

# Two Approaches to Communication

## 1.1 "A good old fiddle"

The BBC radio series *The Ape That Got Lucky* was a series of spoof lectures on human evolution. The first episode concentrated on the origins of language and communication. One point made by the presenter, Chris Addison, is that context is critical to successful communication. He illustrates this with an example of a keen but shy violinist who, while at a party, says to you, "Wait until these people have all gone and I'll dig her out and have a good old fiddle". In the context his meaning seems clear but, Addison points out, the same sentence might be understood very differently if uttered at a funeral.

Jokes of this sort make use of the fact that context is critical to understanding what a speaker means. The example that the psycholinguist Steven Pinker uses to make this point comes from the 1960s American comedy *Get Smart*, in which the central character, Maxwell Smart, was assisted by a humanoid robot called Hymie. One of Hymie's most distinctive traits was that he interpreted commands in a highly literal way. The trademark case was "Give me a hand", in response to which Hymie would remove his own hand and pass it to Maxwell. Others included "Kill the light!" (Hymie uses his pistol to shoot the lightbulb), "Get hold of yourself!" (Hymie grasps each of his arms with the other), and "Knock that stuff off!" (Hymie knocks all the paper off the desk). There is a clear difference between Hymie's interpretation of the command and the meaning that Maxwell intended to communicate - and so when Hymie takes Maxwell's words literally, our expectations are violated in a humorous way.

What these and many other examples show is that there is sometimes a big difference between the *literal meaning* of what is said (sometimes called *sentence meaning* or *linguistic meaning* or *utterance meaning*), and the speaker's *intended meaning* (sometimes called *speaker meaning* or *utterer's meaning*). The literal meaning is the 'decoded' meaning, based upon the dictionary definitions of the words used, and the speaker meaning is the meaning that was actually intended. Sometimes the difference between these two is obvious and large, such as in sarcasm, and other times the difference

is not so big - but it is always there to some degree. Even an utterance as apparently simple as "It's raining" can have multiple different intended meanings, such as 'I don't want to go out after all', or 'Get the umbrella'. In fact, "It's raining" does not even specify where or when there is rain, yet these facts are critical if the utterance is to be relevant. In short, even the most simple, straightforward utterance can be interpreted in different ways, depending on the context. The technical term for this is *underdeterminacy*: the literal meaning of an utterance underdetermines the speaker meaning (Carston, 2002a; Atlas, 2005). In other words, linguistic communication is never just literal meaning. Literal meaning helps us understand speaker meaning, but it is not the same thing.

My thesis in this book is that the origin of language was the consequence of the creation of an evolutionarily novel form of communication, in which underdeterminacy is an inherent and inevitable feature. Underdeterminacy is often seen as a defective quality for a communication system, one that creates ambiguity and misunderstanding. It is true that such vagaries do follow in its wake, but these characteristics are also assets, since they allow communication to be used in incredibly flexible, creative and indeed funny ways. It was, I will argue, the emergence of such a communication system that allowed our ancestors to create the expressively rich form of communication that we call language. This evolutionarily novel form of communication is called *ostensive-inferential communication*, or just *ostensive communication* for short.

In this opening chapter I will describe exactly what ostensive communication is, and how it differs from other types of communication. To do this, I will describe two different ways of thinking about the very possibility of communication. The first is called the *code model*. It is an intuitive way of thinking about communication, and it provides a good description of most and perhaps all animal communication. However, it is unable to handle the underdeterminacy that is inherent to human communication. In contrast, the second way of thinking about communication, called the *ostensive-inferential model* (or just *ostensive model*, for short), was developed to account for these realities. Having described these two approaches, I will then expand on several important preliminary issues, including the precise difference between code and ostension, how they interact with one another, the role of 'codes' in language, and the meaning of meaning, among others. These discussions will set the stage for my discussion, in Chapter 2, of exactly why the transition to ostensive communication was so revolutionary.

## 1.2 The code model

When presented with questions such as 'What is communication?' or 'How does communication work?', most people's intuitions are typically along

the following lines: that communication involves the transmission of information, via some communication channel. Information is encoded into a signal, sent along this channel, and then decoded at the other end. If the algorithms for encoding and decoding are appropriately calibrated to one another, then what is encoded at one end is the same as what is decoded at the other end. The result is that information has been transferred, from the signaller to the receiver. This way of thinking about communication is called the *code model*.

The code model combines at least two metaphors of how communication works (Blackburn, 2007). The first is the idea that signals contain messages that are packaged up and sent along some channel, to be unwrapped at the other end. This metaphor is called the *conduit metaphor* (Reddy, 1979; Figure 1.1). (A conduit is a channel for conveying things from one location to another.) The power of this metaphor is illustrated by the degree to which it pervades our everyday language about communication: ‘Get your message *across*’; ‘His feelings *came through* clearly’; ‘I couldn’t *extract* the meaning from your writing’; ‘I *gave* you that idea’, and so on.

The second metaphor that contributes to the code model is the information-theoretic approach to communication developed by Claude Shannon (1948). Here, signals are seen as strings that are to be transmitted along a communication channel (Figure 1.2). The intellectual challenge comes from devising a way to do this that overcomes the presence of noise and error in the transmission process - particular problems for electrical engineering, telecommunications, and other areas that involve the transmission of digital information.

It is not difficult to see how the conduit metaphor and information theory combine to form the essence of the code model: ‘information’ is packaged up and then transmitted along a conduit, to be unwrapped at the other end. Having said that, the code model is not so much the precise summation of these two metaphors, but rather a general background assumption about how communication works that has been greatly influenced by them.




Figure 1.1 *The conduit metaphor*. In the conduit metaphor, communication involves the packaging of messages, which are then to be passed along a conduit, to be unwrapped at the other end.




Figure 1.2 *The information-theoretic approach to communication.* Information is encoded, and then transmitted as a signal. During the transmission process, the signal may become distorted due to noise. The received signal is then decoded. Information theory provided a way for the processes of encoding and decoding to counteract the effect of the noise.

The code model has had impact on a great many researchers in a diverse range of disciplines. One particularly salient example is in the study of animal communication, where there are many examples of this way of thinking (Rendall et al., 2009). Here are three definitions of communication taken from that literature: “the provision of information from a sender to a receiver” (Bradbury *Sc* Vehrencamp, 1998, p. 2); “...the transmission of information from one animal to another. Information is encoded by one individual into a signal. When received by another animal, this information undergoes decoding...” (Green *Sc* Marler, 1979, p. 73); “the transfer of information via signals sent in a channel between a sender and a receiver” (Hailman, 1977, p. 52). And here is a definition of animal communication from a linguist: “One organism has a message in mind that he or she wants to communicate to another organism. He or she emits some behavior ...that encodes that message. The other organism...perceives the behavior, identifies it in terms of the meaning encoded, and treats the result of that decoding as the meaning of the message” (Anderson, 2006, p. 17). There are many more definitions with a similar flavour.

