

ON VERBAL IRONY*

Deirdre WILSON and Dan SPERBER

Abstract

In this paper, we reconsider an earlier treatment of irony (Sperber and Wilson 1981) in the light of recent developments in pragmatic theory. We argue that, though correct in essentials, this treatment was too restrictive in one respect, and mistaken in another. It was too restrictive because certain types of irony do not fit the analysis of irony as echoic mention proposed in 1981; we argue that they do fit the more general analysis of irony as echoic interpretive use developed in Sperber and Wilson 1986. We were also mistaken, in our earlier paper, in suggesting that the interpretation of non-literal language involves entirely different representational and computational resources from those involved in the interpretation of literal language; we show how a unified treatment is possible within the framework of Sperber and Wilson 1986.

1. Introduction

Some years ago, a referendum was held on whether Britain should enter the Common Market. There was a long campaign beforehand: television programmes were devoted to it; news magazines brought out special issues. At the height of the campaign, an issue of the satirical magazine *Private Eye* appeared. On the cover was a photograph of spectators at a village cricket match, sprawled in deckchairs, heads lolling, fast asleep and snoring; underneath was the caption: 'The Common Market - The Great Debate.'

This is a typical example of verbal irony. As such, it is of interest not only to linguists analysing spontaneous discourse, and to critics analysing literary texts, but also to students of humour. It is curious, though, how little attention has been paid, by linguists, philosophers and literary theorists, to the nature of verbal irony. Theories of metaphor abound. By contrast, while there are many illuminating discussions of particular literary examples, the nature of verbal irony is generally taken for granted. Where theoretical definitions are attempted, irony is still essentially seen as a figure of speech which communicates the opposite of what was literally said. In an earlier paper, 'Irony and the use-mention distinction' (Sperber and Wilson 1981), we drew attention to some problems with this definition, and sketched an alternative account. We would now like to return to some of the issues raised in that paper, and propose some developments and modifications.

2. Traditional accounts of verbal irony

In classical rhetoric, verbal irony is a trope, and as such involves the substitution of a figurative for a literal meaning. Irony is defined as the trope in which the figurative meaning is the opposite of the literal meaning:

Irony is the figure used to convey the opposite of what is said: in irony, the words are not taken in their basic literal sense.

[Du Marsais: Des Tropes, chapter XIV]

Or, as Dr Johnson put it, irony is 'a mode of speech in which the meaning is contrary to the words.'

Modern pragmatic definitions of verbal irony remain firmly in the classical tradition. According to Grice (1975:53), the ironist deliberately flouts the maxim of truthfulness, implicating the opposite of what was literally said. The only significant difference between this and the classical rhetorical account is that what was classically analysed as a figurative meaning is reanalysed as a figurative implication or implicature.

Yet the traditional definition of irony has many weaknesses. In the first place, there are obvious counterexamples to the claim that an ironical utterance invariably communicates the opposite of what is literally said. Here are some illustrations:

A. Ironical understatements

We come upon a customer complaining in a shop, blind with rage and making a public exhibition of himself. I turn to you and say:

- (1) You can tell he's upset.

This is a typical example of ironical understatement. Understatements are traditionally analysed as saying, not the opposite of what is meant, but merely less than what is meant. Though (1) is intuitively ironical, it does not communicate either (2a) or (2b), as the traditional definition of irony would suggest:

- (2) a. You can't tell he's upset.
b. You can tell he's not upset.

Or take Mercutio's ironical comment on his death-wound:

- (3) No, 'tis not so deep as a well, nor so wide as a church door; but 'tis enough, 'twill serve.

[Romeo and Juliet]

Mercutio did not mean to convey that his wound was not deep enough, and would not serve. Ironical understatements thus fail to fit the traditional definition of irony.

B. Ironic quotations

Imagine (4) as said in a cold, wet, windy English spring, or (5), as said in a rainy rush-hour traffic jam in London:

- (4) Oh to be in England
Now that April's there.
[Browning, 'Home thoughts from abroad']
- (5) When a man is tired of London he is tired of life.
[Boswell, Life of Johnson]

Either remark could be ironically intended. To succeed as irony, it must be recognised as a quotation, and not treated merely as communicating the opposite of what is literally said. What (4) would communicate when ironically intended is not - as the traditional definition suggests - a desire to be out of England now that April has arrived, but the idea that the English spring does not always live up to expectations, that the memory of home is not always accurate, that romantic thoughts do not always survive reality, and so on. The point of (5) would be not so much - as the traditional definition suggests - to deny the claim that when a man is tired of London he is tired of life, as to make fun of the sentiments that gave rise to it, the vision of London it was originally intended to convey. Ironic quotations thus fail to fit the traditional definition of irony.

C. Ironic interjections

You have invited me to visit you in Tuscany. Tuscany in May, you write, is the most beautiful place on earth. I arrive in a freak cold spell, wind howling, rain lashing down. As you drive me home along flooded roads, I turn to you and exclaim:

- (6) Ah, Tuscany in May!

My exclamation would almost certainly be ironically intended. Ironic exclamations do not fit the traditional definition of irony. They do not express a complete proposition; hence, they cannot be true or false, and cannot usefully be analysed as deliberate violations of a maxim of truthfulness. Moreover, it is hard to see what the opposite of the interjection 'Ah, Tuscany in May!' would be. Yet verbal irony is clearly present here.

D. Non-ironical falsehoods

So far, we have considered three cases in which irony is present but the traditional definition is not satisfied. In a fourth case, the traditional definition appears to be satisfied, but irony is absent - which suggests that something is missing from the definition.

This example is taken from Grice (1978:124). We are out for a stroll, and pass a car with a broken window. I turn to you and say:

- (7) Look, that car has all its windows intact.

