normative rather than psychological (Saul 2002b). As noted in the previous section, Levinson also holds that pragmatic processes contribute to fixing the truth-conditional content of utterances. However, he holds that these processes are limited to application of the GCI principles (as opposed to context-specific enrichment), and he does not recognize a distinction between explicature and implicature, which, he argues, has no principled basis (2000, pp.194–8).

¹⁴ The searches for explicatures and implicatures should not be thought of as independent of each other. The search for an explicature may be constrained by the need to find an explicit meaning that supports a contextually relevant implicature. For example, if a person replies ‘I’m tired’ when asked ‘Do you want to go out?’, then interpreting their utterance may involve narrowing down the meaning of ‘tired’ to a more specific degree of tiredness suitable to implicate ‘I don’t want to go

Since relevance theory holds that implicature derivation is guided by psychosocial principles, it is in a very broad sense Gricean (in effect, it puts all the weight of derivation on the maxim of Relevance, reinterpreted as the presumption of optimal relevance in the technical sense). However, there are big differences between it and Levinson's neo-Gricean theory. First, Levinson holds that the GCI principles are applied at an early stage in language processing and that the interpretations they yield are the default ones. Relevance theory, by contrast, holds that meanings are processed in order of accessibility, starting with their semantically coded content and enriching it (and deriving implicatures if necessary) until current expectations of relevance are met. As we shall see later, this difference between neo-Griceanism and relevance theory yields conflicting predictions, which have been experimentally tested. A second difference is that neo-Griceanism holds that some implicatures are generalized and context-independent, whereas relevance theory sees implicature derivation as *context-driven* (for example, Breheny et al. 2006). Since optimal relevance is defined in terms of the speaker's abilities and preferences, hearers' expectations of relevance will vary from context to context, and similar utterances may generate an implicature in one context but not in another.

To illustrate this, consider the following examples (taken from Sperber and Wilson 1995, p.277):

(2) *Henry*: If you or some of your neighbours have pets, you shouldn't use this pesticide in your garden.

Mary: Thanks. We don't have pets, but some of our neighbours certainly do.

(3) *Henry*: Do all, or at least some, of your neighbours have pets?

Mary: Some of them do.

out'. Interpretation involves a search for the *combination* of explicit and implicit contents that together makes the utterance optimally relevant (Noveck and Sperber 2012, p.313).

Here, neo-Griceanism predicts that both of Mary's replies should implicate that not all of her neighbours have pets, thanks to an automatic application of the Q-principle. Sperber and Wilson suggest that this is wrong; only her second reply carries that implicature. In the first example, the reading of 'some' as 'some and possibly all' (which Sperber and Wilson assume is the more basic one) is sufficient to satisfy Henry's expectations of relevance. In the context it does not matter whether all of Mary's neighbours have pets. In the second example, by contrast, Henry has made it clear that it is relevant to him to know whether all of Mary's neighbours have pets, and Mary's answer would not meet this expectation on the basic reading of 'some'. Henry therefore engages in further processing, reasoning that Mary did not say that all of her neighbours had pets because she was not in a position to do so, and that she means him to understand that they do *not* all have pets (Sperber and Wilson 1995, pp.277–8).

To sum up, relevance theory holds that there are no general inferential principles involved in the derivation of implicatures (other than the presumption of optimal relevance), and no distinction between generalized and particularized implicatures; in effect, it treats all implicatures as particularized.

2.2 Convention theory

In the literature on implicature recovery, neo-Griceanism and relevance theory are the main players. However, I want to introduce another approach, which draws on a non-Gricean analysis of conversational implicature developed by Wayne Davis (Davis 1998; see also Morgan 1978).

Davis argues that conversational implicatures cannot be calculated, even in principle, by applying Gricean conversational principles. I have discussed some of his arguments in previous chapters (for example, Chapter 3, sections 1.1–1.3, Chapter 4, section 1.3), and I shall introduce another one later in section 3.2 below. In opposition to Grice, Davis argues that particularized implicatures (he calls them 'speaker implicatures') are determined by the intentions of the speaker, and that generalized implicatures ('sentence implicatures') are determined by semantic conventions. It is this latter claim that I want to focus on here.

According to Davis, languages are associated with implicature conventions. It is a convention of English that sentences of the form 'Some F are G' are used to

implicate that not all F are G, that sentences of the form ‘p or q’ implicate ‘Not both p and q’, that sentences of the form ‘p and q’ implicate that p preceded q, and so on.¹⁵ Davis defines a convention as

an arbitrary social custom or practice. More explicitly, a convention is a regularity in the voluntary action of a group that is socially useful, self-perpetuating, and arbitrary. (1998, p.133; italics in original)

And he argues that implicature practices of the sort just mentioned are conventions in this sense. They are socially useful, promoting ‘*cooperative, efficient, polite, and stylish communication*’ (1998, p.174, italics in original). It is often quicker, politer, and more stylish to implicate something than to say it explicitly. They are largely arbitrary; different implicature practices are possible, and it is a historical accident that we have the ones we do.¹⁶ (Davis accepts that there will have been some pre-existing relation between the literal meaning of a sentence and the implicature it has come to carry, which made the practice ‘fitting’, ‘appropriate’, or ‘intelligible’ (he calls this the *Principle of Antecedent Relation*) (1998, pp.183–4). However, he argues this relation is never strong enough to uniquely determine the implicature.) Finally, implicature practices are self-perpetuating; once a

¹⁵ Davis also argues that there are more general implicature conventions, which are not associated with a particular sentence form, but are in effect procedures for generating one-off speaker implicatures. For example, we have conventions of implicating one thing by asserting its denial (as in irony); of implicating a piece of information by making a closely related statement (as in Grice’s example of saying ‘There is a garage round the corner’ to convey where petrol can be bought; or of implicating an affirmative or negative answer by asking a question whose answer is obvious (as in ‘Is the pope Catholic?’) (Davis 1998, pp.148–54; see also Morgan 1978).

¹⁶ Davis also points out that some generalized implicatures are language-specific, citing Wierzbicka’s work on cross-cultural pragmatics (1985, 1987, 1991). For example, in English ‘An X is an X’ implicates that one X is as good as another, but the Polish equivalent implicates that there is something uniquely good about an X (Davis 1998, p.144).

practice has become established, people have reason to continue following it if they wish to communicate successfully.¹⁷

Davis claims that the semantic conventions that fix implicatures are of a different type from those that fix literal meaning (sentence meaning, or conventional meaning in the usual sense). The latter are *first-order*: they are rules for assigning meanings to words and sentences. By contrast, the conventions that determine implicatures are *second-order*: they are rules for assigning further meanings to sentences when they are used with the meanings assigned by the first-order conventions. Thus, first-order conventions dictate that ‘Some politicians take bribes’ means that some politicians take bribes, and a second-order convention dictates that using that sentence with that meaning expresses the *further* meaning that not all politicians take bribes. As Davis puts it:

The first-order rules are conventions for using sentences to *directly* express certain thoughts. The second-order rules are conventions for *indirect* expression, rules for expressing further thoughts by expressing thoughts assigned by first-order rules. (Davis 1998, p.156)¹⁸

Davis holds that languages are defined by their first-order rules, not their second-order ones and that a language’s implicature conventions are not essential

¹⁷ Morgan also argues for the existence of implicature conventions, though from a Gricean perspective (Morgan 1978). Take the use of ‘Can you X?’ to request that the hearer do X. Originally, Morgan suggests, this implicature was a particularized one, and the hearer inferred the connection between the literal and implicated content by applying Gricean principles. But as the use of such indirect requests spread, a convention was established whereby one could request someone to do X by saying ‘Can you X?’, and hearers no longer needed to calculate or even notice the rational connection. The implicature became, as Morgan puts it, ‘short-circuited’. Although Davis denies that implicatures are calculable in the Gricean way, he envisages a similar historical process, in which repeated use of a one-off implicature gradually establishes a conventional connection between literal and implicated content (Davis 1998, pp.164–5).

¹⁸ This distinction corresponds closely to Searle’s distinction between *conventions of language* and *conventions of usage* (Searle 1975; see also Morgan 1978).

to it. He compares implicature conventions to *speech act rituals*, such as saying ‘N speaking’ when answering the telephone or asking ‘How are you?’ when greeting someone. Because implicature conventions are second-order, Davis predicts that second-language learners should take longer to master them than lexical conventions, at least when they are different from those of their first language (Davis 1998, pp.161–2).

Davis also contrasts sentence implicatures with idioms, such as ‘kicked the bucket’ (meaning ‘died’) (1998, pp.162–6). Like sentence implicatures, idiomatic meanings are not derivable from the literal meaning of the component words, though there is some relation between the literal and idiomatic meaning that makes the connection appropriate. However, unlike sentence implicatures, idioms do not depend on the current literal meaning of the words used. ‘Kicked the bucket’ does not mean ‘died’ in virtue of meaning ‘struck the bucket with the foot’. Davis suggests that idioms typically start life as nonce implicatures (metaphors), which later become conventional and finally fossilized into idioms. The literal meanings dropped out of the picture, and the phrases came to express the idiomatic meanings directly. (In effect, second-order conventions became first-order ones.)¹⁹

I shall refer to Davis’s view of generalized implicature as *convention theory*. As will be clear from the previous summary, convention theory is a theory of implicature generation — of what makes it the case that certain utterances carry certain implicatures. There is therefore no *direct* comparison between it and either neo-Griceanism or relevance theory, which are (primarily) theories of implicature recovery. However, as we have seen, issues of generation and recovery are closely interconnected and convention theory does have some implications for implicature recovery. (I should stress that these implications are not identified by Davis himself, who focuses on issues of implicature generation. My remarks here and

¹⁹ As noted in Chapter 2, Grice also holds that some implicatures are conventional. For example, ‘p therefore q’ conventionally implicates that q follows from p (Grice 1989, pp.25–6). However, these implicatures are different from sentence implicatures in Davis’s sense. They are part of the meaning of the words used and depend on first-order conventions. Davis’s claim is that many of what Grice regarded as *non-conventional*, ‘conversational’ implicatures are also conventions, though of a second-order kind (Davis 1998, p.157).

later in this chapter may be thought of as a preliminary sketch for a cognitive counterpart to convention theory.)

First, if convention theory is correct, then recovering a generalized implicature will involve knowing and applying the relevant second-order convention, as opposed to applying a general inferential principle or searching for an optimally relevant interpretation. Exactly what cognitive states and processes are involved in this is of course an open empirical question. Second, like neo-Gricean theory, convention theory holds that some implicature processing at least is not context-driven, but involves applying context-independent rules. Third, convention theory agrees with relevance theory that literal (linguistically coded) meanings are in a sense psychologically more basic than implicated ones. As we have seen, knowledge of literal meanings requires mastery of first-order conventions only, whereas knowledge of implicatures requires mastery of second-order conventions as well. Since it is possible to acquire the former without the latter (as in the case of second language learners), this suggests that knowledge of lexical conventions may be stored and accessed separately from knowledge implicature conventions. As we shall see in section 5, this means that convention theory may offer new ways of interpreting experimental work on scalar implicature.

2.3 *Weak neo-Griceanism*

The fourth approach to implicature recovery that I want to introduce is one that, so far as I know, does not have a name, though it is a possible position and one that is represented in the literature. (I shall consider an example later.) It is an inclusive position, which can be seen as a halfway house between neo-Griceanism and relevance theory.

Neo-Griceanism can be thought of as involving two core claims: (1) some implicatures are derived from the application of broadly Gricean principles, and (2) these principles are applied *by default*. Claim (2) itself can be understood to mean (a) that the principles are applied automatically whenever an appropriate expression is detected, regardless of context, and (b) that they yield the same interpretations of the same expressions each time. (Subclaim (b) follows from subclaim (a); it is because the principles are not sensitive to context that they yield the same interpretations each time.) Claim (2) is important to Levinson's view that

the principles serve to speed up language processing; it is because they provide rapid context-independent enrichments that they save time. But (1) does not entail (2), and the principles might still serve a useful interpretative function even if they are not applied by default.

The view I want to introduce accepts (1) but rejects (or remains neutral about) (2). That is, it holds that we derive some implicatures by applying Gricean principles, such as the Q-principle, but does not claim that we do so by default. It allows that contextual factors may determine when and how the principles are applied, and that the principles may yield different interpretations of the same expressions in different contexts. This view agrees with neo-Griceanism that some implicatures are derived by applying general pragmatic principles but holds that the principles in question are applied in a context-sensitive way — perhaps when a literal interpretation fails to meet relevance expectations. Hence, it agrees with relevance theory that there are no truly generalized, context-independent implicatures. I shall refer to this broad approach as *weak neo-Griceanism* (‘Griceanism’ because its principles are derived from Grice’s; ‘neo’ because it applies Gricean principles to implicature recovery; and ‘weak’ because it is not committed to (2)).

2.4 Back to Levinson

Having outlined these broad alternatives to neo-Griceanism, I shall now turn to the task of assessing neo-Griceanism itself, in the form developed by Levinson. I shall highlight some problems for the theory and indicate how one or other of the alternative approaches might be applied instead. I shall not attempt to adjudicate between the alternatives themselves but merely show that each may better explain some of the cases discussed.

As already noted, Levinson discusses a huge number of examples, often persuasively, and I cannot possibly engage with the range and detail of his analyses. (Nor, indeed, can I consider more than a tiny fraction of the alternative analyses offered by the rival approaches.) But I will focus on some key problem cases.

