



# WHAT GOES ON IN THE HEAD IS NOT MAGIC

AN INTERVIEW WITH DEIRDRE WILSON

by Marianne Vahl and Cathrine Felix.

Deirdre Wilson is professor of Linguistics at University College London. She has won great renown for her groundbreaking work with Dan Sperber in the field of pragmatics. According to Wilson and Sperber, in order for pragmatics to explain linguistic behaviour it needs only one principle - the Communicative Principle of Relevance. In essence, this holds that in any given context, what people communicate is expected to be relevant. By claiming that the search for relevance is a central aspect of human cognition, relevance theory argues that the hearer, when hearing an utterance, automatically looks for an interpretation that satisfies his expectations of relevance. The comprehension process can be seen as a form of reasoning by which the hearer constructs and evaluates hypotheses about explicatures and implicatures – the implicit and explicit parts of the speaker’s meaning – and stops when he finds an interpretation that is relevant in the expected way. We were so lucky to have a talk with her when she was in Oslo to give a speech at the CSMN/ Shared Content workshop on implicatures.

*You argue that the rules Grice thought that people follow in order to communicate efficiently: quantity, quality, relation and manner, can be reduced to one principle only: The Communicative Principle of Relevance. What do you mean by relevance?*

Although I disagree with Grice about the details of the pragmatic principles, the work that I have done with Dan Sperber over the years could not have been done without Grice. What’s important about Grice was his emphasis on the fact that un-

derstanding is a matter of recognizing speakers’ intentions, recognizing a certain sort of overtly expressed intention. His cooperative principle and maxims were his hypothesis about how a rational hearer and speaker might converge on the same intention. What Dan Sperber and I have been trying to do is to find an empirically plausible account of the principles and mechanisms that hearers do exploit when identifying the speaker’s intentions. Grice said he didn’t really know what relevance was. We thought it would be good to try to produce a theoretically acceptable notion of relevance. The main idea is that relevance isn’t something that only matters in conversation (as Grice suggested). If you hear a loud noise, you pay attention, because it could be relevant to you. If you see smoke coming through the door, it could be relevant to you. So the basic notion of relevance has to be defined for cognition in general, and not just for communication. In a nutshell, we claim that information is relevant to you if it combines with your background knowledge to yield what we call a positive cognitive effect – one that makes a worthwhile difference to your representation of the world. The greater the cognitive effect, and the more easily it is achieved, the greater the relevance.

*It is controversial how to define pragmatics. Levinson, for instance, defined it as being «the study of those relations between language and context that are grammaticalized, or encoded in the structure of language». He has been criticised for not telling how we can connect user and grammar, but on the other hand Levinson has criticized the view that explains*

*pragmatics as whatever happens whenever users «do things» with words as being a very broad notion of pragmatics. Can you tell us whereabouts you place yourself in this landscape?*

Levinson in his 1983 text book considered several definitions of pragmatics and found them all unsatisfactory. The one he didn't consider was the view that pragmatics aims to provide a cognitively plausible and empirically testable account of how hearers understand utterances, and in particular how they bridge the gap between the meaning of a sentence and the meaning the speaker intends to convey by uttering that sentence on a particular occasion. Linguists agree that the grammar will tell you what the meaning of the sentence is. However, when the sentence is uttered, it can convey many different meanings on both the explicit and the implicit level. Hearers bridge that gap between sentence-meaning and speaker's meaning extremely easily, and we need an explanatory theory of how that is done. So I don't like either of the two definitions that you quoted. The idea that pragmatics should look only at the grammaticalization of relations between sentence and context is much too narrow, but the idea that pragmatics is a theory of everything that language users do is much too broad. As Chomsky has often pointed out, you can't have an explanatory theory of everything that language users do: there are too many different principles or mechanisms involved. The basic idea of relevance theory is that when someone utters a sentence in a particular situation, there is a mental grammar that will decode the sentence into its rather fragmentary and incomplete linguistic meaning, which can then be enriched in many different ways in different contexts, with each enrichment yielding a different cognitive effect, and hence being relevant in a different way. The hearer's goal is to enrich the fragmentary sentence meaning in the way intended by the speaker. This is something that we are all capable of doing very fast, and I see it as the job of pragmatics to explain how this is done.

