18

Information, influence and inference in language evolution

THOMAS C. SCOTT-PHILLIPS AND SIMON KIRBY

18.1 Introduction

The various chapters that appear in this volume reflect a range of perspectives on a question of contemporary and interdisciplinary debate: the nature of communication. There are several reasons (surveyed elsewhere in this volume) why this issue has arisen at the present time. One of these is the growing research interest in the origins and evolution of human language, which has highlighted the need for a general framework and vocabulary with which to describe communication, since the development of such a framework would assist cross-disciplinary discussion of the transition(s) from non-human primate communication to language (Rendall *et al.*, 2009). In addition, several protagonists in the current discussion have argued that comparisons to language have, depending on your perspective, either hindered or enabled research on animal communication (e.g. Rendall Owren & Ryan, 2009; Seyfarth *et al.*, 2010). Despite these motivations, the discussion has mostly taken place in the animal communication literature, and the contributors to it have mostly been experts in that same field. Consequently, debate can be advanced by a detailed discussion of how linguistics, and the study of language origins in particular, conceives of communication. This chapter will address these needs.

First, we will briefly chart the development of language evolution research, with a particular focus on how the discipline has tended to conceive of communication (Section 18.2). We then describe a distinction, commonly made in pragmatics (the branch of linguistics that deals with language use in context), between two different approaches to communication (Section 18.3), and discuss

Animal Communication Theory: Information and Influence, ed. Ulrich Stegmann. Published by Cambridge University Press.

the implications of this distinction for language evolution (Section 18.4). We will explain why consideration of language highlights the need for discussion about the nature of communication to develop an account that is sufficiently general to describe a communication system that is heavily context-dependent. We then provide a set of definitions that are able to do this, and that hence provide the foundation for a general account of communication (Section 18.5). Our overall objectives are thus to present a unified, consilient account of communication, and to discuss the implications of that account both for language evolution research, and for the more general issues that have motivated the present volume.

18.2 Communication in language evolution research

Those unfamiliar with recent research into the evolution of language are often struck with an obvious question: how is it even possible to study the origins of a behavioural trait that is apparently unique to one species and leaves no direct fossil remains? In other words, where are the data for those interested in language evolution going to come from, and what are the appropriate methodologies that should be employed? Perhaps because for a long time there were no clear answers to these questions, the field has sometimes been viewed as one that leads to unconstrained speculation, giving rise to an abundance of theories with little other than personal taste to help choose between them.

It was in this context that the use of computer simulations was pioneered in the late 1980s (in particular Hurford, 1989, 1991), as a way to tackle the problems of unconstrained speculation. Although we cannot study real evolving populations of communicating hominids, we can create populations in computer-simulated environments, and use these as a test-bed for evaluating different hypotheses about the mechanisms, selection pressures and so on implicated in the origins of language. This general methodology proved highly influential through the 1990s and 2000s, when computational modelling became an established methodology for language evolution, with a large proportion of papers in the regular international conferences and edited collections using such simulations (e.g. Briscoe, 2002; Cangelosi & Parisi, 2002). This foundational research often conceived of communication in a very idealised way, typically for practical, implementational reasons – but these assumptions shaped much of the subsequent thinking in the wider field of language evolution. Because of this, it is worth going into some of the detail of a typical early model in order to see how such models bear on the topic of this volume (for more thorough reviews of the early modelling literature, see Kirby, 2002; Steels, 2003).

In a typical simulation, there would be a set of independent simulated 'agents' implemented using very simple algorithms. The interesting behaviour of these models arose from the consequences of interactions between these agents. For example, the agents might engage in a communication task, taking turns as a producer of a signal and a receiver of a signal. The agents might tackle this task in a range of different ways, each particular response arising either because the agents learned their behaviour from each other or because it was encoded for each as a set of idealised 'genes' in the specification of the agents. Populations of the agents might be static or change over time, and could be used to examine hypotheses about the cultural and biological evolution of communication systems.

The very first simulation models of language evolution focused principally on the biological evolution of signalling systems (e.g. Werner & Dyer, 1992; MacLennan & Burghardt, 1993; Ackley & Littman, 1994). Models that addressed other concerns, such as the social transmission of language through learning, and the consequent cultural evolution of linguistic systems, came later (e.g. Batali, 1998; Kirby, 2001; K. Smith, 2004; Brighton, Smith & Kirby, 2005). In one of these early models, agents are born with a set of genes that specify directly what their behaviour is, and this behaviour does not change throughout their lives (Oliphant, 1996). Agents engage in a communicative task with each other, and their success at this task determines their fitness. In other words, better communicating agents were more likely to pass on their genes to future generations of agents. The particular concern of this piece of research is the way in which fitness could be calculated in these models, and the details of the communication task are not a central concern. Accordingly, a set of simplifying idealisations about how communication works is made for the purposes of modelling. At each time-step of the simulation, one of the agents is prompted to produce a signal for a particular meaning. To do so, it consults a table determining its transmission behaviour, which is specified directly by the agent's genes. This signal is then transferred to another agent, which consults a corresponding table determining reception behaviour. This results in a response to the signal received. If this response is appropriate then communication is deemed to be successful; see Figure 18.1.