When you ask me what on earth I mean, I explain that I was merely trying to draw your attention, in an ironical way, to the fact that the car has a broken window. My remark meets the traditional definition of irony. I have said something patently false, intending to communicate the opposite, namely (8):

(8) That car has one of its windows broken.

Why do you not instantly leap to the conclusion that (8) is what I meant to convey? As Grice points out, though it fits the traditional definition, (7) cannot be understood as ironical in the circumstances. Clearly, something is missing from the definition.

The traditional definition of irony thus fails on the purely descriptive level: some ironical utterances do not communicate the opposite of what is literally said. But there is a more general problem. According to the traditional definition, an ironical utterance communicates a single determinate proposition which could, if necessary, have been conveyed by means of another, purely literal utterance. On this account, the ironical (9) should be pragmatically equivalent to the strictly literal (10):

(9) What a wonderful party.

(10) What an awful party.

Yet (9) and (10) clearly differ in their pragmatic effects. Intuitively, (9) expresses a certain attitude, creates a certain impression in the hearer. Thus even examples that fit the traditional definition are not adequately described by saying merely that they communicate the opposite of what was literally said. Yet one looks in vain, in either classical rhetoric or modern pragmatics, for attempts to deal with the obvious differences in effects achieved by ironical utterances and their strictly literal counterparts.

The traditional definition of irony raises another, more general problem. An adequate account of irony should provide not just descriptions but explanations. We need to know not just what verbal irony is, but why it exists, how it works, and what is its appeal. Now saying the opposite of what one means is, on the face of it, neither natural or rational. Traditional accounts of verbal irony thus suggest a certain sort of explanation: they suggest that irony is a deviation from the norm; that it should not arise spontaneously; that it is governed by arbitrary rhetorical rules or conventions, which may vary from culture to culture. We believe, on the contrary, that verbal irony is both natural and universal; that it can be expected to arise spontaneously, without having to be taught or learned. If this is so, then we need not only a different definition of verbal irony, but one that suggests a different explanation.

3. Irony as echoic mention

In 'Irony and the use-mention distinction', we outlined a new account of irony based on a distinction between use and mention. This distinction was originally developed to deal with the following sorts of contrasts:

- (11) a. Natasha is a beautiful child.
b. 'Natasha' is a beautiful name.
- (12) a. He deliberately provoked controversy.
b. He deliberately mispronounced 'controversy'.

In (11a), the word 'Natasha' is used to refer to a child; in (11b) it is used to refer to a word of English. In (12a) the word 'controversy' is used to refer to a debate; in (12b) it is used to refer to a word. This self-referential use of words or other linguistic expressions is known in the philosophical literature as 'mention'. Thus, in (11a) and (12a) the words 'Natasha' and 'controversy' are used; in (11b) and (12b) they are mentioned.

In written English, as in the above examples, quotation marks are often used to mark off cases of mention. In the spoken language, such clues are rarely available. Sometimes, there is little room for doubt as to whether use or mention was intended: for instance, it is hard to see how a rational speaker could have intended 'Natasha' in (11a) to refer to anything other than a child, or in (11b) to refer to anything other than a word. Sometimes, though, matters are less straightforward.

For instance, compare (14a) and (14b) as answers to the question in (13):

- (13) Peter: What did Susan say?
- (14) a. Mary: I can't speak to you now.
b. Mary: "I can't speak to you now."

In (14b), as in (11b) and (12b) above, quotation marks are used to distinguish mention from use. In (14a), the sentence 'I can't speak to you now' is used to describe a certain state of affairs; in (14b), it is used to refer to a sentence of English - in other words, it is mentioned. Here, either (14a) or (14b) would, on the face of it, be an acceptable response to the question in (13), and in the spoken language some criterion for recognising the intended interpretation is needed. Notice that, despite their linguistic similarities, the two utterances would be understood in very different ways: for example, in (14a) the referent of 'I' is Mary, the referent of 'you' is Peter, and the referent of 'now' is the time of Mary's utterance; in (14b) the referent of 'I' is Susan, the referent of 'you' is the person Susan was speaking to, and the referent of 'now' is the time of Susan's utterance. In order to understand Mary's reply, Peter must be able to recognise whether the sentence 'I can't speak to you now' was being used or mentioned. A general

criterion for resolving this and other linguistic indeterminacies is proposed in section 5 below.

Utterance (14b) is, of course, a direct quotation. In direct quotations a sentence or other linguistic expression is mentioned. In 'Irony and the use-mention distinction', we argued that indirect quotations could be analysed as cases of mention too. Consider (15) as a possible reply to (13):

(15) Mary: She couldn't speak to me then.

This utterance has two interpretations closely paralleling those of (14a) and (14b). On one interpretation, Mary is not reporting what Susan said - Susan may not have said anything at all - but merely explaining why Susan did not speak. This interpretation parallels (14a) above: Mary uses the sentence 'She couldn't speak to me then' to represent a certain state of affairs. On the other interpretation, paralleling (14b) above, Mary is reporting what Susan said. She is not directly quoting Susan's words: for example, Susan would have said 'I', not 'she,' 'can't', not 'couldn't', and 'now,' not 'then'. On this interpretation, (15) is an indirect quotation, an attempt to reproduce not Susan's words but her meaning.

Now because (15), on this interpretation, is not a direct quotation, it cannot be analysed as involving mention of the sentence Susan uttered. Hence, the contrast between the two interpretations of (15) cannot be analysed in terms of a distinction between the use and mention of sentences. In 'Irony and the use-mention distinction', we argued that it should be analysed in terms of a distinction between the use and mention of propositions. On both interpretations of (15), we claimed, the sentence 'She couldn't speak to me then' is used to express a proposition; the difference between the two interpretations lay in whether that proposition was itself mentioned or used. On the interpretation paralleling (14a), it was used to represent a certain state of affairs; on the interpretation paralleling (14b), it was mentioned - that is, used to represent itself. On this account, (14b), a direct quotation, mentions the sentence Susan spoke, whereas (15), an indirect quotation, mentions the proposition she expressed.