*So you think that if you take the Austin approach you might go out too broadly?*

Although Austin raised many important issues about linguistic communication, I think his concentration on speech acts turned out not to be very productive in the long term. One of the reasons is that Austin was a very taxonomic person; he belie-



ILLUSTRATION: CATHERINE FELIX

ved in producing long lists of possible speech acts and their associated felicity conditions. Often, the people who were attracted to his view turned out to be taxonomists too: they were interested in classifying speech acts and setting out the conditions on performing them, but lost sight of the general principles that might be seen as underlying all acts of communication. Austin was also very interested in the institutional aspects of speech, and looked for explanations at the social rather than the psychological level. But times have changed, psychology has moved on – partly under the influence of Chomsky, who is interested in the mental representations involved in knowledge of language. So a cognitively plausible approach to pragmatics would look for cognitive mechanisms that take as input a sentence meaning plus a wide array of background knowledge, and yield as output a hypothesis about the speaker's meaning. This is very different from what Austin was thinking about. That doesn't mean that this is the only thing to do, but it is one thing that needs to be done, and it seems reasonable to call it pragmatics.

*You work in the field both of linguistics and pragmatics. But considering linguistics as a field it is much more static than*

*pragmatics, pragmatics is dynamical, how can these very different disciplines be united?*

Linguistics looks at sentences in abstraction from speakers and contexts, and pragmatics doesn't: it looks at utterances, which are tied to a particular speaker, place and time. I started out as a philosopher – I was trained in Oxford by ordinary-language philosophers – and I was quite disappointed by the behaviourist view of mind that was influential in philosophy at that time. It was Grice talking about the role of intentions in communication that really aroused my interest in philosophy, and it was Chomsky talking about the role of the mind in language acquisition that made me think linguistics is worth studying, and so I went to MIT and did a PhD with him. Then I came back to England and thought: Where do I belong? And the answer was: not quite anywhere. As a specialist in pragmatics, I might be seen as on the edge of the linguistics department, and not close enough to being a philosopher, and not close enough to being a psychologist. But all those subjects interest me, and I think pragmatics is very much an interface subject where you have to know something about philosophy, psychology and linguistics in order to do serious work.

*Pragmatics is often called «the wastebasket of linguistics,» what do you think about that saying?*

I think pragmatics has sometimes been used as the wastebasket of linguistics, mainly by semanticists. Formal semanticists come up against a problem that they can't handle, and they hope that pragmatics will maybe provide an answer, and so it becomes the wastebasket for all the unsolved problems of semantics. But of course there is another way to look at pragmatics, and that is not as the wastebasket for semantics, but as the theory that aims to explain both verbal and nonverbal communication, in the way I just described. For a start, it aims to explain how children can start communicating with their parents before they have language, how children can read the intentions of speakers when acquiring the vocabulary of their language, and how adult hearers use contextual information to enrich fragmentary sentence meanings in the intended way. In other words, pragmatics is a proper theory in its own right, and if we succeed in building a cognitively plausible theory of pragmatics, it should also have implications for semantics. For instance,

recent work in pragmatics has very much changed our view of lexical meaning. Grice started out by saying that semantics doesn't need to account for everything a speaker conveys by using a word, that general pragmatic principles can relieve some of the burden of semantics. The more you look into what pragmatics can do in this area – in dealing with metaphor, approximation, lexical narrowing, neologisms and so on – you find that lexical meaning can be even simpler and more schematic than Grice thought. That is, there's a huge amount of pragmatic enrichment that takes place at the explicit level when you hear a word uttered in a particular situation. You don't necessarily start by taking it literally, as Grice thought: you might broaden its denotation or narrow its denotation, taking it to convey a more general or a more specific sense than the linguistically encoded one. People do this absolutely automatically, and that then brings into question the standard philosophical view that the norm in language use is to speak literally except for decoration or vividness. Speakers aren't necessarily expected to speak literally: they use the language as a clue to the thoughts that they are trying to convey. Very often the relation between language and meaning, even explicit meaning – what's asserted – is much less direct than is widely assumed by linguists, philosophers and psychologists.