The code model is highly intuitive, and can be used as a way to describe many instances of communication in a very simple and general way. As an example, consider bacterial communication. Many species of bacteria communicate by a process known as quorum sensing, in which individual bacterial cells produce small diffusible signal molecules in particular local environments (Schuster et al., 2003; Wagner et al., 2003). These molecules bind to cognate receptors on other bacteria, and in doing so effectively reveal the presence of the bacterium in the local environment. When the concentration of these molecules reaches a threshold level (the quorum), it triggers population-wide behaviour, such as gene regulation (Williams et al., 2007; Ng *Sc* Bassler, 2009). Such interactions are most easily described in terms of a code. The encoding algorithm is something

like 'If in environment  $X$ , produce molecule  $Y$ , and the decoding algorithm something like 'If the concentration of molecule  $Y$  is above the threshold  $t$ , then perform behaviour  $Z$ '

Following this logic, one student of animal communication once compiled a six-signal 'dictionary' for grasshoppers (Moles, 1963, pp. 125-126). It reads:

Signal I: It is fine, life is good.

Signal II: I would like to make love.

Signal III: You are trespassing on my territory.

Signal IV: She's mine.

Signal V: Oh, how nice it would be to make love!

Signal VI: How nice to have made love!

A dictionary like this is, in effect, a description of a code.

Communication of this sort depends upon mechanisms of association. In signallers, certain states of the world are associated with the production of particular signals; and in receivers, the reception of those same signals is associated with particular behaviours. In fact, this associativity is the defining feature of the code model. If communication is made possible by associations, it is, by definition, an instance of code model communication.

These associations can be more or less complex. For example, in many species signal production depends in part on whether an audience is around, and if so, what sort of audience that is. This is called *audience design* (see e.g. Gyger et al., 1986 for a review of audience design in chickens). Sensitivity to the existence and nature of the audience does not, however, imply that the system is not explicable in associative terms. It just means that the associations are nuanced to these environmental cues (i.e. the code might be 'if there is a leopard in the environment, and also members of my group/kin, then produce the relevant alarm call').

Some cases of human communication are explicable in the same associative terms. The ring of pigmented skin that surrounds nipples is called the areola. The areolar glands of new mothers secrete chemical compounds, which signal to newborns the location of the nipple, and hence of milk and its various nutrients (Doucet et al., 2009). In effect, this 'chemosignal' allows the mother's body to communicate with the infant's body. In terms of the code model, the encoding algorithm is 'if having just given birth, secrete chemical  $X$ ', and the decoding algorithm is 'if  $X$  is detected, then move towards it'. Tears are also thought to contain a chemosignal, one that causes sympathy in others (Gelstein et al., 2011). It is likely that there are many more signals, chemical or otherwise, by which human bodies communicate with one another.

Another instance of code model communication in humans is involuntary laughter. Involuntary laughter is often called ‘Duchenne’ laughter, after the French physician who first distinguished the different muscles used in ‘real’ and ‘faked’ smiles. Several strands of research suggest that involuntary laughter serves a communicative function, namely to express social solidarity (Gervais & Wilson, 2005). Evidence also suggests that laughter in humans and laughter in the other great apes are evolutionary homologues of one another; that is, they both descended from the same common ancestral behaviour (van Hooff, 1972; Davila Ross et al., 2009; Bryant & Aktipis, 2014).

The same is likely to be true of several other behaviours such as snorting, smiling, frowning, sighing, pouting, and so on, which all occur in both humans and other animals, and have been hypothesized to serve a communicative function of some kind, typically one that involves managing social relations (Darwin, 1872; Owren & Bachorowski, 2003; Davila Ross et al., 2009). Monkey alarm calls fall into this box too: several lines of evidence, in particular common neural organization, suggest that these calls are homologous to spontaneous human emotional vocalizations, and not speech (Owren et al., 2011). The term ‘gesture-calls’ has been proposed as a label for the human forms of these behaviours, because of the evolutionary relationship that they share with the communicative gestures and calls of our primate cousins (Burling, 1993; 2005). The term has not caught on, but the idea that this class of behaviours should have an identifying label is correct. Regardless, the key point for the present purposes is that we communicate with these behaviours, and this communication is describable with the code model.

The point I am making is not that chemosignals and involuntary laughter operate by similar mechanisms. On the contrary, there are plainly many differences. The point I am making is instead that despite these differences, chemosignals and involuntary laughter (and grasshopper calls, and quorum sensing) do share one basic quality in common, which is that they depend upon *associations* of one sort or another, and are hence describable in the terms of the code model. This will be important shortly, when we come to contrast these examples with language and some other forms of human communication.

There are two final points that I should make about code model communication. The first is that it does not require that encoding and decoding be fully deterministic. Production and reception may be more flexible than this. For example, they could be used probabilistically instead. In this case, the code might read: if X, then do Y with a certain probability, and Z with some other probability (rather than simply ‘if X, then do Y’). Computational models of the evolution of communication show that signalling systems can still emerge and be useful when these probabilistic codes are used (Vogt, 2002; Smith, 2005; Steels, 2011). Indeed, it seems likely to me that many cases of animal communication operate at least to some extent in this way.

Because they are not fully deterministic, such systems have been described as inferential (e.g. Smith, 2005). However they are still, strictly speaking, instances of the code model, because the basic paradigm remains one in which producers are prompted to transmit signals for particular meanings, which receivers then decode. Both production and reception are still associative, albeit in a probabilistic way.

The other point I want to make is that the competent use of code model communication systems is not necessarily a cognitively trivial task. For one thing, probabilistic codes can involve a great deal of complexity. But more than that, this type of communication can, in some cases, also involve a great deal of social intelligence. Of particular relevance to my concerns in this book is the possibility that code model communication systems can be used in an intentional way. I will discuss this possibility further in Chapter 4; here, I simply want to use it to make the point that just because a communication system may be predicated on associations, that does not, *a priori*, make it simple or unsophisticated.

From what I can gather, the label 'code model' appears to have been coined by the cognitive scientists Dan Sperber and Deirdre Wilson, in their seminal book *Relevance: Communication and Cognition* (1995). In that book, not to mention many further publications since, they contrasted the code model with an alternative way of thinking about communication. The next two sections present this alternative.

### 1.3 The expression and recognition of intentions

The code model depends upon associations, but there are some instances of communication that are simply not explicable in associative terms. For example, suppose that a couple have had a number of conversations about, say, a particular TV show, which many people think is brilliant, but which they both agree is terrible. Now suppose that they head to a party one evening, and a friend of theirs starts telling them how great he thinks the show is. The friend then leaves the conversation. One half of the couple wants to acknowledge to the other half that this is just the sort of thing they had been talking about. To do this, she can improvise and use an indefinite variety of behaviours: a raised eyebrow; puffed cheeks; a glance over the shoulder; sticking her tongue out; and many others. There is no code in operation here: none of these behaviours is conventionally associated with her intended meaning, even probabilistically. There is something else going on, something that requires an alternative to the code model. In this section I lay the foundations for the presentation, in the following section, of just such an alternative.

We all mentally represent the world in some way. For example, right now my mental representations include the knowledge that Durham is in England, the belief that tomorrow will be sunny, and my intention to write

another 1,000 words before the end of the day. As this list of examples illustrates, mental representations can take several different forms: beliefs, assumptions, goals, knowledge, and so on. These representations can change, in particular when we are provided with new inputs that suggest we should change our existing representations in some way, or create new ones. For example, if I look up the weather forecast and it predicts rain tomorrow, I may, depending on how much I trust the forecast, update my belief about tomorrow's weather.