In this work, signals and meanings are, respectively, referred to as *observable behaviours* and *states of the world/appropriate responses* (Oliphant, 1997, pp. 15–16), yet in the model there is none of the richness that this might imply. Signals and meanings are simply members chosen from a finite set. Success relies on signaller and receiver having coordinated mappings from meanings to signals, and signals to meanings. This work was influential in the modelling community (as were others like it) in that it set the scene for much future computational

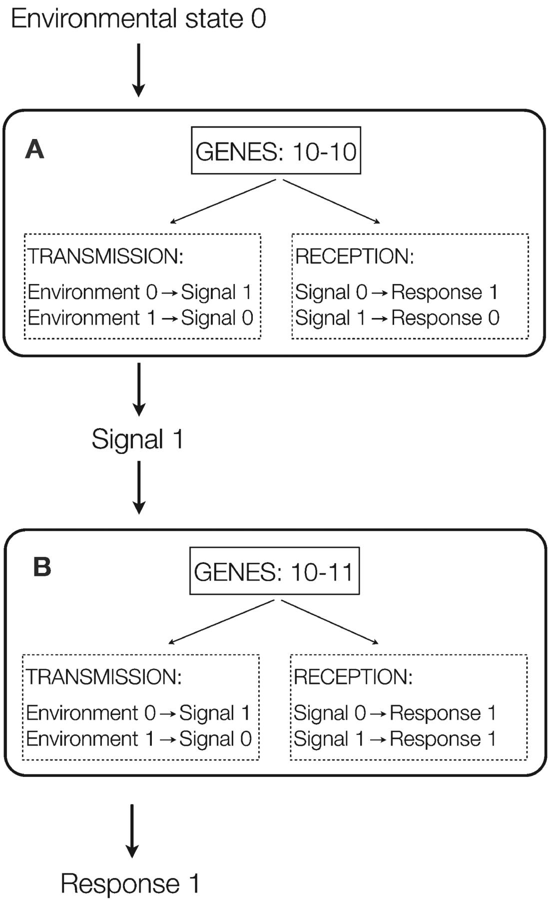

Figure 18.1 A sketch of Oliphant's (1996) model of an innate signalling system. Individuals' communicative behaviour is divided into transmission and reception. Transmission relates distinct states of the environment to signals, whereas reception relates signals to particular responses. (In the literature more broadly, the environmental states and responses are typically both glossed as 'meanings'; we follow this convention in the main text.) These relationships are determined by a string of heritable genes, the first half of which give the transmission behaviour, and the second half of which give the reception behaviour. There are two components to each of these: the first figure describes the output given an input of 0; the second describes the output given an input of 1. In this example, agent A signals to agent B. The environmental state is 0, so agent A signals 1 (the first part of the signalling genes '1 0'). Agent B then receives this input, and so produces response 1 (the second part of the signalling genes '1 1').

work on language evolution. It provided a backdrop which focused on the evolution of mappings between sets of possible signals and possible meanings in which communication was about individuals encoding meanings as transmittable signals and others decoding those signals into meanings.

Later work loosened various parameters of these models. In particular, models were built that added contexts in which to ground symbols, showing how agents can converge upon stable meaning–signal associations using techniques such as cross-situational learning (disambiguation across multiple contexts) and even that communication can be successful when agents do not represent meanings in the same way as one another (e.g. Steels, 1999; Vogt, 2002; Smith, 2005). Although such research can be described as inferential (Smith, 2005) because meanings are not given directly, the basic paradigm remained one in which producers are prompted to transmit signals for particular meanings, which receivers then decode.

More recently, language evolution has seen a rapid growth of experimental approaches, much of it informed by the previous modelling work (Scott-Phillips & Kirby, 2010). In particular, the general framing of communication in terms of a system that encodes and then transmits signals with fixed 'meanings' has persisted in much (but not all) of this experimental work. For example, in one study participants learn to associate strings of syllables (signals) with coloured shapes (meanings). They then go on to produce strings for particular shapes that they are prompted with. The resulting language collated from one participant's output is used to train the next participant in the laboratory. In this way, cultural evolution of language is modelled in a way that is entirely analogous to the computational simulations (Kirby, Cornish & Smith, 2008). The analysis of the results is even conducted in the same fashion, treating the data as a matrix of signal–meaning associations. This way of thinking about communication is sometimes called the *code model*, because it assumes that communication consists of the encoding of a message into a signal, which is then transmitted along some channel, to be decoded (subsection 18.3.1 below).