*You say that the human cognitive system automatically allocates its attention to information which is relevant. Can you explain what you mean by cognitive system?*

I take the cognitive system to include at least the mechanisms for perception, memory and inference. We all know that when we hear a loud noise we automatically pay attention to it; when someone moves towards us suddenly, we notice. In relevance theory, we would say that there are automatic perceptual mechanisms or heuristics for picking out potentially relevant inputs to perception and assigning them a conceptual representation that maximises their relevance. Similarly, there are automatic inferential mechanisms for deriving conclusions that are potentially relevant, and automatic memory retrieval mechanisms for retrieving contextual information that contributes to relevance. According to relevance theory, as a result of constant selection pressures, the cognitive system as a whole has tended to evolve in the direction of increasing cognitive efficiency. In our terms, this means that it tends to operate in a way

that maximises relevance, picking out the most relevant perceptual and conceptual inputs, retrieving the most relevant background knowledge and deriving the most worthwhile conclusions. This is what underpins the relevance-theoretic approach to pragmatics.

*So it doesn't imply that knowledge is an aspect of this cognitive system?*

The cognitive system is geared to picking out the most relevant information, and processing it as relevantly as possible, which means yielding the greatest cognitive effect for the least effort. An efficient cognitive system is not going to waste its time allocating attention to information that has no chances of making a worthwhile difference to the individual's representation of the world – and in general, that means information that is at least evidenced, hence likely to be true. We claim that the guiding principle in both cognition and communication is the search for relevance. If someone speaks to you, offers you information, they are suggesting that the information they are trying to convey will be relevant enough to be worth your attention. Suppose someone attracts your attention by clapping their hands, and then produces an utterance. The utterance has many different possible interpretations, each of which will be relevant in a different way. How should you understand the utterance? The answer suggested by relevance theory is that you will automatically look for an interpretation that satisfies your expectations of relevance, and assume that this is what the speaker was intending to convey. Of course you can't know with certainty what was in the speaker's mind, but this is the interpretation a rational hearer should choose.

*You and Dan Sperber have claimed that you cannot violate The Communicative Principle of Relevance, even if you want to. Do you still say that?*

Yes, look at how the communicative principle of relevance is formulated. It is not an instruction to speakers. Grice's maxim of relevance says: be relevant! That is an instruction to speakers that they can ignore or violate, and the idea that speakers can violate the maxims plays an important role in Grice's theory. The Communicative Principle of Relevance says: every utterance addressed to someone creates an expectation of relevance. This is not an instruction to speakers, and it is not the sort

of thing that speakers can violate. We claim that, whether they want to or not, speakers automatically raise that expectation in their addressees. The Communicative Principle of Relevance is like a scientific law of communication: when you attract someone's attention, you will automatically create an expectation of relevance. Of course, the hearer may immediately see after the first two words that you are not going to be relevant enough, but as long as you go on making them think they are going to get enough relevance, they will go on paying attention.

*So a rational speaker is someone following The Communicative Principle of Relevance?*

A rational speaker should aim to satisfy – or at least seem to satisfy – the hearer's expectation of relevance if she wants to hold the hearer's attention. There will be utterances that do not satisfy this expectation, because the speaker, either accidentally or intentionally, does not succeed in being relevant enough, or is deceiving the hearer and only pretending to be relevant. What the speaker is guaranteeing to the hearer (whether honestly or deceptively) is that she is not requesting his attention and effort for no reason at all. Given this guarantee, when the hearer comes across a part of an utterance that seems to cost more effort than expected, he has to ask himself: was this an accident, or was it intentional? If it is intentional, e.g. if the piece is written by a good stylist or a good poet, the communicator must have been intending to get across more than could have been got across with a simpler utterance. That's how indirect answers to questions work. When a speaker answers a question indirectly, instead of saying directly «yes» or «no» – which would have cost the hearer less effort – she is typically not only answering the question but giving a reason for her answer. For example, if you ask «do you want some coffee?» and the answer is «coffee will keep me awake,» in many situations you are entitled to infer that the speaker is refusing coffee because coffee will keep him awake. In other situations, the same answer could, of course, convey that the speaker is accepting the offer of coffee. How do you decide which answer was implied, or implicated? This is the sort of question that a cognitively plausible account of pragmatics should aim to answer. Relevance theory offers a quite neat answer by looking at the amount of effort the different interpretations would take to construct,



ILLUSTRASJON: SVEIN JOHAN REISANG

and the different implications they would yield, in different situations.

