

# THE DEFINITION AND PROCESSING OF IRONY

by  
Svend Østergaard

Can we give a definition of irony? From the point of view of cognitive linguistics, the category of ironic utterances is fuzzy, and therefore a semantic definition cannot be provided. This paper argues for a minimal pragmatic definition. Irony is defined as an insincere statement, where the speaker intends the listener to perceive it as insincere and where the statement constitutes a misfit with some aspect of the context. The processing of irony depends on a conflict between a contextual meaning and a salient, but contextually inappropriate meaning. The processing will depend on the equilibrium between the two sources of meaning. If the contextual meaning is strong, the irony of the statement will be easy to access; however, if the contextual meaning is weak, the salient meaning of the statement will be foregrounded and the ironic meaning will be more difficult to access.

## *1. Introduction*

With the term *irony* we refer to many different types of verbal expressions, but also to situations that are not necessarily expressed verbally. Irony is therefore not a unitary phenomenon. For instance, in a narrative the protagonist can utter a sentence that is ironic; this is called verbal irony. However, the sequence of events that unfold in the narrative can itself be called ironic relative to the protagonist. In the first case, the protagonist utters a sentence (performs a speech act) which is insincere; for example, it might mean the opposite of what she considers as the veridical situation. In the latter case, the protagonist acts in view of a goal, but in the situation there are

circumstances which, unknown to the protagonist, cause her acts to have the opposite effect. There is a schematic similarity between the two cases, since in both of them, the act of the protagonist is negated by the veridical situation; but there is also a big difference, since the discrepancy between the act and the veridical situation is intended in the verbal case, but not in the non-verbal case. In fact, in the verbal case it is exactly the intended mismatch between the veridical situation (as seen by the speaker) and the meaning of the utterance that qualifies it as ironic, whereas in the ironic event it is the unintended mismatch that makes up the irony.

This amounts to saying that, in accordance with the general theory of categorization developed in cognitive linguistics, cf. Lakoff (1987), the term irony denotes a category with a heterogeneous structure. The difference depicted above holds true at least as regards ironic events versus ironic discourse, but even if we restrict ourselves to verbal irony, there is no reason to believe that the category of ironic utterances has a uniform structure, as assumed in many analyses of irony, such as those offered by Brown and Levinson (1978: 226) and Searle (1979), who suggest that irony is understood by assuming the opposite of the sentence's literal meaning. In Sperber and Wilson (1981), it is claimed that an ironic statement is like an echo reminding the listener of a similar statement which on a previous occasion, has been or could have been uttered; through the echoic form, the speaker expresses her attitude towards the situation. In Clark and Gerrig (1984), verbal irony involves pretense, meaning that the speaker of an ironic sentence pretends to be some other person proclaiming the utterance to an unknown audience. Kumon-Nakamura et al. (1995) make the claim that ironic remarks have an effect by alluding to a failed expectation. Many theories on irony argue that the phenomenon can be defined through necessary and sufficient conditions, but in so far as we are talking about semantic conditions, this seems not to be a fruitful approach.

A cognitive approach to irony will proceed from the following assumptions:

1) Irony is not primarily a rhetorical device or a literary technique; it is found in ordinary language with a relatively high frequency. In the corpus studied in Tannen (1984), irony was found in 8% of all turns; in Gibbs' (2000) study, it was found in 7% of all turns. In many cases, a sentence is understood as ironic without this being the intention (Gibbs and O'Brien 1991). Conversely, the meaning of a sentence is often understood without the sentence being recognized as ironic. All of this points to irony as an integrated part of human communication which works implicitly and automatically. The studies mentioned above show that irony is used in talk among friends to establish group membership by commenting on individuals who are not group members. The case of irony is a bit like what Lakoff and Johnson (1980, 1999) found concerning metaphor. Metaphor is not (just) a rhetorical figure found in literary texts; it is used in ordinary language all the time and it is processed without necessarily being recognized as a metaphor. However, there is also a crucial difference, because the cognitive ability to think in metaphors (and metonymy) is a necessary condition for language use; without these two operations, the lexicon would explode and put too big a load on the semantic memory. Irony, on the other hand, is a pragmatic operator with no influence on the structure of language.