18.3 Two models of communication

18.3.1 The code model

The code model is an idealisation of communication in which signal–meaning mappings are the primary representation: the signaller 'looks up' the correct signal to use for a given meaning and produces it, and the receiver then takes that signal and 'looks up' the associated meaning in order to interpret it. Any approach that models communication in this way is using the code model, even if the representations of meanings may be different in signaller and receiver. As far as we are aware, the term *code model* was coined by Dan Sperber and Deirdre Wilson in *Relevance: Communication and Cognition* (1986), in an effort to distinguish it from the *ostensive-inferential model* that they wished to promote, and which we discuss below. But the code model itself has a much longer, if often unrecognised history, particularly in theoretical linguistics (Blackburn, 2007). Although there is no definitive definition of the code model, perhaps because its use often goes unrecognised, the foundational assumption is that communication consists of the encoding of a message into a signal, which is then transmitted along some channel, to be decoded. The communication system is the set of mappings (i.e. codes) between messages and forms – and these codes exist independently of any given user. When applied to language, this way of thinking about communication leads to the belief that languages are codes that define a correspondence between sound and meaning.

A historical analysis suggests that the code model combines at least two metaphors of how communication works (Blackburn, 2007). The first is the conduit metaphor, in which signals are packaged up and sent along some channel, to be unwrapped at the other end (Reddy, 1979; see Figure 18.2A). This view pervades our everyday language about communication (e.g. “Send me your ideas”; “Get your message across”), and it has been hugely influential in the animal signalling literature, where many textbooks define communication in this way (see Rendall *et al.*, 2009 for a critical review).

The second metaphor is the telecommunications metaphor. It has its roots in Claude Shannon’s information theory (Shannon, 1948; see Figure 18.2B). This was designed to address a particular set of electrical engineering and telecommunication problems, in particular how to transmit digital strings along noisy channels. The metaphor is to treat language as if it operates in the same probabilistic, context-independent way. Yet there is no attempt in Shannon’s work to relate the structure or semantic aspects of any underlying message to the real world (Blackburn, 2007), and Shannon himself is explicit about this: “these semantic aspects of communication are irrelevant to the engineering problem [that this work is concerned with]” (Shannon, 1948, p. 5).

Despite its influence, the code model is plainly an insufficient characterisation of linguistic communication. Much of what is meant is not encoded in what is said (that is, the utterance *underdetermines* the speaker’s meaning). In sarcasm, for example, the ‘encoded’, literal meaning of an utterance is the opposite of the speaker’s intended meaning. More generally, linguistic communication often leaves much implied, and both speakers and hearers lean heavily upon context to ensure that communication is achieved successfully (Grice, 1975; Sperber & Wilson, 1986; Carston, 2002; Atlas, 2005). For example, if Ann asks John, “Where’s Bill?”, and John replies, “There’s a yellow VW outside Sue’s