We went on to argue that verbal irony is a variety of indirect quotation, and thus crucially involves the mention of a proposition. The argument ran as follows. Note first that indirect quotations may be used for two rather different purposes - we called them reporting and echoing. A report of speech or thought merely gives information about the content of the original: in (15), for example, Mary may simply want to tell Peter what Susan said. An echoic utterance simultaneously expresses the speaker's attitude or reaction to what was said or thought: for example, Mary may use (15) to let Peter know not only what Susan said, but how she reacted to Susan's utterance, what she thought or felt about it. Irony, we argued, is a variety of echoic utterance, used to express the speaker's attitude to the opinion echoed.

Echoic utterances are used to express a very wide range of attitudes. Compare (16b) and (17b):

- (16) a. Peter: Ah, the old songs are still the best.
b. Mary (fondly): Still the best.
- (17) a. Peter: Ah, the old songs are still the best.
b. Mary (contemptuously): Still the best!

In both cases, Mary's utterance is echoic. In (16b), her attitude to the thought she is echoing is one of approval; from which it follows that she, like Peter, believes the old songs are best. In (17b), her attitude is one of disapproval. She dissociates herself from the thought she is echoing, perhaps indicating indirectly that she believes the old songs are not the best.

Verbal irony, we argue, invariably involves the expression of an attitude of disapproval, thus falling into the same broad category as (17b). The speaker echoes a thought she attributes to someone else, while dissociating herself from it with anything from mild ridicule to savage scorn. To illustrate, consider the following scenario. Mary has lent some money to Bill on the understanding that she will get it back next day. She wonders aloud to Peter whether Bill will keep his word. Peter replies as in (18), thus reassuring her that Bill is trustworthy:

- (18) Bill is an officer and a gentleman.

Next day, Bill rudely denies all knowledge of his debt to Mary. After telling Peter what has happened, Mary comments:

- (19) An officer and a gentleman, indeed.

This utterance is clearly ironical. Mary echoes Peter's earlier reassurance in order to indicate how ridiculous and misleading it turned out to be. To understand (19) as ironical, all that is needed is a realisation that it is echoic, and a recognition of the type of attitude expressed.

Not all ironical echoes are as easily recognisable. The thought being echoed may not have been expressed in an utterance; it may not be attributable to any specific person, but merely to a type of person, or people in general; it may be merely a cultural aspiration or norm. For example, because the code of an officer and a gentleman is widely held up for admiration, a failure to live up to it is always open to ironical comment; hence (19) could be ironically uttered even in the absence of an explicit reassurance such as (18).

From both descriptive and explanatory points of view, the echoic account of irony compares favourably with the traditional account. On the descriptive level, it deals with the case where the speaker communicates the opposite of what she says, and with the various cases where she does not. What is common to all these cases is that the speaker echoes an implicitly attributed opinion, while simultaneously dissociating herself from it. What differ from case to case are the reasons for the dissociation.

Perhaps the most obvious reason for dissociating oneself from a certain opinion is that one believes it to be false; in that case, the speaker may implicate the opposite of what was literally said, and the utterance will satisfy the traditional definition of irony. However, as is shown by the ironical understatement in (1), a speaker may dissociate herself from an opinion echoed not because it is false but because it is too mild - because only someone dull-witted and imperceptive could put it forward in the circumstances. As is shown by the ironical quotations (4) and (5) and the ironical exclamation (6), she may dissociate herself from an opinion echoed not because it is false but because to hold it or express it in the circumstances would be patently absurd. The echoic account thus deals both with the examples that fit the traditional definition, and with those that do not.

The echoic account also sheds light on the problematic example (7) above, in which all the traditional conditions for irony are met but no irony is present. On the echoic account, what is wrong with this example is that no-one's views are being echoed and made fun of. As soon as the missing condition is supplied, irony appears. Thus, consider the following scenario. As we set off for a stroll, I complain to you that my street is being used as a dumping ground for broken-down cars. You tell me I'm imagining things: the cars all look in perfect condition to you. Just then, we pass a car with a broken window, and I turn to you and say:

(7) Look, that car has all its windows intact.

In the circumstances, this remark would certainly be ironical. I am echoing back to you the opinion you have just expressed, but in circumstances where it would clearly be ridiculous to maintain it. Thus, all that is needed to make (7) ironical is an echoic element and an associated attitude of mockery or disapproval.

Notice how inadequate it would be with this example to say that I was merely trying to communicate the opposite of what I had said. The main point of uttering (7) is to express my attitude to the opinion you have just expressed, and in doing so to imply that you were wrong to disagree with me, wrong to think the world is not going to the dogs, and so on. If I had merely wanted to communicate (8) I would, of course, have expressed this proposition directly.

The echoic account of irony also differs from the traditional account in its explanation of irony. If irony is merely a variety of echoic utterance, then it should arise as naturally and spontaneously as echoic utterances in general, and require no separate rhetorical conventions or training. Since echoic utterances are not normally treated as departures from a norm, there is no reason to treat ironical utterances any differently. In fact, the existence of echoic utterances, and the ease with which they are understood, strongly suggests that there is no norm or maxim of literal truthfulness, as most modern pragmatists believe. Any utterance may be understood in two quite different ways: as expressing the speaker's own opinion, or as echoing

or reporting an opinion attributed to someone else; it is up to the hearer to decide which interpretation was intended.

The echoic account of irony has considerable intuitive appeal. Indeed, it seems to accord with the intuitions of that expert ironist Jane Austen, who has Darcy say to that other expert ironist Elizabeth Bennet:

- (20) ... you find great enjoyment in occasionally professing opinions which ... are not your own.