2) This then raises the question, what is the pragmatics of irony? Again, it is probably not possible to give a unique definition of the pragmatic effects of irony. In the literature, we find different – sometimes contradictory – claims about irony which, however, can all be true. For instance, some argue that irony makes a negative meaning less rude (Dews et al. 1995), whereas others suggest that ironic statements are more rude than literal statements (Colston 1997). Interestingly, Ivanko and Pexman (2003) hypothesize that

in a strongly negative context, an ironic statement is considered more mocking and less polite than a literal statement, whereas in a slightly negative context the opposite is true. We will return to this study in more detail below.

There are also competing theories about the social function of irony, considering it to be either a source of affiliation or of conflict between individuals. Such competing views do not constitute a problem for a cognitive approach to irony, since there is no reason to assume that irony is a unitary phenomenon that can be described using a single set of necessary and sufficient conditions. In fact, it might even be unclear whether a given sentence is meant to be ironic or not. For instance, the first sentence in the following text – taken from a homepage on tennis – might be considered ironic, but will probably not be recognized as such by the majority of readers:

Maybe it wasn't as warm as it appeared. The seven-time Roland Garros champion had made a point of blaming his slow start in his opening two rounds firmly on the cold conditions; yet here he was, once again bogged down in another long first set despite the agreeable mid-afternoon temperatures.

If we claim that irony is pragmatically motivated (Haverkate 1990), it follows that it is not conceptual, i.e. we cannot give a semantic description of irony. In this regard, irony also differs from metaphors, which rely on conceptual mappings between semantic domains (Lakoff and Johnson 1999). Still, there are cases where we can say that irony is motivated by general principles of how information is presented; this will also be discussed below.

## 2. Ironic meaning as the opposite of literal meaning

According to the classical definition of irony, an ironic statement is a figure of speech which conveys the opposite of the meaning it expresses verbally; the prototypical case would be someone saying, *What a lovely weather*, when in fact it is pouring down. The problem with this definition is that it is easy to come up with counterexamples, because in many cases it is not clear what the opposite of the literal meaning is. For instance, if someone says *Thanks* in an ironic way, the meaning is not *No thanks* or, *It is not the case that I thank you*, but rather, *You have done something for which you do not deserve gratitude*.

There seems to be classes of ironic statements where the ironic meaning cannot be expressed by its opposite. For instance, the meaning of ironic understatements can never be the opposite of the statement: *It seems to be raining*, when it is pouring down, obviously does not mean *It does not seem to be raining*. Ironic expressive forms like *Thanks* never convey the opposite meaning. The same is true of ironic questions like, *Is it possible for you to arrive in time just once?* Wilson and Sperber (1992) mention ironic quotations as examples; for instance, if someone in a cold, wet English spring says, *Oh to be in England now that April's there* (a quotation from Robert Browning), the meaning is not a desire to be in England, but to express (contrary to Browning's feeling) that the English spring does not always live up to expectation. Furthermore, sentences where the irony is directed toward presupposed meaning also fall outside the classical definition (cf. Haverkate 1990). For instance, *Jane has stopped organizing her exciting parties*, where *exciting* is ironic. The presupposition is that Jane has previously been organizing parties and now she does not. The meaning can therefore not be rendered by the negation *Jane has not stopped organizing her exciting parties*. So there is a host of different types of ironic expressions that cannot

be described by assuming that the intended meaning is the opposite of what is said.

But even in the simple case where the meaning seems to be the opposite of what is literally expressed, Sperber and Wilson (1981) point to a deeper problem. For instance, *Lovely weather!* – said when it is pouring down – is only ironic if the receiver of the sentence believes that the speaker knows that it is pouring down and that she therefore experiences the weather as awful. It seems therefore counterintuitive to say that the speaker wants to convey the meaning *The weather is awful*, because this is already known by the dialogue participants. So what is the purpose of saying, *Lovely weather!*, when in fact everybody acknowledges that everybody acknowledges that the weather is awful? It is at least clear that the ironic remark does not convey information about the veridical world, because if I say, *Awful weather!*, then, given the situation, I convey 'trivial information' (cf. Beaugrande and Dressler 1981). Normally, trivial information has to be upgraded in the information stream; for instance, *It is awful weather, but I like it because it gives rain to the crops*. If that is not the case, then the trivial information does not serve the function of information transfer, but fulfills other intersubjective functions, such as to remind us of common expectations, attitudes, or history – not an uncommon way of using language. Think of a group of football fans talking about the match afterwards; a lot of utterances will be about facts that are known to all the participants having watched the match, but they still serve intergroup bonding rather than providing information. Or consider two former schoolmates meeting after many years, recalling old episodes from their common history. They are not informing each other about anything, just confirming their common past.

