

(In Jean Khalfa (ed.) *What is Intelligence?* Cambridge University Press (1994), 179-198.)

Dan Sperber
UNDERSTANDING VERBAL UNDERSTANDING

All of us humans speak and understand at least one language, English for instance. How intelligent we are! Well, how intelligent are we? Is mastery of a language truly proof of intelligence?

I take it that intelligence, like beauty, is a property rather than a thing. There is no area of a brain that might properly be called its intelligence. On the other hand, a variety of doings and doers can be called intelligent. Though there probably isn't a satisfactory definition, typical instances of intelligent doings are not too hard to characterize. They involve creative reasoning achieved by bringing together various informations (for instance new and old informations) on the basis of which novel conclusions, insights, or decisions can be reached. More technically, we happily describe as intelligent inferential processes that are not entirely "data-driven." If a process isn't at all inferential, or if it is an automatic reaction to a specific stimulus, then talk of intelligence is much less appropriate.

Notwithstanding the doubts that surround the very notion of intelligence, the view that human language is proof of human intelligence may well be the oldest and most common philosophical cliché. Here is today's quote:

As the soft lips and pliant tongue are taught
With other minds to interchange the thought;
And sound, the symbol of the sense, explains
In parted links the long ideal trains;
From clear conceptions of external things
The facile power of Recollection springs.
Whence REASON's empire o'er the world presides
And man from brute, and man from man divides

(Erasmus Darwin, *The Temple of Nature* 4.
265-72)

The claim that linguistic behaviour shows human superior intelligence can be taken in three ways. It may be the things we say: our words give evidence of our thoughts, and the soundness and creativity of our thoughts show how intelligent we are.

It may also be argued that, without linguistic communication, there would be no literature, no science, no law, in a nutshell, no cumulative building of knowledge, theoretical or practical. Now the cultural

transmission of knowledge, the Darwin Lectures for instance, no doubt contributes greatly to the development of individual intelligence. These two lines of argument have been soundly and creatively pursued by Daniel Dennett in his contribution to this volume, and I won't consider them further here.

According to a third line of argument, nothing better establishes our intelligence than the very existence of verbal communication, whatever we say. As Descartes argued, even fools saying foolish things display a form of intelligence that no other animal possesses. The very ambiguity of the word "understanding," meaning both intelligence and comprehension, is, in this respect, quite telling. This view of the relationship between language and intelligence, however, doesn't square too well with an even more common view, according to which verbal communication is a matter of coding and decoding. For are coding and decoding truly intelligent activities?

Coding and its limits

What do we do when we communicate? What cognitive skill do we thereby exhibit? The common view is that communication is possible just in case interlocutors share a code. A language such as English is seen as a complex code. Whether simple or complex, a code is a device that generates pairs made up of a message and a signal: the Morse code, for instance pairs each letter of the alphabet with a series of short or long beeps; a language pairs linguistic senses and sounds. The pairing of messages and signals generated by the code can be made to work in two kinds of devices: encoders and decoders.

Humans can perform both as encoders of linguistic senses and as decoders of linguistic sound, and that, so the story goes, is how they communicate with one another. Failures of communication occur when encoding or decoding isn't done properly, or when noise damages the sound signal, or, more significantly, when the codes of the interlocutors are not properly matched. Otherwise, such code-based communication is sure to run smoothly. This is a simple and powerful explanation of the successes and failures of communication. However, if this explanation is correct, then the ability to communicate linguistically shouldn't be described as intelligent at all.

The work of an encoding or decoding device is neither inferential nor creative. It is not inferential because the symmetrical relation between a message and a signal is quite different from the asymmetrical relation of premise to conclusion: just as the letter "m" does not logically follow from two long beeps, the meaning of a sentence does not logically follow from its sound. The work of an encoding or decoding device is not creative either: it is an automatic reaction to the input message or signal. Actually, it had better *not* be creative: a creative bout at the encoding or decoding end would jeopardize the

symmetry between the two processes and hence the success of code-based communication.

However, the common view of verbal communication is false and, as I will try to show, coding and decoding are just ancillary components of what is essentially a creative inferential process.

Take an ordinary sentence such as:

It's late

As an English speaker, you know, don't you, what this means? In a situation where you wanted to convey that meaning to another English speaker, you would say "It's late" and be understood, and that's all there is to it. Or is it?