3. Assessing the Q-principle

This section looks at some problems relating to Q-implicatures. It looks first at a core example, and then discusses some more general theoretical concerns about the neo-Gricean treatment of scalar implicatures.²⁰

3.1 'An X'

Since Neo-Griceanism is derived from Grice's account of generalized implicature, I will begin by returning to the example Grice uses to introduce the notion (which he describes, in a characteristically cautious way, as one that he 'hope[s] may be fairly noncontroversial'; Grice 1989, p.37). This example, which I discussed briefly in Chapter 2, concerns the indefinite article. Grice notes that when a speaker uses an expression of the form *an X*, they typically implicate that the X in question is not closely connected to them. For example, 'I met a man' implicates that the man was not my close relation or friend, 'I found a cat' implicates that the cat was not mine or one known to me, and so on. Grice proposes that this is due to the first maxim of Quantity:

the implicature is present because the speaker has failed to be specific in a way in which he might have been expected to be specific, with the consequence that it is likely to be assumed that he is not in position to be specific. (Grice 1975/1989, p.38)

Grice explains that when a person or object is familiar to the speaker, it will usually be informative to indicate that it is, since our interactions with familiar persons and objects are typically very different from our interactions with unfamiliar ones. If a

²⁰ One common area of application for the Q-principle is in relation to number words. Neo-Griceans typically hold that these literally specify only minimum amounts ('three' literally means 'at least three') and that the exact meanings they commonly carry ('exactly three') are the product of implicatures generated by application of the Q-principle. However, number words present special problems (see, for example, Carston 1998; Levinson 2000, p.88), and I shall not focus on them here.

speaker does not indicate that the person or item in question was familiar (for example, by using ‘my’ rather than ‘an’), they implicate that it was not familiar.

There is a problem with this, however. For, as Grice notes, in some cases the implicature does not hold. ‘I have been sitting in a car all morning’ does not implicate that the car was not my own. And in some cases the opposite implicature holds. ‘I broke a finger yesterday’ implicates that the finger was mine. (Both examples are Grice’s.) Since there is no explicit cancellation in these cases, how can ‘an X’ carry a generalized implicature of lack of connection?

Levinson revisits this case and proposes a slightly different and more comprehensive account. He points out that the implicature from ‘an X’ to ‘not my X’ cannot be a simple Q-implicature, since possession and indefiniteness are different types of relation, and there is therefore no scale $\langle my, a \rangle$. Rather, he argues, it is a two-stage process, involving a Q-implicature followed by an I-implicature. First, there is an entailment scale $\langle the, a \rangle$ since definite and indefinite reference are similar relations, so ‘an X’ Q-implicates ‘not the X’. In using the indefinite article, we implicate that we do not mean to refer to some definite, unique X. Second, ‘the X’ I-implicates ‘the salient X’ — that is, the one I am familiar with or closely connected to in some way.

Grice’s examples are thus indirectly explained: when one says ‘I went into a house’ one Q-implicates ‘I didn’t go into the house’, where the definite suggests (I-implicates) my very own house. (Levinson 2000, p.92)

This is neat. But how is ‘I broke a finger’ to be explained? According to this account, it should implicate that the speaker cut someone else’s finger, not their own, but that would not be the usual interpretation. Levinson agrees but argues that this is not a counterexample. He explains (taking ‘I cut a finger’ as his example):

[W]hen I say ‘I cut a finger’ I merely implicate that it was no unique, otherwise salient finger (say, the one I cut before) that suffered, and that interpretation is compatible with the assumption that it was my own

finger (which in turn is a more stereotypical reading than one that involves the chopping of other peoples' fingers). (Levinson 2000, p.92)

The idea is that the Q-implicature to 'not *the* finger' goes through as usual, indicating that the speaker did not mean to pick out some specially salient finger, such as the one that they had been planning to cut or had cut the day before. But this does not rule out the finger being his or her own.

There are problems with this response, however. First and most obviously, although the implicated message is compatible with the claim that the finger in question was one of the speaker's own, it is also compatible with the claim that the finger was someone else's. If a manicurist were talking, for example, it would be natural to read the utterance as implicating that they had cut a client's finger — as Levinson himself acknowledges (2000, p.17). Yet the implicated message is not that the finger *might* have been my own, but that it *was* my own. The parenthesis in the passage just quoted suggests that Levinson would respond by appealing to I-principle. If the speaker is not a manicurist, then the stereotypical reading will be one on which the finger was their own. But this cannot be right. For the I-principle (we are assuming) tells us to draw on *general* background knowledge, not on knowledge of the specific context. This would include knowledge that manicurists use sharp objects on other people's fingers, whereas non-manicurists rarely do this, but it would not include knowledge that the speaker himself is a manicurist. That is context-specific knowledge and hence irrelevant to GCIs. So the implicature that the finger cut was someone's else cannot be a *generalized* one. At most there is a generalized implicature that it was not the salient one — whichever that might be.

A second problem for Levinson's account is that it does not explain other very similar cases. In English, 'I cut a finger' implicates that the cut finger was the speaker's own, but the formally and semantically similar 'I cut a nose' typically carries the opposite implicature — that the nose was *not* the speaker's own. On Levinson's account it is hard to see why this should be so. Use of 'a nose' should implicate 'not *the* nose' — the unique, otherwise salient nose — and that is compatible with the assumption that it was the speaker's own. Moreover, an appeal to stereotypicality reinforces this. We are more likely to cut our own noses than to cut other people's, just as we are more likely to chop our own fingers than those

of others. It seems that other factors or principles must be in play here, beyond those mentioned by Levinson. In fact, the uniqueness or otherwise of the bodily part in question seems to be crucial here. ‘I hurt an X’ implicates that the X is the speaker’s own if the speaker has more than one X (for example, finger, toe, ear, breast, testicle), but implicates ‘The X was someone else’s’ if the speaker has just one X (for example, head, nose, chin, penis, vagina, etc.). We might call this the *Uniqueness Principle*. Why should this hold? It might be suggested that since ‘an X’ implicates ‘not the X’, it cannot refer to a unique feature of the speaker. Compare ‘A wheel of the car is loose’, where the use of ‘a’ rather than ‘the’ indicates that the reference is to a (non-unique) road wheel rather than the (unique) steering wheel (example adapted from Levinson 2000, p.155). The cases are not parallel, however. A speaker would not refer to their own nose as ‘the nose’, and ‘nose’ unlike ‘wheel’ is not ambiguous between unique and non-unique features. Besides, this does not get to the heart of the matter. The implications we need to explain are not of uniqueness or otherwise, but of *possession*. Why should ‘a nose’ implicate ‘not my nose’ while ‘a finger’ implicates ‘my finger’? It is not obvious that this can be explained with the resources Levinson has to offer, and a more plausible explanation may be that it is simply a convention of English usage, in line with convention theory.²¹

It turns out then, that Levinson’s account of the ‘an X’ case, does not fare much better than Grice’s. If GCIs are genuinely context-independent, then the same expression-form should give rise to the same implicature in every context (and indeed every language), but ‘an X’ does not, and neither Grice nor Levinson can fully explain why. Given that this was the example with the notion of generalized implicature was originally introduced into the literature, this is a problem for

²¹ Informal discussions with speakers of other languages suggest that the Uniqueness Principle is indeed a convention of English. In many languages it is not applicable, since it is not felicitous to use a ‘a leg/nose’ in this context without explicitly indicating whose it is, by use of a pronoun or reflexive verb. However, in languages where the phrase is not infelicitous, the principle does not always hold. In Finnish, for example, the reference would be to the speaker’s own X, whether unique or not, and in Greek the reference would be ambiguous. (My thanks to the friends and correspondents who have shared their intuitions on this topic with me.)

Gricean approaches. Even if there are generalized implicatures associated with use of the indefinite article, they do not appear to be derivable from general principles and may depend on language-specific conventions.

3.2 *Scalar implicatures*

I turn now to some more general considerations about Q-implicatures, and, specifically, scalar implicatures. These are perhaps the clearest examples of GCIs, and Levinson's account of them builds on a rich body of pre-existing work on the topic, including Gazdar 1979, Horn 1972, 1984, 1989, and Hirschberg 1985. However, even here there are reasons for thinking that the neo-Gricean account may not be the best. I shall begin with a general worry about the Gricean approach to scalar implicature, raised by Wayne Davis (2014).

As Davis points out, the idea behind Quantity implicatures is that we recover the implicated message by reference to what is *not* said. He quotes Levinson's own (1983) account of the implicit reasoning involved:

- (i) S has said p
- (ii) There is an expression q , more informative than p (and thus q entails p), which might be desirable as a contribution to the current purposes of the exchange (and here there is perhaps an implicit reference to the maxim of Relevance)
- (iii) q is of roughly equal brevity to p ; so S did not say p rather than q simply in order to be brief (i.e. to conform to the maxim of Manner)
- (iv) Since if S knew that q holds but nevertheless uttered p he would be in breach of the injunction to make his contribution as informative as is required, S must mean me, the addressee, to infer that S knows that q is not the case ($K\sim q$), or at least that he does not know that q is the case ($\sim Kq$).

(Levinson 1983, p.135)

But of course, in any exchange there are *many* things that are not being said. As an example, Davis takes 'Some athletes smoke'. This implicates the denial of the stronger claim that all athletes smoke. But that is far from being the only stronger

relevant statement of roughly equal brevity the speaker could have made and did not. Davis illustrates this with a range of scales on all of which ‘Some athletes smoke’ figures as the weakest element:

<All athletes smoke, Nearly all athletes smoke, Most athletes smoke, Many athletes smoke, Several athletes smoke, Some athletes smoke>

<100% of athletes smoke, At least 90% of athletes smoke, At least 50% of athletes smoke, At least 10% of athletes smoke, At least 1% of athletes smoke, Some athletes smoke>

<Some athletes smoke constantly, Some athletes smoke regularly, Some athletes smoke often, Some athletes smoke occasionally, Some athletes smoke>

<Some athletes, maids, and cops smoke, Some athletes and maids smoke, Some athletes smoke>

<Some athletes smoke filterless Marlboros, Some athletes smoke Marlboros, Some athletes smoke>

<Everyone knows some athletes smoke, I know some athletes smoke, Some athletes smoke>

< $n\%$ of athletes smoke ($0 < n < 100$), Only some athletes smoke, Some athletes smoke>

(Adapted from Davis 2014)

By Levinson’s reasoning, Davis argues, ‘Some athletes smoke’ should implicate the denial of all of the stronger statements. Yet in fact it implicates the denial of only one of them:

Among the infinity of statements stronger than ‘Some athletes smoke,’ ‘All athletes smoke’ is highly unusual in that people typically implicate its denial. (Davis 2014)

Now Levinson has an answer to this. In the passage quoted above he mentions that the stronger statement must be ‘of roughly equal brevity’ and ‘desirable as a contribution to the current purposes of the exchange’. And in later work he sets

out two more precise constraints that an entailment scale must meet in order to support Q-implicatures. First, the stronger items in the scale must be lexicalized to at least the same degree as the weaker ones. That is, the stronger items must consist of as few or fewer words than the weaker ones, so that, for example, if the weakest item is monolexemic, then all the other elements are monolexemic too. Second, all the items in the scale must be ‘about’ the same semantic relations and thus ‘in conceptually salient opposition’. So for example, the scale *<regret, know>* does not support Q-implicatures, since ‘regret’ involves a conceptual element not present in ‘know’ (Levinson 2000, p.80). And these conditions rule out most of Davis’s examples. None of his scales except the first meets the lexicalization constraint, and the penultimate one at least fails the aboutness constraint.

The lexicalization constraint also gives Levinson a response to what would otherwise be a serious objection. He holds that ‘and’ is typically strengthened by application of the I-principle to indicate temporal or causal sequence (2000, p.37–8). So ‘They got married and had a child’ I-implicates that the marriage preceded the child’s birth. But it seems that the Q-principle could also be applied here to produce precisely the opposite implicature. Since the speaker chose ‘and’ rather than the informationally stronger ‘and then’, we might infer that they were not in a position to assert that the events took place in that order described and thus that they do not know that they did and perhaps know that they did not. Since Q-implicatures take precedence over I-implicatures, we should therefore take the utterance to implicate that the child’s birth did not take place after the marriage (Davis 1998, p.52–3). However, the scale involved here, *<and then, and>*, does not satisfy the lexicalization constraint, since the stronger element is less lexicalized than the weaker one, and it does not therefore support Q-implicatures (Bezuidenhout 2002, p.264–5).²²

²² We might press the objection, pointing out that some speakers use ‘then’ (incorrectly, according to grammarians) as a coordinating conjunction, as in ‘They got married, then had a child’, which suggests that in their idiolect at least there is a legitimate scale *<then, and>* which does support the Q-implicature. However, it may be that in these cases the comma before ‘then’ serves as a coordinating conjunction, so the real scale is *<[comma] then, and>*, which, arguably, does not meet the lexicalization condition.

But why should the Q-principle be restricted in this way? (The fact that it saves Levinson's account from a serious objection is not itself a reason, unless one is already convinced that the account is correct.) If we typically assume that what is not said isn't the case, then why doesn't saying that some athletes smoke implicate that it is not the case that some athletes, maids, and cops smoke? What justifies the two constraints? Levinson suggests answers. He justifies the lexicalization constraint on the grounds that where stronger items are *less* lexicalized (more wordy), then any Q-implicature would be undercut, since the hearer might think that the speaker had avoided the stronger term simply because they were avoiding being 'clumsy and prolix' (following Grice's maxim of Manner), rather than because they were not in a position to assert it (Levinson 2000, pp.79–80). I assume the aboutness constraint is justified in a similar way, by reference to the maxim of Relevance. Where a stronger term would have introduced a different kind of information, the hearer may take the speaker to have avoided it simply in order to remain relevant, rather than because they were not in a position to assert it, thus undercutting any potential Q-implicature it might have supported. These justifications are not unreasonable, and plainly the Q-principle (if it is a principle) would have to be restricted in some way — otherwise it would produce endless implicatures from every utterance. However, we might wonder if an appeal to Gricean maxims is sufficient to justify the conditions in the strict form Levinson proposes. The assumption that a speaker is following the maxims of Manner and Relevance doesn't require us to suppose that they would have avoided even *slightly* longer phrases or introduced *any* new information. Moreover, even if we accept the conditions, some problems remain.