Turning to *awful's* ironic counterpart, *lovely weather*, we have a non-trivial sentence which in fact contradicts our common assumptions

about the situation. Beaugrande and Dressler call this 'third order' information, and when we encounter this in an information exchange it has to be downgraded to normal information; this could for instance be, *It is lovely weather today, because the crops really need water*. Here, the unexpected sentence is justified according to our general schematic knowledge about the world. However, if the sentence stands alone and is not downgraded, then it separates itself from the information stream and fulfills the same function as its trivial counterpart, inasmuch as it only serves intersubjective bonding, by reminding the discourse participants of the shared expectations and attitudes. It is generally acknowledged that speakers (and texts) avoid trivial information that cannot be upgraded (Beaugrande and Dressler 1981 – which could be one of the decisive reasons for using the ironic version. Another and probably more important reason is that the ironic remark opens a direct window on the speaker's expectation about the world, while leaving the veridical situation implicit. However, as the veridical situation is already directly accessible to the discourse participants, more is gained pragmatically by directly expressing the implicit expectation.

The above reasoning holds in all cases where the literal meaning of the ironic phrase is in direct opposition to the contextual information; in such cases, the utterance constitutes third order information which conflicts with our expectations of the situation. If it is not downgraded, it cannot be informative and must necessarily serve the discourse participants' sharing of common attitudes; alternatively, it marks a discord between the speaker and the referent of the utterance in the case where it is directed toward a discourse participant.

However, as mentioned previously, not all irony consists in a negation of contextual facts. Compare Haverkate's (1990) example *I love people with good manners*, said in a context where a misbehaving person has caught the attention of the discourse participants.

Again, the ironic meaning comes from the fact that the sentence cannot be informative if it stands alone – in which case the flow of informativity would require an upgrading, making it clear to the listener why the speaker is saying this (for instance, *I love people with good manners, but in this case the rude behavior is justified*). So, regardless of whether the speaker is presenting trivial information or third order information, the ironic meaning comes from its lack of upgrading or downgrading in the information flow.

The above is not to say that irony cannot appear in a discourse flow, but this will typically be a local exchange where irony is answered with irony, as in the following example from Gibbs:

Anne: By the way, were our wonderful guests still here when you came out and ate lunch?

Dana: I had a sandwich and...

Anne: Isn't it so nice to have guests here?

Dana: Totally!

Anne: I just love it, you know, our housemates. They bring in the most wonderful guests in the world and they can totally relate to us.

Dana: Yes, they do.

Anne: (laughs) Like I would just love to have them here more often (laughs) so I can cook for them, I can prepare (laughs)

Dana: to make them feel welcome?

Anne: Yeah, isn't this great, Dana? Like today I was feeling all depressed and I came out and I saw the guests and they totally lightened up my mood. I was like the happiest person on earth.

Dana: Uh huh.

Anne: I just welcome them so much, you know, ask them if they want anything to drink or eat (laughs). (Gibbs 2000: 6)

In this example we see that Anne is making an ironic statement and then maintains it throughout the discourse. Dana is not contributing to the irony, but is confirming Anne's statements.

Two things are worth mentioning concerning this exchange:

1) Anne refers to the guests with the ironic words *wonderful guests*, and then she elaborates on this scenario: how the guests delighted her, how she would like to prepare the food etc. But all of this is imaginative and ironic, expressing her attitude toward the guests, and it does not provide Dana with any new information. This is a general characteristic of irony; it expresses an attitude toward what should generally be recognized in the situation, and as long as the ironic mode is maintained, no new information is given.

2) Then the question arises again, why not express the attitude to the guests by referring to them as awful? This is probably due to the dynamics of the discourse. The ironic version is much more flexible and less committed than is the direct mentioning of the real attitude. Firstly, if Dana had not responded positively, it would have been easier to drop out of the ironic mode without being committed to an opinion the discourse participant does not have. Secondly, imagine Anne had started by saying, *By the way, were our terrible guests still here when you came out and ate lunch?* Then the following

elaboration, preparing the food etc., would not have been possible. Positive statements are more open for possibilities; if the weather is wonderful, one can ride a bike, take a walk, go on a picnic etc. These dynamic possibilities are probably one of the reasons for using irony and one of the reasons that the majority of ironic statements have a positive literal meaning.