[Austen: *Pride and Prejudice*]

However, we would like to modify one aspect of our original treatment. In the next section, we will argue that the use-mention distinction is merely a special case of a more general distinction, which is needed to account for the full range of echoic utterances, and of ironical utterances in particular.

4. Irony as echoic interpretation

In 'Irony and the use-mention distinction,' we noted that the traditional definition of irony fails to explain the very close links that exist between irony and parody. The Princeton Encyclopaedia of Poetry and Poetics (Preminger 1975), defines parody as 'the exaggerated imitation of a work of art':

Like caricature, it is based on distortion, bringing into bolder relief the salient features of a writer's style or habit of mind. It belongs to the genus satire and thus performs the double-edged task of reform and ridicule.

[Princeton Encyclopaedia:600]

If parody is exaggerated imitation, and irony is saying one thing and meaning the opposite, it is hard to see what the two can have in common. On the echoic account of irony, their similarities and differences can be brought out. Roughly speaking, parody is to direct quotation what irony is to indirect quotation: both involve an echoic allusion and a dissociative attitude, but in parody the echo is primarily of linguistic form; in irony, as we have seen, it is of content.

However, while both irony and parody intuitively involve echoic allusion, it is hard to see how parody can strictly speaking be analysed as a case of mention. Consider the following, from a parody of the later Henry James:

- (21) It was with the sense of a, for him, very memorable something that he peered now into the immediate future, and tried, not without compunction, to take that period up where he had, prospectively, left it.

[Beerbohm: A Christmas Garland:3]

There are clear echoes here of James's style, but in what sense is (21) a mention? Mention, we have seen, involves identical reproduction of an original; but (21) is not a reproduction of anything James wrote: it merely resembles what he wrote. While direct quotation involves mention in the strict sense - the exact words of the original are reproduced - parody is typically based on looser forms of resemblance.

According to the Princeton Encyclopaedia, parody may be directed not only at style, but also at content, or 'habits of mind'. The following, from a parody of Galsworthy's Forsyte Saga, illustrates this aspect of parody:

- (22) Adrian Berridge paused on the threshold, as was his wont, with closed eyes and dilated nostrils, enjoying the aroma of complex freshness which the dining-room had at this hour... Here were the immediate scents of dry toast, of China tea, of napery fresh from the wash, together with that vague, super-subtle scent which boiled eggs give out through their unbroken shells. And as a permanent base to these there was the scent of much-polished Chippendale, and of bees' waxed parquet, and of Persian rugs...

Just at that moment, heralded by a slight fragrance of old lace and of that peculiar, almost unseizable odour that uncut turquoises have, Mrs Berridge appeared.

[Beerbohm: A Christmas Garland:110-11]

Clearly, there are echoes here of both form and content: of the sort of thing Galsworthy said and the way he said it; but in what sense is (22) a mention? It is not an identical reproduction of anything Galsworthy wrote: it merely resembles what he wrote. Strictly speaking, then, neither parody of form nor parody of content can be analysed in terms of mention.

In parody, as the Princeton Encyclopaedia says, an element of exaggeration is often involved. The same is true in many standard examples of verbal irony. One such example is treated in Paola Fanutza's excellent dissertation 'Irony in Jane Austen's Emma' (Fanutza 1985), in which a wide variety of ironical utterances are insightfully discussed. Emma is playing with her sister's child. Mr Knightley comments:

'If you were as much guided by nature in your estimate of men and women, and as little under the power of fancy and whim in your dealings with them, as you are where these children are concerned, we might always think alike.'

To which Emma replies:

'To be sure - our discordancies must always arise from my being in the wrong.'

[Fanutza: 47-8]

What Emma ironically echoes back to Mr Knightley is a caricature of the opinions he has just expressed. If mention involves identical reproduction of an original, then where irony involves an element of exaggeration or caricature, an analysis in terms of mention is too narrow.

In fact, what is true of ironical echoes is true of all indirect quotations. Reports of speech are not always identical reproductions of the content of the original: they may be paraphrases or summaries; they may be elaborations, spelling out some assumptions or implications that the original speaker took for granted, or that struck the hearer as particularly relevant. In such cases, the content of the indirect speech report resembles the content of the original without, however, being an identical reproduction of it; and the analysis of indirect speech in terms of mention is too restrictive.

In our book Relevance (Sperber and Wilson 1986: chapter 4, sections 7 & 9) we therefore replaced the notion of mention by a notion of interpretive resemblance, or resemblance of content. In the appropriate circumstances, we argued, any object in the world can be used to represent any other object it resembles. A uniformed doll can be used to represent a soldier, an arrangement of cutlery and glasses can be used to represent a road accident, a set of vertical lines to represent the heights of students in a class. Such representations are used in communication for two main purposes: to inform an audience about the properties of an original, and for the expression of attitude. I may show you a uniformed doll so that you can recognise a soldier when you see one; I may communicate my attitude to soldiers by, say, kicking the doll.

Utterances, like other objects, enter into a variety of resemblance relations. It is not surprising, therefore, to find these resemblances exploited, and for just the same purposes, in verbal communication. Onomatopoeia is based on resemblances in sound, verbal mimicry on resemblances in phonetic and phonological form, direct quotation and parody on resemblances in syntactic and lexical form, translation on resemblances in logical form, and paraphrase and summary on resemblances in propositional content. Where resemblance of propositional content is involved, we talk of interpretive resemblance; we reanalyse echoic utterances as echoic interpretations of an attributed thought or utterance, and verbal irony as a variety of echoic interpretation. In other respects, the account of verbal irony developed in 'Irony and the use-mention distinction' remains unchanged.

What does it mean to say that one thought or utterance interpretively resembles another? Resemblance in general involves a sharing of properties: the more shared properties, the greater the resemblance. Interpretive resemblance, or resemblance in propositional content, we argue, is best analysed as a sharing of logical and contextual implications: the more shared implications, the greater the interpretive resemblance. It is possible for two propositions to share all their implications; when one of these is interpretively used to represent the other, we say that it is a literal interpretation of that other proposition. On this account, literalness is just a special case of

interpretive resemblance. However, one representation may interpretively resemble another when the two merely have implications in common.