First, the constraints do not rule out all the problematic scales. For example, consider the scale <*several, some*>. This meets the two constraints, but does not support GCIs. 'Some athletes smoke' does not imply that it is not the case that several athletes smoke. Levinson might reply that 'several' forms part of a larger scale that continues up to 'all' and that it is only the strongest element on the scale whose denial is implicated (Levinson 2000, p.77). However, even if this further qualification is added, it is arguable that exceptions remain. For example, consider:

<*is a cardiologist, is a physician*>

<commit murder, commit a crime>

<likes baseball, likes sports>

Again, these scales meet Levinson's constraints but do not support generalized scalar implicatures. 'X is physician' does not generally implicate that X is not a cardiologist; 'Some cops commit crimes' does not generally implicate that no cops commit murder, 'Y likes sports' does not generally implicate that Y does not like baseball (though there might be specific contexts in which those implicatures would hold).

Second, there are entailment scales which do not meet the constraints but *do* plausibly support GCIs. Consider, for example:

<got a distinction, passed>

<right up to, near>

<got a good look at, saw>

These scales do not meet the lexicalization constraint. Yet my intuition is that they support scalar implicatures. 'Amy passed' implicates that Amy did not get a distinction; 'Bob went near the edge' implicates that Bob didn't go right up to the edge; 'Cal saw the robber' implicates that Cal didn't get a good look at the robber.²³

Perhaps further constraints could be added to deal with these exceptions, but they would begin to look ad hoc. Moreover, if the Q-principle were supplemented with even more constraints, it would no longer look like a rule that could be applied automatically at an early stage of processing. Applying it would be a complicated business, which would involve checking that multiple conditions hold, and it might slow down communication rather than improve its efficiency. If there are such tight constraints on the application of the Q-principle — with the result that the

²³ This is confirmed by the recognized diagnostics for scalar implicatures, summarized by Levinson (2000, p.81). If a scale, <S, W>, supports scalar implicatures, then the following cancelling and suspending phrases should be permissible: 'W and even S', 'Not only W, S', 'W in fact/indeed S', 'W or possibly/even S' and 'W if not S'. This is the case with the three scales mentioned.

exceptions to it hugely outnumber the cases to which it applies — then why posit the principle in the first place? Wouldn't it be more economical to appeal to conventions of use rather than the general principles? Perhaps it is a convention of English (and other languages) that when we use 'some', with its basic meaning of 'at least one', it is understood to implicate 'not all'. Thus, in order to recover the implicature, hearers would simply need to know the convention and recognize that it applies in this case. This is a simpler hypothesis than supposing that speakers have to access a relevant scale, check that the scale meets multiple constraints, and then apply an inferential principle.

This conclusion is reinforced, I think, by another point made by Davis (Davis 2014). If someone asks 'Do any athletes smoke' the response 'Some do' will carry the implicature that not all athletes smoke, but the answer 'Yes' will not. Yet, Davis points out, in the context 'Yes' is logically equivalent to 'Some do' and is a no less cooperative response. Since the two exchanges are informationally equivalent we should expect them to produce the same implicatures, if general principles are at work. If the Q-heuristic produces a scalar implicature in the first case, then it should do so in the second too. The fact that no implicature arises in the second case strongly suggests that a general principle is not involved in either case. Again, it seems more appropriate to appeal to a convention of use, which is associated with particular expressions. There may be a convention of use that 'some' (in the right context) implicates 'not all', but there is of course no convention that 'yes' implicates 'not all'.

3.3 Reducing scalar GCIs to PCIs

There is another way of looking at scalar implicatures, which stresses their continuity with particularized implicatures. Levinson allows that there may be other types of Q-implicature, in addition to those based on entailment scales and clausal contrasts (see Levinson 2000, pp.98–103, from where the examples below are taken). For example, Q-implicatures can be based on *non-entailment* scales, such as <*succeed, try*> (saying that John tried to reach the peak implicates that John did not succeed, even though succeeding does not entail trying). Q-implicatures can also be based on sets of *alternatives*, where the choice of one alternative implicates that the others do not apply (for example, 'The flag is white'

implicates ‘The flag is not white and red’), and on *levels of specificity*, where the use of a more general term implicates that the speaker cannot be more specific (for example, ‘I just saw a horrid animal in the larder’ implicates that the speaker is not sure what sort of horrid animal it was). Levinson notes that many of the inferences underlying these implicatures are weak unless contextually reinforced, and thus lie at the border between GCIs and PCIs (2000, p.103).

Following Fauconnier (1975), Levinson also notes that scalar implicatures can be generated by contingent scales, which depend on our beliefs about the world rather than on the meanings of the terms involved. He offers this example (Levinson 2000, p.104):

- (4) He can drive small trucks
Implicature: He can’t drive big ones.

The Q-implicature here depends on the scale *<driving big trucks, driving small trucks>* which is based on knowledge of truck-driving rules and skills (people licensed to drive big trucks are also allowed to drive small ones, but not vice versa).

Q-implicatures can also depend on contextually given *nonce* scales. Levinson quotes the following example (Levinson 2000, p.105, quoted from Hirschberg 1985, p.50):

- (5) A: Did you get Paul Newman’s autograph?
B: I got Joanne Woodward’s.
Implicature: I didn’t get Paul Newman’s.

Here the speaker assumes a scale of autograph prestige *<Newman, Woodward>*, and by affirming that they secured the lower value item, they implicate that they did not secure the higher-value one.

Julia Hirschberg proposes a systematic treatment of scalar implicatures which includes such contingent and context-dependent ones (Hirschberg 1985; for

discussion, see Levinson 2000 pp.104–8).²⁴ According to this, scalar implicatures are supported by *orderings* (Levinson calls them ‘Hirschberg scales’) constructed from a contextually salient set of *values* (expressions) and an *ordering relationship* of some kind (it can be any relation that is salient). For example, the values ‘oak’, ‘maple’, ‘tree’ and the ordering relation *is-a-kind-of* would give the following ordering:

<{*oak, maple*}, *tree*>²⁵

Note that this ordering is a *partial* one, in that it does not apply to every pair of items in the set. ‘Oak’ and ‘maple’ are both ordered with respect to ‘tree’, but not with respect to each other. In Hirschberg’s account, scalar implicatures require only partial orderings to support them. Similarly, the ordering relationships *has-parts*, *has-attribute*, and *has-prior-stage*, might yield the following orderings (examples quoted in Levinson 2000, pp.106-7):

<{*book, {chapter 1, chapter 2, ...}*>

<*Greek, {Greek-speaking, Greek relatives, Greek residency, Greek ancestry}*>

<*marriage, engagement, going-steady, dating*>

The rules for scalar implicatures are then as follows. Affirming a lower expression in an ordering (to the right) implicates either that the speaker doesn’t believe that a higher expression applies or that they do not know which, if any, does. Thus ‘It’s a tree’ implicates that the speaker does not know which kind of tree; ‘I’ve read Chapter 1’ implicates that the speaker hasn’t read the whole book, and ‘I’ve Greek relatives’ implicates that the speaker is not Greek. By contrast, denying a higher

²⁴ Hirschberg is a computer scientist, and her aim is to develop a formal framework for the representation and calculation of scalar implicatures that could be implemented computationally, rather than to identify the psychological mechanisms involved in human implicature recovery. The framework is detailed and complex and only its broad outlines are relevant here.

²⁵ Example from Levinson 2000, p.106.

item implicates that the speaker believes that a lower one applies, or may do so. Thus, ‘It’s not an oak’ implicates that it is a tree, ‘I haven’t read the whole book’ implicates that the speaker has read some of the chapters, ‘We’re not married’ implicates that the speaker may be engaged or dating, and so on.

Hirschberg holds that scalar implicatures can also be generated by *unordered* sets of alternatives, such as {*chapter 1, chapter 2 ...*}. Here the rule is that affirming one expression implicates that the others do not apply or are not known to apply, and denying one expression implicates that one of the others may apply. So, for example, ‘I’ve read Chapter 1’ implicates that the speaker has not read Chapter 2, and ‘I’ve not read Chapter 1’ implicates that the speaker may have read Chapter 2 (Levinson 2000, p.106).

Crucially, Hirschberg extends this treatment to scalar implicatures based on entailment scales. Given the contextually salient expressions ‘all’ and ‘some’ and the relation of entailment, we can form the Hirschberg scale $\langle all, some \rangle$. Applying the rules, ‘Some came’ implicates that not all came, and ‘Not all came’ implicates that some came. Since these implicatures too are generated from contextually salient orderings, it follows that there is no sharp distinction between PCIs and GCIs, and that Hirschberg’s approach reduces GCIs to PCIs.²⁶

Levinson rejects this conclusion, of course. Though he concedes that Hirschberg offers a neat treatment of particularized scalar implicatures, he argues that it does not tend to undermine the distinction between GCIs and PCIs, since we can still draw a clear-cut distinction between context-independent scalar implicatures that are based on contrasts in meaning (GCIs) and context-dependent ones that are based on contrasts salient in particular contexts (PCIs):

The GCI theorist is simply claiming that speakers carry their lexicons on their backs, as it were, from context to context, and it is mutual knowledge of this fact that elevates the Q-heuristics to a default mode of inference.

(Levinson 2000, p.108)

²⁶ Hirschberg writes: ‘the traditional distinction between generalized and particularized implicature is a false one, an artefact of the inventiveness of analysts — or lack thereof’ (Hirschberg 1985, p.42).

This is questionable, however. Levinson assumes that some expressions (such as ‘some’, ‘sometimes’, ‘possibly’) will evoke the same set of contrasting values and the same ordering relation (namely, entailment) in all contexts, thus supporting context-independent scalar implicatures. That is, the scales that support GCIs will be salient in all contexts. But this is questionable. Consider these exchanges, for example:

(6) A: Is it true that she bought ten pairs of shoes yesterday?

B: She bought some.

(7) A: Do you visit her as often as you used to?

B: We visit her sometimes.

In the context of (6) ‘some’ does not evoke the entailment scale *<all, some>* but the nonce specificity scale *<ten, some>*, and by using ‘some’ B implicates that she doesn’t know if the more specific figure of ten is correct. Similarly, in (7) the entailment scale *<always, sometimes>* is not salient, but instead the nonce scale *<as often as we used to, sometimes>* is. Of course, in many contexts where ‘some’ is used, the familiar entailment scale *<all, some>* would be salient, but as the examples just given show, there are contexts in which it would not be. Given this, it is more appropriate to think of scalar implicatures as lying on a continuum, from relatively particularized ones, which depend on orderings that are salient only in a few contexts, to relatively generalized ones, which depend on orderings that are salient in many contexts. But since this difference is a matter of degree and there will be a full range of intermediate cases, this view tends to undermine the sharp neo-Gricean distinction between PCIs and GCIs, and, with it, Levinson’s case for existence of a distinct level of utterance-type meaning.²⁷

²⁷ This is not to deny, of course, that there are general principles at work in scalar implicature, on Hirschberg’s account. There are the rules that affirming a lower-ranked element in an ordering implicates the denial of stronger one, and that denial of a stronger element implicates affirmation

Levinson also objects that Hirschberg's account will overgenerate implicatures, since it imposes no constraints on scalehood (echoing Davis's complaint against him, discussed above) (Levinson 2000, p.107). But this, I think, mistakes Hirschberg's aims. In fact, it is a virtue of her account that it imposes no such constraints. For, as Hirschberg stresses, given the right ordering, *any* expression can generate a scalar implicature.²⁸ Consider again the examples Davis gives. There is no general implicature from 'Some athletes smoke' to 'It is not the case that some athletes smoke Marlboros', but there are contexts where it arises:

- (8) A: Do some athletes smoke Marlboros?
B: Some athletes smoke.

The implicature depends on the contextual salience of the expression 'smoke' and 'smoke Marlboros' and the ordering relation *is-a-specific-form-of*. Affirming that some athletes engage in the general activity of smoking implicates that they do not

of a weaker one. But these rules by themselves do not generate implicatures, even when combined with the lexicon. A contextually salient ordering must also be given.

²⁸ Hirschberg illustrates this with the following example. Each of B's replies generates a different scalar implicature, based on a different implicit ordering or set of alternatives:

- A: Did the girl in the red dress spill a diet coke?
- a. B: She spilled a diet *pepsi*.
 - b. B: She spilled a *regular* coke.
 - c. B: She spilled a *glass of tomato juice*.
 - d. B: *Jane* spilled a diet coke.
 - e. B: The girl in the red *slacks* spilled a diet coke.
 - f. B: The girl in the *green* dress spilled a diet coke.
 - g. B: The girl in the *green slacks* spilled a diet coke,
 - h. B: The *boy* in the red dress spilled a diet coke.
 - i. B: The girl in the red dress *will spill* a diet coke.
 - j. B: The girl in the red dress *drank* a diet coke.
 - k. B: The girl in the red dress spilled *the* diet coke.

engage in the more specific Marlboro-involving form of smoking (or are not known to do so).²⁹

This suggests another way of looking at the constraints on scalehood that Levinson proposes. I have already suggested that these are inadequate, and I think we can now see why. If Hirschberg is right, then our intuitions about which scales do and do not support scalar implicatures are based on our judgements about contextual salience. We judge that the entailment scale <*all, some*> does support implicatures and that the specificity scale <*ten, some*> does not because in most contexts the former would be salient and the latter would not. But, as we have seen, there are contexts in which the reverse is the case. We cannot hope to provide general rules of scalehood since there are no truly general, context-independent scalar implicatures. Any rules would have potentially unlimited context-specific exceptions.