Often, we wish to change the mental representations of others. Right now, I wish and intend to change your mental representations about mental representations (or perhaps create new ones, if you have never thought about them before). When preparing dinner, I may wish and intend to create in my partner a belief that I need some help chopping the vegetables. To try and satisfy such intentions, we provide evidence for them, in the expectation that the audience will be able to infer what our intentions are, and act to satisfy them. Exactly what the evidence will consist of will depend both on the change the signaller is trying to cause, and the targeted individual's existing representations. In the case of the vegetables, one way to make my partner believe that I need help might be to push a large pile of unchopped vegetables, and a knife, in her general direction. Note that for this to work, she needs to already have a mental representation about how knives are used, and one that vegetables are normally chopped before cooking. Armed with these representations, and others, she is able to make appropriate inferences about what I intended by my actions. In this case, her representations (about whether I needed help) were easy to change, but in other cases (e.g. political views) they may not be.

An *informative intention* is a signaller's intention that the receiver change their representation of the world in response to the signaller's behaviour. I have an informative intention that my partner recognize that I need help chopping the vegetables. If you ask someone in the street for directions, you have an informative intention that the passer-by believe that you are lost, and need help. Here is a third example: I am in a coffee shop, I catch the eye of the waiter, and I tilt my coffee cup in a particular, somewhat stylized way. The waiter then comes over and refills my cup. Here, I have an informative intention that the waiter understands that I would like a refill. And so on. The content of an informative intention is, in colloquial terms, the information that it provides. More specifically, it is the changes that the signaller wants to make to the receiver's mental representations.

Informative intentions are not the only sort of intention involved in this sort of communication. The tilt of my coffee cup expresses my informative intention, but it also expresses something just as important: the very fact that I wish to communicate with the waiter at all. After all, I could be tilting the coffee cup incidentally. This is something that happens all the time, such as when chatting to friends. How does the tilt reveal to the waiter that it is a signal? In other words, how does it signal its own signalhood,



and how does the waiter recognize as much (Scott-Phillips et al., 2009)? These questions highlight that in addition to an intention to make it clear to the waiter that she wants more coffee, I must also make it clear to the waiter that I am trying to communicate with him at all. My intention to do this - that is, my intention to create in my audience a representation of the fact that I have an informative intention - is called a *communicative intention*. This intention is expressed when I establish eye contact with the waiter and tilt my cup in a particular, *ostensive*, way.

If I did not make eye contact, and did not tilt my in an ostensive way, I would only be expressing an informative intention - and the waiter would likely never know that this is what I am doing. In short, if the expression of an informative intention is not accompanied by the expression of communicative intention too, then the intended audience would never know that they are the intended audience, and communication would not occur.

There are, then, two types of intention involved in this type of communication. One, the informative intention, is an intention that one's audience recognize *what* one is trying to communicate (for example, that I want more coffee). The other, the communicative intention, is an intention that one's audience recognizes *that* one is trying to communicate; in other words, that one has an informative intention. Another way to say this is to say that the content of a communicative intention is an informative intention. Signals that express communicative intentions, and hence informative intentions, are called *ostensive* signals.

Typically, both types of intention are expressed in one and the same behaviour, such as the tilt of the coffee cup. Here is another example (from Sperber, 2000). Mary is eating berries. She wants Peter to know that she thinks that the berries are very tasty, so she eats them in a somewhat exaggerated, stylized way, and pats her tummy as she does so. This reveals two things to Peter: (i) that Mary thinks the berries are tasty (this is the content of her informative intention); and (ii) that Mary wants to communicate this fact to Peter (this is the content of her communicative intention). If Mary simply ate the berries enthusiastically, but did not do so in any stylized or exaggerated way, Peter would still be able to infer that they are tasty, but not because Mary had expressed either an informative or a communicative intention. There would be no communication in that case. (I will define communication and other, related terms in the next chapter.)

Other cases illustrate that sometimes the content of the informative intention can be so transparent from the context that it does not need to be expressed by the signal itself. Then the signal has to do little more than make it apparent that communication is taking place i.e. it only has to express a communicative intention. The couple who communicate their views about a TV programme with a non-specific behaviour like puffed cheeks are an example. The puffed cheeks express the woman's communicative intention, but only very loosely express the content of her informative intention. This is because the context (the just concluded conversation,

combined with the memory of the couple's previous private conversations) does so much work that explicit expression of the content of the informative intention is unnecessary.

In this respect, this example is different to, say, the coffee shop example. In that context, an ostensive signal directed at the waiter could potentially have one of several possible meanings (e.g. 'Can I have the bill?', 'I'd like an extra slice of carrot cake', and so on). Consequently the signal has to not only express a customer's communicative intention, but also specify the content of the informative intention. It is this latter concern that explains why it takes the particular form that it does. A different ostensive signal (e.g. a mime of signing a cheque) would express a different informative intention. In contrast, the puffed cheeks could have been replaced by one of several other possible behaviours, so long as the behaviour was produced ostensively, and the meaning would have remained more-or-less the same - because the content of the informative intention was mostly derived from the context.

A common response when presented with this account of what is involved in communication is to suggest that it is overblown, or over-theorized. It seems to be a complex description of something that is instinctively much simpler. In particular, the embedding of an informative intention inside a communicative intention might be seen as theoretically unnecessary, since on an intuitive level an informative intention alone seems sufficient. I will address these complaints in detail in Chapter 3, but let me here make one short point: that just because a formal description of a phenomenon is complex, this does not mean that we necessarily experience it as complex. A formal description of what is involved in, say, vision, is a mighty complex thing, but our conscious experience makes it seem far more straightforward than that. So there is no *a priori* reason to think that the fact that there is a disparity between our instinctive, everyday experience of communication and this formal description is problematic.

We are now ready and able to set out an alternative to the code model of communication. It is called the *ostensive-inferential model*.

## 1.4 The ostensive-inferential model

Ostensive-inferential communication is the expression and recognition of informative and communicative intentions (Sperber & Wilson, 1995). This expression and recognition is achieved by the appropriate use of evidence: signallers provide evidence for their intentions, and receivers interpret it (*ibid.*). This evidence can come in many forms, such as points, shrugs, vocalizations (including language) and indeed potentially any physical act that the signaller can perform. The receiver must then take this evidence and draw an inference about the signaller's informative and communicative intentions, and hence about the meaning the signaller intends to convey. Signallers and receivers thus perform distinct but complementary jobs

in successful communication: the former provide evidence, and the latter interpret it. This is why the phrase 'ostensive-inferential' is double-barrelled: ostension is the provision of evidence; inference the interpretation. Nevertheless, 'ostensive-inferential' is a cumbersome term, and so from hereon I will often use the shortened terms *ostensive communication* and *ostensive model*, especially in later chapters.