### 3. *Irony as a violation of Grice's maxims and the processing of irony*

The Gricean approach (Grice 1975, 1978) is similar to the one discussed in the previous section. The only difference is that in the classical approach, an ironic statement semantically means the opposite of what it literally says. For Grice, this is not necessarily the case; instead, the opposite meaning is inferred pragmatically. This is because an ironic utterance violates the maxim of quality. The process leading the listener to the right interpretation is as follows: 1) The literal meaning is processed. 2) The listener understands that the literal meaning does not apply to the situation. 3) The listener infers that the intended meaning is the opposite of the literal one. Of course, we find the same problems here as in the classical approach: many ironic statements do not involve the opposite meaning; also, this theory does not provide a plausible explanation of why people use irony instead of literal statements.

However, Grice's theory is interesting in this connection since it refers to the above-mentioned reasoning process. The stages in the Gricean reasoning are philosophically motivated, but do they have anything to do with the psychological processes at work when people hear and understand irony? Several studies have addressed this question. There are two main results that seem to contradict each other. One is a confirmation of the 'direct access hypothesis'. This hypothesis states that the context in which an ironic statement is

made can suppress the literal meaning of the words used, so that the ironic meaning is immediately accessed. In contrast to this, we have the 'graded salience hypothesis', which states that irrespective of the context, the salient meaning of the words will always somehow interfere with the contextually determined meaning.

#### 3.1. The graded salience hypothesis

In support of this hypothesis, behavioral experiments were performed measuring the reading times of single words; in addition, response times were measured for lexical decision tasks. Here, words referring to the ironic as well as the literal meaning were presented to the participants having to decide whether the words in question were words or non-words. In Giora et al. (2007), it is shown that in reading a target sentence like *It is terrific news*, where *terrific* can either be ironic or literal, it took the participants longer to read the final word in the ironic context. Moreover, in the lexical decision task – done after the reading of the text – responses to literally related probes were faster than to ironically related probes. All of this suggests that it is faster to derive the literal interpretation of a statement than the ironic one. Note that this does not in any way confirm the standard pragmatic account of how the listener gets to the contextually appropriate meaning of the statement, as described above. In a dynamic model of language processing, the established meaning is always a competition between different possible outcomes, also in so-called literal language. Normally, the context will be a determining factor in establishing a stable meaning, but in the case of irony, there will be a conflict between the salient, but contextually inappropriate meaning of the statement and the contextually appropriate meaning. We can view the two forms of meaning as coexisting in a dynamic process, which will eventually stabilize in the contextually determined meaning (that is, if the listener or reader understands the irony).

In Colston and Gibbs (2002), subjects are presented with stories of the following type:

You are a teacher at an elementary school.

You are gathering teaching supplies with your assistant teacher.

Some of the scissors you have are in really bad shape.

You find one pair that won't cut anything.

You say to your assistant, "This one's really sharp". (Colston and Gibbs 2002: 58)

This is compared to a story of the same length about a clever student ending with the same target sentence, *This one's really sharp*. In the former case, the sentence is ironic while in the latter, it is metaphoric. In the experiment, the participants were reading the stories on a computer screen, one line at a time. The participants were instructed to press a particular key when they had read and understood the line. In this way, the reading time could be measured. One of the conclusions in the article is that it takes a longer time to process the ironic target sentence. The explanation given is that metaphoric expressions are understood descriptively, because they are about states in the world, whereas ironic expressions are meta-representations, since they are representations of 'pretend thoughts', thoughts about thoughts.

Using a metaphor provides information about states of affairs in the world. In the example above, we were dealing with a coded metaphor, since being intelligent is one of the meanings of *sharp* listed in the dictionary. The context will therefore ensure that this particular sense of *sharp*, which is already one of the attractors, will

be stabilized in the mind of the reader. The ironic version is not coded; we therefore get a conflict between the salient, but contextually inappropriate meaning(s) of *sharp* and the context (as in Giora et al. 2007). The coded lexical meanings will constitute a resistance to the contextually appropriate meaning, which therefore will stabilize more slowly. This is logical: the reader has already established a mental model where the scissors are blunt; the ironic version is a simulation of a counterfactual, desirable world where the scissors are sharp; the reader is therefore confronted with two contrasting mental spaces and has to organize them so that the last model expresses an attitude toward the first one. If the protagonist had said, *This one is really blunt*, this would have fit into the already established mental model: no further processing is necessary.