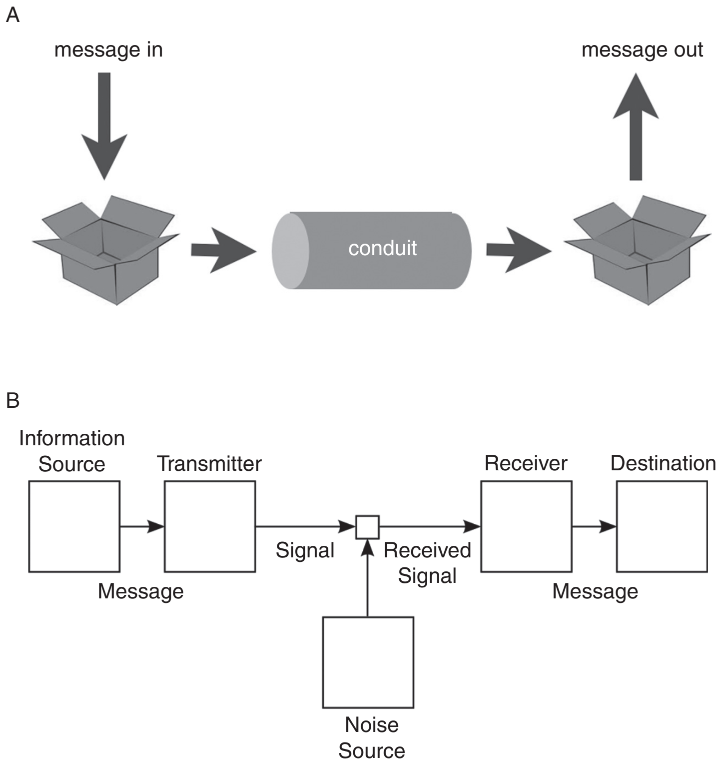

Figure 18.2 Two metaphors for communication. **A**, The **conduit metaphor**, in which information is packaged up and sent along a channel, from signaller to received. **B**, The **telecommunications metaphor**, in which communication is conceived of as the transmission of digital strings along noisy channels (from Shannon, 1948). While both approaches are intuitively appealing, it is important to keep in mind that they are metaphorical, and should not be taken as a complete description of what is involved in communication.

house", then John has not explicitly stated that he thinks (or at least has reason to believe) that Bill is inside Sue's house; this is only implied (example from Levinson, 1983). Furthermore, important contextual information is required if Ann is to comprehend John's utterance appropriately. Specifically, Ann must know that Bill drives a yellow Volkswagen, and perhaps that Bill knows Sue, and hence it is plausible that he is visiting her house. If Sue does not know this or, worse, believes otherwise, then it becomes difficult for her to make sense of Bill's utterance. The study of how such contextual factors affect meaning is called pragmatics.

This underdeterminacy is ubiquitous: there is *always* a distinction to be made between what a speaker says and what they mean (Carston, 2002). As such, both speakers and hearers must take contextual information into account if communication is to be successful. This makes the code model inherently unsuitable as a description of language use in context. Consequently, a large part of the development of pragmatics as a discipline has been the creation and refinement of an alternative to the code model (Austin, 1955; Grice, 1975; Sperber & Wilson, 1986). We now describe that alternative.

18.3.2 The ostensive-inferential model

In the *ostensive-inferential model*, signals do not encode the signaller's intended meaning; rather, they provide evidence for it. The signal and its various properties (in particular, the conventional, 'decoded' meaning of the utterance) are part of that evidence, but so too is the situation in which it is produced. Similarly, the receiver does not decode the signal; rather, they take the evidence presented to them and must then infer the signaller's intended meaning – a task that can only be done when context is taken into account. As such, ostensive-inferential communication involves the expression and recognition of intentions, and the ostensive-inferential model highlights that the signal always and necessarily underdetermines the speaker's intended meaning. Coding and decoding are not (indeed, cannot be) all there is to linguistic communication.

The term *ostensive-inferential* is double-barrelled to reflect the two tasks involved: speakers must produce ostensive stimuli, rich enough for their purposes, and listeners must draw context-dependent inferences from those stimuli (Sperber & Wilson, 1986). Note that although the model is typically presented in terms of vocal linguistic communication (with 'signaller' and 'receiver' replaced by 'speaker' and 'listener'), that need not be the case; the model applies to any forms of ostensive communication, vocal, linguistic or otherwise. A classic example of non-vocal, non-linguistic yet ostensive behaviour is pointing (see e.g. Tomasello, 2008).

Because it emphasises the fact that both parties are doing some work in this process, the ostensive-inferential model contrasts with the conduit metaphor, in which listeners are characterised in a more passive way, as recipients of prepackaged pieces of information. Nevertheless, ostensive-inferential is a wordy term, and so *ostensive communication* or *inferential communication* are sometimes used as shorthand. One question that follows from this discussion is how these different ways of thinking about communication (the code model and the ostensive-inferential model) inform the current debates about the nature of animal communication. In Section 18.5 we will discuss how the ostensive model relates to a view of communication predicated on influence. Before then, we wish to highlight the importance of ostension and other pragmatic factors for the study of language evolution itself.

18.4 The relevance of pragmatics for the study of language origins

Although pragmatics is an active branch of linguistics, and indeed forms part of any typical undergraduate curriculum, there are many other areas of core linguistics, such as syntax, phonology and so on, that can and do operate largely without consideration for the role of ostension, inference and other pragmatic concerns. Moreover, it would not be inaccurate to say that for many linguists pragmatics falls at the periphery of the discipline. However, although it may not be part of the core of linguistics, pragmatic factors are of critical importance for the study of language origins (Hurford, 2007; Tomasello, 2008). Language is the archetypal example of ostensive communication, and hence any comprehensive explanation of the origins of language must, in our view, account for why this form of communication has apparently emerged in only one species. We consider this to be a central question for current language evolution research (Scott-Phillips & Kirby, 2010).