This provision and interpretation of evidence contrasts with the associative processes of encoding and decoding employed in the code model. The puffed cheeks example illustrates the point most clearly. Suppose, quite reasonably, that the couple have never previously used that behaviour to communicate with one another - or even that they had never previously used it to communicate with anybody at all. If so, there is no way that it could be *associated* with the meaning the woman wishes to express. Furthermore, even if the puffed cheeks had become associated with something like lack of approval, this does not encode all that is relevant here: it does not, for one thing, encode what it is that is being disapproved. As such, the puffed cheeks cannot work as a piece of code model communication. Nevertheless, they could and indeed are likely to still be successful as a signal. This is because what such cases ultimately depend on is not the existence and use of any pre-existing associations, but rather the abilities of signallers and receivers to reason about each other's intentions and other mental states.

There is, then, a clear difference in the cognitive abilities required for ostensive-inferential communication, and those required for code model communication. The latter require the ability to form associations. This can be a simple task. Indeed, in its most basic form it is so simple that even bacteria can do it. Ostension and inference, on the other hand, are more complex. This is because the sort of evidence required for a particular intended meaning is specific to the current time, place, and intended audience, among other things (Sperber & Wilson, 1995). This means that to construct appropriate signals, signallers must take into account not just what it is they wish to communicate, but also their knowledge and beliefs about what their audience already knows or believes, given the current context. Receivers must do something similar, but in reverse: they must take into account not just the form that the signal takes, but also the context and their knowledge and beliefs about the signaller's general motivations and intentions (*ibid.*).

To illustrate, consider again the example of Mary eating berries in a way that makes it apparent to Peter that she thinks that the berries are tasty. Suppose now that Mary knows that Peter is curious about whether the berries are tasty or not, and Peter knows that Mary knows this. Now, Mary's intended meaning (that the berries are tasty) is the same as previously, but her knowledge of Peter's goals is different. Consequently, she produces a different behaviour: because she knows that Peter wants to know how the berries taste, all she does is nod approvingly; the exaggerated mime of pleasure that she performed in the previous scenario is unnecessary. As

for Peter, he knows that Mary knows that he is curious about the tastiness of the berries, and so he is able to interpret her nodding correctly - *but he would not be able to do this if he did not have the knowledge that Mary knows about his curiosity*. What this shows (and many similar examples could be created or observed) is that exactly what sort of evidence is appropriate for a given intended meaning is not fixed, or even probabilistic; instead, it depends upon a range of contextual factors - and these include, in particular, the prospective signaller's beliefs and knowledge about the beliefs and knowledge of the intended conversation partner. The same is true of interpretation as well. As such, successful ostensive-inferential communication depends upon abilities to reason about others' minds, beliefs, knowledge, and so on. In short, it is inherently *metapsychological*: it is only possible between individuals able to think about each other's thoughts, to reason about each other's reasons, and so on.

This is, then, the crucial difference between ostensive-inferential communication and code model communication. They are made possible by different internal mechanisms. That is, *the mechanisms that are causally responsible for the existence of each type of communication are different in each case*: associations on the one hand, metapsychology on the other. Put in different but equivalent terms, the difference between these two types of communication is that in one, the signaller encodes the content, while in the other she provides evidence that she has a communicative intention to express some specific content.

Correspondingly, while the evolutionary function of both types of signal is to cause changes in the receiver's behaviour, the more immediate function is to do this in quite different ways: the production of ostensive stimuli is designed to change mental states, while the production of encoded stimuli is designed to trigger associative responses. (This difference between evolutionary and more immediate functions is typically called the ultimate/proximate distinction. I shall say more about it in §1.6.)

Another way to think about the difference between code and ostension is that it is analogous to the difference between mathematics and science. In mathematics, the correct answer to a question is determined using logical operators (e.g. +, -, x, ÷, etc.), and as such when we use these operators we know, with certainty, that the answer arrived at is the correct answer to the question posed, so long as the inputs are correct and the operators have been correctly applied. In contrast, in science we can only amass evidence for the answer to a question, and so we can never *know*, with certainty, whether a proposed answer is correct. Instead we must make inferences about the most likely correct answer, given the evidence. Indeed, this is the very essence of what science is.

Similarly, because the code model is built on associations, then if a receiver assumes that these associations are correctly applied then she can fully (i.e. with 100 per cent certainty) determine what a particular signal means. (In the probabilistic cases, the receiver knows with certainty

what the probability is that a particular signal has a particular meaning.) However in ostensive-inferential communication the audience is only ever supplied with evidence for an intended meaning, and so the best she can do is make inferences based upon that evidence - and to make these inferences, she must reason, just as the scientist must, about both the data she receives (the form the signal takes), and the conditions under which it was observed (what she already knows about the signaller's intentions, beliefs, knowledge, and other mental states). Furthermore, the signaller must also create the right sort of evidence, and this too requires reasoning about the other party's mental states. As such, these acts of ostension and inference are analogous to, respectively, the production and interpretation of scientific data.

Is there a third type of communication, with a different foundation still? None presently exists ("There just is not to this day ...a third type of explanation of the very possibility of communication" (Origgi & Sperber, 2000, p. 149)). Unless and until philosophical research is able to generate a third account of how communication can even exist, these are our only two options. This is, incidentally, not an empirical question, but a conceptual/philosophical one. The question is: how is communication even possible? The code model and the ostensive-inferential model provide two possible answers to this question. No other answers presently exist.

### 1.5 Natural codes and conventional codes

Where does linguistic communication fit in to this distinction between these two different models of communication? The typical intuition is that it operates according to the code model. After all, there are clearly reliable associations between linguistic signals and their meanings: the word 'cat' is reliably associated with feline animals, for example. Many linguists consider this characterization at least sufficiently accurate to operate as a background assumption to their research (Reddy, 1979; Blackburn, 2007). Yet it is clearly not the whole story. We have already observed, in the opening passages to this chapter, how language is replete with metaphors, allusions and other figurative expressions that express far more than the literal, 'decoded' meanings of what is said. Furthermore, language is not the only piece of evidence that helps listeners determine a speaker's intended meaning. Other aspects of production, such as intonation and body language, are important too. For example, even an utterance as simple as "I'll see you later" can have a wide range of quite different speaker meanings, depending on the manner of its delivery, which could express, say, conviviality (such as from one friend to another, for example), a threatening attitude (from a bully to a victim), a prediction (from one work colleague to another), a request (from a suitor to his beau), or any one of numerous other dispositions. To determine between these readings, and indeed to use them

appropriately, we must reason about our conversational partner's mental states. Linguistic communication clearly involves some use of ostension and inference, in addition to the linguistic code.

There is, then, a question about the relationship between the linguistic code on the one hand, and the role of ostension and inference, and pragmatics more generally, on the other. One way to conceive of this relationship is as follows: the development of widely shared associations between signals and meanings - codes - is what makes language possible, and the human capacity for ostension and inference then makes it especially powerful i.e. able to express an extremely wide range of propositions. This conception of the role of pragmatics in communication is the dominant one in mainstream, contemporary linguistics, although it is rarely if ever stated quite so explicitly. It is reflected in, for example, the textbooks, encyclopaedia entries and other material designed to provide a basic statement of the discipline's domain of enquiry: in the vast majority of cases, those aspects of language that involve the code (in particular, but not only, syntax and semantics) are presented as the central, core concerns, while pragmatics and other phenomena that are not easily accountable in terms of a code are presented as peripheral topics. The implicit assumption is that languages are, when stripped to their most basic features, extremely rich coding schemes. Thus, the code lies at the heart of the standard conception of language. Ostension and inference are bonus add-ons.