### 3.2. The direct access hypothesis

Other experiments support the direct access hypothesis. In Gibbs (1986), it is demonstrated that it did not take the participants longer to process an ironic expression such as *He is a fine friend* (meaning he is a bad friend) than it did for the same expression in a literal context, or the non-ironic equivalent *He is a bad friend*. This suggests that the logical steps in the Gricean implicature do not necessarily correspond to the psychological processing of ironic meaning.

Gibbs explains the finding that ironic statements sometimes are processed as fast as or even faster than are literal meanings by referring to the context constituting an ironic situation, i.e. a situation where there is a contrast between what is expected and the frustrating reality. For instance, the ironic statement, *This sure is an exciting life*, should be easier to process if it is clear from the context that the speaker had a previous expectation of an exciting life which had been frustrated, than if the speaker had no such prior expectations.

This makes sense: if the protagonist's expressed expectation of an exciting life has been foregrounded at some previous point in the reading process, then it will be easier for the reader to access it as a non-veridical statement and thereby understand it as irony. Although the irony is not coded directly here, we have a cognitive model of irony that implies a frustrated expectation; hence, the sheer mention of this expectation will prompt an ironic reading.

It was mentioned above that Giora et al. (2007) provided experimental evidence for the graded salience hypothesis, such that the salient (literal) meaning will be accessed despite contextual information that makes the literal meaning inappropriate. However, the experiments also provided evidence for the claim that an expectation of irony did not facilitate the reading of an ironic statement, where expectation was manipulated by embedding an ironic speaker *in vivo* in the experiment. This study was a response to the study by Ivanko and Pexman (2003) mentioned in the introduction, who aim to support the direct access principle by presenting a series of complicated results. The material used in their study was of the following form:

Sam agreed to pick Christopher up after school. *Sam never arrived to pick up Christopher and never called to say why/Sam arrived an hour late and apologized/Sam and Christopher talked about the dance on Friday.* The next day Christopher is explaining to Jodi what happened. Christopher says:

Ironic statement: Sam is a nice friend

Literal statement: Sam is a rotten friend

Wrap up sentence: Christopher and Jodi were walking home from school. (Ivanko and Pexman 2003: 276)

This story was presented on a screen, word for word, so that the reading time could be measured. The three sentences in italics represent three different contexts: a strongly negative context, a slightly negative context, and a neutral context. For each of the contexts, the participant was presented with either the ironic statement or the literal statement, and in all cases the story was followed by the wrap-up sentence.

The idea is to show that the experience of irony is dependent on a discrepancy between statement and context. The discrepancy can be manipulated by varying the statements, but in this particular experiment, the ironic statement (*Sam is a nice friend*) is fixed while the context is varied. This variation gives interesting results when it comes to reading time. In the strongly negative context, the reading time of the fifth word in the statement is significantly longer in the case of the ironic statement than in the literal one; also the reading times of the sixth, seventh, and eighth words in the wrap-up sentence are longer in the strongly negative context. For the slightly negative context, the inverse seems to be true, as the literal statement takes a longer time to read than the ironic statement. The sixth, seventh and eighth words in the wrap-up sentence take about the same time. These results were tested against other experiments on the same material, which for instance showed that in the strongly negative context, statements rated as more sarcastic showed a slower reading time, whereas in the slightly negative context, the inverse was true: the more sarcastic rated statements being those with a faster reading time. Finally, the participants rated the ironic statements in the strongly negative context as more mocking than the literal statements; in the slightly negative context, the inverse was true, the ironic statements being perceived as more polite than were the literal statements. These results add up to the following, likely prediction: In a strongly negative context, an ironic statement will be perceived as more negative than a literal statement, but in a

slightly negative context, it will be perceived as more positive than the literal statement.

When we read a text, each statement opens a set of possibilities for its continuation; the more predictable a continuation, the faster it will be read. From this we can conclude that in strongly negative contexts, the readers will predict a literal statement (something which also transpires from the participants' reports, as presented in Ivanko Pexman's article). Apparently, we only use irony in mildly negative contexts; if things get really serious, we tend to be literal.

Following the logic mentioned above, it may come as a surprise that in the slightly negative context, readers would predict an ironic statement. However, negative statements take a longer time to process than positive statements; also in lexical decision tasks, responses to negative words like *useless* were found to take a significantly longer time than did positive words like *useful*. One reason for this could be that our default expectations of the world are positive; any negative statement represents a deviation from this and will therefore require more processing time. In Ivanko and Pexman (2003), the literal statement was always negative, as in the example above *Sam is a rotten friend*; encountering such a strong negative predicate as *rotten* in a situation that is only slightly negative is perhaps not what one would expect. This could be part of the explanation of why the literal statement took a longer time to read in this context.