To begin with, in saying "It's late" you might intend to convey not just the explicit information that it is late, but also something implicit: say, that it is time to go home. The existence of such implicit content has long been recognised, and its study, inspired by the work of the philosopher Paul Grice, has become the main focus of pragmatics. Not everything that is communicated is wholly encoded: that much is generally agreed. I would like to make a stronger claim: nothing that is communicated is wholly encoded. In other words, even the explicit part of communication is never fully explicit, not by a wide margin, actually.

Consider the following dialogue:

Peter: When does the train arrive?

Mary: It's late!

Here Mary is using "It's late" in a meaning different from the one you probably had in mind a minute ago. Whereas in the earlier utterance, the "it" did not refer to anything at all and was just a syntactic filler, in Mary's utterance, the "it" refers to the train. In yet other utterances of the same sentence, "it" might refer to an apology, a payment, a chemical reaction, or whatnot. Analysing the word "it" in the hope of finding out what, if anything, it refers to would be as unhelpful as staring at a pointed finger in the hope of finding out what, if anything, it is pointing at. What must be taken into consideration in order to determine what "it" may refer to is, of course, the context.

An utterance must be taken together with a context. How is that done? It is often supposed that there must exist some system of rules that applies to an utterance and its context taken together, to yield the intended interpretation. However, this presupposes that the context is somehow given and, together with the utterance itself, provides well circumscribed initial data for the interpretation process. This presupposition is quite mistaken.

"It's late," taken without contextual information, suggests that the "it" is a syntactic filler without a referent. Knowing just that "It's late" was a reply to the question: "When does the train arrive?" suggests that "it" referred to the train. But further contextual information might reverse this interpretation:

Mary: I'm tired. I want to go home. We don't even know that Johnny will be on the train.
Peter: When does the train arrive?
Mary: It's late! Forget the train! Let's go home!

Here, the "it" in "It's late" is once more the usual syntactic filler. Further contextualisations would again reverse the preferred interpretation. As such examples show, the context is not a given on the basis of which comprehension might proceed. Rather, deciding what constitutes the pertinent context is part and parcel of the interpretation process.

Whether "it" refers, and to what, is by no means the only source of indeterminacy in "It's late." Nothing is late in and of itself. To be late is to be late with respect to some expectation, schedule or timetable. But, of course, several such schedules may enter into consideration. Even referring to specific objects such as trains doesn't determine a single relevant schedule. Thus, I might say

The late train to Cambridge is hardly ever late

without contradicting myself: my two uses of "late" would merely relate to two different schedules: lateness in the day and lateness on departure or on arrival. Deciding which schedule is intended by the speaker again depends on a context that must be discovered as part of the interpretation process.

Suppose, however, that we know that "it" refers to the train to Cambridge, and that the train is said to be late with respect to its expected time of arrival, 5:25 p.m.. Surely, in that case, we know what "It's late" means. Or do we? Here we are, let us imagine, on the platform at Cambridge railway station at 5:20. Three people are about to utter the same sentence, "It's late," referring to the same train and with the same schedule in mind and yet they will mean quite different things.

First Peter, who believes the worst of British Railways, asserts at 5:21:

Peter: No doubt it's late again!

However at 5:24 the train is heard, and Mary says to Peter:

Mary: You're a real seer! Indeed it's late!

Our third character is a railway controller looking at his stopwatch. He sees that the train stops at 5:26, and says sternly:

The controller: It's late!

Peter and the controller are speaking in earnest, while Mary is speaking ironically. Peter and Mary use "late" rather loosely, while the controller uses it so strictly as to see sixty seconds late as late enough to mention. Because of these subtle differences, their three utterances, though the words, the train, and the schedule are the same, convey quite different thoughts.

Yes, as an English speaker, you know the linguistic sense, or rather senses, of "It's late," but these senses are very incomplete affairs; they are not, by themselves, the meanings you might want to convey when you utter the words. As a hearer too, merely knowing the linguistic senses of "It's late" does not tell you what a speaker means in uttering these words. The same is true of every sentence in every human language: the thought we intend to convey can never be fully encoded, and linguistic decoding is only a first step in understanding a speaker's meaning.

