

Grades of communication

Ronald de Sousa

16.1 Introduction

The chapters in this volume illustrate the astounding variety and pervasiveness of communication in the living world. From individual alleles through bacteria and social species to human speech, communication is everywhere. The means of communication also include all channels, chemical, electrical, visual, tactile, auditory, and olfactory, whereby organisms acquire any sort of information at all. For some purposes, one might want to limit the relevant sense of ‘communication’ to transactions involving conspecifics: so far, when we write articles and books, we intend them only for human readers. But that is only a special case: we also need to take account of cases where individuals of one species affect the behaviour of another in such a way as to affect their own reproductive success and thereby their own genome. In such cases, the individuals of the second species are part of what Dawkins has called ‘extended phenotypes’ of the first (see also Chapter 10). Typical examples involve parasites that manipulate their hosts, or predator and prey that have co-evolved in the course of an arms race. But a concept may lose some of its usefulness if its application is excessively broad. True extended phenotypes must be carefully distinguished from the endless variety of mere effects that individuals of one species can have on another, without being reflected in the former’s genome (Dawkins 1982, 2004). We need to be able to say what, in the interactions of cells, organs, or individuals, is *not* communication. What exactly, then, do all those phenomena have in common which may legitimately fall under the concept of ‘communication’?

Let me start with the obvious. Where there is communication, there must be a sender *S*, communicating something, *X*, to a targeted receiver, *R*. This simple truism raises several questions: first, about the potential partners *R* and *S*; second, about the nature of *X*; and third, about the constraints on and the point of the whole process.

On the first question, our initial intuition might be that both *S* and *R* must be individuals. They need to be sufficiently separate for one to be in possession of information that the other lacks. The concept of a biological individual is a complex and elusive one, however (de Sousa 2005), and it is soon apparent that *S* and *R* don’t have to be individuals in any strong sense. As several of the chapters in the present volume make clear, the entities engaged in communication can be part of a kind of ‘super-organism’, such as colonies of ants or bees (see also Chapters 2, 7, and 10) and a surprising number of cases of *intra-organismic* **conflict** and cooperation testify to the fact that parts of a single organism, or even of a single gene (see Chapters 12 and 13), can communicate among themselves to cooperate or compete.

On the second question, we can start by assuming that what *X* stands for at its most general is *information*; but the meaning of that word is in need of elucidation. What has come to be known as ‘Shannon information’ (see Box 16.1) affords a useful measure of *quantity* of information (Rheingold 2000, Chapter 6), but it is, as we shall see, notoriously inadequate for providing an assessment of its *quality* in terms of what communicators might be interested in.

As to the third questions, we should ask whether information counts as communication any time

some X is transmitted from S to R , or whether further constraints should be imposed if the concept is to be of any interest. To understand the point of communication we should require, for example, that the participants in a process of communication derive some benefit from the process. But if so, must benefit accrue to both the sender and the receiver, or only to the sender? And what further conditions apply?

Once the issue of benefit is raised, the issues of reliability, error, and deception loom large. Language, in a famous quip attributed to Talleyrand, was given to the human race in order to enable us to conceal our thoughts. As if to echo that saying, Crespi (Chapter 13) notes that we use language to 'manipulate the thoughts of others'. And it is obvious from surveying instances of non-human mimicry in nature that it is not only humans whose messages are not invariably veridical. Intelligence, it is sometimes said, is an arsenal of weapons in an arms race, an essentially Machiavellian tool (Dunbar 1993), the real point of which is rarely the simple conveying of information, but rather the manipulation of others' responses.

This last observation challenges us to understand how, in the case of organisms to which we

are not tempted to ascribe conscious intentionality, we might make sense of the application of the idea of 'manipulation': is it a mere metaphor or does it have literal application? What suggests itself here is that we need to allow for different grades of **signalling**, corresponding to different positions on a continuum of degrees of intentionality. At one end there may be a kind of signal that we can, without metaphor, ascribe to unicellular organisms, or even to their component parts. At the other end, there will be the fully fledged intentional communication of explicit language. The former offer examples of special purpose *functions*, while the latter have collected, over the long stretch of evolution, increasingly sophisticated devices serving an unbounded variety of individual purposes. Many of the special tricks of language are doubtless unique to it; yet we should not assume that we have lost the non-verbal aspects that marked the communication styles of our simpler ancestors. On the contrary, they may remain to constitute a mainstay of the pragmatics of language, with connotation, innuendo, irony, and other figures of speech and exploitation of context responsible for shifts or elaborations of meaning.

In Section 16.2 I begin with an outline of an influential account of 'non-natural meaning'; as it

Box 16.1 Shannon information

A quantitative measure of information introduced by Claude Shannon (1948). Intuitively, if a message is considered as a series of random variables selected from a finite set, the information afforded by each variable is measured by the number of yes/no questions that must be answered to guarantee that the value has been identified. Hence the *Shannon entropy* of a message is the minimum average message length, in binary units or 'bits' (using base-2 logarithms), that must be sent to communicate the true value of the random variable to a recipient.

More formally, the information entropy of a discrete random variable X , that can take on possible values $\{x_1 \dots x_n\}$ is

$$\begin{aligned} H(X) &= E(I(X)) \\ &= \sum_{i=1}^n p(x_i) \log_2(1/p(x_i)) \\ &= -\sum_{i=1}^n p(x_i) \log_2 p(x_i) \end{aligned}$$

where $I(X)$ is the information content or self-information of X , which is itself a random variable, and $p(x_i) = \Pr(X = x_i)$ is the probability mass function of X .