That the last words in the wrap-up sentence take a longer time to read in the ironic statement than in the strongly negative context does not seem strange. The situation in the strongly negative context is rather dramatic, and one would expect that the following sentence somehow elaborates on this; however, such is not the case, but this becomes clear only at the end of the sentence.

All of Ivanko and Pexman's (2003) results can thus be explained by the information theoretical assumption that in processing language, the brain is always occupied with predicting possible continuations of the discourse, written or verbal. The processing time will then depend on the extent to which the continuation fits into a predictable scenario. For the authors, these results disconfirm the graded salience hypothesis – that salient lexical meaning cannot be suppressed – and even support the hypothesis that ironic meaning can be accessed directly.

We thus have very contradictory results concerning the discussion: graded salience versus direct access. This should not pose a problem, however. Since irony is not a unitary phenomenon, we should not expect a unitary irony processing in the brain. Rather, every situation is different, and the various studies use very different material. So why should we expect processing in one case to be replicated in exactly the same way in another case? For instance, *This is really sharp* (said ironically about a pair of scissors) invokes a very different situation from *Sam is a nice friend* (said ironically about some other person); hence, our access to the ironic meaning might be very different in the two cases. Still, the fact that there are strong cases in favor of some kind of processing of the salient meaning, should not come as a surprise either, as the salient lexical meaning is necessary for the ironic meaning to be grasped. It is not as if we processed the literal meaning, discarded it, and concluded that the speaker intended to convey the opposite meaning. In reality, the speaker does not intend to convey the opposite meaning; he or she intends to be ironic. In accordance with this the hearer – via the salient meaning – understands the ironic intention; depending on the situation, this process may be slower or faster.

#### 4. Irony as echoic mention

The theory of irony as echoic mention was first proposed by Sperber and Wilson in 1981 and then further developed in many of their later articles. Their approach is cognitive rather than semantic, as we can see from the following quotation:

Because ... it seems to be a mistake to take IRONY itself as the object of investigation, and to limit one's attention to its more standard cases. There is a whole range of utterance-types that can be more or less loosely called irony. (Sperber and Wilson 1981: 298)

This means that even if we have a category of utterances called ironic, we cannot find a semantic definition of this category in terms of necessary and sufficient conditions. This is Grice's main criticism; he sees the violation of truthfulness as both a necessary and a sufficient condition for ironic interpretation. But as we have seen earlier, in Section 2, truthfulness is not a necessary, and clearly not a sufficient condition, since not every false statement is a case of irony. What is missing, according to Sperber and Wilson, is that an ironic statement is a case of mention, meaning that the ironic statement is one that has been (or could have been) uttered by someone previous to the time of speech. Consider the trivial statement *Wonderful weather!*, said during a picnic while the rain is pouring down. Let us suppose that this has been uttered previously, maybe as an argument for going on a picnic. Then the utterance would be a standard example of an echoic mention. However, even if it has not been said by anyone, it is still the kind of cultural expectation we have when going on a picnic, and therefore it is still an echo: namely of a norm.

A more sophisticated example could be of someone saying, *Have a piece of mind*, while offering some food at a lunch table to a

philosopher who has just been talking about the identity of mind and matter. This is indirectly an echo of the philosopher's thesis; of course it is meant ironically. In many cases, however, it is not immediately clear that an ironic statement should be an echo of a previous statement or of an unspoken norm or thought. For instance, take the example from *Pride and Prejudice* – mentioned in Wilson and Sperber (1992) – where Elizabeth Bennet remarks that she began to appreciate Fitzwilliam Darcy, when she first set eyes on his magnificent estate of Pemberley. What is this an echo of?

In any case, echoic mention seems to be a kind of prototypical irony (if we assume that the category of ironic remarks has a graded structure with prototypical members and fuzzy borders), and repeating an utterance seems to attract ironic perception. This is exploited by Shakespeare in *Julius Caesar* (Act III, Scene ii), when Mark Antony makes the statement, *But Brutus says he was ambitious; And Brutus is an honourable man*. The statement is repeated six times altogether. In the play, Mark Antony is allowed to give a speech at Caesar's funeral on the condition that he does not say anything negative about Brutus. What he does instead is to say something positive, but repeats it. The first time he says *Brutus is an honourable man*, it is not perceived as ironic. But echoing this statement again and again changes the perception, and in the end it is deeply ironic. But why does repetition have this effect? To repeat a statement that has been uttered previously in the discourse constitutes trivial information; it can therefore not be part of the information flow, and this changes its status into a non-seriously meant contribution. So when Antony repeats *Brutus is an honourable man*, the listeners can no longer believe him to be seriously committed to the veridical content of the statement. Echoic mention is therefore strongly connected to insincerity, which is part of the last theory of irony we will briefly consider.