*You are part of the core group at CSMN (Centre for The Study of Mind in Nature). What are you planning and hoping to do at CSMN?*

CSMN is a wonderful opportunity to bring together philosophers and linguists and scientists working on moral, rational and linguistic norms. I am a part of the linguistic agency project; we have three subprojects, and the one that I am mostly concerned with is on metarepresentation. A metarepresentation is a higher-order representation with a lower-order representation embedded inside it. Attributions of utterances or thoughts are the most typical cases of metarepresentation. For instance, “John said that the earth is round” is a metarepresentation of an utterance, and “Sue thinks that the earth is round” is a metarepresentation of a thought. Pragmatic interpretation involves a huge amount of metarepresentation. The pragmatic interpretation process itself takes a metarepresentation of an utterance as input (e.g. “The speaker said P”), and gives as output a metarepresentation of the speaker’s meaning (e.g. “The speaker meant Q”). So both the inputs to and the outputs of the pragmatic interpretation process are metarepresentations. One of the things that relevance theory has been used for is to look at the development of pragmatic metarepresentational abilities in children and the breakdown of these abilities, in particular in people with autism. There are certain things that are typically said about the comprehension abilities of people with autism: for instance, that they are very literal, they can’t deal with metaphor or irony, they don’t do approximation, they are not very good at Gricean implicatures – in other words, they are a sort of literalists. But if you look more closely, you can see that people with autism are quite capable of some pragmatic inferences: they can do certain types of disambiguation and reference assignment, for example. Similarly, children up to the age of three are said to have limited pragmatic abilities and to be literalists in some sense, yet they are heavily involved in non-verbal communication and also capable of drawing some pragmatic inferences, even though they make characteristic mistakes. One of the standard examples you find in the literature is that children under the age of three don’t always disambiguate – choose the intended meaning of an ambiguous word – in the way that adults would. For instance, the English word “hair/

hare” is ambiguous (when pronounced) between the hair on one’s head and the hare – an animal like a large rabbit. When adults are asked to draw a “hair/hare” running across a field, they typically draw a large rabbit. When a child of three, who has shown that she knows both meanings, is given the same instruction, she typically draws a hair on the head with added legs. Why is that? This “hair on the head” sense is the most frequently used, and therefore the first to come to mind. It seems that children up to the age of three have a tendency not to reflect on whether the first remotely plausible interpretation that occurs to them could reasonably have been intended by the speaker. If they can make some sense of it, as in a fairy tale, for instance, they will go ahead and assume that this was the intended meaning. This is also what autistic people have a tendency to do, and that’s why they are sometimes seen as over-literal or insufficiently attentive to contextual information. One of the things we want to do in the metarepresentation project is look at how the philosophy, the psychology and the pragmatics of metarepresentations fit together, aiming to produce a fully adequate account.

*Do you consider your method a hypothetical deductive method concerning the inferential nature of comprehension?*

That is a really tricky question. Our idea is that there is a spontaneous inferential system that takes as premise the fact that the speaker has produced a certain utterance, and yields as output a hypothesis about what the speaker meant. That makes it look like a hypothetical deductive process. But, first, the hypothesis formation process is very constrained: we claim that there is an actual mental mechanism that follows the path of least effort in looking for relevance, and goes straight to the right interpretation in the simplest cases. So the hearer is not just sitting there with a bunch of premises and doing repeated applications of modus ponens in order to get to conclusions. There is a quite complex automatic comprehension heuristic which produces a hypothesis, a tentative interpretation, that then has to be evaluated. And second, in spontaneous comprehension, the hearer doesn’t set about testing hypotheses very much. The cognitive system is monitoring for contradictions, so if a contradiction between interpretation and background knowledge is immediately apparent, the hearer may reject the first hypothesis and look for another one. If the system doesn’t immediately come up with a contradiction, the assumption is that this interpretation

is the best available one, and there is no point in looking further. So the comprehension system does construct hypotheses, but under severe constraints, and doesn't do a huge amount of in-depth testing.