Equivalently, the Shannon entropy is a measure of the average information content the recipient is missing when he does not know the value of the random variable.

Partially adapted from the article on 'Shannon entropy' in Wikipedia, at http://en.wikipedia.org/wiki/Shannon_information.

applies to human intentional communication, followed, in Section 16.3, by a brief summary of the best current philosophical solution to the problem of how to analyse teleology without intention. In Section 16.4, I give a very informal sketch of the notion of Shannon information, and explain why it is both indispensable and insufficient for a satisfactory account of communication. In Section 16.5, I address some of the conditions under which it is useful to describe a given transmission of information as a reliable signal, in terms of the different ways in which the process has a distinctive biological function. Finally, in a short conclusion, Section 16.6, I raise a couple of tentative and speculative questions.

16.2 Natural and non-natural meaning

In an influential 1957 article, the philosopher Paul Grice proposed an analysis of the notion of ‘non-natural meaning’ (Grice 1989). He started by noting a number of contrasts between (1) ‘Those spots meant measles’ and (2) ‘Those three bells meant that the bus is full’. In the case of (2), but not (1), one can consistently infer that *someone meant to convey* something. One can speak of the *content* of what is conveyed in (2), and place it between quotation marks, as in ‘three bells meant “the bus is full”’. But it would be nonsensical to say ‘those spots meant “measles”’, as we might say ‘In French, “varicelle” means “measles”’. Furthermore, in the case of (2) one could go on to say: ‘but it was a mistake, as the bus wasn’t full’, but it would seem odd to add to (1), ‘but it was a mistake, as it was not measles’. (The mistake in (1), we might say, is made not by the sender but by the receiver.) Grice referred to the meaning in (1) as ‘natural meaning’, and to the kind of meaning alluded to in (2) as ‘non-natural meaning’. The former involves inferences from perceptions of facts or events to correlated facts or events, including causal antecedents or consequences. The latter, on the other hand, involves an intention to communicate. In human language, that intention is crucial, and in later work by Grice and others the analysis of the role of intention reached truly daunting levels of sophistication and complexity. A representative

sample of an intermediate level is the following definition (Grice 1989, pp. 99–100):

‘U meant something by uttering x’ is true if and only if (for some A and for some r):

- (a) U uttered x intending
 - (1) A to produce r
 - (2) A to think U to intend (1)
 - (3) A’s fulfillment of (1) to be based on A’s fulfillment of (2)
- (b) there is no inference-element E such that U uttered x intending both
 - (1’) that A’s determination of r should rely on E and
 - (2’) that A should think U to intend that (1’) be false.

Although further analysis found this characterization inadequate (Schiffer 1972), the example is sufficient to give some idea of the complex nexus of intentions involved in an unadorned case of someone meaning something by an utterance.

Quite obviously this isn’t the sort of thing that can be involved when we speak of communication among bacteria, genes, or neurons. Nor can it account for birds or mammals ‘signalling’ to mates or predators. If we are to speak meaningfully of information and communication among organisms not suspected of being capable of formulating conscious intentions we need to find a way of cashing out what in such contexts can only appear as metaphors.

To do this requires two closely related but importantly different tasks. The first is to explain how teleology can be brought under the aegis of ordinary causality. The second is to show that a concept of teleology thus explicated is adequate to provide a theory of communication, spanning both the most basic types and the more elaborate forms of conscious intentionality.

16.3 Objective teleology

In the past 50 years philosophers of biology have successfully accomplished the first task. The work of Larry Wright (1973), refined and elaborated by many others and especially Ruth Millikan (1984, 1993), has established that the core concept of function or teleology can indeed be explicated in terms of an *aetiological schema*, applicable equally well, with minor adjustments, to the functions of human actions (including communicative acts such as

uttering a meaningful sentence), of artefacts, and of biological organs. The aetiological schema is so named after the Greek word for 'cause', and its aim is to reduce teleological notions such as purpose, goal, or function to purely causal notions (Box 16.2). The idea is that the function of an act, object, organ or token of behaviour is distinguished among its many potential effects as the one that *causally explains its presence*.

In the case of an intentionally produced object or behaviour, this explanation typically refers to an intentional agent's pre-existing plan or purpose. In the case of an organism or part of an organism, where the functionality in question cannot be attributed to any conscious intention, biologists tend to be somewhat casual about cashing out talk of function or purpose. This is perfectly sensible, given the heuristic fruitfulness of talk of purpose and design in nature. Given a piece of anatomy or a mode of behaviour, the first thing to ask about it is what it is *for*, even though we are clearly not expected to infer that any actual purpose or design is involved. But talk of 'purpose' or 'design' needs to be cashed out. To effect such a cashing out by reducing it to causal terms is the point of the aetiological analysis.

In regard to non-intentional teleology, the aetiological analysis adverts to the effect causally responsible for the relative reproductive success of some lineages over others. The vertebrate heart, among other effects, both produces rhythmic sounds and circulates the blood. To say that the latter is its function, while the former is not, is to say that present-day vertebrates have hearts as a result of the advantage afforded in ancestral hearts by the circulation of the blood, and not by any advantage conferred by rhythmic sounds. It is important to note that this analysis does not require that everything that serves the interests of a particular organism on a specific occasion must be held to be an adaptation. As Sober (1984) has shown, the factors that confer selective advantage on a type of organism are those that the organism is selected *for*, but there will be many others that happen to go along with that without in themselves conferring a **fitness** advantage. These last are *selected*, but not *selected for*. What was previously selected without being selected for—as well as traits that merely resulted from random genetic drift—can become functional if, in changed circumstances, they have effects that confer new fitness advantage, and thereafter begin to be preserved by **natural selection**.