To this end, comparative research has begun to identify the similarities and differences between the social cognition of humans and other primates (Tomasello *et al.*, 2005), and has suggested some of the selection pressures that may explain these patterns (Humphrey, 1976; Byrne & Whiten, 1989; Dunbar, 1998; Dunbar & Schultz, 2007). Thus one of the present challenges is to focus on the social cognition of ostensive-inferential communication in particular, in order to make this account more precise. What are the specific cognitive mechanisms involved, and how did they evolve? Recognition of the importance of this question is reflected in recent trends in the comparative literature, which has seen a move towards studies focused on the psychological and cognitive mechanisms involved in non-human primate communication, rather than their surface form (e.g. Cartmill & Byrne, 2010).

Indeed, one of the main criticisms made of the informational approach is that it borrows constructs from language and attempts to shoehorn them into descriptions of animal communication (Rendall *et al.*, 2009; Owren *et al.*, 2010). For example, the development of the notion of *functional reference* can be seen as a way to bring the idea of reference, developed in the study of human communication, into animal (and particularly non-human primate) communication. In this light, it is worthwhile to point out that when we consider the role of pragmatics, it becomes clear that the conduit metaphor is an insufficient characterisation of even human communication, where, as discussed above, reference and other factors that contribute to literal meaning are only part of the story. Context-dependent factors are also critical. This is why the trend towards investigation of the underlying cognitive mechanisms required for ostensive communication is to be welcomed. The picture that appears to be emerging is that non-human primates possess and use some but not all of these mechanisms (Hurford, 2007). For example, evidence of intentionality (widely defined as behaviour that is directed towards another, with the apparent objective of obtaining a goal, and which is employed flexibly, depending on the audience response) has been observed in the communication of several non-human primate species, particularly in the gestural modality (e.g. Pika, 2008; Tomasello, 2008; Cartmill & Byrne, 2010; Holbaiter & Byrne, 2011).

Such results present a significant theoretical challenge. Specifically, they suggest that ostension and inference are not all-or-nothing phenomena. The characterisation of human communication as ostensive and non-human communication as not ostensive may be too simple. Rather, it seems that there are shades of grey here; there may be forms of communication that are 'somewhat' ostensive and inferential. Therefore an important research problem for language evolution (and indeed for pragmatics) is to develop a theoretical framework with which to characterise such systems, in order to compare them against one another. This is a highly interdisciplinary problem. One way in which it might be approached is in the context of a more general account of what communication is, at its most basic level.

18.5 A general framework for communication?

The code model has been a dominant framework not only for the study of language and linguistics, but also for studies of animal communication (Rendall *et al.*, 2009; Owren *et al.*, 2010). The strengths and weaknesses of this approach have been recently debated (Rendall *et al.*, 2009; Seyfarth *et al.*, 2010), and are a central concern of the present volume. What the preceding discussion highlights is that regardless of its utility for the study of animal communication, the code model is certainly an insufficient description of human ostensive communication, including language, and therefore cannot be sufficient to describe communication in the most general sense. This section presents an alternative approach, one predicated on functionality, that is able to describe both coded and ostensive communication. The implications and applications of this framework are then discussed.

18.5.1 Definitions

We adopt functional definitions of communication and associated terms. We are inspired by definitions developed in evolutionary biology (Maynard Smith & Harper, 2003; Scott-Phillips, 2008), but here we express them in more general terms, so that they can be applied whether we are concerned with ultimate or proximate sources of design. A signal is any behaviour that satisfies the following three conditions: (i) it causes a reaction in another organism; (ii) its function is to cause the reaction; and (iii) the reaction's function is to be caused by the action. If these conditions are all satisfied, then the action is a signal; the reaction is a response; and the overall interaction is communicative. If only conditions (i) and (ii) are satisfied, the action is coercive; and if only conditions (i) and (iii) are satisfied, the action is a cue. (Interestingly, a similar view, expressed in terms of manipulation and exploitation rather than actions and reactions, was developed by some of the pioneers responsible for the early computational models of the evolution of communication and language (e.g. di Paolo, 1997; Oliphant, 1997). However, a detailed discussion of the way in which different definitions of communication have been used in the literature is beyond the scope of this chapter.)