Of course this means that the notion of contextual salience has a lot of work to do in Hirschberg's account, and Levinson suggests that this is a major problem for the account:

All implicatures are made dependent on the contextually salient ordering relation, so we have no account of implicature generation without an account of how this is arrived at. (Levinson 2000, p.107)

In fact, Hirschberg devotes a whole chapter to the discussion of contextual salience, identifying some of the cues that make an ordering salient to a speaker and hearer (syntactic, intonational, semantic, pragmatic, and communication dynamical) and proposing ways in which assignments of salience might be

²⁹ Other examples from Davis's list would generate implicatures based on unordered sets of alternatives:

A: Is it true that athletes, maids, and cops smoke?

B: Some athletes smoke.

Here the salient set of alternatives is {*athletes, maids, cops*}, and the affirmation of one element implicates the denial of the others.

formally represented in a computational model (Hirschberg 1985, Chapter 6). It is true that this does not amount to (and is not intended to be) a complete theory of contextual salience (indeed, providing such a theory would be a major achievement for psycholinguistics). But that does not undermine the case for thinking that we need such a theory, and Levinson has not shown that we don't need one.

According to Hirschberg, then, there are genuine scalar implicatures, whose recovery involves the application of general principles (affirmation of a weaker scalar term implicates denial of a stronger one, and denial of a stronger one implicates affirmation of a weaker one). But these principles do not yield default, context-independent interpretations. The *orderings* to which the principles are applied are contextually determined, and the same expression might evoke a different ordering, and hence a different scalar implicature, in different contexts. Thus, Hirschberg's approach (at least as I have interpreted it) is an example of the class of views I called *weak neo-Griceanism*.³⁰

3.4 Q-implicature and T-implicature

There is another set of considerations that pose a challenge for Levinson's account of Q-implicature. Levinson's principles, like Grice's maxims and the Cooperative Principle from which they follow, are rooted in the assumption that the aim of communication is the efficient sharing of information. Thus, the Q-principle tells us to assume that speakers will give as much (relevant) information as they can. But as several writers have noted, conversation often has other aims besides the communication of information. Sometimes it is more important to be polite than to be informative. Geoffrey Leech has formalized this idea, proposing a *Politeness Principle*, 'Minimize (other things being equal) the expression of impolite beliefs', which he breaks down into a series of maxims, of Tact, Generosity, Approbation, Modesty, Agreement, and Sympathy (Leech 1983, p.132). Leech notes that

³⁰ The particularized implicature carried by Mr Bronston's reply 'The company had an account there', discussed in Chapter 1, might be regarded as a scalar implicature, dependent on the ad hoc scale <me, my company>, where the ordering relation is something like *seriousness of holding a bank account in the name of*.

communicative exchanges frequently involve a trade-off between the Cooperative Principle and the Politeness Principle, and he explores in detail the complex pragmatics of politeness. Here I shall focus on a narrow range of cases, in order to highlight a potential problem for neo-Griceanism.

In British English at least it is common to use understatement to convey information or instructions that will be unwelcome to the hearer. We might see this practice as obeying what Leech calls the *Tact maxim*: ‘Minimize the expression of beliefs which express or imply cost to other’ (Leech 1983, p.132). For example, a manager might tell a subordinate, ‘There is a problem with your report’, to convey that they are in fact seriously displeased with it. We might call this a *Tact implicature*, or *T-implicature*. Here are some more examples:

(9) I might not be able to do that.

T-implicature: I won’t do that.

(10) I would like to see your passport.

T-implicature: You must show me your passport.

(11) Someone’s eaten the icing off the cake.³¹

T-implicature: You have eaten the icing off the cake.

Note that these implicatures are not M-implicatures. The expressions used are not untypical or unusual. Indeed, these sentences could be used without generating the implicatures. If (9) were uttered by a friend who was clearly anxious to help as much as they could, or (10) by someone known to be interested in the design of passports, the implicatures would not arise. To this extent, then, these are context-dependent, particularized implicatures. However, they are not *highly* contextualized, and the examples given can be easily understood without any background information.

In these cases, the hearer applies what we might call a *Tact principle*: interpret the affirmation of a weaker, less unwelcome statement or request as implicating a

³¹ Example borrowed from Leech 1983, p.80.

relevant stronger, more unwelcome one. What isn't said, is, we might say. This is, of course, the opposite of Levinson's Q-principle, which says that affirmation of a weaker claim implicates the denial of a relevant stronger one. Moreover, the same utterance may potentially generate both a Q-implicature and a T-implicature. In (9), 'I might not be able' should Q-implicate that it is not the case that the speaker definitely won't be able to do the thing requested — which is, of course, at odds with the T-implicature that the speaker won't do it. Here are some more examples:

(12) It is possible that the train will be delayed.

Q-implicature: It is not probable that the train will be delayed.

T-implicature: It is probable that the train will be delayed.

(13) Some of the staff you sacked are angry.

Q-implicature: Not all the staff you sacked are angry.

T-implicature: Many or even all of the staff you sacked are angry.

(14) I think you dropped this.

Q-implicature: I am not sure that you dropped this.

T-implicature: You dropped this.

Given the right context, the T-implicatures here would take precedence over the Q-implicatures. (Imagine (12) said by a grim-faced railway employee.) And this poses a problem for Levinson's account, which holds that Q-implicatures are generalized and take priority over other implicatures. It is true, as I mentioned earlier, that Levinson allows that Q-implicatures can be cancelled if they conflict with entailments of what is said or with background assumptions, or if they are obviously irrelevant. But it would be a major concession to allow that Q-implicatures can also be overridden by T-implicatures, which are context-dependent, play a social role rather than an informational one, and may even be culturally determined.

This is only tentative, of course, but it tends to support the earlier suggestion that scalar implicature is much more context sensitive than Levinson allows. I

suspect this point could be reinforced by considering other manifestations of the Politeness Principle.

3.5 *Scalar implicature or explicature?*

I shall close this section by introducing another alternative to the neo-Gricean treatment of scalar implicature, this time from a relevance theory perspective. (The following draws on Noveck and Sperber 2012.)³²

The view in question is that some supposed scalar implicatures are not in fact implicatures but *explicatures*. As explained earlier, an explicature of an utterance is (roughly) a pragmatically enriched version of its linguistically coded meaning, with ambiguities resolved, references identified, gaps filled, and so on. It can be thought of as the speaker's explicit meaning (as opposed to any distinct implicated meaning).³³ Now, one of the central processes in explicature is *narrowing*, in which the meaning of an expression is narrowed to express a more specific meaning, often an ad hoc, contextually determined one. Noveck and Sperber give the following example:

- (15) Henry: Do you want to go on working, or shall we go to the cinema?
Jane: I'm tired. Let's go to the cinema.³⁴

'Tired' can be used to express a wide range of physical and mental states from boredom and mild weariness through to outright exhaustion. In the context, however, it is clear that Jane's utterance of 'I'm tired' is relevant only if she means something like 'tired enough to prefer going to the cinema to going on working',

³² Bezuidenhout has made a similar proposal (Bezuidenhout 2002). Unlike Noveck and Sperber, however, she focuses mainly on number terms.

³³ For detailed discussion of explicature, see Carston 2002, 2004a, 2012. As noted earlier, Levinson denies that the distinction between explicature and implicature can be drawn in a rigorous way (Levinson 2000, pp.194–8). However, this does not affect the coherence of the position described in the text. We can agree that narrowing is a real phenomenon, whether or not we think of it as contributing to a distinct process of explicature.

³⁴ Noveck and Sperber 2012, p.312

and Henry will narrow down the meaning of the term to express that ad hoc concept.

Noveck and Sperber maintain that many cases of supposed scalar implicature are in fact cases of narrowing of this type. As an example, they give the following sentence, which we are to imagine being uttered in the context of a discussion of scientific literacy in America:

- (16) Most Americans are creationists and some even believe that the Earth is flat.³⁵

Here the hearer will narrow down the basic, semantically coded meaning of ‘some’ (which Noveck and Sperber take to be ‘at least two and possibly all’) in the search for a relevant interpretation (that is, one that has contextually useful, easily processed implications). Given the context, it would obviously not be a useful contribution to utter (16) if there were only two Americans who believe the Earth is flat. The utterance is contextually relevant only if a significant number is meant. Moreover, it is common knowledge that not all Americans believe the Earth is flat, so that information has little value, and the contrast with ‘most’ makes it clear that the speaker means a smaller number than the number of creationists. In this way, the meaning of ‘some’ is narrowed down at both ends to mean something like ‘a number large enough to be relevant to the discussion, but smaller than the number of creationists’.

In cases like this, ‘some’ is interpreted as having a meaning that is narrowed down *at both ends of the scale*. This narrowed-down meaning will, of course, *entail* the more limited top-end narrowing that the Q-principle would have produced, but it is a much richer and more contextually useful one. Noveck and Sperber give other examples. If Henry is preparing dinner and Jane tells him ‘Some of the guests are arriving’, a vague reading of ‘some’ as ‘more than one and less than all’ will be sufficient to render her utterance optimally relevant (having various contextual implications about what Henry should do next). Similar points, Noveck and Sperber note, apply to other scalar terms. For example, ‘possible’ may

³⁵ Noveck and Sperber 2012, p.313.

be narrowed down to indicate a modest probability, excluding both certainty *and* tiny probability.³⁶

This view provides an alternative, and perhaps more plausible, analysis of some of the cases discussed in the previous subsection. For example, the implicatures in examples (6) and (7), which I suggested would generate implicatures based on nonce Hirschberg scales, may be better thought of as yielding interpretations involving contextually narrowed meanings of ‘some’ and ‘sometimes’. Similarly, the ‘tact’ implicatures I discussed in the previous subsection might be re-interpreted as resulting from narrowing-down of the key concepts rather than an application of a general Tact principle. In (12) and (13), for example, ‘possible’ is narrowed to mean ‘highly probable’ and ‘some’ to ‘many and even all’, since these are the meanings from which most contextual implications can be drawn (concerning what actions the hearer should take).

Noveck and Sperber do not claim that we never draw scalar implicatures. They hold that where there is an implicit or explicit question as to whether a stronger term applies, we typically do draw one. If Henry had explicitly asked Jane whether all the guests had arrived, then her utterance would have prompted him to derive the implicature ‘not all’ (Noveck and Sperber 2012, pp.314–5). (Example (3) in section 2.1 above would be another example of what Noveck and Sperber would regard as a genuine scalar implicature.) But they argue that such cases are much rarer than neo-Griceans believe:

From the point of view of relevance theory, then, the classical neo-Gricean theory of scalar implicatures can be seen as a mistaken generalisation of the relatively rare case where a weaker claim genuinely implicates the denial of a stronger claim which is under consideration in

³⁶ Noveck and Sperber also give an example in which the meaning of ‘some’ is *broadened*. Henry has agreed to go and pick up dessert as soon as the dinner guests start arriving, and Jane calls to him ‘Some of the guests are arriving’. Here the relevance of Jane’s utterance does not depend on how many guests are arriving, and Henry will understand ‘some’ as compatible with any number of guests arriving, from one to all — which is a broadening of the basic meaning of ‘some’, as Noveck and Sperber understand it (Noveck and Sperber 2012, pp.313–4).

the context, to the much more common case where the denotation of an expression is narrowed to exclude marginal or limiting instances with untypical implications. (Noveck and Sperber 2012, p.315)

If this analysis is correct, then it is problematic for Levinson. In cases where considerations of relevance narrow down the meaning of scalar terms, automatic application of the Q-principle will be at best redundant, slowing down the comprehension process rather than speeding it up. If such cases are common, it is hard to see why there would have been pressure for the Q-principle to be applied by default.³⁷

It might be suggested that the cases Noveck and Sperber highlight should be regarded as falling under the I-principle, which tells us to enrich interpretations of utterances in the light of background knowledge. I don't think this is a promising suggestion, however, since in these cases the enrichment would depend heavily on contextual factors that have no role in a process of default interpretation (see section 4 below for more on this). Moreover, it is unclear why the Q-principle would not be applied in these cases, and if it is, then, on Levinson's account, the implicatures it produces should override ones produced by the I-principle.

³⁷ Levinson does allow that considerations of relevance (in the everyday sense) may affect how a scalar term is interpreted. He gives the following example:

A: Is there any evidence against them?

B: Some of their identity documents are forgeries.

(Levinson 2000, p.51)

Here he argues, 'some' is not interpreted as meaning 'not all', since the stronger claim is irrelevant to the speaker's communicative goal (establishing that there is evidence against the people in question). However, on his view this does not mean that the implicature is not derived, but only that is derived and then cancelled.

4. Assessing the I- and M-principles

This section looks at Levinson's other two principles, the I-principle and the M-principle. Again, I shall argue that it is doubtful that these principles support the existence of a level of utterance-type meaning.

4.1 Stereotypes and defaults

The I-principle tells hearers to enrich the content of utterances by drawing on knowledge of relevant stereotypes. Levinson notes that this is a powerful heuristic, which 'allows an interpreter to bring all sorts of background knowledge about a domain to bear on a rich interpretation of a minimal description.'

When introducing this principle, I noted a possible objection. Since the principle tells us to draw on background knowledge to interpret an utterance, it does not look like one that yields default interpretations, associated with utterance types. I suggested that Levinson would reply that the relevant background knowledge can be applied without considering the context of the utterance in question. This is a plausible reply when an expression reliably evokes a single stereotype. Levinson's examples are, arguably, of this kind: 'secretary' is interpreted as 'female secretary', 'road' as 'hard-surfaced road' (though the former might be considered problematic). But, as Anne Bezuidenhout points out, the same expression can evoke different stereotypes in different contexts (Bezuidenhout 2002). As an example, she takes the utterance 'Susan turned the key *and* the engine started.' (an example used by Levinson; 2000, p.117). Here, the I-principle tells us to enrich the interpretation of 'and' by drawing on background knowledge (conjunction buttressing). But (as Levinson himself notes; 2000, p.117), several different enrichments are possible, corresponding to temporal sequence, causal sequence, and goal:

- (a) GCI: Susan turned the key and then the engine started.
- (b) GCI: Susan turned the key and as a result the engine started.