Going from linguistic sense to speaker's meaning is not a matter of further decoding, not even of context-sensitive decoding: it is a matter of inference. What then is the form of inference involved in verbal understanding?

Inference

In any instance of verbal understanding, there is an initial premise and a goal. The premise is the information that a certain person uttered a certain sentence. The goal is to discover what that person meant in uttering that sentence. The premise and the conclusion of a successful comprehension process are both complex in the same way, but not to the same degree.

The first premise in comprehension might be something like this:

Carol says: "It's late"

The conclusion might be something like this:

Carol means that it is time to go home

Here, premise and conclusion both involve a "meta-representation": that is, they contain a representation of a representation. The premise is about an utterance ("It's late") and directly quotes that utterance. The conclusion is about a thought (that it is time to go home) and

indirectly quotes that thought. On closer analysis, however, it turns out that while the premise is a first-order meta-representation, the conclusion is a higher-order meta-representation.

Just like the attribution of an utterance, the attribution of a belief, for instance

Carol believes that it is time to go home,

may be a simple meta-representational affair. There is an important relationship between meaning and believing but it is not a simple one.

When John concludes that, in saying "It's late," Carol means that it is time to go home, he need not at all attribute to her the belief that it is time to go home: after all, she might be insincere and attempting to communicate what she does not believe. What John must be attributing to Carol is an intention rather than a belief: the intention that *he* should believe that it is time to go home.

An intention is a mental representation of a desired state of affairs. What state of affairs does a speaker desire? A state of affairs in which some information becomes represented in her hearer's mind as a result of her utterance. We call such intentions "informative intentions." Simple informative intentions are first-order meta-representations. Thus the content of Carol's intention is the following meta-representation (I will use different lines and indents to separate the different representational levels):

John should believe

that it is time to go home

Thinking that someone has an informative intention is entertaining a second-order meta-representation. Thus John's understanding of Carol's utterance contains the following second-order meta-representation:

She intends

me to believe

that it is time to go home

I will argue later that full-fledged comprehension involves reaching a more complex conclusion than this. It is plausible, however, that young children at a certain stage in communicative development do not go beyond such second-order meta-representations and that adults too do not systematically attend to the further meta-representational tiers involved in communication. Still, reaching this already complex attribution of an informative intention may, in many cases, provide adequate understanding, as I will now describe.

Going from the premise:

Carol says: "It's late"

to the conclusion:

She intends me to believe that it is time to go home,

is a case of inductive inference. The conclusion goes beyond the information contained in the premise in two respects: it goes from the ambiguous and incomplete linguistic sense of "It's late" to the proposition that it is time to go home; and it goes from first-degree meta-representational attribution of an utterance to second-degree meta-representational attribution of an informative intention.

Let us look first at the meta-representational aspect. Whereas the initial premise of this inductive inference merely attributes a behaviour to the speaker, namely the production of a certain utterance, the conclusion attributes an intention to her. Computers are coming close to being able to recognise a speaker's utterance, and that is of course a remarkable achievement. However, recognising a speaker's intentions requires far more intelligence than computers have been equipped with so far.

In general, behaviours can be conceptualised as bodily movements or as realising intentions. Conceptualising voluntary behaviours as realising intentions is far more economical, more explanatory, and of greater predictive value than merely conceptualising them as bodily movements. However, in order to conceptualise behaviours in terms of underlying intentions, an organism needs the ability to entertain meta-representations. Very few animals have any meta-representational ability. There is some experimental and anecdotal evidence suggesting that chimpanzees and possibly other non-human primates do possess some such ability in a rudimentary form.

What about humans? Do they have a meta-representational ability? Do birds fly? Do fish swim? Humans can no more refrain from attributing intentions than they can from batting their eyelids. The only issue regarding humans is developmental: at what age and through what stages do meta-representational abilities develop, if they are not there from the start? There is plenty of current research on the issue, much of it assuming that the ability to attribute beliefs and intentions appears well after the development of rich verbal abilities. I should point out that this is inconsistent with the picture of verbal understanding that I am proposing here. If the picture I am presenting is correct, some ability to attribute intentions precedes the ability to communicate verbally.