Box 16.2 Objective teleology and the aetiological analysis of function

There are two principal varieties of teleology: *goal* or *purpose*, and *function*. One can say of a tool that it has a function rather than a goal, but it was with the goal of serving such a function that the tool was designed. A goal, then, will commonly be a certain state of affairs, while a function will more likely be identified with a specific means of achieving that state of affairs. Where intentional actions and artefacts are concerned, functions are relative to the goals and interests of agents, and goals differ from one agent to another. But in biology, there are no agents, and so no real goals. We can, however, identify *replication* as a metaphorical 'goal' of nature. Whether a particular gene or set of genes is or is not more successful than another is a matter of fact which—however

difficult it might be to ascertain—is not relative to any agents or interests. We can therefore identify purely objective teleological properties such as functions providing we analyse them in strictly causal terms. That is the point of the aetiological analysis (AA) which explicates the intuition that the function *F* of an element *X* (an organ or part or an organism) can be identified with the specific effect *because* of which *X* currently exists:

(AA) An existing element *X* has the direct proper function *F* if and only if:

1. *X* results from the reproduction of an antecedent element, ancestral *X*;
2. Ancestral *X* effected *F* in the past, in virtue of properties reproduced in *X*.

Such traits are first what S. J. Gould called 'spanrels', and when further shaped by natural selection they become 'exaptations' (Gould and Vrba 1982). Good examples are the signals involved in **sexual selection**, which may begin either as random individual preferences by females, favouring inheritance by their offspring of both (through the male) the character preferred and (through the female) the preference itself. Or else they can begin with a trait actually correlated to fitness, particularly where it is subject to allometric development (Cronin 1991, pp. 183–204). The first, which Helena Cronin calls the 'good taste' variant, exacerbates a trait that did not in itself have a selective advantage before it became the object of female preference. The latter, which Cronin calls the 'good sense' variant, originally functioned as a perceptible indication of the presence of a desirable trait. It therefore acted merely as a piece of useful information guiding mate choice, rather than a fully fledged signal. Once it becomes enhanced by sexual selection, however, it can be seen as a genuine signal.

The aetiological theory is well equipped to explain apparent outcomes of design where no intentionality or even mentality is involved. Many well-known examples of mimicry are of this kind. And so is camouflage, though as we shall see below there is reason to think that camouflage, though functional, should not count as a true signal. The function of mimicry and camouflage is to deceive, although no intention can be ascribed at all. The deceptive traits are merely the outcome of reproductive lineages of organisms among whom a higher degree of resemblance to a poisonous species or resemblance to the background afforded a fitness advantage. In such cases, the mimicry does not depend on any representation of the current situation in the organism in question (see also Chapter 4).

When there is such a representation which responds to changes in the immediate environment, we can speak of a ground level of intentionality. It presupposes no sentience in the usual sense of the word, but only its simpler ancestor, a capacity for detection. Higher levels of intentionality, more plausibly attributed to conscious mental states, are involved in the sort of communication made possible by second- and third-order representation, culminating in the process of human

communication Grice describes. Thus children and some animals are said to acquire a 'theory of mind' which allows them to respond to their second-order representation of what is represented in the minds of others. Primatologists have found clear evidence that low-ranking animals can be aware of what can be seen by higher-ranking animals and modify their behaviour accordingly (Cheney and Seyfarth 1990). They have also found apparent cases of active deception which presupposes awareness of what the other would discover if placed in a position to do so (Tomasello and Call 1997). The anthropologist Robin Dunbar has suggested that humans are characterized by a fourth level of intentionality, involving the capacity to mention a reference made by an individual to a reference by a second individual to what a third thinks about a fourth (Dunbar 2004). It's not clear that such an achievement requires, as Dunbar intimates, more conceptual resources than are afforded by a clear capacity to go to a third level of representation. As we are able to speak of the content of another's mind, it would seem that one is *ipso facto* enabled to envisage higher orders of reference. But while that is merely a plausible supposition when what is in question are only non-intentional functional capacities, it is certain that once language is available to codify such representations and embed them in iterable syntactic structures, there is no clear theoretical limit to the number of iterations that become possible (see also Chapter 14).

16.4 Information: quantity and quality

So far the notion of 'information', as the X that gets transmitted from S to R , has been taken for granted. It is time to look at it more closely.

Let us again begin with a truism: what I already know conveys no information. This suggests a first approach to the characterization of information, as a *measure of surprise*. Learning something highly unlikely is maximally surprising and therefore maximally informative. One might be tempted, then, to identify information simply with the inverse of probability: the lower the prior probability of p , the more informative it is to learn that p is true. But there are several problems in the way before one can make this into a usable idea.

The first problem is that there are various ways of mathematizing the raw intuition. One little-noticed problem with the idea that information is inversely proportional to probability is that it appears to entail a paradoxical consequence. From the point of view of purely epistemic rationality, concerned exclusively with information and truth in abstraction from any other values (see Box 16.3), believing any proposition and believing its negation will come out to be equally rational.