(c) GCI: Susan turned the key with the goal of bringing it about that the engine started.

(Bezuidenhout 2002, p.266).³⁸

Now, in any given context, we will pick out one of these as the relevant stereotype. But plainly the words of the utterance itself cannot determine which is the relevant one, since they are the same in all three cases. In order to access the relevant stereotypical information, it seems we must draw on contextual information. As Bezuidenhout puts it:

Thus hearers will need to rely on the information made accessible in the wider context, such as information from prior discourse context (i.e., the mutual linguistic context), from the mutual physical environment, or from other shared sources of knowledge. (Bezuidenhout 2002, p.266)

As another example, Bezuidenhout gives ‘Professor White’s book is on the table’. Here, the I-principle tells us to enrich the possessive (via narrowing this time) by treating it as the stereotypical person-book relation. But again there is no single relation of this kind. The book might be one the professor owns, bought, borrowed, wrote, and so on. And which one the hearer chooses will be determined by the wider context. If the conversation is between assistants in a bookshop, the hearer will probably take the speaker to mean the book the professor wrote; if they are librarians processing requests from academics, they will probably take them to mean the one the professor requested, and so on (Bezuidenhout 2002, p.267–8).

Many more examples could be given. And, of course, similar cases will arise with the M-principle. What counts as the relevant *non-stereotypical* reading of an expression will also vary with context. ‘Bill caused to car to stop’³⁹ implicates that

³⁸ And, as Bezuidenhout notes, this does not exhaust the possible enrichments. ‘And’ can also implicate relations of temporal inclusion (‘He went to London and he saw the Queen’; co-occurrence (‘She likes to ride her bike and listen to her Walkman’; enabling (‘I forgot to hide the cake and the kids ate it’), and more (examples from Bezuidenhout 2002, p.272).

³⁹ Levinson’s example (Levinson 2000, p.39).

Bill didn't stop the car in the usual way but leaves open a wide range of options, from which the hearer will choose, depending, for example, on whether Bill was the driver, a passenger, a bystander, a policeman, and so on. Indeed, the range of possible M-implicatures will be wider than that of I-implicatures since there are many more ways of being non-stereotypical than of being stereotypical.

All this undermines Levinson's claim that I- and M-inferences are default ones, supporting a level of meaning associated with utterance *types*. If the default interpretation is the one that is most easily accessible, then, as Bezuidenhout notes, expressions will have *many* defaults, varying with context (Bezuidenhout 2002, p.272). Bezuidenhout concludes that Levinson faces a trilemma. If the I-principle produces multiple interpretations of the same expression in every context, then it does not serve the function of speeding up language processing. If it produces different interpretations of an expression in different contexts, then it is not part of a system of default interpretation. And if it produces the same interpretation in every context, then it will often hinder processing, since in many cases this interpretation will have to be overridden and corrected (Bezuidenhout 2002, p.274).

One way of resolving this would be to adopt a weak neo-Gricean view. We might say that hearers employ a general principle which tells them to read unmarked expressions in a stereotypical way (and marked ones in a non-stereotypical way), but that *which* of the many available stereotypes (or alternatives) they settle on will be determined by contextual factors. On this view, the I- and M- principles would guide interpretation, but would not yield default interpretations.

4.2 A deeper problem

There may be a deeper general problem with the I-principle. Levinson holds that the function of the Q-, I-, and M-principles is to overcome the bottleneck in human communication caused by the relatively slow speed of our articulatory processes. But he does not, of course, think that these principles exhaust our pragmatic competence; he assumes that they supplement a system of context-driven pragmatic processing, which processes PCIs and utterance-token meaning

generally.⁴⁰ The three principles speed up communication by providing automatic enrichments of certain types of utterance, pre-empting or reducing the need for context-driven pragmatic processing.

Now, it is crucial to this function that they are *formal* principles. Each is triggered by the presence of some formal property (for example, a certain expression type), which can be detected at an early stage of processing. The principle then specifies a formal procedure that can be applied to enrich the content of the utterance in reliable ways. In the case of the Q-principle, the triggering condition is the presence of a lower-ranked expression from a suitable scale, such as ‘some’, ‘possibly’, ‘may’, and the procedure is the replacement of the weaker expression with the negation of a stronger one from the same scale. As we have seen, there are some problems for the Q-principle, but in outline at least it looks like a feasible and effective strategy of default enrichment.

In the case of the I-principle, however, the picture is less clear. The triggering condition here is the presence of an unmarked expression — that is, one that is simple, brief, and familiar. This in itself is problematic, since this condition is the default one: people generally use unmarked expressions, unless they have some reason not to. Given this, it would seem more cost-effective in processing terms to look out for situations in which the condition *doesn't* hold than for ones in which it does. Second, the procedure to be applied is simply an instruction to draw on background knowledge to interpret the expression in the standard way. But this is, presumably, what would have happened anyway, thanks to the context-driven pragmatic processes. In effect, the I-principle says to check if an utterance is

⁴⁰ He writes, for example:

In the composite theory of meaning, the theory of GCIs plays just a small role in a general theory of communication. In this regard, GCI theory is not in direct competition with holistic theories like Sperber and Wilson's theory of Relevance, which attempts to reduce all kinds of pragmatic inference to one mega-principle — GCI theory is simply not a general theory of human pragmatic competence. Instead it attempts to account for one relatively small area of pragmatic inference. (Levinson 2000, pp. 21–2)

linguistically normal, and if it is, to process it in the normal way. But this does not pre-empt or reduce pragmatic processing; it is simply permission for it to go ahead, and explicitly applying it would be more likely to delay the process than speed it up. What is the point of checking if a condition holds, unless you're going to do something differently if it does? In short, the I-principle appears to be simply redundant, even from a weak neo-Gricean perspective.⁴¹

There is a related worry about the M-principle. Here the trigger is the presence of a marked (prolix, unusual) expression, and the procedure is to look for a nonstandard interpretation. This may seem more effective. When a marked expression is detected, we skip the usual processing and jump straight to a nonstandard interpretation. There is a problem, however. How can we tell what a nonstandard interpretation might be until we know what the standard one is? We cannot set aside the standard interpretation until we have identified it. It seems that in order to execute the M-principle, we (or rather, our cognitive systems) will have to let the normal pragmatic processes run until they reach the standard interpretation, *and then let them run further*, looking for alternative, less obvious interpretations.

This does not, however, mean that the M-principle is redundant. It does not pre-empt or reduce context-driven pragmatic processing and does not yield default interpretations. But it does guide how pragmatic processing is conducted and when it terminates. From a weak neo-Gricean perspective, the M-principle may still have a role to play, even if the I-principle does not.

5. Experimental evidence

Neo-Gricean theories were developed by drawing on linguistic intuitions and analyses rather than by experiment. However, since they involve, or at least imply,

⁴¹ It might be suggested that processing costs could be reduced by using a single detector, sensitive to only marked expressions, to implement both the I- and M- principles. If the detector is triggered, the M-principle procedure is executed, if it is not triggered, then the I-principle procedure is. However, if executing the I-principle simply involves letting normal pragmatic processing run, this would in effect eliminate the I-principle, since when the detector isn't triggered, nothing different happens.

claims about the mental processes involved in implicature recovery, they are open to experimental testing, and in recent years relevant work has been done in the growing field of *experimental pragmatics* (Noveck and Reboul 2008; Noveck and Sperber 2004). In particular, there are many experimental studies of scalar implicature, designed in part to test whether neo-Griceanism or relevance theory gives a better account of it. In this section I shall briefly survey some of this work and assess its findings.

5.1 Reaction-time studies

A key difference between neo-Griceanism and relevance theory concerns the order in which interpretations of scalar terms are processed. According to Levinson, scalar inferences are made automatically by default, and processing a literal interpretation of a scalar term will involve cancelling the scalar inference. According to relevance theory, by contrast, the initial interpretation is typically the linguistically coded one, and pragmatically enriched interpretations are derived only if needed to meet current expectations of relevance. Thus, in the case of ‘some’, for example, Levinson’s view predicts that the pragmatic meaning ‘some but not all’ is derived first, whereas relevance theory predicts that the basic meaning ‘some and possibly all’ is. (The latter meaning is sometimes referred to as the *logical* one, since it corresponds to the existential quantifier of predicate logic). Thus, if Levinson is right, pragmatic interpretations of ‘some’ should take less time to process than the logical one, and if relevance theory is right, the opposite should be the case.

In a pioneering study (conducted in French), Lewis Bott and Ira Noveck sought to test these predictions (Bott and Noveck 2004). They focused on what they call ‘underinformative’ sentences, such as ‘Some giraffes have long necks’, which make a claim that is true on a logical reading of ‘some’ but false on a pragmatic one (such sentences are said to be *pragmatically infelicitous*). Bott and Noveck used a sentence verification task, in which participants were presented with sentences of the form ‘Some/All F are G’ and asked to classify each as true or false. A sixth of the sentences were underinformative ‘some’ sentences, the rest

were control sentences that were straightforwardly true or false.⁴² There were two sessions. In one, the participants were told to treat ‘some’ logically, as meaning ‘some and possibly all’; in the other they were told to treat it pragmatically, as meaning ‘some but not all’. Bott and Noveck reasoned that if underinformative sentences generate scalar inferences by default, then participants should take longer to respond to such sentences when told to treat them logically than when told to treat them pragmatically, since in the former case the default pragmatic inference would have to be cancelled before the logical reading could be derived. In fact, the opposite happened. Participants responded to underinformative sentences more quickly in the logical condition than the pragmatic one, taking around 800 ms in the former and nearly 1400 ms in the latter. (They also responded more quickly to control sentences in the pragmatic condition, though the difference was not as great.) Participants also gave fewer incorrect answers when instructed to adopt a logical reading (90% correct as opposed to 60% in the pragmatic condition), suggesting that they found it easier to apply the logical interpretation.⁴³

In a variant of the experiment, Bott and Noveck allowed the participants to interpret ‘some’ as they wished. Those participants who classified the underinformative sentences as true were assumed to have adopted the logical interpretation and those who classified them as false were assumed to have adopted the pragmatic one. Again, there was a significant difference in response time, with

⁴² The control sentences consisted of equal numbers of true ‘some’ sentences (for example ‘Some mammals are elephants’), false ‘some’ sentences (for example ‘Some elephants are insects’), and three sets of ‘all’ sentences produced by substituting ‘all’ for ‘some’ in underinformative, true, and false ‘some’ sentences.

⁴³ A possible weakness in the experiment is that the underinformative sentences called for a positive response in the logical condition and a negative response in the pragmatic one. If participants were quicker to confirm a sentence than to deny it, this might explain the difference in response times. To control for this, Bott and Noveck ran a second experiment in which the underinformative sentences called for the same response in both conditions. (They achieved this by asking participants to assess a second sentence that expressed a true/false verdict on the original one and switching the value of the verdict between the conditions.) The results were in line with those of the first experiment.

those who adopted the logical reading responding more quickly than those who adopted the pragmatic one (2700 ms as opposed to 3300 ms).⁴⁴

In a final variant of the experiment, Bott and Noveck manipulated the time participants were given to respond to the sentences presented to them. There were two experimental conditions, Long and Short. In the Short condition participants were allowed 900 ms to respond, in the Long condition they were allowed 3000 ms. Bott and Noveck found that participants were more likely to classify underinformative sentences as true in the Short condition than in the Long condition (72% 'true' responses in the former, versus 56% in the latter). In other words, forcing subjects to respond more quickly (and thus limiting the cognitive resources available for producing their response) increases the likelihood of their treating 'some' as meaning 'some and possibly all' and reduces the likelihood of their drawing a scalar inference 'not all'. (The claim that the availability of cognitive resources affects implicature processing has been confirmed in another study (Pouscoulous et al. 2007). Pouscoulous et al., showed that by using a simpler task, with fewer distracting factors and more basic terms, implicature processing improved across the board from age 4 to adult.)

Bott and Noveck conclude that their studies provide evidence against the neo-Gricean view that scalar inferences are automatic and default and support for the relevance theory view that scalar implicatures take time and effort to process and are derived only when contextually required. The data do certainly indicate that scalar implicatures are not default interpretations, and they thus pose a problem for neo-Griceanism. However, this does not leave relevance theory as the only option.

First, the data are compatible with weak neo-Griceanism. The data suggest that the Q-principle is not applied automatically, before logical interpretations are processed; but it might be applied later and with more effort, if the context makes the logical reading unsatisfactory. That is, the Q-principle may be responsible for scalar implicatures if they are derived, even though they are not derived by default.

⁴⁴ Another study found an even greater difference, with pragmatic responders taking nearly twice as long as logical ones; see Noveck and Posada 2003.