Let us illustrate these ideas with an example. Mary says to Peter:

(23) a. I met an agent last night. b. He can make me rich and famous.

As we have seen, an utterance such as (23b) has two possible interpretations, (24a) and (24b):

(24) a. He can make me rich and famous, I believe.
b. He can make me rich and famous, he says.

On interpretation (24a), Mary's utterance is a straightforward assertion. On interpretation (24b), it is either an echoic utterance or a report of speech, and must therefore bear some degree of interpretive resemblance to what the agent said.

Suppose that what the agent said was actually (25):

(25) I can make you rich and famous.

Then Mary's utterance would be a literal interpretation of what the agent said: the propositions expressed by the two utterances would be identical, and hence share all their implications in every context. In that case, it is quite reasonable to see Mary as having mentioned the proposition the agent originally expressed.

Suppose, however, that what the agent said was actually (26):

(26) I can do for you what Michael Caine's agent did for him.

Then Mary's utterance would be a less than literal interpretation of what the agent said, and it would not be reasonable to claim that Mary had mentioned the proposition the agent originally expressed. It may be common knowledge, though, that Michael Caine's agent made him very, very rich and famous. In a context containing this assumption, (26) would contextually imply (23b). The report in (23b) thus interpretively resembles the agent's utterance in (26): the propositions expressed by the two utterances have implications in common. Many reports of speech, and many echoic utterances, are based on this looser form of resemblance.

We propose, then, to analyse indirect speech reports, echoic utterances and irony not as literal interpretations (i.e. mentions) of an attributed thought or utterance, but simply as interpretations, literal or non-literal, of an attributed thought or utterance. This change corrects an over-restrictive feature of our earlier account.

5. The recognition of irony.

Wayne Booth tells of a puzzling encounter with a graduate student, a sophisticated reader who was arguing that the whole of Pride and Prejudice is ironic. This student expressed a dislike of Mr Bennet, and when asked to explain said, 'Well, for one thing, he's really quite stupid, in spite of his claims to cleverness, because he says towards the end that Wickham is his favourite son-in-law.' Booth comments:

He retracted in embarrassment, of course, as soon as we had looked at the passage together: "'I admire all my three sons-in-law highly,' said he. 'Wickham, perhaps, is my favourite; but I think I shall like your husband [Darcy] quite as well as Jane's.'" How could he have missed Mr Bennet's ironic joke when he was in fact working hard to find evidence that the author was always ironic? [Booth 1974:1]

Such failures, even in sophisticated readers, are quite common. Walter Scott, for example, is notorious for having missed the irony in Elizabeth Bennet's remark that she began to appreciate Darcy when she first set eyes on his magnificent estate at Pemberley (see Southam 1976:155, 159 and 165, footnote 8). The subtler the irony, the greater the risks.

There is no such thing as a fail-safe diagnostic of irony. All communication takes place at a risk. The communicator's intentions cannot be decoded or deduced, but must be inferred by a fallible process of hypothesis formation and evaluation; even the best hypothesis may turn out to be wrong. The standard works on irony (e.g. Booth 1974, Muecke 1969) provide good surveys of the sort of clues that put alert readers or hearers on the track of irony; but the clues themselves have to interact with more general interpretation procedures. In our book Relevance, we outline a general criterion for the resolution of linguistic indeterminacies which, we suggest, is used in every aspect of utterance interpretation, including the recognition of irony. This criterion is justified by some basic assumptions about the nature of relevance and its role in communication and cognition, which we can do no more than sketch briefly here. (For further details, see Relevance; for summary and discussion, see Sperber and Wilson 1987.)

Human information processing, we argue, requires some mental effort and achieves some cognitive effect. Some effort of attention, memory and reasoning is required. Some effect is achieved in terms of alterations to the individual's beliefs: the addition of contextual implications, the cancellation of existing assumptions, or the strengthening of existing assumptions. Such effects we call contextual effects. We characterise a comparative notion of relevance in terms of effect and effort as follows:

Relevance

- (a) Other things being equal, the greater the contextual effect achieved by the processing of a given piece of information, the greater its relevance for the individual who processes it.
- (b) Other things being equal, the greater the effort involved in the processing of a given piece of information, the smaller its relevance for the individual who processes it.

We claim that humans automatically aim at maximal relevance: that is, maximal contextual effect for minimal processing effort. This is the single general factor which determines the course of human information processing. It determines which information is attended to, which background assumptions are retrieved from memory and used as context, which inferences are drawn.

To communicate is, among other things, to claim someone's attention, and hence to demand some expenditure of effort. People will not pay attention unless they expect to obtain information that is rich enough in contextual effects to be relevant to them. Hence, to communicate is to imply that the stimulus used (for example, the utterance) is worth the audience's attention. Any utterance addressed to someone automatically conveys a presumption of its own relevance. This fact, we call the principle of relevance.

The principle of relevance differs from every other principle, maxim, convention or presumption proposed in modern pragmatics in that it is not something that people have to know, let alone learn, in order to communicate effectively; it is not something that they obey or might disobey: it is an exceptionless generalisation about human communicative behaviour. What people do have to know, and always do know when they recognise an utterance as addressed to them, is that the speaker intends that particular utterance to seem relevant enough to them to be worth their attention. In other words, what people have to recognise is not the principle of relevance in its general form, but the particular instantiations of it that they encounter.

Speakers may try hard or not at all to be relevant to their audience; they may succeed or fail; they still convey a presumption of relevance: that is, they convey that they have done what was necessary to produce an adequately relevant utterance.