Endorsing an improbable belief would, of course, be unlikely to pay off; but that would be offset by the large gain in information that would accrue if we turned out to be right. Rather as a perfectly fair bet has the same expected utility as not betting at all, the case of belief would work like this. Supposing I'm wondering whether to believe X , which has a probability of p . By definition $Not-X$ has a probability of $(1-p)$. But if $(1-p)$ also the epistemic value of X and p the epistemic value of $\sim X$, then the expected desirability of believing X [EU(BX)] is precisely the same as the expected desirability of believing $Not-X$ [EU(B $\sim X$)], namely zero, or the desirability of believing neither:

$$EU(BX) = [p \times (1-p)] + [(1-p) \times -p] = EU(B\sim X) = 0.$$

In the first term, p is X 's probability and $(1-p)$ is its information value if true. In the second, the two parameters are simply reversed.

To avoid this unwelcome result, the measure of information generally prescribed for Shannon information is not the inverse of probability, but its logarithm base 2. That is not a merely arbitrary dodge designed to avoid the awkward result just

mentioned. The log base 2 can be thought of as the number of questions required to arrive at a solution if one is attempting to identify a single item in a structured set. Thus 20 binary questions will suffice to zero in on a single number between zero and 1 million. Provided that the state space of possibilities is suitably structured, then the number of binary units (or 'bits') of information represented by the identification of that one number in a million is determined by the maximum number of binary partitions needed to home in on it. This gives an *objective quantitative measure* of information.

Furthermore, this conception of information is linked to both the physics of thermodynamics and statistical theory by means of the notion of *entropy*. This can be illustrated in terms of the classical model of an ideal gas. Imagine two containers of equal size, linked by a passage currently blocked by a gate. In the left container there is a volume of gas, while the other is a pure vacuum. Now consider what happens when the gate is opened, from the three different perspectives of physics, statistics, and information theory.

From the thermodynamic point of view, the gas in the left container has been exerting a certain pressure on the gate. According to the classical corpuscular theory of gases, that pressure depends on the temperature of the gas, and is actually equivalent to the vector that represents the mean molecular momentum of the particles that constitute the gas exerting pressure on the gate. The pressure represents usable energy: a piston placed between the containers might use it to effect some work (Fig. 16.1, A). But if we simply open the gate, the gas will

Box 16.3 Epistemic rationality

A rational strategy maximizes the probability of success. But what kind of success is relevant? Practical or economic rationality looks to gains and losses, and in biology those get cashed out in terms of fitness. But in the context of information and knowledge, we can think of success exclusively in terms of the likelihood of being *right*, or *believing truly*. That is the point of view of *epistemic rationality* (Levi 1967).

Practical and epistemic rationality can be thought to conflict. It can arguably be of practical benefit to have a false but encouraging belief in one's strength, health, or ability, or to flee from what is falsely believed to be a predator on the maxim 'better safe than sorry'. For a recent discussion of when it might be practically rational to be epistemically irrational, see Stephens (2001).




Figure 16.1 Movement of particles of an ideal gas between containers. See text for details.




Figure 16.2 Movement of a single particle between containers. See text for details.

gradually diffuse into the second container (Fig. 16.1, B) until the pressure in both containers is equal (Fig. 16.1, C). At which point, in accordance with the second law of thermodynamics, there is no further possibility of using, from within the closed system described, the energy embodied in the motion of the particles of gas. From the physical point of view this final state is known as the state of maximum entropy.

Switching to the statistical point of view, consider a single particle positioned right between the two containers, at the location of the now open gate (Fig. 16.2). Since particles move randomly, it has an equal chance of heading left or right. But since all the particles are originally in the left container, the initial probability of the particle passing from the left into the right container is 0.5, while the probability of a particle going from right to left is zero. As more and more particles end up in the right container, the probability of a random particle going from right to left will increase, in exact proportion to the ratio of particles in the right to those in the left container. At the end of the process, for purely mathematical reasons, the probability of a particle passing from left to right will be precisely equal to the probability of its going the other way. That constitutes the highest point of statistical entropy. This is equivalent to saying that no state of the whole set-up is more probable.

This probabilistic interpretation of entropy provide a handy way to think of the second law, loosely paraphrased as ruling that the passage from order to disorder is always to be expected, as a mere consequence of the mathematical tautology that less probable states are likely to give way to more probable ones. The notion of probability provides the link to the informational point of view, in terms of the basic intuition about surprise with which I started. The state in which all the gas is in the left container (Fig. 16.1, A) is intuitively a state of ‘order’ as opposed to the state of ‘disorder’ represented by the fully diffused state of the gas (Fig. 16.1, C). That original state of maximal order, if it were to result from random fluctuations alone, would be highly surprising by dint of its being the least probable outcome of the random motions of the individual molecules constituting the gas. Just as maximal disorder, or maximal entropy, is equivalent to the most probable distribution of particles in an enclosed area, so minimal entropy can be identified with maximal information.

The foregoing considerations bear on communication in two ways. The first concerns the trade-off between the length of a message and what we might call its *informational density*. The second is that further constraints on the notion of information are needed before it can be of practical use in understanding concrete cases of communication between parties that have different interests and different states of prior information.