Second, the data are also compatible with convention theory (or rather its cognitive counterpart). As noted earlier, convention theory treats literal meanings as more basic than generalized ('sentence') implicatures, since it holds that they depend on first-order semantic conventions rather than second-order ones. Moreover, convention theory, I suggest, predicts (at least tentatively) that scalar implicatures will require more time and effort to process than literal interpretations. According to convention theory, deriving literal interpretations involves applying first-order semantic rules only, whereas deriving sentence implicatures involves applying *both* first-order *and* second-order semantic rules. (Second-order semantic rules are rules for expressing further meanings by using sentences with their basic first-order meaning, so a second-order rule cannot be applied until the relevant first-order meaning has been processed. Otherwise, we would be dealing with an idiom rather than an implicature.) This suggests that sentence implicatures should take more time and effort to process than literal meanings. This is only a tentative prediction, of course; to make firm predictions we would need a theory of how knowledge of semantic conventions is stored and accessed. But it is a plausible initial one. *Prima facie*, then, convention theory fits the experimental data quite well.⁴⁵

⁴⁵ Other methods are also being used to test theories of scalar implicature. In one of the first studies of its kind, Bezuidenhout and Morris used eye movement monitoring to detect how long participants took to read different regions of a sentence, indicating the different processing demands each region made (Bezuidenhout and Morris 2004). Their aim was to compare Levinson's view (they call it the Default Model, DM) on which scalar terms automatically trigger Q-implicatures, and models such as those discussed in section 3.5 above, where expressions such as 'some' are semantically underspecified and undergo a process of contextually cued pragmatic enrichment (Bezuidenhout and Morris call this the Underspecification Model). Participants were asked to read passages such as the following, in which a 'some' sentence is followed by a sentence explicitly cancelling the supposed 'not all' implicature:

Some books had colour pictures. In fact all of them did, which is why the teachers liked them.

5.2 Developmental studies

Experimental work has also been done on the development of implicature processing in children. In a pioneering study, Ira Noveck ran a series of experiments to test competence with scalar implicature in French children and adults (Noveck 2001).

In one experiment, participants were presented with underinformative ‘some’ sentences (such as ‘Some elephants have trunks’) and control sentences, and asked to say whether they agreed with them. Noveck found that, whereas most adults rejected the underinformative sentences, the majority of the children accepted them (89% of eight-year-olds and 85% of ten-year-olds accepted them, as opposed to 41% of adults), suggesting that most of the children were adopting the logical reading of ‘some’ and not deriving the implicature ‘not all’. (The children correctly evaluated control sentences.) Noveck obtained similar results using the scalar terms ‘might’ and ‘must’. When asked to assess a claim that something *might* be the case (for example, ‘There might be a parrot in the box’) in a condition in which they knew it *must* be the case, children were much more likely than adults to accept the sentence as true (such sentences were accepted by 80% of seven-year olds, 69% of nine-year-olds, and 35% of adults). Again, this indicates that children tend to adopt a logical interpretation of the modal term, treating ‘might’ as meaning

Bezuidenhout and Morris reasoned that if participants don’t make the scalar inference by default when reading the first sentence, but simply start searching for the most contextually appropriate enrichment of ‘some’ (in line with the UM), then they should spend more time on the word ‘all’, since it is a strong clue that ‘some and possibly all’ is the appropriate enrichment. On the other hand, if participants automatically make the scalar implicature to ‘not all’ (as the DM model predicts), then they will not be pulled up by ‘all’ since they already have an interpretation of ‘some’, and it is not until they reach ‘them did’ that it becomes clear that this interpretation is wrong (‘all’ might have governed some other predicate). They should thus spend more time processing the words ‘them did’, which indicate the need for reinterpretation. The results favoured UM rather than DM. In comparison with control sentences in which (for example) ‘The books’ was substituted for ‘Some books’, participants spent more time processing ‘all’ and actually spent less time processing ‘them did’.

‘possibly and perhaps necessarily’, whereas most adults adopt a pragmatic reading, taking the affirmation of possibility to implicate the denial of necessity (Noveck 2001). Noveck concluded that logical interpretations of scalar terms are developmentally primary and that children are, in a sense, *more logical than adults*.

Noveck’s findings have been replicated by other researchers (see, for example, Guasti et al. 2005, Experiment 1; Papafragou and Musolino 2003, Experiment 1; Pouscoulous et al. 2007, Experiment 1). Children, it seems, do not spontaneously make scalar inferences. There is evidence, however, that they do have the *ability* to make them, given suitable prompting. After confirming five-year-olds’ apparent lack of sensitivity to scalar implicature, Papafragou and Musolino went on to see if they could improve the children’s performance by training them to detect pragmatic infelicity. They prepared the children by telling them stories in which a character said ‘silly things’, which were true but inappropriate (for example, describing a dog as ‘a little animal with four legs’) and asking how the character might ‘say it better’. They also changed the experimental task itself (which involved assessing descriptions of acted-out stories) to make it clear that it was relevant to know whether or not the stronger statements were true.⁴⁶ The result was that a much higher proportion of the children rejected underinformative ‘some’ statements (52.5% of five-year-olds as opposed to only 12.5% in the previous experiment).⁴⁷ Moreover, the children who rejected them justified their answer by pointing out that the stronger term was applicable (Papafragou and Musolino 2003).

⁴⁶ For example, the children would hear about a character Mickey, who had been challenged to put all his hoops round a pole, and, after trying hard, had succeeded. They would then hear Minnie respond to a question about how Mickey had done by saying ‘Mickey put some of his hoops round the pole’. The children would then be asked if Minnie had answered well (Papafragou and Musolino 2003, p.271).

⁴⁷ The children were also tested on the scales <*finish, start*> and <*three, two*>, and showed similar increases in pragmatic responding (47.5% vs 10% on <*finish, start*>, and 90% vs 65% on <*three, two*>. The experiments were conducted with Greek-speaking children.

Other studies have confirmed this. Feeney et al. found that in pragmatically rich contexts (using storyboards and photographs to tell a story and asking participants to assess claims made by one of the characters) only 21% of seven-to-eight-year-olds adopted the logical reading of ‘some’, as opposed to 57% on a simple sentence verification task (Feeney et al. 2004, Experiment 2). Similarly, Guasti et al. found that in a realistic conversational setting where all the relevant evidence was easily accessible, seven-year-olds derived scalar implicatures at adult levels (Guasti et al. 2005, Experiment 4). Guasti et al. note, however, that the same does not go for younger children. In tests, only half of five-year-olds rejected underinformative statements, even when the statements were presented in a natural way (although the ones that did so, did so consistently) (Chierchia et al. 2001; Papafragou and Musolino 2003). Guasti et al. suggest that at that age some children simply lack the knowledge or ability to derive scalar implicatures, either because the weaker expression does not activate the contrasting stronger one or because the inference from the affirmation of the former to the denial of the latter isn’t made (Guasti et al. 2005, p.694).⁴⁸

A possible weakness in some of the experiments reviewed is that they required young children to make metalinguistic judgements (judgements about how well a character had described a situation). These tasks may have been too demanding for younger children, hiding their pragmatic competence. In an ingenious experiment, Yi Ting Huang and Jesse Snedeker sought to get round this problem by using pictures and eye-tracking (Huang and Snedeker 2009). They presented five-year-old children with a range of pictures showing four characters, two boys and two girls, each of whom had a number of items, either socks or soccer balls (but not both). While looking at one of these pictures, the children then heard an instruction of the following form:

⁴⁸ However, in another study where the task was naturalistic and informational demands clear, children of four-to-five years made scalar inferences at a high level. This extended to particularized scalar implicatures, dependent on nonce scales. For example, if a character was asked whether it had wrapped two presents and replied that it had wrapped one of them, 90% of children detected the implicature that it had not wrapped the other (Papafragou and Tantalou 2004).

Point to the girl/boy that has some/all/two/three of the socks/soccer balls.

Their eye movements were recorded as they listened to the instruction and looked for the intended target.

The pictures were designed in such a way that when ‘all’, ‘two’, or ‘three’ were used in the quantifier position, the identity of the target could be inferred from gender and quantifier alone, and the children tended to look to the target, even before processing the final words (‘socks’ or ‘soccer balls’). When ‘some’ was used, however, the identity of the target remained uncertain *unless* ‘some’ was interpreted pragmatically (for example, the relevant options might be a girl with two socks and a girl with *all* the balls). The final words of the instruction then resolved the ambiguity in favour of the pragmatic reading. Huang and Snedeker reasoned that if children made the scalar inference when they processed ‘some’, then they would look to the correct target before the end of the sentence. In fact, they did not, but delayed looking to the target until they heard the disambiguating words at the end, suggesting that they had not derived the implicature.

Variants of the experiment confirmed this. When the pictures were adjusted so that ‘some’ identified the correct target whichever reading was adopted (when, for example, the relevant options were a girl with a subset of the socks or a girl with no socks at all), the children looked to the correct target before the end of the sentence. In a final version of the experiment, ‘some’ was ambiguous, and the pragmatic reading of it now indicated the *wrong* target (for example, a girl with a subset of the balls, when in fact the correct target was a girl with *all* the socks). Huang and Snedeker predicted that if children were drawing the scalar inference, they would take longer to look to the correct target at the end, since they would have to correct an original misidentification. In fact, the adjustment made no difference; the children looked to the correct target just as quickly when it was inconsistent with the scalar inference as when it was not. In all of the experiments, the results from adult participants showed the opposite tendency, indicating that they were drawing the scalar implicature.

As Huang and Snedeker note, this evidence for children’s lack of competence with scalar implicature presents a puzzle (Huang and Snedeker 2009, p.1737). For young children are very good at certain kinds of pragmatic processing, in particular

at learning new words through interpreting speakers' communicative intentions. Why then are they so slow to master scalar implicature? Huang and Snedeker suggest that it is because pragmatic processes play different roles in word learning and scalar implicature. In word learning children infer meanings directly from non-linguistic evidence of the speaker's intentions, such as pointing. The pragmatic process is a top-down one, from intentions to meanings, and it can proceed without any prior semantic processing. In the case of scalar implicature, the process is bottom up. The child must start with an analysis of word meaning and move from that to an implicated meaning — which is a much more demanding task. Huang and Snedeker note that other tasks that are hard for young children, such as interpreting irony and metaphor, are also of this bottom-up kind:

In each case, pragmatic success requires listeners to calculate an interpretation that builds upon but goes beyond the initial linguistic meaning. These postsemantic processes may be particularly difficult, because they require that some feature of the child's initial analysis be revised. (Huang and Snedeker 2009, p.1737)

To sum up then: The experimental data strongly suggest that (a) children initially adopt a logical reading of scalar terms, (b) pragmatic competence increases with age, and (c) children (of seven years and up at least) can derive scalar implicatures if they are provided with suitable contextual assistance.

These results are not what neo-Gricean theory would predict. If the language comprehension system applies the Q-principle automatically, then children should find the pragmatic reading of scalar terms natural and should not need contextual help to derive scalar implicatures. Instead, it should be the logical reading that they find hard to master, since its application will involve cancelling the default pragmatic one. Levinson might reply that the Q-principle is initially applied in a slow and effortful way and only later becomes automatized, but this does not fit well with his view that it is an evolutionary adaptation designed to alleviate the articulatory bottleneck.⁴⁹

⁴⁹ Levinson writes:

Relevance theory, on the other hand, predicts the experimental results. If logical interpretations of scalar terms are more accessible than pragmatic ones, and if pragmatic interpretations are derived only when required to satisfy expectations of relevance, then we should expect children to derive scalar implicatures less often than adults. For children are typically less aware of informational demands and opportunities than adults and have more limited cognitive resources, which means that they will have lower expectations of relevance and will find implicatures more costly to process and hence less relevant. Likewise, relevance theory predicts that children will draw more implicatures if the context is adjusted to raise their informational expectations and make implicature derivation easier, as they in fact do.

Again, however, there are other options besides neo-Griceanism and relevance theory, not typically considered in the experimental literature. First, the data are broadly compatible with convention theory. It would not be surprising if children learn the first-order rules that govern literal meaning before the second-order rules that govern sentence implicatures. (As noted earlier, Davis predicts that second-language learners will be slower to learn second-order rules, and it may be that children are slower to acquire them in their first language; Davis 1998, p.159.)⁵⁰ However, children who have not mastered the conventions for scalar implicature might still be able to work out individual scalar implicatures in a particularized way, drawing on contextual clues and theorizing about the speaker's intentions. This would explain why children derive more scalar implicatures when tested on

Now it is quite clear that ... intelligent agents with the asymmetrical abilities in thinking and speaking I have just elucidated, would find a way around the articulatory bottleneck (just as, as a matter of fact, evolution has). The essential asymmetry is: inference is cheap, articulation expensive, and thus the design requirements are for a system that maximizes inference. (Levinson 2000, p.29)

⁵⁰ It should be stressed that Davis himself does not commit to the second claim and does not rule out the possibility that first-language learners can master a language's implicature conventions simultaneously with its literal meaning conventions (personal communication).

statements produced in realistic conversational settings where contextual clues are available.

In this context, it is interesting to note that there is evidence that word choice affects derivation of scalar implicatures in children. Pouscoulous et al. found that French nine-year-olds were more likely to draw scalar implicatures when ‘quelques’ was used for ‘some’ instead of ‘certains’, even though their responses on control problems showed that they understood the meaning of ‘certains’ (Pouscoulos et al. 2007). (42% adopted the logical reading when ‘certains’ was used, as against 0% when ‘quelques’ was; the change made no significant difference to adults’ responses.) This result is difficult for neo-Griceans to explain, since the words have the same meaning and should both automatically trigger application of the Q-principle. Pouscoulous et al. suggest that ‘certains’ is a more complex word semantically, which uses up extra processing resources, leaving fewer free for implicature processing — an explanation that fits well with relevance theory. But convention theorists might offer another explanation, suggesting that the implicature conventions governing the two words are different and that those associated with ‘certains’ take longer to learn.

Second, the data are compatible with weak neo-Griceanism. The evidence indicates that scalar inferences are not drawn automatically whenever scalar expressions are processed, but this is compatible with the view that they are drawn in a more effortful way, when contextually cued. Indeed, there is evidence that even very young children of three-to-four years can draw scalar inferences from contextually salient orderings of non-linguistic stimuli (Stiller et al. 2011; see also Papafragou and Tantalou 2004).