Evolutionary approaches have, for the most part, also adopted this assumption, although here too it is almost never expressed explicitly. Instead, it is a background assumption, of which we can observe only the surface manifestations (Scott-Phillips & Kirby, 2013). One such manifestation is the language used to describe linguistic communication. Let me quote a passage from a seminal paper in the field. It is littered with the terminology of the code model in general, and information theory in particular: "the vocal-auditory *channel* has some desirable features as a medium of communication: it has a *high bandwidth*... however it is essentially a *serial interface* ...the basic tools of a *coding scheme* employing it are an *inventory* of distinguishable symbols and their *concatenation*" (Pinker & Bloom, 1990, p. 713, italics added).

A second manifestation of the assumption that languages are at their core very rich coding schemes is the research techniques used, and the questions that are pursued. For instance, there is a small industry of mathematical and computational models of the origins and evolution of signalling systems, which frequently extend their methods and claims to encompass language. Indeed, such models comprise a significant proportion of all research that purports to study language origins (see Kirby, 2002; Steels, 2011 for reviews of the computational models; and Skyrms, 2010 for detailed discussion of mathematical models). It is worthwhile to enter into a short discussion of these models to illustrate the general approach employed here.

A typical model frames communication as follows. There is a pre-defined set of possible states of the world. Which one actually applies at any given time is determined at random. This state then triggers the production of one member of a similarly pre-defined set of signals, which in turn triggers the production of one of a set of responses, also pre-defined. As such, the signaller's strategy is a set of mappings, from the states of the world to the signals they produce in each case; and the receiver's strategy a set of mappings from signals to responses.

The simplest possible model of this sort has two states (*A* and *B*), two signals (1 and 2), and two responses (*a* and *b*), with payoffs that are maximized if and only if state *A* leads to response *a* and state *B* to response *b*. If we then assume, for simplicity, that signallers will produce different signals for different states, and that receivers will produce different reactions for each different signal, then there are just two possible strategies for each player:

- signaller strategy 1 (S1): state *A*  $\rightarrow$  signal 1  
state *B*  $\rightarrow$  signal 2
- signaller strategy 2 (S2): state *A*  $\rightarrow$  signal 2  
state *B*  $\rightarrow$  signal 1
- receiver strategy 1 (R1): signal 1  $\rightarrow$  response *a*  
signal 2  $\rightarrow$ \* response *b*
- receiver strategy 2 (R2): signal 1  $\rightarrow$  response *b*  
signal 2  $\rightarrow$  response *a*

It is not hard to see that in such a game, the pairs  $\langle S1, R1 \rangle$  and  $\langle S2, R2 \rangle$  are equilibria. In these cases, state *A* leads to response *a*, and state *B* to response *b*, and hence payoffs are maximized. That is not the case with the pairs  $\langle S1, R2 \rangle$  and  $\langle S2, R1 \rangle$ . As such, if this set of strategies is subject to natural selection, one or the other of  $\langle S1, R1 \rangle$  and  $\langle S2, R2 \rangle$  will be the outcome (Skyrms, 2010). This basic game can be made more complex in numerous different ways, most obviously by manipulating the number of possible states, or signals, the probabilities of the different states, the payoffs associated with each outcome, and other component parts.

However, we need not enter into such details, because this basic sketch is sufficient to make the following, important point: that such models operate according to the code model. We know this because the associations that exist between states and signals, and between signals and responses, are associations that make communication *possible*; remove either of them, and there is no communication. This is the very definition of the code model. A great deal of research in language origins uses such models, to address topics such as the conditions under which shared codes can emerge. I am thinking here not only of mathematical models like the one above, but also

the many computational models that have for a long time been a staple of the discipline (see Kirby, 2002; Steels, 2011 for reviews).

Yet any assumption that codes are what make linguistic communication possible in the first place, and ostension and inference are what make it powerful, is wrong. It is upside down. Instead, ostension and inference make communication possible, and the development of widely shared associations between signals and meanings is what makes it linguistic, and hence expressively powerful (Sperber & Wilson, 1995; Origgì & Sperber, 2000). We know this because if we consider what is actually entailed by these two alternative accounts of how linguistic communication works, it soon becomes apparent that only the ostension-makes-it-possible-and-codes-make-it-powerful account is plausible. The codes-make-it-possible-and-ostension-makes-it-powerful account is not.

To see why, recall the notion of underdeterminacy: the fact that linguistic meaning underdetermines speaker meaning (Carston, 2002a; Atlas, 2005). The existence of underdeterminacy is widely acknowledged. The most obvious examples are referring expression (e.g. pronouns, demonstratives), and figurative uses of language, such as sarcasm and irony, but there are a great many others, such as understatement, similes, and indirectness, and I gave some examples of these in the opening to this chapter. What is less appreciated is how far this underdeterminacy extends. We have already seen how even an apparently transparent utterance such as “I’ll see you later” can have multiple different interpretations, and further examples are not difficult to provide. “Higher!”, “It’s raining”, “Yes, please”: all are open to multiple interpretations, depending on context. In other words, the literal meaning of each of these is more-or-less clear, but this is insufficient to determine the speaker’s intended meaning. In fact, when we begin to study linguistic communication in this light, what we see is that there is underdeterminacy in *any and every* utterance you care to produce (Carston, 2002a; Atlas, 2005).

The problem is not the tractable one that literal meanings may correspond to more than one of a still finite number of speaker meanings, and that we must choose between them. It is far more serious than that. For any literal meaning, there is an *infinite* range of possible speaker meanings. To illustrate, consider again “It’s raining”. Two possible speaker meanings are, as I mentioned earlier, ‘I don’t want to go out after all’ and ‘Get the umbrella’. But here is a third: ‘Even the weather can’t lighten my mood’. And a fourth: ‘I told you so!’ I could go on, and that is the point. Here, then, is the challenge for the skeptic who doubts the ubiquity of underdeterminacy: describe an utterance and a corresponding list of possible speaker meanings that is *wholly* exhaustive, in the sense that there are no other possible speaker meanings whatsoever. This task is, I submit, impossible: given an utterance with literal meaning, and a list of possible speaker meanings that the literal meaning might be used to express, it is *always* possible to add another possible speaker meaning to the list.



One consequence of this fact is that the linguistic code cannot make linguistic communication possible. The existence and ubiquity of underdeterminacy means that the meaning that a speaker intends when they produce an utterance cannot ever be determined by analysis of the literal meaning of the utterance alone. In short: linguistic utterances are, on their own, unable to specify speaker meaning, and hence they logically cannot make any sort of communication possible.

To put the point another way: *as codes*, languages are very defective indeed. In fact, they are wholly ineffectual (Origi & Sperber, 2000). Here is a simple illustration. Consider the following dialogue:

MARY: Would you like to join us for dinner?

PETER: I ate earlier.

The important thing to notice about Peter's response is that it does not, *on its own*, tell Mary anything at all that is relevant to her enquiry. It is only with the addition of context and shared knowledge (for instance, the knowledge that people do not eat two meals in one evening) that Peter's meaning becomes clear. The linguistic code is insufficient for communication to take place, but communication will still succeed - so it cannot be the linguistic code that makes linguistic communication possible. If Hymie the humanoid robot really could only process literal meaning, then it would not be only metaphors and similes that caused him problems. Every utterance would.