Relevance, we said, is a matter of contextual effect and processing effort. On the effect side, it is in the interest of hearers that speakers offer the most relevant information they have. However, speakers have their own legitimate aims, and as a result may choose to offer some other information which is less than maximally relevant. Even so, to be worth the hearer's attention, this information must yield at least adequate effects, and the speaker manifestly intends the hearer to assume that this is so. On the effort side, there may be different ways of achieving the intended effects, all equally easy for the speaker to produce, but requiring different amounts of processing effort from the hearer. Here, a rational speaker will choose the formulation that is easiest for the hearer to process, and manifestly intends the hearer to

assume that this is so. In other words, the presumption of relevance has two parts: a presumption of adequate effect on the one hand, and a presumption of minimally necessary effort on the other.

As we have seen, the linguistic form of an utterance grossly underdetermines its interpretation. Direct quotations, indirect quotations, echoic utterances and irony are not recognisable from their linguistic form alone. Various pragmatic theories appeal to complex sets of rules, maxims or conventions to explain how this linguistic indeterminacy is contextually overcome. We claim that the principle of relevance is enough on its own to explain how linguistic form and background knowledge interact to determine verbal comprehension.

In a nutshell, for an utterance to be understood, it must have one and only one interpretation consistent with the principle of relevance - one and only one interpretation, that is, on which a rational speaker might have thought it would be relevant enough to be worth the hearer's attention. The speaker's task is to see to it that the intended interpretation is consistent with the principle of relevance; otherwise, she runs the risk of not being properly understood. The hearer's task is to find the interpretation which is consistent with the principle of relevance; otherwise, he runs the risk of misunderstanding it, or not understanding it at all.

To illustrate these ideas, consider how Peter might set about interpreting Mary's remark in (19) above:

(19) An officer and a gentleman, indeed.

As we have seen, this remark has two possible interpretations, (27a) and (27b), corresponding to what we originally called use and mention of the proposition expressed:

- (27) a. Bill is an officer and a gentleman, I believe.
b. Bill is an officer and a gentleman, you said.

Suppose that interpretation (27a) is the first to occur to Peter, and thus the first to be tested for consistency with the principle of relevance. To be consistent with the principle of relevance, the utterance, on that interpretation, must achieve adequate contextual effects, or at least have been rationally expected to do so. To achieve contextual effects, it must either have contextual implications, strengthen an existing assumption, or contradict and eliminate an existing assumption. Now the hypothesis that, in the circumstances described, Mary might genuinely believe that Bill is an officer and a gentleman contradicts known facts; rather than eliminating existing assumptions, it is likely itself to be rejected. In the circumstances, Mary could not rationally have expected her utterance, on this interpretation, to achieve adequate contextual effects, and interpretation (27a) must be rejected as inconsistent with the principle of relevance.

Now consider (27b). This could be understood as either a report of speech or an echoic interpretation of an attributed thought or utterance. Suppose Peter decides to test the hypothesis that it is a straightforward

report of speech. To be consistent with the principle of relevance on this interpretation, Mary's utterance must achieve adequate contextual effects - for example, by adding contextual implications, or by strengthening an existing assumption - or must at least have been rationally expected to do so. But, unless Peter's memory is defective, he will be able to remember his earlier remark, and will need no reminding of it. Hence, the hypothesis that Mary's utterance was intended as a report of speech is inconsistent with the principle of relevance.

The only remaining possibility is that Mary's utterance was intended as echoic: that is, she was echoing Peter's earlier utterance in order to express her attitude to it. What attitude was she intending to express? The hypothesis that her attitude was one of approval can be ruled out for reasons already given: in the circumstances, the idea that Mary could genuinely believe that Bill is an officer and a gentleman contradicts known facts. Hence, the only possible hypothesis is that Mary was echoing Peter's utterance in order to dissociate herself from the opinion it expressed.

Is Mary's utterance, on this interpretation, consistent with the principle of relevance? Would it achieve adequate contextual effects for the minimum necessary effort, in a way that Mary could manifestly have foreseen? It is easy to see how it might achieve adequate contextual effects: for example, it draws Peter's attention to the various ways in which Bill's behaviour has fallen short of the ideal, and to the fact that he has made a mistake, is possibly responsible for Mary's loss of money, is unlikely to be trusted so readily again in his assessment of character, and so on; moreover, these are effects that Mary might easily have foreseen. As long as no other utterance would have achieved these effects more economically, this interpretation would also be satisfactory on the effort side, and would therefore be consistent with the principle of relevance.

In *Relevance*, we show that having found an interpretation consistent with the principle of relevance, the hearer need look no further: there is never more than one. The first interpretation tested and found consistent with the principle of relevance is the only interpretation consistent with the principle of relevance, and is the one the hearer should choose.

6. The communication of impressions and attitudes

What do ironical utterances communicate? While rejecting the traditional claim that they invariably communicate the opposite of what was literally said, we have, as yet, offered no alternative account. Ironical utterances, we said, communicate a certain attitude, create a certain impression in the hearer; but how are attitudes and impressions to be dealt with in a theory of communication?

At the end of 'Irony and the use-mention distinction', we were rather sceptical about the possibility of dealing with the communication of impressions and attitudes within the framework of what we called 'logical pragmatics', in which utterance interpretation was seen

primarily as an inferential process involving the construction and manipulation of propositional (conceptual) representations:

An ironical utterance carries suggestions of attitude ... which cannot be made entirely explicit in propositional form. In this respect, a logical-pragmatic model does not provide a better description ... than a semantic model. On the other hand, our analysis of irony ... crucially involves the evocation of an attitude - that of the speaker to the proposition mentioned. This attitude may imply a number of propositions, but it is not reducible to a set of propositions. Our analysis thus suggests that a logical-pragmatic theory dealing with the interpretation of utterances as an inferential process must be supplemented by what could be called a 'rhetorical-pragmatic' or 'rhetorical' theory dealing with evocation.