If, for example, an individual deliberately pushes a colleague from her chair, and she then falls to the floor, then this satisfies criteria (i) (the push caused the colleague to fall) and (ii) (the fall was the intended outcome of the push). Criterion (iii) is not satisfied, however, and the interaction is *coercive*. If the pusher's boss, unbeknownst to the pusher, observes the incident, and subsequently acts on it (by punishing the pusher, for example), then that interaction between the pusher and the observer is a *cue*: it satisfies criteria (i) (the push caused the boss to punish the pusher) and (iii) (the function of that punishment was to take account of what had been seen), but not (ii) (since the pusher did not intend the boss to see the incident). Now suppose that there is one further colleague, and that the pusher had actually intended this second observer to see the incident, perhaps as part of a practical joke. Then criterion (ii) is satisfied, and the push is a *signal* to this colleague. Note that these examples also illustrate how the same event can involve several different interactions, each with a different communicative status (see Figure 18.3).

In these examples, the required functionality comes from the intentions of the individuals involved. However, we have in mind a more general notion of functionality (see Millikan, 1984). Specifically, a trait or behaviour's function is the task that it performs that at least in part explains why it is reproduced from one generation to the next (Millikan, 1984). For example, hearts make noise, contribute to body weight and pump blood, amongst other things. Yet it is only the last of these that is their function, since it is their capacity to pump blood, and the contribution this makes to survival, that explains why hearts are reproduced from one generation to the next. This notion of functionality applies to any adaptive system, including (amongst others) human cognition, and natural selection, whose dynamics produce traits and behaviours whose function is to maximise an organism's inclusive fitness (Grafen, 2002, 2006). Hence these definitions work across multiple domains in the sense that they accurately

Figure 18.3 Functional definitions of communication.
This illustration shows one man (in the centre of the image) in three different interactions, each of which is of a different type *vis-à-vis* communication. In one, he is pushing a colleague from her chair. This is **coercion**. In another, he is seen pushing his colleague by a second colleague (on the right of the image), and he intended for this to be the case. This is **communication**; the push is a signal, and the laugh of his colleague is a response. In a third interaction, his boss (on the left of the image) has also seen him pushing his colleague, but this was not the man's intention; indeed, he does not know that his boss has seen him. The pushing is a **cue** for the boss; it guides the boss's future action, but this was not the function of the push. See main text for further discussion.

capture various *prima facie* cases of communication and non-communication, whether these cases occur in human communication (as in Figure 18.3, above), animal communication (see Maynard Smith & Harper, 2003; Scott-Phillips, 2008 for discussion and examples of how this definition applies to animal communication) or elsewhere (e.g. computer-to-computer communication, where the source of the functionality is the software engineer that designed the program). The result is a set of definitions that make no commitments regarding the mechanisms involved in any particular instance of communication, and as such allow for discussion of communication in general terms.

18.5.2 Communication is a matter of effects

Can the above definitions be used to inform the present debates about how best to conceptualise communication? Our definitions, which are centred on the effects that signals have upon receivers, plainly align with the view that communication should be seen, at the most basic level, as a matter of influence (Dawkins & Krebs, 1978; Rendall, *et al.*, 2009). However, this does not mean that information should be rejected from the study of communication (linguistic or otherwise). On the contrary, it is possible to integrate information into our approach (Carazo & Font, 2010). In fact, we suspect (but do not know) the following: that once the functionally interdependent relationship between signal and response described above (subsection 18.5.1) is established, then a necessary corollary will be a correlation between some properties of the signal and some properties of the world (e.g. there is a negative correlation between a deer's size and the formant dispersion of its roar; Fitch & Reby, 2001). If this is correct, then it will be possible to identify something we may wish to term information (others have termed this *functional information*; Carazo & Font, 2010).

Our point is only that functional effects are what lie at the heart of communication, by which we mean: it may be possible to observe and/or quantify information transfer, but we can only do this in a post-hoc way, after we have specified what the effects of a signal are (Scott-Phillips, 2008). Indeed, this is a general point about communication, be it animal communication or human language. First and foremost, signals *do* things. Only once we know what they do can we identify information, conventional meaning and other associated phenomena – since these things simply do not exist until there is functional symbiosis between signals and responses. Effects are methodologically prior.

We will not labour this point, since it has been discussed at length previously, in a range of different contexts (di Paolo, 1997; Blackburn, 2007; Scott-Phillips, 2008; Rendall, *et al.*, 2009), and is discussed further elsewhere in this volume. What we do wish to emphasise is that our insistence that effects are a general, fundamental property of communication aligns with the point made by the discipline of pragmatics that linguistic communication is about more than conventional meaning. Utterances are behaviours that cause others to do things; conventional meanings are simply part of the story of how these effects are achieved (Origgi & Sperber, 2000; Scott-Phillips, 2010a). Indeed, one of the founding texts of pragmatics is entitled *How To Do Things With Words*, with the emphasis very much on the *do* (Austin, 1955). Moreover, there have been recent calls in the pragmatics literature for *all* types of linguistic utterances to be reconceptualised as acts of social influence (Reich, 2010). One of the arguments presented in favour of this view is that it is particularly compatible with evolutionary theory, and hence that it allows for comparison with non-human communication. This emphasis within pragmatics on the social, interactive aspect of linguistic utterances in turn aligns with the ostensive-inferential model of communication, in which utterances have effects only by virtue of the inferences they cause others to make.

