

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.