### 5. Irony as 'allusional pretense'

The theory of irony as allusional pretense is presented in Kumon-Nakamura et al. (1995) and is intended to provide a more precise description of irony than the echoic mention theory does. It consists of two claims: 1) In an ironic statement, there is an allusion to an expectation or a norm that has been violated. 2) The statement is pragmatically insincere. A critical feature is here pragmatic insincerity. This is also the case in Grice's definition, where irony is a violation of the maxim of truthfulness, but Kumon-Nakamura et al. (1995) extend it to other speech acts such as offerings, questions, expressives, etc. For instance, *Have another piece of cake* said to someone who has just gobbled up most of the cake is clearly a violation of a social norm and an insincere offer. As to questions, we can imagine the case of a professor, who unexpectedly displays a lack of knowledge, which thus gives rise to the insincere question, *You are a professor, are you not?* These examples have been used in the literature to refute the idea that irony only concerns the maxim of truthfulness, as mentioned above, but the new thing in Kumon-Nakamura et al. (1995) is that a defining feature of these different ironic speech acts is their **insincerity**.

It could well be the case that insincerity is the primary defining feature of irony. Any ironic remark is clearly an insincere speech act, but is any insincere statement where the speaker intends the statement to be understood as insincere, also ironic? In a very general understanding of irony we can postulate that this is indeed the case. Consider the following example. X is driving in a taxi with two famous chess players, Mikhail Tal and Bobby Fischer. X is a chess player himself, but not very well-known. The taxi driver drives very hazardously and X says, *If we die now, one can read in the newspaper tomorrow that X died in a car accident together with two other passengers*. This is a case of self-irony, but it is not clear what expectation has been

violated here; it is evident, though, that the statement is meant to be insincere. The misfit is between the statement and the situation, so one could say that **irony is an insincere statement which is meant to be understood as insincere and which constitutes a misfit with the situation**. However, it is an open question whether an insincere statement, that is intended by the speaker to be perceived as insincere, by definition constitutes a misfit.

### 6. Summary

There are three main questions related to our understanding of irony: its definition, its function, and its processing. A major part of the classical literature deals with the definitional part. Here, we have presented four definitions: the classical view according to which an ironic statement is a figure of speech that means the opposite of what it says; Grice's pragmatic theory, where irony is a violation of the maxim of quality (in particular) truthfulness; Sperber and Wilson's echoic mention theory; and finally the allusional pretense theory of discourse irony propounded by Kumon-Nakamura et al. Given the variety of ironic statements, we have argued that it is not possible to give a semantic definition in terms of necessary and sufficient conditions. Any definition of irony has to be pragmatic, and we have suggested that insincerity, as proposed by Kumon-Nakamura et al. (1995), is indeed a critical feature; hence our preferred definition: **An ironic statement is an insincere speech act which is intended by the speaker to be understood as insincere and which constitutes a misfit with some aspect of the context**.

The second question related to irony concerns its purpose. Here, too, it is not possible to give a unitary definition because of the variety of situations in which irony is being used. One of its obvious function is to be humorous; other possible functions are to be

derogative, to alleviate social discomfort, to provoke a reaction, etc. To these pragmatic purposes, one could add an information theoretical advantage contained in irony. The use of irony happens in a situation where there is strong contextual information, whereas the ironic remark refers to a counterfactual situation. In this way, the ironic remark represents a different window on reality than the one presented by the context; in this way, more information is provided by the ironic remark than by a literal reference to the context. Of course, this 'more information' has to be understood pragmatically as being insincere, as stipulated above.