How are intentions identified? In many cases, the attribution of a specific intention involves a simple pattern of inference: the behaviour of an individual is observed to have a certain desirable effect; the individual is assumed to have intended this very effect. A man shoots an arrow and kills a deer: an observer infers that it was the man's intention to kill the deer. Though often successful, this pattern of inference will sometimes fail: a behaviour may not achieve the intended effect; or a behaviour may produce a desirable effect that had not been foreseen, and therefore not been intended either. A more reliable inference pattern will see as intentional not the actual desirable effect of a behaviour but an effect that the agent may have seen as desirable and as made more probable by his behaviour. Humans easily perform inferences of this more complex kind: a man shoots an arrow and the arrow comes close to hitting a deer: all the same, an observer has no trouble inferring that it was the man's intention to kill the deer.

Capable of attributing mental states to others, humans are also capable of forming an intention to change those mental states. Peter wears the tie that Mary has given him. A knowledgeable observer recognises Peter's intention to please Mary. Bobby tiptoes to his mother and suddenly shouts "Boo!" An observer recognises Bobby's intention to startle his mother.

Consider a behaviour intended to change the mental state of an individual, and imagine that that target individual herself observes this behaviour and recognises the underlying intention. Will this affect the chances of the intention being fulfilled? Well, it may sometime hinder, and sometime help the fulfilment of the intention. If Bobby's mother sees him coming and understand his intention to startle her, she won't be startled. On the other hand, if Mary realises that Peter intends to please her, she may be even more pleased.

Informative intentions are typically helped by being recognised by the intended audience. Communication in particular is a means of fulfilling an informative intention by making one's audience recognise it.

It is relatively easy to recognise a behaviour as intentionally informative. Verbal and other communicative behaviours are typically attention-catching, and they have to be. I can't communicate with you if you don't pay attention. Verbal behaviour, moreover, is largely specialised for communication.

Communicative behaviour also typically calls ideas to the mind of the audience. This may be done by non-verbal as well as by verbal means. Carol might establish eye contact with John, lean her head sideways, and close her eyes for a second, in other words, she might mime falling asleep. This would bring to John's mind the idea of sleep, or of

Carol sleeping. Saying to John "It's late" would, in the same situation, bring to his mind the senses of the sentence, which all include some idea of lateness. In either scenario, John would, at this point, recognise that Carol was intending to inform him of something, but he would still have to discover what she wanted to inform him of. The ideas she has brought to his mind fall quite short of representing definite information.

We have focused so far on the meta-representational dimension of comprehension. We are now brought back to the semantic dimension: how does John, for instance, go from the ambiguous and incomplete linguistic senses of "It's late" to the well-understood proposition that it is time to go home? I will describe three interpretation strategies that John might adopt, a naive optimistic, a cautious optimistic, and a sophisticated strategy.

Naive optimism

Suppose John is very trusting indeed and takes for granted that Carol is behaving both benevolently and competently. Then John can take for granted two further things: that the information Carol wants to convey to him is information worth his attention, and that the means she is using to convey that information should make it as easy as possible for him to retrieve it.

Deirdre Wilson and I have argued in our book *Relevance* that information worth one's attention is information which brings about significant cognitive effects, that is, information which, when taken together with what the individual already knows or assumes, allows inferences which would not have been possible otherwise. Such information is *relevant*, and the greater the cognitive effects it brings about, the more relevant it is. We have argued, on the other hand, that the greater the mental effort needed to acquire and process relevant information, the less relevant it is. This being so, when a hearer assumes that a speaker is benevolent and competent, he can take for granted that the information conveyed by her is relevant enough to be worth his attention, and that retrieving it should not cause unnecessary effort, for this would pointlessly diminish relevance.

In particular, John may take for granted that the linguistic senses and ideas that Carol's utterance has brought to his mind have not been evoked in vain, and that they provide an optimal starting point for retrieving the information Carol intended to convey. Given the situation and the way in which these ideas have been evoked, one linguistic sense is more salient than the others, some ways of fleshing it out come more easily to mind, some contextual information is more easily invoked, and thus the possible interpretations are ranked in John's mind in order of ease of access. In other situations, of course, the various possible interpretations of the same utterance

would be ranked differently.