To understand the trade-off between length and density, note that while Shannon information is conveniently measured by counting binary units, it does not require to be packaged in such units. Implementation in the form of bits is indeed the obvious solution for information processed in computers, in which every basic unit can be regarded as either on or off. The simplicity of the basic vocabulary, which we can then conceive of as being made up of just two elements, 1 and 0, comes at the price of long messages. Thus 999 takes just three elementary signs in the decimal system, but takes up 10 elementary signs in the binary system, where it is represented as 1111100111. Adding the 26 letters of the English alphabet plus a space to the decimal digits 0–9 makes a total of 37 elementary signs, which affords the same amount of information

in just two characters, and is, in fact, to an order of magnitude comparable to the number of phonemes in human spoken languages, which ranges between a dozen and a hundred (see also Chapter 14). At the other end of the spectrum from machine language, written Chinese language provides a striking example of a system capable of minimizing message length at the cost of requiring mastery of a very large number of distinct elementary signs. The first Chinese dictionary, commissioned by the Emperor Kang Xi, contained about 47,000 distinct characters, of which a literate but not erudite Chinese might be expected to know some 10,000. There is a trade-off between the memory storage required to distinguish 10,000 characters and the economy this allows in the length of each message. The way this trade-off plays out relates to informational density: the 0 or 1 of machine language represents single bit. A Chinese character, by contrast, compresses 16 bits of information, which is why two 8-bit 'words' or bytes are required to specify it in your word processor. Since an English character takes up just one byte, it is not surprising that a Chinese text invariably takes up much less space on the page than its English translation. On the other hand, the English alphabet can be memorized in an hour, whereas it takes years to master 10,000 Chinese characters. Since oral language is probably a more significant indicator than written language of the constraints under which our brains function, this suggests that the most efficient point at which the trade-off between memory and message length tends to settle requires somewhere between a dozen and a hundred basic elements.

Of course, complete messages are not built directly out of those elements. The power of language largely derives from the way that one or two intermediate levels of structure intervene between the atoms (phonemes) and complete messages (stories, arguments, speeches, poems, etc.), in the form of words, themselves articulated into phrases and sentences.

Here then is one hypothesis that can be derived from these abstract considerations about the notion of information. If a sender is to communicate information to a receiver, energetic factors will be involved, but they will be secondary to a sender and receiver's capacity to discriminate and

produce a number of distinct elementary signs, and to their capacity to process signals of a given length and density. Humans naturally 'chunk' information when it involves more than a few elements that need to be held in immediate memory (Miller 1956), and we would find it extremely difficult fluently to read an English sentence transcribed into machine language, with a string of eight ones and zeros replacing each letter of the alphabet and punctuation. The brains and other storage devices of other communicating organisms will necessarily be subject to similar constraints.

Implicit in this discussion has been the assumption that most animal communication is based on a digital system of representation. That implication was carried in the very notion of a signal's density, which presupposes that the information capacity of any single element is finite. A repertoire of possible signals, like the set of phonemes or letters, typically constitutes a finite set, into which any variant along some continuum would be slotted as belonging to one or another of the elements in question. Clearly, however, there are some dimensions of animal signals that may vary in analogue rather than discrete degrees. The orientation and velocity of a bee's dance comes to mind, as representing the continuous factors of direction and distance (Michelson 1993). But it is noteworthy that neither human language nor the 'language' of the genes could function unless they were organized as digital systems. Both genes and language have a comparable density at the lowest level of analysis, comprising four bases (two bits) or 10 to 100 phonemes (four to seven bits) respectively. Both are then chunked into a larger but still relatively small number of 'words' specifying one of some two dozen amino acids in one case and a few hundred thousand words in the other, which in turn become the components of a huge number of possible protein strings or sentences. It would be interesting to know whether these two examples of genes and language, lying as they do about as far from one another on the scale of mentality, exemplify a very general requirement of signalling, favouring digital systems of representation.

I come now to the shortcoming of the notion of information just detailed. Intuitively, Shannon information fails to connect with our intuitive

understanding of ‘information’, because it fails to relate to any of either the sender’s or receiver’s interests, and so fails to connect with the very idea that led to its characterization in terms of surprise. In short, Shannon information in itself *means* nothing to anyone. If I display for you the number 4,987,654,294,997, you are seeing something the prior probability of which was $1/10^{13}$ and so have just acquired 43 bits of information. You should be surprised! Your surprise is damped, however, by the fact that this ‘information’ is completely meaningless. But what does ‘meaningless’ mean? In itself, what is transmitted at any given synapse or between any two bacteria is also quite properly ‘meaningless’. It is only its role in some larger process that has meaning. We need, therefore, to look more closely at the specific conditions that make it possible and profitable for senders and receivers to exchange information.

16.5 Reliable signals

Any organism equipped with some sort of sensor is capable of acting as an information receiver. But not all information acquired by such an organism should count as a signal. Every cell is capable of some form of tropism, but not all are the effect of signalling. Strategic issues arise only when there is interaction. Even then, the acquisition by one entity of information about another may not constitute communication. This is the case of information acquired by simple perception. If I want to drink, and see water, that’s useful information, but it isn’t any sort of communication. Perceptual information about other organisms is frequently of that sort. Take, for example, humans’ ability to infer information relevant to mate choice from the texture of a person’s skin, described by Craig Roberts (see Chapter 9), and contrast it with the human ability to infer socially relevant information from a person’s posture and movement. The second appears more likely than the first to be an evolved form of functional signal, in the sense of being in some sense manipulated by the sender, though both may be likely to act as triggers for a certain behaviour. In the terms elaborated by Marc Hauser (1996, p. 24), the former type of signal typically pertains to the sending organism’s ‘resource holding potential’

(RHP): the important point about it is that the RHP cannot, while an optional signal generally can, be manipulated at will to mislead. Thus a larger animal will obviously seem to be larger, which may well be crucial information. A smaller animal can send a misleadingly formidable image of its size only to a very limited extent, if it is able to spread its plumage or its fur, or inflate like a blowfish. But even then, as we shall see in a moment, it isn’t clear that what is being sent is a signal in the full sense, rather than something like camouflage, which there is reason to think is not yet quite a signal.