18.5.3 Applications and implications

In this subsection we highlight three ways in which the functional approach to communication outlined above has already provided insights into various questions and topics about communication, and speculate about how it may do so in the future.

First, this approach has allowed microbiologists to investigate the communicative nature of bacterial social behaviour (Diggle *et al.*, 2007). Many bacterial cells use a process known as quorum sensing (QS), in which individual cells produce small diffusible molecules, and once the concentration of these molecules reaches a threshold level they bind to their cognate receptors, causing changes in the expression of QS-dependent genes. This allows the population of cells to behave in a coordinated way. The functional approach described above has allowed microbiologists to determine when behaviour associated with QS should be considered communication, and when it should be considered a cue, or coercive. These differences matter, because they make different predictions about how bacterial cells should behave. In addition to this direct application to the study of bacterial social behaviour, the functional approach also allows evolutionary biologists to integrate two disparate literatures: one on the evolution of animal signals, the other on the social behaviour of microbes (Diggle *et al.*, 2007).

Second, a functional account brings the question of how communication systems emerge into focus. Signals and responses depend upon one another to explain their existence, and this interdependence produces a chicken-and-egg problem: how does such mutual dependency emerge in the first place? A recent mathematical model that explicitly adopts the functional approach describes the different ways in which this problem can be overcome, and explains their relative frequency in nature (Scott-Phillips *et al.*, 2012). As such, the functional view has been shown to provide a comprehensive explanation of the ways in which communication systems can emerge.

Third, the functional approach has been used to choose between different theoretical approaches to pragmatics itself. Like any discipline, there are several different theoretical proposals around the central questions – in this case, how listeners are able to interpret utterances in context (and how speakers are able to construct context-appropriate utterances) – and these accounts compete with one another to varying degrees (e.g. Horn, 1984; Sperber & Wilson, 1986; Levinson, 2000; Bara, 2010). One criterion that we can use to choose between them is compatibility with evolutionary theory, and the functional approach has been used to pursue such a project (Scott-Phillips, 2010b).

Finally, we wish to return to one of the major desiderata for future research identified above: the development of a theoretical framework for pragmatics that is richer than a simple ostensive/not-ostensive dichotomy. The reasons why this is desirable were discussed in Section 18.4: although it is still at an early stage, the evidence from the comparative literature on cognition and pragmatics suggests that some non-human primates possess at least some of the cognitive foundations of ostensive communication. Thus if we wish to understand how ostensive communication evolved, we need a theoretical framework and a vocabulary with which to discuss different possible grades of pragmatic competence. We believe that the functional approach described here is well suited to this task, because it is expressed in general terms, and hence makes no particular commitments regarding the mechanisms involved in communication. This means that researchers in this area are able to develop such frameworks unencumbered by terminology that is weighted towards one or another conclusion. In contrast, the conduit metaphor, and other related approaches to communication, can, by virtue of the language they adopt, lead researchers to think about animal communication in terms of representations and other aspects of human communication when the empirical data do not necessarily warrant this (Rendall *et al.*, 2009).

18.6 Concluding remarks

We started this chapter by surveying the history of computational thinking about the evolution of language. Primarily as an approach to simplifying the problem of modelling the complex adaptive systems underpinning language emergence, modellers naturally gravitated to a form of the code model to implement communication among simulated agents. Given the central role that simulations played in the development of evolutionary linguistics, it is not unreasonable to say that the information-centric code model is taken as a default idealisation by many within the field, extending beyond modelling work and into experimental approaches. Whilst we believe that much important progress has been made with this approach, we hope to have demonstrated that it is plainly incomplete, and that there is much to be gained from thinking not merely about *information* but also *influence* and *inference* in the evolution of language. It is likely that future breakthroughs in the field will come when the theoretical work on ostension and inference in language evolution inform modelling and experimental frameworks. This aligns with arguments that at the most basic level, communication is fundamentally a matter of influence, in which signals and responses are functionally symbiotic. This does not mean that we should seek to remove all talk of information – but neither should we confuse information transfer with a complete description of what communication is about.

Acknowledgements

T.S.P. acknowledges financial support from the Leverhulme Trust.