5.3 Tentative conclusions

Three tentative conclusions can be drawn from the experimental research surveyed. First, the data do not support neo-Griceanism as developed by Levinson. Evidence from reaction-time studies and developmental studies suggests that the default reading of scalar terms is the logical one, and that scalar implicature processing is relatively effortful. Second, the data are compatible with the relevance theory. Third, the data are also compatible with other theories of

implicature recovery such as convention theory and weak forms of neo-Griceanism.⁵¹

6. Conclusions

This chapter has looked at the Gricean framework from the perspective of implicature recovery and linguistic analysis, focusing in particular on the neo-Gricean case for the existence of a class of generalized implicatures derived by the default application of certain general inferential principles. The survey and discussion has necessarily been selective, but it has been sufficient to raise doubts about the neo-Gricean project, at least in the strong form proposed by Levinson. Our examination of Q- I- and M- principles suggests that supposedly generalized implicatures are in fact much more context-sensitive than neo-Griceans suppose, and that default inferences would often need to be cancelled — slowing down, rather than speeding up, the interpretation process. Moreover, the results of experimental work on scalar implicature are at least *prima facie* incompatible with neo-Griceanism.

The chapter also briefly introduced some alternative approaches to implicature recovery, including relevance theory, convention theory (in a cognitive form), and what I called *weak neo-Griceanism*. I have argued that each of these alternatives has some advantages over neo-Griceanism, but I have not advocated one of them in particular. (I shall return to this topic briefly in the final chapter and suggest that elements of different alternative approaches might be combined.)

I noted at the beginning of this chapter that neo-Griceanism can be seen as offering a theory of implicature generation, with utterances being understood to possess the implicatures they would be interpreted as having according to the GCI principles. Thus, in assessing neo-Griceanism, we have, in effect, also been assessing this simplified and restricted version of the Gricean framework. If the concerns raised in the course of this chapter are sound, then this assessment must be largely negative. If the GCI principles did typically guide our interpretation of utterances in the way neo-Griceans claim, then it would be plausible to give them

⁵¹ For further discussion of the experimental literature on scalar implicature, and exploration of its connections with ‘dual-process’ theories of reasoning, see Frankish and Kasmirli 2010.

this sort of normative status. We could treat them as regularizations of normal practice, and use them to judge specific cases. However, the examples we have considered suggest that the principles do not play the role claimed for them. Even if they do have some role in interpretation (especially, perhaps, the Q-principle), they are not applied by default, in a context-independent way. Details of context and speaker intention can colour implicature recovery even in supposedly generalized cases. Moreover, in so far as there are general patterns of implicature associated with some expressions, these may be better explained as arising from conventions of use rather than GCI principles. So even this limited, simplified version of the Gricean framework looks unpromising. Of course, we could still give the principles a normative status and try to revise our practice to bring it in line with them, but it is hard to see why we should accept such artificial norms, which do not reflect our actual practice or the psychological processes underlying it.

Neo-Griceanism is a bold and elegant theory, but (like the Gricean framework that inspired it) it is too ambitious. Even apparently generalized implicatures can be messy and context-dependent, and they resist simple codification.

Chapter 6

Taking stock and looking forward

This short final chapter takes stock and makes some suggestions for future work. The first section is retrospective, briefly revisiting the major themes from the thesis and giving tentative answers to some of the questions posed in the opening chapter. The second section builds on these conclusions to offer some speculations about the function and ethics of implicature.

1. Taking stock

1.1 The Gricean framework

The main focus of this thesis has been Grice's account of how implicatures are generated ('the Gricean framework'), which I introduced in Chapter 2. I raised a number of objections to the account, but the key ones centred around Grice's claim that implicatures can be calculated from general conversational principles. This is an elegant and attractive idea, but I argued that it is wrong to think that there is this kind of rational connection between utterances and what they implicate. The relevant arguments in Chapters 3 and 4 can be seen as presenting a dilemma for Grice. If his account is construed as a descriptive one, which aims to explain our intuitions about what implicatures utterances possess, then it gives the wrong results, predicting implicatures that we do not take to be there and denying the existence of ones we do. If construed as a normative theory, which aims to establish speaker-independent norms of implicature, then it avoids many of the previous objections but ultimately fails on its own terms. The discussion of this latter option linked up with another major theme of the thesis, the normativity of implicature.

1.2 Normativity

A key question addressed in this thesis was whether it is possible to provide speaker-independent norms of implicature, which hold for all implicatures, including particularized ones. Following Saul, I argued that Grice's account is best understood as aiming to provide such norms, and I proposed various revisions to it to make it more consistent with this aim. However, in Chapter 4, I went on to

argue that the account nevertheless fails. Grice identifies what an utterance implicates with the supposition required to preserve the assumption that the speaker is being cooperative, but I argued that there is no way to determine what this supposition is without drawing on information about the speaker, including their background beliefs, intentions, and values.

I then proposed an intention-centred account of implicature which retained a normative element. According to this, what an utterance implicates is (roughly) what a typical audience would take the speaker to intend it to convey. Since this may be different from what the speaker actually intended it to convey, this gives implicature a normative dimension and allows for the possibility that an utterance may implicate something that the speaker does not in fact want to convey. I noted, however, that this account does not provide *speaker-independent* norms of implicature, since how a typical audience would interpret an utterance may depend on facts about the particular speaker, and similar utterances may generate different implicatures when produced by different speakers. I strongly suspect that speaker-independent norms of implicature are not in fact available and that this weak normative conception is the strongest we can hope for.

1.3 Speaker intentions

The role of speaker intentions in implicature has been another theme of the thesis. It surfaced first in Chapter 2, where we saw that there is a tension in Grice's work on implicature. On the one hand, it is plausible to think that Grice regards what a speaker implicates as one aspect of what they mean, from which it follows that implicatures must be backed by appropriate communicative intentions of the sort Grice takes to be involved in speaker meaning. Yet there is no mention of speaker intentions in Grice's definition of implicature itself. In Chapter 3 I proposed that we could resolve this tension by making a distinction between what an utterance implicates and what a speaker implicates, where an utterance implicates that q if the Gricean conditions for implicature are met, and a speaker implicates that q if they produce an utterance that implicates that q and also themselves mean that q .

The issue arose again in Chapter 4, where I argued that Gricean theory should acknowledge a greater role for speaker intentions in implicature. The theory identifies what an utterance implicates with what the speaker must be supposed to

believe in order to make sense of their utterance as a cooperative conversational contribution. But, I argued, this supposition cannot be calculated without drawing on information about the speaker's attitudes, including their intentions with regard to the appropriate background assumptions to use in making the calculation. Thus, what an utterance implicates may depend indirectly on the speaker's intentions. I argued that this undermined the idea that the Gricean framework provides speaker-independent norms of implicature, and, given this, I suggested that there was no reason to deny speaker intentions a more direct role in determining what utterances implicate.

By giving intentions such a role, I argued, we can resolve many of the problems considered in Chapter 3, where a Gricean calculation fails to fix a determinate implicated content. I argued that giving speaker intentions this role need not involve adopting a Humpty Dumpty theory of implicature, since we can retain a normative condition concerning how a typical audience would interpret the utterance.

1.4 An intention-centred account

Developing the idea just mentioned, I sketched an intention-centred account of implicature that retained a normative element. I proposed that an utterance implicates *q* (where *q* is not its literal meaning) if it *makes q available*, where this means that a typical audience would identify *q* as the speaker's intended meaning, inferring this from an open-ended range of evidence. This view is intention-centred, since what an utterance makes available is what the hearer thinks the speaker intended to convey. However, it also retains a normative element, since a speaker may give inadequate or misleading evidence of their intentions, with the result that what their utterance makes available differs from what they intended to convey.

1.5 Generalized implicatures

Since my aim in this thesis was to assess the Gricean framework as a global theory of implicature, I focused heavily on particularized implicatures, which present the hardest cases for the framework. However, this left open the possibility that Gricean principles might explain a more limited class of generalized, context-

independent implicatures. We looked at this possibility in Chapter 5, which discussed neo-Gricean theories, according to which hearers derive generalized implicatures by applying simple interpretative principles closely related to Grice's maxims. (As I explained, these theories can also be construed as theories of implicature generation, which make predictions similar to Grice's own.) The subject is a complex and technical one, but I raised a number of concerns about the approach, arguing that there are many exceptions to the implicature patterns predicted by neo-Griceanism and that it is unlikely that we apply neo-Gricean principles automatically and by default. A review of work in experimental pragmatics offered support for this conclusion, suggesting that implicature recovery is a context-driven process. At best, I suggested, a weak neo-Griceanism may be defensible, according to which neo-Gricean principles play a role in interpretation but contextual factors determine when and how they are applied.

Chapter 5 also considered the possibility that some generalized implicatures arise from language-specific conventions of use. There may be conventions within a language community that certain expressions, used with their literal meaning, convey something else. I suggested that this may offer a more economical explanation of some cases (such as 'an X' implicatures), but other points made in the chapter indicated the need for a cautionary approach to the convention view. I argued that scalar implicatures are attractively explained as due to context-sensitive applications of the Q-principle, in the way Hirschberg proposes, rather than to conventions. And, of course, the numerous exceptions to supposedly generalized implicatures pose as much a problem for convention theory as for neo-Griceanism. If there are implicature conventions, then, it seems, we do not follow them strictly and without regard for context. A proper assessment of convention theory would, however, require detailed work in historical and cross-cultural linguistics.

Under scrutiny, then, the view that there is a clearly defined class of generalized implicatures breaks down, and, with it, the idea that there are any genuinely context-independent norms of implicature. Rather, we find a continuum of cases from more to less particularized, differing in the relative roles played by contextual factors and general principles or conventions.

1.6 Implicature recovery

This thesis did not attempt to survey and assess the range of theories of implicature recovery. The discussion in Chapter 5 was focused on neo-Griceanism, and rival theories were discussed mainly in order to provide contrast with it. The conclusion of the chapter was that neo-Griceanism was an unpromising approach, at least in a strong form, but I did not attempt to adjudicate between the rival theories, nor did I discuss the recovery of particularized implicatures. However, some points fall out naturally from the wider discussion, and I shall make them briefly here.

First, given my emphasis on the role of context in implicature and my scepticism about the existence of generalized implicatures, relevance theory would be the natural complement to the view of implicature generation I have advocated. I take it that relevance theory is consistent with the idea that hearers attempt to detect speakers' communicative intentions, as on the intention-centred account I sketched. If speakers intend to provide their hearers with optimally relevant inputs (as relevance theory says they should), then in searching for the optimally relevant interpretation of an utterance, the hearer is in effect trying to detect the speaker's intentions. (In fact, this seems to be the view that relevance theorists take; see, for example Wilson and Sperber, 2004.)¹ This need not, I suggest, exclude a role for implicature conventions or even general principles, if applied in a context-sensitive way. Knowledge of conventions could feed into the search for optimally relevant interpretations just as knowledge of word meanings and idioms does. And if considerations of relevance dictate it, general principles might be applied to derive enriched interpretations (as in the Henry and Mary example in section 2.1 of Chapter 5). While I do not wish to make a positive commitment to a theory of

¹ They write, for example:

Understanding is achieved when the communicative intention is fulfilled — that is, when the audience recognizes the informative intention ... According to relevance theory, use of an ostensive stimulus may create precise and predictable expectations of relevance not raised by other inputs. ... we will describe these expectations and show how they may help to identify the communicator's meaning. (Wilson and Sperber, 2004, p.611)

implicature recovery at this point, I feel that a pluralistic relevance-driven approach of this kind would be a natural path to explore, for those persuaded by the arguments in this thesis.

2. Looking forward

In Chapter 1 I promised to say something about the ethics of implicature, and I shall close by doing this. What follows is tentative and speculative, though it derives from the earlier discussion.

I will begin with some remarks about the social function of implicature, which is a topic that has, I think, received too little attention in the literature. Theorists often write as if the function of implicature is purely communicative. For Grice, implicating something is a way of making a cooperative contribution to a conversational exchange, and a message is implicated only if it can be interpreted as cooperative.² Neo-Griceans such as Levinson hold that generalized implicatures increase the efficiency of communication. And relevance theorists assume that speakers aim to produce optimally relevant utterances, which convey as much information as they are able and willing to provide. But is this the whole story? We certainly do use implicature to communicate, but why do we sometimes choose to implicate a message rather than speaking literally? What is the *distinctive* function of implicature?

If the distinctive function of implicature were to *improve* communication, then it does not seem very well suited to it. As we have seen, particularized implicatures often depend on subtle contextual cues and knowledge of the speaker's attitudes and habits, and even relatively generalized implicatures require more than simple application of interpretative maxims. If anything, use of implicature would be likely to impede communication, increasing demands on hearers and creating many new opportunities for misunderstanding and confusion. If effective

² It is true (as noted in Chapter 2), that Grice recognizes the existence of a class of *nonconversational* nonconventional implicatures, which are generated by other maxims, including aesthetic, social, and moral ones (Grice 1975/1989, p.28). However, he says very little about such implicatures, and the overwhelming focus in the Gricean literature is on conversational implicature.

communication were our sole aim, it would surely be advisable to avoid implicature.

What other functions might implicature have? I noted Davis's claim that we can use implicature to be polite and stylish (Davis 1998, p.174), and I illustrated how considerations of tact might promote the use of a certain implicature practice (Chapter 5, section 3.4). I have no doubt that implicature does serve these purposes, but I want to make a further suggestion, inspired by the conclusions reached in this thesis.