On the other hand, ostension and inference do make possible a form of communication onto which we can add linguistic codes, and in doing so provide our ostensive communication with a great deal of expressive power. We have already seen several examples of how ostension and inference can be used to communicate without language. All could be augmented by the use of sounds and/or gestures that are reliably, typically associated with certain referents. The phrase 'More coffee, please' could be said as I tilt my coffee cup; the word 'Idiot!' could be added to the puffed cheeks; Mary could say 'Yum!' to emphasize how tasty the berries are. These codes provide a way for us to say what we would otherwise have to show, and the consequence is that ostensive-inferential communication can be used far more precisely, and more expressively, than it otherwise would be (Wharton, 2003b). In this sense, linguistic communication is simply a very important special case of ostensive communication, one in which that ostensive communication is augmented by the linguistic code. In other words, ostension and inference are logically prior to the linguistic code.

In fact, the ambiguity that makes languages such ineffectual codes here becomes a huge asset, because it makes it possible for individual words to be efficiently used to express a diverse - in fact infinite - range of meanings, and in an efficient way too (Piantadosi et al., 2012). There is, then, a critical

difference between the codes used in code model communication, and the codes used in language. The former makes a type of communication *possible*, albeit one that is restricted to a finite set of more-or-less fixed meanings. Consequently, whatever ambiguity there is in the system is limited to a finite set of alternatives, some of which will be more probable than others. Linguistic codes, however, are quite different. They do not make communication possible - because on their own, they have an unlimited range of things they might refer to, and hence they create unresolvable ambiguities. Instead, they are added to a foundation of ostension and inference, and in doing so they make that type of communication incredibly *powerful*. I can nonverbally but ostensively point to any of the objects in this room, but with language I can refer to any object in the world. I can make a request of others by pushing unchopped vegetables in their direction, but with language I can make requests about things remote in time and space. Languages make ostensive communication special.

In short, there is a qualitative difference between the codes used in the code model, and the linguistic code. Put simply, one makes a type of communication possible, the other makes a different type of communication expressively powerful. (In philosophical terms, they have different ontologies.) This difference is widely unappreciated, not only in evolutionary approaches to language, but in fact in the study of language and communication more generally.

Consequently, we invite misunderstanding if we use the same term, 'code', to describe both the linguistic code, which forms part of an ostensive-inferential communication system, and the code used in communication that is based upon the code model. To call both simply 'code' invites the misplaced assumption that these are two instances of the same sort of communication, which differ only in their degree of complexity. Instead, our terminology should recognize the qualitative difference between them.

Evolution is one area in which a failure to make this distinction can have serious consequences. An important evolutionary question is the extent to which various instances of animal communication are similar, or not, to linguistic communication, but we risk inadvertently pre-judging the matter if we use the same label to describe both. This is not to say that no other species communicates in an ostensive way. Whether they do or not is an empirical question, to which I will turn in Chapter 4. The point is instead that we should not pre-judge the matter by using the same term ('code') to describe both animal and linguistic communication. A failure to do this can lead, in turn, to a presumption that it is parsimonious to assume that the two types of code are evolutionarily related to one another: "we may see in [monkey] alarm calls a skeletal version of our own shared codes" (Hurford, 2007, p. 260). Now, this claim may or may not be true. I shall argue in later chapters that it probably is not, but that is not the point here. Instead, what I want to emphasize is that we cannot simply assume that the essential difference between the linguistic code and the code used in

non-human primate communication is a difference of degree, even a large degree. It may instead be a difference of *kind*. And besides, there is certainly a difference of kind somewhere: much non-human communication, such as bacterial quorum sensing, does operate according to the code model, whereas linguistic communication does not. Our terminology should reflect this difference.

I will do this with the labels *natural codes* and *conventional codes* (adopted and partially adapted from Wharton, 2003a). Natural codes are those we observe in systems that operate according to the code model itself. Bacterial quorum sensing, mother-infant chemosignals, grasshopper sounds, gesture-calls, and the various other examples I discussed in §1.2 are all natural codes. In short, natural codes make code model communication possible. In contrast, conventional codes are those codes that enhance and make more expressive and precise an already existing system of ostensive-inferential communication. Languages are conventional codes, but there are other conventional codes too, such as Morse code. I will often simply use the term *conventions* as shorthand for conventional codes. The question for comparative research is: are the codes used in non-human communication natural codes, or conventional codes?

In the UK, one convention is that we drive on the left hand side of the road. We could in principle change this to the right side of the road, if we so wished, so long as we *all* changed our driving behaviour, and at the same time as each other. (Both Sweden, in 1967, and Iceland, in 1968, have done this, changing all the road signs over during the middle of one night.) In the specific context of communication, a conventional code is a reliable association between signal and meaning that holds by virtue of the fact that every member of the community agrees that it holds, and not because of any intrinsic association between the word and the meaning (Lewis, 1969; Clark, 1996). For example, 'dog' refers to canine animals (and not, say, feline ones) only because we all agree that it does. We could in principle all agree to change the word for 'cat' to something else, such as 'tac' (the same sounds, in reverse), and in theory this would not cause any confusion - so long as we all did this, and at the same time as one another.

Conventional codes exist at all levels of linguistic analysis: semantic (the meaning of words), syntactic (the grammar of a language); phonological (the organization of sounds in a language), phonetic (the sounds that are used in a language), and even pragmatic (e.g. in British English, the politeness convention to use indirect requests such as 'Would you mind...', rather than direct requests, wherever possible). It is not hard to see how the expressive potential of ostensive communication explodes with the development of these conventional codes.

Ask 100 linguists what a language is and you'll get 120 different answers. Depending on factors such as their theoretical persuasions, their disciplinary background, and their specific areas of interest, they will emphasize or downplay different aspects of it, such as universal properties of

languages, the cognitive and psychological foundations of language, its role in thought, its use as a means of communication, and so on. My answer is this: a language is the rich, structured collection of conventional codes that augment ostensive-inferential communication within a given community. Tilting your coffee cup is ostensive but not linguistic, since the tilt does not (really) have a conventionalized meaning, but the words 'Can I have some more coffee please?', and the gestures used in sign languages, both do have conventionalized meanings, and as such are both linguistic. This view of a language as a set of conventional codes that augments ostensive communication recognizes both the pragmatic foundations of linguistic behaviour, and the importance and nature of the conventions that make languages different to other, simpler cases of ostensive-inferential communication, such as points, non-linguistic vocalizations, nods of the head, and so on.