When the issue arises of the relative value of information to a sender and a target of transmission, it raises strategic and economic issues. Some critics of sociobiology—including some biologists—have complained that economic concepts are inappropriate imports in biology, and introduce an ideological bias into our conception of the natural world. Margulis and Sagan, for example, have charged that ‘vogue words like “competition”, “cooperation”, “mutualism”, “mutual benefit”, “energy costs”, and “competitive advantage” have been borrowed from human enterprises and forced on science from politics, business, and social thought’ (Margulis and Sagan 2002, p. 16). But the mere fact that concepts can bridge two domains says nothing about whether their claim is illegitimate in one or in the other. In fact, as Maynard Smith foresaw, economic and game theoretic concepts have proved to be of enormous importance in evolutionary biology (Maynard Smith 1984, p. 2000). In fact, we could say that economic concepts and game theory actually apply *literally* only to biology. Their application to human choices require psychological assumptions, notably that humans are utility maximizers, which idealize human motivation out of recognition. By contrast, when economic concepts are applied to biology, none of those questions arise. Fitness provides an objective equivalent of ‘interest’ or ‘advantage’ requiring no intermediate psychological assumptions.

Information is unique among all goods that can be acquired or exchanged, in that it alone does not share the ‘zero sum’ characteristic of other desirable things such as food, territory, or building materials. Unlike the proverbial cake, I can consume information, and give it to another, and still have

it too. The strategic issues raised by information transmission are therefore distinctive in certain ways: when one entity dispenses information, the cost to the former does not amount to the loss of the information in question. But that doesn't mean it is cost-free. On the contrary, in their book-length study of animal signalling, Maynard Smith and Harper (2003) focus principally on the dependence of a signal's reliability on its cost. They first define a signal as follows:

A signal is any act or structure which [1] alters the behaviour of other organisms, which [2] evolved because of the effect, and which [3] is effective because the receiver's response has also evolved....the requirement that a signal evolved *because* of its effect on others distinguishes a signal from a 'cue', [defined as] any feature of the world, animate or inanimate, that can be used by an animal as a guide to future action.

Maynard Smith and Harper (2003)

(compare Zahavi in Chapter 1).

The difference between a cue and a signal is illustrated in terms of the following example. Reichert (1978) studied contests between funnel-web spiders, *Agelenopsis aperta*, over web sites. She found that if there was a difference in weight between two spiders of 10% or more, the smaller spider retreated without risking a fight. A spider can perceive its weight relative to that of an opponent because the contests take place on the web. The spiders signal by vibrating the web, transmitting information about their size: a smaller spider can be converted into a winner by attaching a weight to its back. Thus size itself is not a signal by our definition. It did not evolve *because* of its effect on other spiders. However the act of vibrating the web *is* a signal if, as seems plausible, it evolved because of its effect on the behaviour of an opponent through the information it provides about size (Maynard Smith and Harper 2003).

As I understand it, the crucial difference between the vibration of the web in this case and the blowfish's inflated size is that the receiver of the spider's message, but not a predator watching the blowfish, has had to evolve or learn the meaning of what it perceives. Size is size, ever where it is deceptive; but vibration is merely an indicator of size. On the other hand, most species of spiders are not equipped

with a sufficiently effective visual sense to perceive the other's size, so that the lyriform organ in their feet is the most direct sensory channel available to them to sense the size of anything on the web (David Hughes, private communication). The case cited by Reichert must therefore count as a signal *only if* it is the outcome of an arms race in which spiders have evolved the capacity to amplify their apparent size by vibrating the web in a misleading way, whilst also having refined their capacity to infer real size from vibration.

Following Zahavi and Zahavi (1997), Maynard Smith argues that the costliness of a display can provide a warrant of its reliability. That isn't always the case. Merely being costly and providing information to a receiver does not suffice to make a piece of behaviour count as a signal. *Camouflage* is a case in point, and we can see now why it doesn't count as a fully fledged signal. It is presumably costly to produce, and was selected for. It therefore meets the first two conditions of Maynard Smith's definition. But it fails the third: for the effect that camouflage has on the predator—making the latter less likely to eat it—was not itself selected for. Similarly, if a hare outruns a fox, that may be costly, and the necessary capacity has evolved in the hare. But the insufficient speed of the fox wasn't an evolved response. If, however, the fox has acquired a disposition to economize its own strength by not bothering to undertake the chase, then the hare's speed can count as a signal to the fox. Stotting in gazelles or the alarm calls of vervet monkeys in the presence of predators are costly both in terms of their intrinsic energetic expenditure and owing to the fact that they draw attention to themselves in such a way as to put them at increased risk. And, here again, they will count as genuine signals provided that there is some *evolved* effect on the receiver, causing the latter to save energy by abandoning the chase before it is even begun, as opposed to giving chase in response to the prey's apparent availability.

Once signals evolve to be potentially misleading, the game-theoretical perspective comes into its own. At the ground level, the fact that a signal is not necessarily reliable creates a measure of uncertainty. The frequency of its indicative and misleading occurrences could be assessed by a receiver

equipped with the right kind of memory, and there could be different decision outcomes depending on the stakes implicated in different contexts. But where a certain probability of its being misleading gives the receiver the option of ignoring it, an arms race will give rise to second-order signals, in which some additional element is included to warrant reliability. As is obvious from the experience of ordinary conversation, merely asserting 'This is really true!' before reiterating a dubious proposition does not add credibility. Hence Zahavi's 'handicap principle', which applies particularly to predator-prey communication and to mating signals. The handicap principle posits that an animal can warrant the reliability of the signals by which it advertises health and fitness by indulging in an additional and costly display (Zahavi and Zahavi 1997). The cost of the display is itself an additional signal, providing a kind of warranty insofar as it proves that the animal displaying it has strength and resources to spare.