[Sperber and Wilson 1981:317]

The suggestion was that the representational and computational resources of 'logical pragmatics' would not be adequate to deal with expressions of attitude, which would have to be handled by entirely different mechanisms.

We would not now draw such a sharp distinction between logical and rhetorical pragmatics. We believe that the communication of impressions and attitudes can be handled in much the same terms as the communication of more standard implicatures. In this section, we will suggest how this might be done.

In Relevance (chapter 1), we argued that communication involves an intention to modify the audience's cognitive environment. The cognitive environment of an individual is a set of assumptions that are manifest to him; an assumption is manifest to an individual at a given time if and only if he is capable at that time of representing it conceptually and accepting that representation as true or probably true. Manifest assumptions may differ in their degree of manifestness: the more likely they are to be entertained, the more strongly manifest they are. To modify the cognitive environment of an audience is to make a certain set of assumptions manifest, or more manifest, to him. The intention to modify the cognitive environment of an audience we called the informative intention.

Consider now how utterance (28) might be handled in this framework:

(28) Mary to Peter: I can't stay to dinner tonight.

Mary's utterance modifies Peter's cognitive environment by making manifest to him a variety of assumptions. Peter's task, in interpreting (28), is to recognise Mary's informative intention: that is, to decide which set of assumptions she intended to make manifest, or more manifest, to him. In recognising Mary's informative intention, he is guided by the criterion of consistency with the principle of relevance: that is, he looks for an interpretation on which (28) might rationally

have been expected to achieve adequate contextual effects for the minimum necessary effort.

Among the assumptions made strongly manifest to Peter by Mary's utterance will be (29):

(29) Mary has said to Peter that she can't stay to dinner that night.

It is easy to see how (29) might achieve contextual effects in a context easily accessible to Peter. For example, by assuming that Mary is a trustworthy communicator, Peter can infer that she is unable to stay to dinner; from this, together with other assumptions, he can infer that some of his plans for the evening will have to be abandoned; depending on the relationship between them, the effort he has gone to in preparing the meal, and the reason for her refusal, further implications would follow. In recognising Mary's informative intention, Peter is entitled to assume that she intended to make manifest, or more manifest, to him enough of these implications to make her utterance worth his attention. These will be the implicatures of her utterance.

Manifestness, we said, is a matter of degree. Among the assumptions made manifest to Peter by Mary's utterance, some will be more strongly manifest than others; moreover, Mary's intentions concerning these assumptions will be more strongly manifest in some cases than in others. In the case of (28), for example, it is hard to see how Mary could have expected her utterance to be relevant enough to be worth Peter's attention if it did not make assumption (29) manifest to him. Let us say that when a communicator makes strongly manifest her intention to make a certain assumption strongly manifest, then that assumption is strongly communicated. Then in the circumstances described, (29) will be strongly communicated by (28).

However, not all the speaker's intentions are so easily pinned down. It may be clear, for example, that in saying (28), Mary intended to make manifest to Peter that she couldn't stay to dinner, and fairly clear that she intended him to infer from this that he would have to change his plans for the evening - by inviting someone else, say, by preparing less food, or by abandoning the meal and going out. It may not be so clear, however, that she expected him to follow any particular one of these courses of action, or to carry out the chosen course of action in any particular way. Thus, a wide array of assumptions is made marginally more manifest by Mary's utterance; as the chains of inference grow longer, and the set of possible conclusions wider, Mary's informative intentions become correspondingly less manifest, to the point where they are no longer manifest at all. We might describe this quite standard situation by saying that strong communication shades off into something less determinate, where the hearer is encouraged to think along certain lines, without being forced to any definite conclusion.

Let us say that when the communicator's intention is to increase simultaneously the manifestness of a wide range of assumptions, so that her intention concerning each of these assumptions is itself weakly

manifest, each of these assumptions is weakly implicated. Then, by saying (28), Mary might weakly implicate a range of assumptions having to do with changes in Peter's plans for the evening, changes in her relationship to Peter, and so on. The less strongly manifest her intentions concerning such assumptions, the weaker the communication will be.

Most recent approaches to pragmatics have concentrated on strong communication. One of the advantages of verbal communication is that it allows the strongest possible form of communication to take place: it enables the hearer to pin down the speaker's intentions about the explicit content of her utterance to a single, strongly manifest candidate, with no alternative worth considering at all. On the other hand, what is implicitly conveyed in verbal communication is generally weakly communicated. Because all communication has been seen as strong communication, the vagueness of most implicatures and of non-literal forms of expression has been idealised away, and the communication of feelings, attitudes and impressions has been largely ignored. The approach just sketched, by contrast, provides a way of giving a precise description and explanation of the weaker effects of communication.

Suppose, for example, that in saying (28) Mary speaks sadly, thus making manifest to Peter assumption (30):

(30) Mary has spoken sadly.

The effects thus created can be analysed as weak implicatures. Assumption (30) makes manifest, or more manifest, to Peter a wide array of further assumptions. Why is Mary sad? Is it because she can't stay to dinner? Does Peter want her to be sad? How sad is she? Would she cheer up if he invited her for another evening? Would she stay if he cooked dinner immediately? If he offered to drive her home afterwards? If he lit a fire? If he served fish instead of meat? By processing (30) in a context obtained by answering these and other questions, Peter can increase the contextual effects of (28). On the assumption that Mary intended to make (30) manifest to him, Peter is entitled to conclude that she also intended to make manifest to him enough of these effects to make (30) worth processing.