The final question concerns the processing of irony. Also here we have to acknowledge that, since irony can appear in so many forms and in so many different contexts, there is no unitary way of processing it. According to the direct access theory, contextual information can override the literal meaning of the ironic statement; as a consequence, ironic statements can be processed as fast as (or maybe even faster than) literal statements can. Counter to this, we have the graded salience hypothesis – or 'literal meaning first' hypothesis – which states that even in a highly ironic situation, the salient meaning of the words cannot be suppressed. This will presumably lead to a longer processing time for ironic statements than for literal ones. In accordance with the assumption that there is no unitary processing related to ironic meaning, the different experimental results that have been reviewed here do not allow us to unambiguously decide which of the two theories is the right one. Some results suggest that ironic statements are understood as fast as literal statements, while some suggest the opposite. However, there seems to be a slight tendency in favor of the graded salience hypothesis. This is also what we would expect from a logical point of view, because the ironic meaning can only be understood relative to the literal meaning. This does not entail that the literal meaning is dealt with fully before the ironic meaning is processed. Rather, we can understand irony as the

result of a conflict between the salient meanings and the contextual meanings – a conflict in which where the two are co-present. The time of processing will then depend on the relative strength of the two sources of meaning: if the contextual meaning is weak, it will be difficult to 'get' the irony, but if the contextual meaning is strong, the irony might be more or less immediately accessible.

*Institut for Æstetik og Kommunikation  
Aarhus Universitet  
Jens Chr. Skous Vej 2  
DK-8000 Aarhus C  
semsvend@hum.au.dk*

## References

- Beaugrande, Robert-Alain de & Wolfgang Dressler. 1981. Introduction to Text Linguistics. London: Longman.
- Brown, Penelope & Stephen Levinson. 1978. Universals in language usage: Politeness phenomena. In E.N. Goody (ed.) Questions and politeness: Strategies in social interaction. Cambridge: Cambridge University Press. 56-290.
- Clark, Herbert & Richard Gerrig. 1984. On the pretense theory of irony. *Journal of Experimental Psychology: General* 113. 121-126.
- Colston, Herbert L. 1997. Salting a wound or sugaring a pill: The pragmatic functions of ironic criticisms. *Discourse Processes* 23. 24-45.
- Colston, Herbert L. & Raymond Gibbs. 2002. Are irony and metaphor understood differently? *Metaphor and Symbol* 17. 57-80.
- Dews, Shelley, Joan Kaplan & Ellen Winner. 1995. Why not say it directly? The social functions of irony. *Discourse Processes* 19. 347-367.

- Gibbs, Raymond. 1986. On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General* 115. 3-15.
- Gibbs, Raymond W. & Jennifer O'Brien. 1991. Psychological aspects of irony understanding. *Journal of Pragmatics* 16. 523-530.
- Gibbs, Raymond W. 2000. Irony in talk among friends. *Metaphor and Symbol* 15. 5-27.
- Giora, Rachel et al. 2007. Expecting irony: Context versus salience-based effects. *Metaphor and Symbol* 22. 119-146.
- Grice, Paul H. 1975. Logic and conversation. In P. Cole & J. Morgan (eds.) *Syntax and Semantics, Vol. 3: Speech Acts*. New York: Academic Press. 41-58.
- Grice, Paul H. 1978. Further notes on logic and conversation. In P. Cole (ed.) *Syntax and semantics, Vol. 9: Pragmatics*. New York: Academic Press. 113-127.
- Haverkate, Henk. 1990. A Speech act analysis of irony. *Journal of Pragmatics* 14. 77-109.
- Ivanko, Stacey L. & Penny M. Pexman. 2003. Context incongruity and irony processing. *Discourse Processes* 35. 241-279.
- Kumon-Nakamura, Sachi, Sam Glucksberg & Mary Brown. 1995. How about another piece of pie: The allusional pretense theory of discourse irony. *Journal of Experimental Psychology: General* 124. 3-21.
- Lakoff, George & Mark Johnson. 1980. *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, George. 1987. *Women, fire, and dangerous things*. London, Chicago: The University of Chicago Press.
- Lakoff, George & Mark Johnson. 1999. *Philosophy in the flesh*. New York: Basic Books.
- Searle, John R. 1979. Metaphor. In A. Ortony (ed.) *Metaphor and Thought*. New York: Cambridge University Press. 92-123.
- Sperber, Dan & Deirdre Wilson. 1981. Irony and the use-mention distinction. In P. Cole (ed.) *Radical Pragmatics*. New York: Academic Press. 295-318.
- Tannen, Deborah. 2005. *Conversational style: Analyzing talk among friends*. New York: Oxford University Press.
- Wilson, Deirdre & Dan Sperber. 1992. On verbal irony. *Lingua* 87. 53-76.