At first sight, this signal serves the prey rather than the predator, since it might deter the latter from pursuit. But Zahavi claims it also serves the predator in saving him a pointless pursuit, as in the case of the fox just mentioned. The situation is worth looking at more closely, however. For there is an important asymmetry between the two. Applying the aetiological analysis of function discussed above, it seems that the handicap serves a *strategy* in the case of the prey, but functions merely as a natural signal or *cue* in the case of the predator. The reason is apparent if we compare two counterfactuals suppositions. For the prey, the point of the handicap lies in the effect it has on the predator. If it did not influence the predator, the handicap would not have been selected. The effect on the predator therefore explains the presence of the handicap in the usual way that the presence of a trait is explained by its function. But from the point of view of the predator, there is nothing more to the signal than its capacity to provide useful information, indicating that the chase is not worth it. So while it may seem to be functional for the predator, it is so only in sense of being informative. The asymmetry resides in the fact that one party gets information from the other while the other actually manipulates the information.

To illustrate the point, consider the difference between the information provided to a human being by another's expression of emotion and that provided by a bruise. From a bruise, I can make the inference that the world affords a blunt object in the vicinity and that the person in whom the bruise is seen has had an encounter with it. From the expression of anger, I can similarly infer that the person in question has suffered some injury. (Let us set aside the further information afforded by the observation that the anger is or is not directed at me, in a way that can't be made sense of in the case of the bruise.) The crucial difference between the bruise and the expression of emotion is that it is reasonable to suppose that the indicative function of the latter, but not that of the former, has been selected for. In other words, the expression of emotion exists *so as to provide information*. A bruise, by contrast, is just an effect of burst capillaries, resulting in blood collecting close enough to the skin to be visible. Because it is visible, it can afford information; but it is unlikely that its visibility was selected for. Most likely, then, a bruise has no evolved signalling function. The bodily manifestations of emotion, by contrast, have functions: they are supposed to tell us something.

16.6 Some remaining empirical and theoretical questions

In contrast to the other chapters in the present book, what I have presented here is highly abstract. Little more, I fear, can be expected from a philosopher, whom a scientist can plausibly regard as a kind of scientific poacher, a free-rider who waits in his proverbial armchair for others to conduct serious research in the lab so that he will have something to talk about. But perhaps I can make amends by suggesting, in conclusion, yet another couple of questions the answer to which is not obviously available to simple cogitation, and to which, as far as I can see, answers have not explicitly been provided by the chapters in the present book. Both bear on the question of the potential practical uses to which the findings of animal communication might be put.

The first question is sparked by a recent report in *New Scientist* about a possible strategy for

neutralizing viruses by trapping them in cells that lack the machinery the viruses need to reproduce (Ginsburg 2007). Viruses typically reproduce by hijacking the DNA in the cells they invade for their own reproduction. They respond to particular molecules on the surfaces of cells. These molecules must first act as cues as to the availability of the necessary reproductive machinery inside (though they can be regarded as genuine signals if they result from the fact that cells of that kind have, through a past process of natural selection, been manipulated by the virus). But if these same molecules can be attached to cells that do not contain the machinery in question, they will become *misleading* signals, by which a virus could be lured into a dead end. Since red blood cells have no DNA, a virus trapped in a blood cell will replicate no further. In confirmation of this idea, experiments using genetically modified blood cells bearing glycans on their surface succeeded in trapping a virus into complete extinction *in vitro*, and with notable if incomplete success *in vivo* (Asher *et al.* 2005). Finberg and his colleagues, according to Ginsburg, are now working on synthetic traps that would not require the genetic modification of blood cells. Would this be, in effect, a case of artificial mimicry? That depends on whether glycans (or other similar molecules attaching to the surfaces of cells) already have a function in communication. That isn't known for certain, but it seems possible that glycans on blood cells might indeed function to trap and neutralize viruses. Is this a case of communication or a case of mere camouflage? In its use as a virus trap, the DNA-free cell benefits only the sender; to the receiver it acts as a cue. Yet it can clearly be described as a type of non-intentional level deception. If we are able to use the effect in the manufacture of artificial 'virus traps', it might spark a further step in an arms race. That would happen if it subsequently modifies the behaviour of the virus. And that, in turn, is to be expected, since only mutant versions of the virus that avoid that particular marker will leave copies of themselves. In a case of this sort, perhaps it doesn't matter whether a genuine signal is involved or not. For the purposes of our manipulations, the causal properties of the molecules in question are what counts, not their original functions. But that view

may be hasty. For if we are indeed able to classify such cases under the general heading of 'signals', or merely 'camouflage', rather than merely causal mechanisms, might it suggest analogies with other cases of the exploitation of signalling mechanisms that could yield a broader range of applications?