*In contemporary philosophy it is taken for granted that "context" is something obvious that everybody understands. What actually is this concept of context?*

There is a notion of context that I think is the most appropriate for pragmatics, and that is a mentally represented set of assumptions that are brought to bear in understanding an utterance, or more generally in interpreting inputs to cognitive processes. These contextual assumptions can be derived from perception, from memory or inference, or from the imagination – for instance, if the speaker is presupposing something which you don't believe but that you have to use in order to understand an utterance. So a context is, if you like, a subset of the hearer's assumptions about the world which is brought to bear in interpreting an utterance, or other cognitive input. That is very different from the standard formal account of context, where it is a set of parameters, or the standard linguistic or philosophical accounts, where it is the physical situation of utterance, or the prior discourse, or even the common ground between the speaker and the hearer – mutual knowledge or something like that. Mutual knowledge, or common ground, are not appropriate notions for pragmatics, because very often you understand an utterance which presupposes something you didn't share with the speaker; moreover, the set of assumptions you share with the speaker is much bigger than the actual set of assumptions you use in interpreting an utterance. So what you want is a much more fine grained and cognitively plausible notion of context.

According to relevance theory, the mind develops spontaneous methods for constructing a context of assumptions that are likely to contribute to relevance. The example that I use in lectures is based on the idea that contextual information is needed for cognition as well as communication, and here most of the philosophical and linguistic notions of context don't help. Imagine that you are walking down the street towards your house, and as you get there you see smoke coming out of the window. That fact is relevant to you. In order to understand it, in order to see its implications, you have to use your background knowledge abo-

ut what causes smoke, what the consequences of smoke are, and so on. In fact, you are immediately going to think, something is on fire there, I must call the fire brigade (which requires contextual information about how to use the phone, and so on). There are of course many other possible alternative contexts that you might use when you see smoke coming from your window: you might think that it makes a nice pattern, it would be nice to film it, perhaps you should go to film school and study the films of Spike Lee, etc. From a logical point of view, there would be nothing to stop you starting on such a train of thought: if you are not concerned about your own safety, if you are an artist, that may be a perfectly reasonable way for you to go. However, for almost everyone seeing smoke coming out of their own house, the most relevant implications depend on the spontaneous retrieval of the contextual information that smoke is caused by fire, which can damage a house, that you put out a fire by calling the fire brigade, and so on. These spontaneous context construction processes are working in communication too, and they should be the starting point for any adequate communicative notion of context.

*It is hard to grasp how this spontaneity is capable of dealing with all of our knowledge?*

What goes on in the head is not magic. Your knowledge isn't scattered around in your head like porridge; it is organised in some way – presumably in a way that tends to contribute to cognitive efficiency, i.e. to relevance, perhaps stored under headings, activated by the sort of spreading activation systems discussed in psychology. So if you see smoke, it activates your knowledge related to what causes smoke and what are the consequences of smoke; if this is highly activated, the cognitive system will explore its implications first. So the organisation of the memory and the spreading activation system should contribute to overall relevance. When I say that relevance theory aims to be a cognitive plausible theory, it means that it should be testable by experimental methods, and make use of the evidence that psychologists have been accumulating over the last hundred years. Does this sound too much like magic to you?

*No, not magic. It is exciting that linguistics is quite a new field, and that linguistic pragmatics and implicature is such*

*a dynamic science. New things happen all the time, so if you don't think of doing magic, but try some fortune-telling instead, what do you see around the next corner?*

The photograph on the front of our book *Relevance* is a picture of a fortune-teller, because fortune-tellers are interpreting clues, coming up with hypotheses that go way beyond the evidence, and that's what we think happens in communication. One thing that is very clear is that there is a big movement towards experimental work in pragmatics, and if you look at what has happened over the last twenty years there has been a big shift away from purely social approaches to pragmatics. Levinson did very interesting work on the interaction between sociology and pragmatics, and there are important things to say about the interaction of linguistic, social and cognitive factors in communication. At the moment, though, there is a very clear trend towards a cognitive approach to pragmatics, and I think that's likely to go on, which means that

the relevant methods of testing are likely to be experimental. Very interesting experimental work is being done on whether implicatures are derived by default inference, or whether they are derived by the same processes as more particularized implicatures. I have to say that relevance theory is coming out pretty well on those experimental tests. There is also interesting work being done on metaphor, on irony and on figurative language in general, and there are implications worth pursuing for the study of style and literary communication. Dan Sperber has been looking beyond pragmatics to issues of mental architecture, the interaction between cognitive and social factors in communication, the role of metarepresentation in the construction of social and moral norms, and the relation between understanding and believing. It's all very well to know how people understand utterances, but once you have understood an utterance, you need to know whether to believe it or not. So there is plenty of more work to do!



MONTASJE: DILER QARADAKI