All John has to do now is follow the path of least effort in constructing an interpretation of Carol's behaviour and stop when he reaches an interpretation that provides him with information relevant enough to be worth his attention. The interpretation he will reach in this way will be the one intended by Carol. Why, you may ask, should the most easily reached relevant interpretation be the true one? Why shouldn't the true interpretation be arcane and utterly boring? From a logical point of view, relevance and ease of access and have nothing to do with truth. From a psychological point of view, however, the situation is different: a competent and benevolent communicator will see to it that the information she wants to convey is indeed relevant to her audience, and is more easily retrieved than other otherwise plausible interpretations. So, the first accessed, relevant enough interpretation is, by these very properties, confirmed as the intended one.

On hearing Carol say "It's late," John may not have any specific referent for "it" that springs to mind, so the non-referring interpretation of "it" would be favoured. The first interpretation of "late" he can think of might be late with respect to the time at which they had promised the baby-sitter to be back home. Together with easily accessible contextual premises, for instance the knowledge that their excellent baby-sitter might not come back if promises made to her are not kept, Carol's utterance implies straightforwardly that it is time to go home. This implication makes the utterance relevant enough, and therefore John is led to conclude that Carol intends him to accept this implication, that is, to believe that it is time to go home.

The naive interpretation strategy I have just described will yield adequate comprehension whenever the speaker is indeed benevolent and, above all, competent enough to realise what is relevant and salient for her audience at the time. However, speakers are not always that competent. Suppose for instance that, unbeknownst to Carol, John has just been worrying about a delivery that should have been made that very day. The first relevant interpretation of Carol's utterance "It's late" to come to John's mind might be that the delivery is late, and the simple inference pattern I have described would cause him to accept -- wrongly -- that interpretation as the intended one. Such errors of comprehensions do occur. Young children in particular easily believe that one is talking about what happens to be foremost in their mind (and conversely that what they want to talk about is foremost in the minds of their listeners). Still, most of the time, these errors are avoided. This suggests that competent communicators have a more powerful interpretation strategy at their disposal.

Cautious optimism

What this more powerful interpretation strategy might be is not too

hard to imagine. It is just a special case of competent attribution of intentions. A competent observer can infer that the hunter intended to kill the deer even though the arrow flew just over the animal. The observer's inference pattern consists in seeing as intentional not the actual effect of an action but an effect that the agent might have desired and expected. Similarly, a more appropriate interpretation strategy in matters of comprehension consists in attributing to the speaker an interpretation that she might have thought would be relevant enough and most easily accessed, rather than what actually happens to be the most accessible relevant interpretation. In other words, a competent hearer allows for the possibility that the speaker might have misjudged what would be most accessible and relevant to him.

I have now introduced two optimistic strategies: a naive and a more cautious one. In the naive strategy, the speaker is assumed to be benevolent and competent, and the inference pattern consists in going uncritically where the ideas suggested by the linguistic sense of the utterance will take you: look for easy relevance; assume that it was intended. The conclusion of such an inference is a second-order meta-representational attribution of a first-order meta-representational intention. However, in that inference pattern, attributions of intentions do not serve as premises. Their complex logical structure is derived but not exploited.

In the second, more cautious strategy, the speaker is assumed to be benevolent, but not necessarily competent. She may not know what is on her hearer's mind. She may therefore fail to convey relevant information, or fail to make the relevant information she intends to convey more accessible than any other possible interpretation. As in the naive strategy, the hearer should follow the path of least effort, but he should stop not at the first relevant enough interpretation that comes to mind, but at the first interpretation that the speaker might have thought would be relevant enough to him.

Suppose the interpretation that first occur to the hearer is relevant enough to him. His next step will be to evaluate this interpretation in the light of what he knows about the speaker. Could she have expected this interpretation to occur to him? Would she have seen it as relevant enough to him? Only if the answer to both questions is yes will this interpretation be retained. Otherwise, the next accessible interpretation will be tested in the same way.

This time, second-order meta-representations may serve not just as conclusions, but also as premises. John may reject the conclusion that Carol intends to inform him that the expected delivery is late, because he believes that

she does not know

that I am wondering

whether the delivery took place

He will then go to the next possible interpretation, the interpretation on which it is time to go home. This second interpretation will pass the evaluation test: yes, Carol could have expected it to occur to him and she may have found relevant it enough to be worth communicating. John, using this inference pattern, will not be misled by his own obsession with the possibly late delivery. He will understand Carol correctly.