However that may be, the difficulty of deciding whether one is dealing with a genuine signal in these cases may be merely empirical, rather than theoretical. More thorough knowledge and technological refinements will answer the question. But some difficulties entailed by the mere complexity of the systems involved may be more resistant to empirical testing. I have in mind the question of whether the aetiological analysis endorsed above will scale up smoothly to large *networks* of communicating animals. It seems reasonable to assume that it will. But we may not be able to arrive at a precise formula to describe it, just as we can solve the three-body problem only by simulations and not by mathematical equations. We should perhaps also envisage the possibility that additional effects and properties might be emergent in very large networks of communication. Such very large networks have been discussed in the present volume, at both the 'low' and the 'high' ends of the living world. An example of the former is the quorum sensing of bacteria discussed by Diggle *et al.* (Chapter 2). At the other end of the spectrum are the social networks described in the work of Matessi *et al.* (Chapter 3). In such very large networks, it is conceivable that emergent phenomena might arise that are not clearly predictable on the basis of the teleosemantic approach I have endorsed here. Some of Stuart Kauffman's work on self-organization in complex systems (Kauffman 1995) suggests that such systems are subject to surprisingly strong constraints on a purely mathematical level. Similarly, the work of Albert-László Barabási and others leads us to expect certain emergent properties in the structure of very large unplanned networks. The degree of their connectivity, for example, seems 'naturally' to organize itself in conformity with power laws (Barabási 2002). At that level of complexity, we may find the sort of large effects from small causes that are characteristic of chaos. We may then need to use different techniques for understanding and predicting the spread of information in networks.

Summary

Philosophers interested in meaning have tended to look at the extremes of mere causality on the one hand and fully fledged ‘non-natural meaning’ in human language on the other. But the former (though not simple, as attested by the long and largely vain attempt of philosophers to analyse it) is too simple to count as information, while the complexity of the latter places it far beyond many other forms of genuine communication found in the living world, from bacteria to mammals. Those other forms of communication involve ‘Shannon information’ but aren’t wholly captured by that notion. In this chapter I look at some of the work that biologists have done to construct a coherent concept of information able to span a wide spectrum of communication from such phenomena as ‘quorum sensing’ among bacteria to sophisticated infra-linguistic signalling in primates.

References

- Asher, D.R., Cerny, A.M., and Finberg, R.W. (2005). The erythrocyte viral trap: transgenic expression of viral receptor on erythrocytes attenuates Coxsackievirus B infection. *Proceedings of the National Academy of Sciences of the USA*, **102**, 12897–12902.
- Barabási, A-L. (2002). *Linked: The New Science of Networks*. Perseus, Cambridge, MA.
- Cheney, D.M. and Seyfarth, R.M. (1990). *How Monkeys See the World*. Chicago University Press, Chicago.
- Cronin, H. (1991). *The Ant and the Peacock: Altruism and Sexual Selection from Darwin to Today*. Cambridge University Press, Cambridge.
- Dawkins, R. (1982). *The Extended Phenotype: the Gene as Unit of Selection*. Oxford University Press, Oxford.
- Dawkins, R. (2004). extended phenotype—but not too extended: a reply to Laland, Turner and Jablonka. *Biology and Philosophy*, **19**, 377–396.
- Dunbar, R.I.M. (1993). Co-evolution of neocortex size, group size and language in humans. *Behavioral and Brain Science*, **16**, 681–735.
- Dunbar, R.I.M. (2004). *The Human Story: a New History of Human Evolution*. Faber, London.
- Ginsburg, J. (2007). Weapons of mass deception. *New Scientist*, **196**, 43–45.
- Gould, S.J. and Vrba, E.S. (1982). Exaptation: a missing term in the science of form. *Paleobiology*, **8**, 4–15.
- Grice, P. (1989). Meaning. In: *Studies in the Way of Words*, pp. 213–223. Harvard University Press, Cambridge, MA.
- Hauser, M.D. (1996). *The Evolution of Communication*. MIT Press, Cambridge, MA.
- Kauffman, S. (1995). *At Home in the Universe: the Search for the Laws of Self-Organization and Complexity*. Oxford University Press, Oxford.
- Levi, I. (1967). *Gambling with Truth*. Alfred A. Knopf, New York.
- Margulis, L. and Sagan, D. (2002). *Acquiring Genomes: a Theory of the Origins of Species*. Basic Books, New York.
- Maynard Smith, J. (1984). Game theory and the evolution of behavior. *Behavioral and Brain Sciences*, **7**, 95–126.
- Maynard Smith, J. (2000). The concept of information in biology. *Philosophy of Science*, **67**, 177–194.
- Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press, Oxford.
- Michelson, A. (1993). The transfer of information in the dance language of honeybees: progress and problems. *Journal of Comparative Physiology A*, **173**, 135–141.
- Miller, G.A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, **63**, 81–97.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge, MA.
- Millikan, R.G. (1993). *White Queen Psychology and Other Essays for Alice*. MIT Press, Cambridge, MA.
- Rheingold, H. (2000). *Tools for Thought: the History and Future of Mind-Expanding Technology*. MIT Press, Cambridge, MA.
- Riechert, S.E. (1978). Games spiders play: behavioral variability in territorial disputes. *Behavioral Ecology and Sociobiology*, **3**, 1–15.
- Schiffer, S.R. (1972). *Meaning*. Oxford University Press, New York.
- Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, **27**, 379–423, 623–656 (available online at <http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf>).
- Sober, E. (1984). *The Nature of Selection*. MIT Press, Cambridge, MA.
- de Sousa, R. (2005). Biological individuality. *Croatian Journal of Philosophy*, **54**, 195–218.
- Stephens, C. (2001). When is it selectively advantageous to have true beliefs? Sandwiching the better safe than sorry argument. *Philosophical Studies*, **105**, 161–189.
- Tomasello, M. and Call, J. (1997). *Primate Cognition*. Oxford University Press, Oxford.
- Wright, L. (1973). Functions. *Philosophical Review*, **82**, 139–168.
- Zahavi, A. and Zahavi, A. (1997). *The Handicap Principle: a Missing Piece of Darwin’s Puzzle*. Oxford University Press, New York.