I do not want to enter into a lengthy review of the similarities and differences between this and other views of what a language is, but it is worthwhile to briefly contrast my view with views that focus on structural features. By way of illustration, consider the controversy around Pirahã, a language spoken by an indigenous tribe of hunter-gatherers in the Amazon basin. Some researchers have claimed that Pirahã does not possess *recursion*<sup>2</sup>, a property that other researchers have suggested is critical to language (Everett, 2005). (A language is recursive if some of its linguistic objects are embedded inside versions of the same type. For example, *I like the woman with the hat* has one noun phrase (*the hat*) embedded inside a larger noun phrase (*the woman with the hat*.) Others dispute the claim that Pirahã lacks recursion (e.g. Nevins et al., 2009). What I want to bring attention to is that whichever way this debate is resolved, nobody claims or will claim that Pirahã is not a language. In other words, even if it does lack recursion, that would not be enough, on anybody's terms, to conclude that Pirahã is not a language. Thus, recursion cannot be what makes a language a language. The same goes for all the various structural properties that languages possess: none of them define what a language is, because if we found a language that lacked them, we would not conclude that the observed language is not a language. Indeed, on some analyses there are *no* structural properties that are universal to all languages (Evans & Levinson, 2009).

<sup>2</sup> In that case, any attempt to define what a language is by reference to such properties is doomed. What *is* universal, however, is the fact that languages are rich, more-or-less stable collections of conventional codes that are used to augment ostensive-inferential communication.

Languages are not, however, the only way in which ostension and code can interact with one another. There are other ways too. In production, we can use natural codes in an ostensive way. An example is voluntary laughter. As discussed above, involuntary laughter is a natural code, but we can also exercise voluntary control over laughter (called non-Duchenne laughter), and we often do so in an ostensive way. One way of expressing sarcasm, for example, would be to laugh at a friend's supposedly funny joke

in a highly- and possibly overly-stylized way that mimics genuine laughter, but which at the same time is obviously not genuine laughter. Such behaviour is ostensive: it expresses my intention that my friend knows that I do not think his joke was funny. The other gesture-calls provide similar examples: snorts, smiles, frowns, sighs, pouts and others can all be used in an ostensive way, in addition to the roles they play as natural codes (Wharton, 2003a).

What examples like this show is that, from an evolutionary perspective, the difference between natural codes and conventional codes is critical. What happens when we combine pragmatics and natural codes is quite different to what happens when we combine pragmatics and conventional codes. In the former case we get the use of natural codes, such as laughter, grunts, and so on, in ostensive ways; in the latter case, we get language.

This difference has a number of important implications for the study of the origins of language. The first is that it simply makes clear exactly what it is that we must explain, namely the origins and evolution of the capacity for ostensive communication, and the emergence and continued use of stable, conventional codes. The second implication is that research into the origins and evolution of codes is only of direct relevance to the origins and evolution of language to the extent that those codes being studied are indeed conventional codes, and not natural codes. Since language is a conventional code, then to study the evolution of natural codes and to use the results of that research to draw direct conclusions about the evolution of conventional codes is a category error. One area where this has occurred is comparative research on the communication systems of at least some non-human primates. I will discuss this research, and the important question of whether non-human primate communication is ostensive or not, in Chapter 4.

If there is one point that underpins everything else in this book, it is the one I have made in this section, namely that the common assumption that the linguistic code makes linguistic communication possible is simply false. Instead, linguistic communication is a type of ostensive-inferential communication, made possible by metapsychology. What the linguistic code does is make it expressively powerful. This insight is certainly not original to me. It has been most clearly developed by Sperber and Wilson, but versions of it run through the whole history of the philosophy of language. What I will do in the rest of this book is to describe the implications that it has for the origins and evolution of human communication and language.

## 1.6 Two meanings of meaning

There is one final preliminary topic that must be addressed before we proceed. It is the thorny matter of *meaning*, something that is legendarily the stuff of abstruse philosophical enquiry. It is a central concept for

communication, yet its connotations differ depending on whether it is used within the context of the code model or the ostensive-inferential model of communication. In this section I describe and defend two ways in which the term can be coherently used. One, with its roots in pragmatics, applies to the ostensive-inferential model. The other, with its roots in evolutionary biology, applies to the code model. We shall see that despite several differences, there are also key features that these two approaches to meaning have in common, and these tell us something important about the nature of communication itself.

Let me start with meaning in pragmatics. One of the founding documents of pragmatics as a discipline is an essay, entitled simply 'Meaning', by the philosopher of language Paul Grice (1957). In it, Grice distinguished between natural and non-natural meaning. He did not give formal definitions of these two terms; rather, he illustrated them by way of examples. *Natural meaning* is meaning that derives from a reliable relationship between one thing in the world and another: those clouds 'mean' rain; those spots 'mean' measles. In contrast, *non-natural meaning* is, roughly, the meaning that we have so far referred to as speaker meaning or intended meaning: Mary 'means' that the berries are tasty; I 'mean' that I would like more coffee. It was this idea of non-natural meaning that Grice developed further; he said little more about natural meaning. Incidentally, the terminological link between natural meaning and natural codes is not coincidental: natural codes are what we get when things with natural meaning are used in code model communication (Wharton, 2003a).

Grice's elaboration of non-natural meaning proceeds in three stages. The first is to note that for me to mean something in this sense I must intend to manipulate or add to my audience's mental representations in some way. For example, perhaps I intend to create in my audience the representation that I would like more wine, or that these vegetables need chopping. The second stage is to note that my audience must recognize that I have these intentions. Suppose that I ask my friend to pour me another glass of wine and that, as I expected, he obliges. I expected this not simply because I had asked for more wine, but in fact because, as a result of my asking, my friend *understood* that I wanted another glass of wine. In other words, my request was directed towards the manipulation of my friend's mental states, and he poured me another glass of wine because he recognized this. If I had requested more wine and my friend had not heard me, but poured the wine anyway, simply because he is a good host, then my intentions would *not* have been satisfied.

These two conditions might seem sufficient to say that I 'meant' for my friend to pour me some wine, but, Grice argued, it is not. To see why, suppose that, instead of asking for more directly, I simply move my empty wine glass to a location where it is likely to be noticed by my friend, but I do not bring attention to the fact that the glass is empty in any way. Instead, I simply leave it there. Nevertheless, my friend sees that I have done this,

and understands why I did so. Here, the same two facts as above apply - I still intended that my friend believes that I would like more wine, and my friend still recognized this - but Grice was not willing to grant that there is non-natural meaning in such cases. He argued, and I agree, that "deliberately and openly letting someone know", or "getting someone to think" is not quite enough for (non-natural) meaning (all quotes p. 382). However if, on the other hand, I had made eye contact with my host and simultaneously tilted my wine glass, or expressed my intentions in some other conspicuous way, then we would certainly say that I 'meant' that I would like more wine. The difference between these two cases is this: only in the latter do I *show* the fact that I intended that my friend believes that I would like more wine. Then, *as a result of this showing*, my friend understands what I meant. This is the third stage, and it is the heart of Grice's account of meaning: to mean something, I should intend that my audience believes it, *and they should believe it at least in part because they recognize that this was my very intention*. This is the meaning of (Gricean, non-natural) meaning. It is also the reason why ostensive-inferential communication can be glossed as *intentionally overt* communication.

This account has been much discussed, and in some cases reformulated (e.g. Strawson, 1964; Schiffer, 1972; Neale, 1992). However, in all reformulations one fact remains central, namely the auto-deictic character of the Gricean account; in other words, the fact that ostensive stimuli are effectively pointers to the very intentions that triggered their production in the first place. Indeed, this is arguably what makes an analysis 'Gricean' in the first place: analyses that include this quality are Gricean in spirit, even if they are not in letter; and those that leave it out are less reformulations of Grice, and more fundamental revisions. Perhaps appropriately, given its philosophical origins, this pragmatic notion of meaning is clear and well-developed.