I have argued that implicature is context-dependent to a greater degree than is recognized within the Gricean tradition. In particular, an implicature may depend on facts about the particular speaker, and may be designed for a specific class of hearers. Even supposedly generalized implicatures, I have argued, may depend on awareness of contextual factors, such as recognition of a contextually salient ordering or a shared convention of use. For a purely communicative practice, this would be a disadvantage, but perhaps it hints at an important function of implicature. Perhaps a function of implicature is to establish a bond between speaker and hearer that goes beyond that of simple information exchange. In making use of implicature, a speaker invites their hearer to engage with them in a more intimate way than is required for literal communication, drawing on shared experiences, assumptions, values, and conventions. The speaker invites the hearer to *read between the lines* — to become, as it were, a partner in the communicative act, completing it for themselves. And in doing this, it may be that they are signalling something about their attitude to the hearer — that like them, share their values, feel a rapport with them (or want to establish one).

In employing implicature, speakers may also make an implicit offer. If implicature use demands more from the hearer, then perhaps it signals that the speaker is offering more, too — that they will be more open, confiding, honest. In order to interpret an implicature the hearer must get on the same wavelength as the speaker, and once they have tuned in, they may expect to receive a special message. In these and other ways, implicature use may be seen as an invitation to *trust* the speaker.

If this is right, then I think it casts a new light on the ethics of implicature. The duties of the implicature user are not simply those of the cooperative

communicator (to be truthful, informative, relevant, and so on) but also those of someone who offered a sort of intimacy and invited trust. A person using implicature to mislead is not merely being uncooperative, but, in a minor way, betraying a trust. And this brings us back at last to Mr Bronston. Mr Bronston used implicature to mislead the court. (I take it that he did this intentionally; if the real Mr Bronston did not, then assume I am talking about a fictional one who did.) He invited the lawyer questioning him to read between the lines of what he said, as if they trusted each other, and he exploited the lawyer's willingness to accept his invitation. But it was not only Mr Bronston who was at fault. In court there is no place for the trusting communicative relationship that implicature creates. Chief Justice Burger was right; it was the questioning lawyer's duty to challenge Mr Bronston's answer, and in accepting it, they were negligent.

I am suggesting, then, that in order to understand the role of implicature and to address the ethical issues it raises, we need to take a broader view of the communicative situation and consider the personal relation that implicature use establishes between speakers and hearers. I hope to explore this approach in future work.

References

- Anon. (1999). Criminal Law. Perjury. Sixth Circuit sustains perjury conviction for answer to question with mistaken premise. *United States v. DeZarn*, 157 F.3d 1042 (6th Cir. 1998). *Harvard Law Review*, 112(7), 1783–1788.
- Atlas, J. (1993). *The implications of conversation: The 1990 Leuven lectures*. Ms.
- Bach, K. and Harnish, R. M. (1979). *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.
- Bezuidenhout, A. (2002). Generalized conversational implicatures and default pragmatic inferences. In J. K. Campbell, M. O'Rourke, and D. Shier (Eds.), *Meaning and truth: Investigations in philosophical semantics* (pp.257–283). New York, NY: Seven Bridges Press.
- Bezuidenhout, A. L. and Morris, R. K. (2004). Implicature, relevance and default pragmatic inference. In I. A. Noveck and D. Sperber (Eds.), *Experimental pragmatics* (pp.257–282). Basingstoke: Palgrave Macmillan.
- Bott, L. and Noveck, I. A. (2004). Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of Memory and Language*, 51(3), 437–457.
- Breheny, R., Katsos, N., and Williams, J. (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.
- Bronston v. United States*, 409 U.S. 352 (1973).
- Brown, G. and Yule, G. (1983). *Discourse analysis*. Cambridge: Cambridge University Press.
- Carroll, L. (2009). *Alice's adventures in Wonderland and Through the looking-glass and what Alice found there*. Oxford: Oxford University Press.
- Carston, R. (1998). Informativeness, relevance and scalar implicature. In R. Carston and S. Uchida (Eds.), *Relevance theory: Applications and implications* (pp.179–236). Amsterdam: John Benjamins.

- Carston, R. (2002). *Thoughts and utterances*. Malden, MA: Blackwell.
- Carston, R. (2004a). Explicature and semantics. In S. Davis and B. Gillon (Eds.), *Semantics: A reader* (pp.817–845). Oxford: Oxford University Press.
- Carston, R. (2004b). Relevance theory and the saying/implicating distinction. In L. R. Horn and G. Ward (Eds.), *The handbook of pragmatics* (pp.633–656). Malden, MA: Blackwell Publishing.
- Carston, R. and Hall, A. (2012). Implicature and explicature. In H.-J. Schmid (Ed.), *Cognitive pragmatics* (pp.47–84). Berlin: Walter de Gruyter.
- Castleman, D. (2004). Has the law made liars of us all? bepress Legal Series, Working paper 380. Retrieved from <http://law.bepress.com/expresso/eps/380>
- Cummings, L. (2009). *Clinical pragmatics*. Cambridge: Cambridge University Press.
- Davis, W. A. (1998). *Implicature: Intention, convention, and principle in the failure of Gricean theory*. Cambridge: Cambridge University Press.
- Davis, W. A. (2007). How normative is implicature. *Journal of Pragmatics*, 39, 1655–1672.
- Davis, W. A. (2014). Implicature. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2014 ed.). Retrieved from <http://plato.stanford.edu/archives/fall2014/entries/implicature/>
- Dennett, D. C. (1986). *Content and consciousness* (2nd ed.). London: Routledge and Kegan Paul.
- Fauconnier, G. (1975). Pragmatic scales and logical structure. *Linguistic Inquiry*, 6, 353–375.
- Feeney, A., Scafton, S., Duckworth, A., and Handley, S. J. (2004). The story of some: Everyday pragmatic inference by children and adults. *Canadian Journal of Experimental Psychology*, 58(2), 121–132.
- Frankish, K. (2004). *Mind and supermind*. Cambridge: Cambridge University Press.

- Frankish, K. and Kasmirli, M. (2010). Saying one thing and meaning another: A dual process approach to conversational implicature. In K. Manktelow, D. Over, and S. Elqayam (Eds.), *The science of reason: A festschrift for Jonathan St. B. T. Evans* (pp.259–281). Hove, U.K.: Psychology Press.
- Gazdar, G. (1979). *Pragmatics: Implicature, presupposition, and logical form*. New York, NY: Academic Press.
- Green, M. S. (2002). Review of “Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory” by Wayne A. Davis. *Philosophy and Phenomenological Research*, 65, 241–244.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3), 377–388.
- Grice, H. P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society, Supp. vol.*, 121–153.
- Grice, H. P. (1968). Utterer’s meaning, sentence-meaning, and word-meaning. *Foundations of Language*, 4(3), 225–242.
- Grice, H. P. (1969). Utterer’s meaning and intentions. *The Philosophical Review*, 78(2), 147–177.
- Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp.41–58). New York, NY: Academic Press.
- Grice, H. P. (1978). Further notes on logic and conversation. In P. Cole (Ed.), *Syntax and semantics 9: Pragmatics* (pp.113–127). New York, NY: Academic Press.
- Grice, H. P. (1981). Presupposition and conversational implicature. In P. Cole (Ed.), *Radical pragmatics* (pp.183–198). New York, NY: Academic Press.
- Grice, H. P. (1982). Meaning revisited. In N. V. Smith (Ed.), *Mutual knowledge* (pp.223–243). New York, NY: Academic Press.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Grice, H. P. (2001). *Aspects of reason*. Oxford: Oxford University Press.

- Guasti, M. T., Chierchia, G., Crain, S., Foppolo, F., Gualmini, A., and Meroni, L. (2005). Why children and adults sometimes (but not always) compute implicatures. *Language and Cognitive Processes*, 20(5), 667–696.
- Hall, A. (2008). Free enrichment or hidden indexicals? *Mind & Language*, 23(4), 426–456.
- Harnish, R. M. (1991). Logical form and implicature. In S. Davis (Ed.), *Pragmatics: A reader* (pp.316–364). New York, NY: Oxford University Press.
- Hirschberg, J. B. (1985). *A theory of scalar implicature*. PhD thesis, University of Pennsylvania. Retrieved from <http://semantics.uchicago.edu/kennedy/classes/w14/implicature/readings/hirschberg85.pdf> (Reprinted by Garland Publishing, 1991.)
- Horn, L. R. (1972). *On the semantic properties of logical operators in English*. University of California, LA.
- Horn, L. R. (1984). Toward 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, DC: Georgetown University Press.
- Horn, L. R. (1989). *A natural history of negation*. Chicago: University of Chicago Press.
- Horn, L. R. (2004). Implicature. In L. R. Horn and G. Ward (Eds.), *The handbook of pragmatics* (pp.3–28). Oxford: Blackwell.
- Huang, Y. T. and Snedeker, J. (2009). Semantic meaning and pragmatic interpretation in 5-year-olds: Evidence from real-time spoken language comprehension. *Developmental Psychology*, 45(6), 1723–1739.
- Leech, G. N. (1983). *Principles of pragmatics*. London: Longman.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.

- Martinich, A. P. (1991). A theory for metaphor. In S. Davis (Ed.), *Pragmatics: A reader* (pp.507–518). New York, NY: Oxford University Press.
- Morgan, J. L. (1978). Two types of convention in indirect speech acts. In P. Cole (Ed.), *Syntax and semantics 9: Pragmatics* (pp.261–280). New York, NY: Academic Press.
- Neale, S. (1990). *Descriptions*. Cambridge, MA: MIT Press.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15(5), 509–559.
- Noveck, I. A. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition*, 78(2), 165–188.
- Noveck, I. A. and Posada, A. (2003). Characterizing the time course of an implicature: An evoked potentials study. *Brain and Language*, 85(2), 203–210.
- Noveck, I. A. and Reboul, A. (2008). Experimental pragmatics: A Gricean turn in the study of language. *Trends in Cognitive Sciences*, 12(11), 425–431.
- Noveck, I. A. and Sperber, D. (Eds.) (2004). *Experimental pragmatics*. Basingstoke: Palgrave Macmillan.
- Noveck, I. and Sperber, D. (2012). The why and how of experimental pragmatics: The case of “scalar inferences.” In D. Wilson and D. Sperber (Eds.), *Meaning and Relevance* (pp.307–330). Cambridge: Cambridge University Press.
- Papafragou, A. and Musolino, J. (2003). Scalar implicatures: Experiments at the semantics-pragmatics interface. *Cognition*, 86(3), 253–282.
- Papafragou, A. and Tantalou, N. (2004). Children’s computation of implicatures. *Language Acquisition*, 12(1), 71–82.
- Peters, P. (2004). *The Cambridge guide to English usage*. Cambridge. Cambridge University Press.
- Pouscoulous, N., Noveck, I. A., Politzer, G., and Bastide, A. (2007). A developmental investigation of processing costs in implicature production. *Language Acquisition*, 14(4), 347–375.

- Recanati, F. (1989). The pragmatics of what is said. *Mind and Language*, 4(4), 295–329.
- Recanati, F. (1993). *Direct reference: From language to thought*. Oxford: Blackwell.
- Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy*, 25(3), 299–345.
- Sadock, J. M. (1991). On testing for conversational implicature. In S. Davis (Ed.), *Pragmatics: A reader* (pp.365–376). New York, NY: Oxford University Press.
- Salmon, N. (1989). Illogical belief. *Philosophical Perspectives, Vol. 3, Philosophy of Mind and Action Theory*, 243–285.
- Saul, J. M. (2001). Critical studies: Wayne A. Davis, *Conversational implicature: Intention and convention in the failure of Gricean theory*. *Noûs*, 35(4), 630–641.
- Saul, J. M. (2002a). Speaker meaning, what is said, and what is implicated. *Noûs*, 36(2), 228–248.
- Saul, J. M. (2002b). What is said and psychological reality; Grice’s project and relevance theorists’ criticisms. *Linguistics and Philosophy*, 25(3), 347–372.
- Saul, J. (2010). Speaker-meaning, conversational implicature and calculability. In K. Petrus (Ed.), *Meaning and analysis: New essays on Grice* (pp.170–183). Basingstoke: Palgrave Macmillan.
- Searle, J. (1975). Indirect speech acts. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp.59–82). New York, NY: Academic Press.
- Soames, S. (2009). *Philosophical essays, volume 1: Natural language: What it means and how we use it*. Princeton, NJ: Princeton University Press.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Oxford: Blackwell.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23, 391–434.

- Sterelny, K. (1982). Against conversational implicature. *Journal of Semantics*, 1(2),
- Stiller, A., Goodman, N. D., and Frank, M. C. (2011). Ad-hoc scalar implicature in adults and children. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*. Boston, MA.
- Tiersma, P. M. (1989-1990). The language of perjury: Literal truth, ambiguity, and the false statement requirement. *Southern California Law Review*, 63, 373–432.
- Turner, K. (2001). Review of Wayne A. Davis, *Implicature: Intention, convention, and principle in the failure of Gricean theory*. *Journal of Linguistics*, 37(2), 451–462.
- United States v. Bronston 453 F.2d 555 (2nd Cir. 1971).
- Wierzbicka, A. (1985). Different cultures, different languages, different speech acts. *Journal of Pragmatics*, 9, 145–178.
- Wierzbicka, A. (1987). Boys will be boys: “Radical semantics” vs. “radical pragmatics”. *Language*, 63, 95–114.
- Wierzbicka, A. (1991). *Cross-cultural pragmatics: The semantics of human interaction*. New York, NY: Mouton de Gruyter.
- Wilson, D. and Sperber, D. (1991). Inference and implicature. In S. Davis (Ed.), *Pragmatics: A reader* (pp.377–393). New York, NY: Oxford University Press.
- Wilson, D. and Sperber, D. (2002). Truthfulness and relevance. *Mind*, 111(443), 583–632.
- Wilson, D. and Sperber, D. (2004). Relevance theory. In L. R. Horn and G. Ward (Eds.), *The handbook of pragmatics* (pp.607–632). Oxford: Blackwell.