This second, cautious pattern of inference will yield adequate comprehension whenever the speaker is benevolent, even though she might not be competent enough to realise what is relevant and salient for her audience at the time.

But now let me tell you in confidence: speakers are not always benevolent, and communication is not always a nice affair. I you doubt me, just look at the amount of communicative energy spent on trying to convince young people that smoking will make them glamorous! Very competent, of course, but hardly benevolent. So, a truly sophisticated hearer does not assume that every communicator is benevolent. And, for that matter, a truly sophisticated communicator is not always benevolent.

Imagine the following scenario. The baby-sitter usually leaves at midnight. This time, however, Carol, thinking that the party she and John are invited to will be great fun, has asked the baby-sitter to stay until one. Carol does not enjoy the party after all, but John does. At 11:30, she says to him: "It's late," expecting him to think that it they have to go home for the baby-sitter. Carol does not know that John knows about the special arrangement she has made with the baby-sitter. Now, if John took for granted that Carol was benevolent, he would be bound to misunderstand her: she could not intend him to believe something she knows to be false and that therefore does not constitute relevant information. He would then look for another interpretation, and assume for instance that she means late with respect to some other schedule or expectation.

You might say: given what he knows, John will, in this case, drop the assumption that Carol is benevolent. True enough, but very puzzling. Remember that the two interpretation strategies I have described so far both presuppose this very assumption. If John assumes that Carol is benevolent, he will not find the right interpretation and he will have no reason to doubt Carol's benevolence. If John drops this assumption he will not find any interpretation at all. When you take benevolence for granted, you cannot recognise bad faith or lies. When you don't take benevolence for granted you cannot use optimistic interpretation

strategies that presuppose benevolence. Hearers who are capable of recognising lies must be using yet another strategy.

Sophisticated understanding

Competent hearers (who are usually also competent speakers) realise that speakers use communication to pursue their own ends, which may correspond in some respects to the ends of their audience, and differ in others.

Communication is a special way of fulfilling an informative intention. An informative intention can be fulfilled in many ways. One way is to provide evidence, genuine or spurious, of the information one wants someone to accept as true. Thus Carol might have pretended to fall asleep at the party; John might be fooled by this behavior and come to believe that it is time to go home, without ever realizing that this is exactly what Carol intended him to believe.

Another, often more practical way of fulfilling an informative intention is to provide evidence for the fact that one has that very same informative intention. Instead of pretending to fall genuinely asleep -- and hurting her hosts' feelings as a side effect -- Carol might, as I mentioned, mime falling asleep in a manifestly intentional way. This ostensive behaviour would give John evidence of her intention: she intends, by her behavior, to make him realise that it is time to go home. John might well, as a result of grasping Carol's intention, end up believing that it is time to go home. After all, if someone you trust shows you that she wants you to believe something, then that gives you a reason to believe it.

Compare three ways in which Carol might try to convince John that it is time to go home: she might deceitfully pretend to fall asleep; she might overtly mime falling asleep, or she might say "It's late." In terms of outward appearances, the deceitful pretence and the overt mime are much more similar to one another than either is to the linguistic utterance. In terms of procedure, however, the overt mime and the utterance fall together. They both consist in trying to fulfil Carol's informative intention by giving John genuine evidence of that very intention. The deceitful pretence, on the other hand, stands alone; it consists in giving false evidence of the information Carol wants John to accept, and no evidence whatsoever of Carol's intention.

To fulfil an informative intention by making it known is, properly speaking, to communicate. Communication is not an accident that happens when people have informative intentions and these informative intentions somehow become known to their intended audience. Communication is generally intentional. An intention to communicate, or, as we call it, a "communicative intention," is a second-order informative intention: the intention to make a first-order

informative intention known. A simple informative intention is a first order meta-representation. Knowing that someone has an informative intention is entertaining a second-order meta-representation. Having a communicative intention -- that is, intending someone to know that one has an informative intention -- is entertaining a third-order meta-representation. Attributing a communicative intention to someone is therefore entertaining a fourth-order meta-representation. Thus, if John is aware of Carol's intention, not just to cause him to believe, but to *communicate* to him, that it is time to go home, then he is entertaining a thought of the following complexity:

She intends
 me to know
 that she intends
 me to believe
 that it is time to go home

When an agent is known to have some main intention, it can be inferred that he also intends whatever he believes is necessary for his main intention to be fulfilled. The hunter intending to kill the deer intends his arrow to hit the deer in a vulnerable spot and with enough strength to penetrate it.