In contrast, in evolutionary biology, and specifically in animal signalling theory, the term 'meaning' has been used in a more casual way, not subject to the detailed analysis that the Gricean notion has been. Instead, it has mostly been used in a colloquial, mostly anthropomorphic way i.e. *as //animals have the same intentions as humans do*. A related fact is that the vast majority of research on animal communication (perhaps all of it) implicitly adopts a code model of communication - not because there has been extensive consideration of whether the code model or the ostensive-inferential model is more appropriate, but rather because the model is very intuitive, and works well as a description of most (and perhaps all) animal communication. Indeed, I suspect that many animal communication researchers are not aware of the fact that they have adopted this model, or that there is an alternative.

The set of background assumptions that are embedded into the code model are rarely problematic for the study of non-human communication. However, they can become so when we wish to view human communication,

which does not operate according to the code model (§1.5), through the lens of evolutionary biology, as we do when studying language origins. There is here the potential for misunderstanding and confusion. In order to avoid this outcome, some researchers think that we should avoid the term 'meaning' altogether, unless we do explicitly want to invoke the Gricean sense described above (e.g. Rendall et al., 2009). While sympathetic to the agenda behind this suggestion, I think it is too restrictive. Here I will try to explain how meaning can be profitably used within the code model. I will then discuss how that usage relates to the Gricean idea of (non-natural) meaning discussed above.

The first thing we must do is recognize an important distinction, central to evolutionary theory, between what are called ultimate and proximate explanations of biological phenomena (Scott-Phillips et al., 2011). Biological evolution is change in gene frequencies (Futuyma, 2005). There are two classes of question we can ask about these changes: why do they occur?; and what do the genes that do exist actually do (i.e. what proteins and hence organs do they create, how do those organs work, what behaviours do they lead to, etc.)? Answers to the first question are called ultimate explanations; answers to the second question, proximate explanations. In short, ultimate explanations are concerned with why a behaviour exists, and are correspondingly expressed in terms of the function of the behaviour in question (or the lack of any function); proximate explanations are concerned with how it works, and are correspondingly expressed in terms of the various causal mechanisms (be they physical, physiological, psychological, social, chemical, etc.) involved (Mayr, 1961; Scott-Phillips et al., 2011).

For example, an ultimate explanation of infant crying is that it elicits care and defence from mothers and other caregivers - this is the function of infant crying. Proximate explanations include the external triggers of crying, such as physical separation from the caregiver, cold, or a lack of food, and also the internal mechanisms, such as the limbic system and the endogenous opioids involved in the cessation of crying. As this example makes clear, these two different types of explanation are distinct and complementary, and neither is more important than the other. (As such, the term 'ultimate' may be unfortunate, since it might suggest that this type of explanation is the more important of the two. This is not the case, but the term is well-established, and as such it would be counter-productive to suggest an alternative.) The ultimate/proximate distinction is central to any evolutionary analysis of behaviour (Davies et al., 2012). This is because it is not simply a heuristic that researchers can choose to adopt, or not, as a way to organize their enquiries, but instead a logical consequence of the very definition of biological evolution (see above).

The headline point about meaning and animal communication is that here, 'meaning' is an ultimate level explanation of the signalling behaviour, applicable across all different types of communication system. In



other words, when we talk about the meaning of an animal signal, we are making a statement about its ultimate function (Krebs & Dawkins, 1984; Scott-Phillips, 2008a). (In contrast, the Gricean account is a proximate explanation, applicable only to systems that operate with the mechanisms of metapsychology described above.) For example, the function of mating calls is to cause in others a willingness to mate, the function of marking territory is to cause others to not encroach upon that territory, and so on and so on. As such, when animal communication researchers use the term 'meaning', they typically do so to describe the command that would trigger these effects, *as if* this were an instance of linguistic communication. In other words, mating calls 'mean' 'come and mate with me', territory marking 'means' 'do not encroach upon this territory', and so on. This is not so much a theoretically developed notion of meaning, but rather a description of how the term seems to be used in practice.

However, when such research is conducted or interpreted through the lens of comparison with human language, the danger of misunderstanding is significant. This is not only due to the fact, already mentioned, that the evolutionary approach to meaning differs from the pragmatic one described above. It is also possible because with such interdisciplinary research a significant proportion of the intended audience is likely more familiar with other approaches to meaning (in particular the Gricean approach discussed above), and may hence interpret claims about the meanings of animal signals in terms quite differently to those intended. For example, when a researcher states that, say, a monkey call 'means' 'eagle!', she may only have in mind the idea that the ultimate function of the call is to elicit a response that is appropriate for the presence of eagles. She may not mean to suggest that monkey communication is ostensive, and hence that the monkey in some sense intends to change the mental representations of other monkeys, but there is however a danger that some audiences will interpret her comments in this way. Equally, it is also possible that she *does* think that monkey vocal calls are Gricean, but this too may not be clear. In short, there is significant potential for ambiguity and misunderstanding here. At the same time, both the Gricean approach to meaning, and this more intuitive, coded sense of the term, are coherent and useful. As such, I will employ them both, and where there is any ambiguity about which is being used, I will state which I have in mind.

These two approaches to meaning have at least one important feature in common, which is this: that both are, at bottom, about how signals *do* things to an audience, and indeed how they are *designed* to do so (Scott-Phillips, 2010a). In one case, what signals do is change behaviour, and the design comes by virtue of natural selection, which produces organisms that behave in goal-directed ways (Dennett, 1995; Gardner, 2009). In the other case, what signals do is change mental states, and the design comes by virtue of human intentions (notwithstanding the fact that the capacity for this is of course itself a product of natural selection). Indeed, one of the

seminal texts in the history of pragmatics is entitled *How to Do Things with Words*, with the emphasis very much on the *Do* (Austin, 1955). In sum, then, 'meaning', in either sense of the word, is about doing things to others in a designed way. Of course, the broader idea that meaning derives ultimately from how it is used in the world has a long philosophical history that pre-dates pragmatics, most famously in the work of Ludwig Wittgenstein (1953). In the next chapter we shall see how the fact that the function of signals is to do things to others has profound consequences for how communication systems emerge in the first place.

\* \* \*

*Chapter 1. There are two models of the very possibility of communication: the code model and the ostensive-inferential model. The former is built upon mechanisms of association; the latter upon mechanisms of metapsychology. Most and perhaps all animal communication can be described with the code model. It is commonly assumed that linguistic communication operates according to the code model, but it does not. It is instead an instance of ostensive-inferential communication that is augmented by a set of shared communicative conventions. For this reason, it is important to distinguish between these conventional codes, which augment ostensive-inferential communication, and natural codes, which make code model communication possible. The notion of meaning plays out differently in each of the two models of communication, but in both cases it is about doing things to others in a functional, designed way.*

*Chapter 2. What does animal communication theory tell us about how communication systems emerge? Do ostensive-inferential systems emerge any differently to code model systems? Why is human language the only communication system in the natural world that makes widespread use of meaningful combinatorial communication? Did ostensive-inferential communication evolve from coded communication, or from some other source?*