What we are suggesting is that the assumptions made manifest to Peter by Mary's utterance include not only the contextual effects of the proposition she has expressed, but also those of various descriptions of her utterance - her tone of voice, facial expression, accompanying gestures, and so on. Some subset of these may form part of the intended interpretation of her utterance, this subset being selected, as usual, by the criterion of consistency with the principle of relevance. The resulting communication will, of course, be weak, but it will not be different in kind from the communication of quite standard implicatures. In either case, the interpretation process will involve the inferential processing of newly presented information in the context of assumptions supplied by the hearer. What makes communication weak is merely the fact that a very wide array of assumptions is made manifest, or more manifest, so that, in forming hypotheses about the

speaker's informative intentions, the hearer has a very wide range of contexts and contextual effects to choose from.

Let us return, in the light of this suggestion, to our original example: the magazine cover with the caption 'The Common Market - The Great Debate' printed across a photograph of spectators asleep at a village cricket match. How should this cover be understood?

The caption and photograph would, between them, have made a variety of assumptions manifest to contemporary readers. The caption would give them access to encyclopaedic information about the Common Market, including the information that a referendum on Britain's entry to it was shortly to be held, and that the referendum issue had been repeatedly referred to by politicians and journalists as the 'Great Debate'; the photograph would give them access to a range of assumptions about the length and uneventfulness of village cricket matches, the lack of excitement normally felt by spectators, and so on. What set of assumptions was this cover intended to make manifest to contemporary readers? That is, on what interpretation might it have been intended to achieve adequate contextual effects for the minimum necessary effort?

Here, some hypotheses can be automatically eliminated as inconsistent with the principle of relevance. These would include the hypothesis that the utterance was literally intended and that the designers of the cover merely intended to make manifest, or more manifest, the assumption that the debate on the Common Market would be exciting, together with some subset of its contextual effects. A communicator who merely wanted to achieve these effects could have achieved them without putting readers to the unnecessary effort of processing information about village cricket matches: hence, the use of this cover to achieve these effects would be inconsistent with the principle of relevance.

Consider now the hypothesis that the description 'The Great Debate' was echoically used - a hypothesis that would have come easily to contemporary readers. Clearly, for reasons just given, the attitude being expressed to the opinion echoed cannot have been one of approval. By contrast, the hypothesis that it was one of dissociation or disapproval is strongly confirmed by the accompanying photograph. This photograph conveys an impression of stupefying boredom. On the assumption that the Common Market Debate resembles the village cricket match in relevant respects, readers can infer that this debate too is one of stupefying boredom; that to call it a 'Great Debate' is ridiculous; that it is not, in fact, a great debate. On the assumption that the cover designers intended to make these assumptions manifest to the audience, their behaviour would be consistent with the principle of relevance. This is the interpretation, then, that the audience should choose.

On this account, the magazine cover would achieve a combination of strong and weak communication. It would strongly communicate that it was ridiculous to call the Common Market debate a 'Great Debate', that this debate was very boring, and that it was not, in fact, a great debate. It would weakly implicate a wide array of contextual effects

derivable from these assumptions, in terms of which readers would be able to create for themselves an impression of just how ridiculous the media descriptions were, and just how boring the debate was likely to be.

7. Concluding remark

In this paper, we have analysed irony as a variety of echoic interpretive use, in which the communicator dissociates herself from the opinion echoed with accompanying ridicule or scorn. The recognition of verbal irony, and of what it communicates, depends on an interaction between the linguistic form of the utterance, the shared cognitive environment of communicator and audience, and the criterion of consistency with the principle of relevance. This approach to irony, which appears to offer both better descriptions and better explanations than traditional accounts, has one surprising consequence which is perhaps worth mentioning here.

It is tempting, in interpreting a literary text from an author one respects, to look further and further for hidden implications. Having found an interpretation consistent with the principle of relevance - an interpretation (which may itself be very rich and very vague) which the writer might have thought of as adequate repayment for the reader's effort - why not go on and look for ever richer implications and reverberations? If we are right, and considerations of relevance lie at the heart of verbal communication, such searches go beyond the domain of communication proper. Though the writer might have wished to communicate more than the first interpretation tested and found consistent with the principle of relevance, she cannot rationally have intended to. Relevance theory thus explains how irony is (fallibly) recognised, and sets an upper limit to what the ironist can rationally expect to achieve.

8. Notes

* Deirdre Wilson is grateful to the British Council, the Associazione Italiana di Anglistica and the University of Turin for their hospitality at the AIA conference on *Le Forme Del Comico*, October 1985, at which an earlier version of this paper was presented.

9. References

- Beerbohm, M. (1950) A Christmas Garland. London, Heinemann.
- Booth, W. (1974) A Rhetoric of Irony. U. of Chicago Press, Chicago.
- Fanutza, P. (1985) Irony in Jane Austen's Emma. Tesi di Laurea (Lingue e Letterature Straniere), Universita Degli Studi Di Cagliari.

Grice, H.P. (1975) 'Logic and conversation.' In P. Cole and J. Morgan (eds) Syntax and Semantics 3: Speech Acts. New York, Academic Press, 41-58.

Grice, H.P. (1978) 'Further notes on logic and conversation.' In P. Cole (ed.) Syntax and Semantics 9: Pragmatics. New York, Academic Press, 113-28.

Muecke, D. (1969) The Compass of Irony. Methuen, London.

Preminger, A. (ed.) (1974) Princeton Encyclopaedia of Poetry and Poetics. Macmillan, London.

Southam, B.C. (1976) Sense and Sensibility, Pride and Prejudice and Mansfield Park: A Casebook. Macmillan, London.

Sperber, D. and Wilson, D. (1981) 'Irony and the use-mention distinction.' In P. Cole (ed.) Radical Pragmatics. Academic Press, New York, 295-318.

Sperber, D. and Wilson, D. (1986) Relevance: Communication and Cognition. Blackwell, Oxford; Harvard UP, Cambridge, Mass.

Sperber, D. and Wilson, D. (1987) 'A Precipis of Relevance' and 'Presumptions of Relevance', in Behavioral and Brain Sciences 10, 697-754.