In the same way, when someone is known to have a communicative intention, further intentions of hers become inferable. To have a communicative intention is to hope to fulfil an informative intention by making it known to one's audience. For this to succeed, there is an obvious condition that must be fulfilled: the audience must believe that the information the communicator wants to convey is relevant enough (and is true to the extent that truth is necessary for relevance), or else the audience won't pay attention, or won't accept the information as true. It is manifest therefore, to communicator and audience alike, that the communicator intends her audience to assume that what she intends to communicate is relevant enough. In other words, every act of communication, and in particular every utterance, conveys a presumption of its own relevance. This universal fact is what we call *the principle of relevance*.

As I have argued, if a hearer can postulate that the speaker is benevolent and competent, and therefore that her utterance is relevant to him, then he has a good and simple interpretation pattern at his disposal. However the initial postulate is often false. If a hearer can postulate that the speaker is at least benevolent, and therefore that she believed that her utterance would be relevant to him, he has a good, somewhat more complex interpretation pattern at his disposal. However, this postulate may still be false. If a hearer now attributes to

the speaker a full-fledged communicative intention, then he need not *postulate*, he can *infer* that she intends her utterance to seem relevant to him; if he is right in assuming that she is communicating to him, he cannot be wrong in making this further assumption. Does this assumption provide him with a good interpretation pattern? Yes it does.

In the third, sophisticated strategy, the speaker is not assumed to be benevolent or competent. She is merely assumed to intend to seem benevolent and competent. The beauty is that now even if that intention fails, interpretation may succeed. As in the other interpretation strategies, the hearer should follow the path of least effort, but he should stop, not at the first relevant enough interpretation that comes to mind, nor at the first interpretation that the speaker might have thought would be relevant enough to him, but at the first interpretation that the speaker might have thought would *seem* relevant enough to him.

Suppose the first interpretation the hearer comes up with is not relevant to him at all, because he knows that it is false, however it would be relevant enough if were true. The hearer will evaluate this interpretation: could the speaker have expected him to come up with it? Would she have seen it as likely to seem relevant enough to him? If the answer to both questions is yes, this interpretation will be retained. Then the hearer may well wonder: did the speaker assume that this interpretation would appear relevant to him because she herself believes it is true and relevant, or was she trying to mislead him? Either way, the speaker's *informative* intention will have failed: the hearer does not accept the intended interpretation as true. But *the communicative intention will have succeeded*: the hearer has correctly retrieved the intended interpretation.

Thus John will correctly understand that Carol intended him to believe that it was time to go home because of the baby-sitter. However he will disappoint Carol by not believing what she intended him to believe, and moreover he will correctly recognise her disingenuousness.

Full-fledged communicative competence involves, for the speaker, being capable of having at least third-order meta-representational communicative intentions, and, for the hearer, being capable of making at least fourth-order meta-representational attributions of such communicative intentions. In fact, when irony, reported speech, and other meta-representational contents are taken into consideration, it becomes apparent that communicators juggle quite easily with still more complex meta-representations. This does not imply that communicators are conscious of the complexity of their mental representations. What it does imply is that every tier of these representations may play a role in inference. Much of everyday communication takes place between people who are benevolent to one

another and who know one another well enough. In such circumstances, cautious, and even naive optimism can serve as "default" interpretation strategies, and the higher level meta-representational tiers involved in sophisticated understanding may play no role at all. Still, when the optimistic strategies fail, a competent hearer resorts to the sophisticated strategy, and performs the complex meta-representational inferences it involves, without the slightest difficulty.

The ability to use such complex meta-representations is ignored in psychology, where the study of much simpler meta-representation is a relatively new topic. The logic involved in meta-representational inferences is also hardly ever studied. Are humans intelligent enough to know how intelligent they are?