References

Ackley, D. H. & Littman, M. L. (1994). Altruism in the evolution of communication. In R. Brooks & P. Maes, eds., *Artificial Life IV*. Cambridge, MA: MIT Press, pp. 40–48.

Atlas, J. D. (2005). *Logic, Meaning and Conversation: Semantic Underdeterminacy*. Oxford: Oxford University Press.

Austin, J. L. (1955). *How to Do Things With Words*. Oxford: Oxford University Press.

Bara, B. (2010). *Cognitive Pragmatics: The Mental Processes of Communication*. Cambridge, MA: MIT Press.

Batali, J. (1998). Computational simulations of the emergence of grammar. In J. R. Hurford M. Studdert-Kennedy & Knight, C., eds., *Approaches to the Evolution of Language: Social and Cognitive Bases*. Cambridge: Cambridge University Press, pp. 405–426.

Blackburn, P. (2007). *The Code Model of Language: A Powerful Metaphor in Linguistic Metatheory*. SIL International e-books.

Brighton, H., Smith, K. & Kirby, S. (2005). Language as an evolutionary system. *Physics of Life Reviews*, **2**, 177–226.

Briscoe, E. J. (2002). *Linguistic Evolution Through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press.

Byrne, R. W. & Whiten, A. (1989). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*. Oxford: Clarendon Press.

Cangelosi, A. & Parisi, D. (2002). *Simulating the Evolution of Language*. London: Springer Verlag.

Carazo, P. & Font, E. (2010). Putting information back into biological communication. *Journal of Evolutionary Biology*, **23**, 661–669.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Cartmill, E. A. & Byrne, R. W. (2010). Semantics of primate gestures: intentional meanings of orangutan gestures. *Animal Cognition*, **13**, 793–804.

Dawkins, R. & Krebs, J. R. (1978). Animal signals: information or manipulation? In J. R. Krebs & N. B. Davies, eds., *Behavioural Ecology*, 1st edn. Oxford: Blackwell, pp. 282–309.

di Paolo, E. A. (1997). An investigation into the evolution of communication. *Adaptive Behavior*, **6**, 285–324.

Diggle, S. P., Gardner, A., West, S. A. & Griffin, A. S. (2007). Evolutionary theory of bacterial quorum sensing: when is a signal not a signal? *Philosophical Transactions of the Royal Society of London Series B*, **362**, 1241–1249.

Dunbar, R. I. M. (1998). The social brain hypothesis. *Evolutionary Anthropology*, **6**, 178–190.

Dunbar, R. I. M. & Schultz, S. (2007). Evolution in the social brain. *Science*, **317**, 1344–1347.

Fitch, W. T. & Reby, D. (2001). The descended larynx is not uniquely human. *Proceedings of the Royal Society of London B*, **268**, 1669–1675.

Grafen, A. (2002). A first formal link between the Price equation and an optimization program. *Journal of Theoretical Biology*, **217**, 75–91.

Grafen, A. (2006). Optimisation of inclusive fitness. *Journal of Evolutionary Biology*, **238**, 541–563.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan, eds., *Syntax & Semantics III: Speech Acts*. New York: Academic Press, pp. 41–58.

Holbaiter, C. & Byrne, R. W. (2011). The gestural repertoire of the wild chimpanzee. *Animal Cognition*, **14**, 745–767.

Horn, L. R. (1984). Towards a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin, ed., *Meaning, Form, and Its Use in Context: Linguistic Applications*. Washington, DC: Georgetown University Press, pp. 11–42.

Humphrey, N. (1976). The social function of the intellect. In P. Bateson & R. Hinde, eds., *Growing Points in Ethology*. Cambridge: Cambridge University Press, pp. 303–317.

Hurford, J. R. (1989). Biological evolution of the Saussurean sign as a component of the language acquisition device. *Lingua*, **77**, 187–222.

Hurford, J. R. (1991). The evolution of the critical period for language acquisition. *Cognition*, **40**, 159–201.

Hurford, J. R. (2007). *Origins of Meaning*. Oxford: Oxford University Press.

Kirby, S. (2001). Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, **5**, 102–110.

Kirby, S. (2002). Natural language from artificial life. *Artificial Life*, **8**, 185–215.

Kirby, S., Cornish, H. & Smith, K. (2008). Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences USA*, **105**, 10681–10686.

Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.

Levinson, S. C. (2000). *Presumptive Meanings*. London: Longman.

MacLennan, B. & Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, **2**, 161–187.

Maynard Smith, J. & Harper, D. G. C. (2003). *Animal Signals*. Oxford: Oxford University Press.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.

Oliphant, M. (1996). The dilemma of Saussurean communication. *Biosystems*, **37**, 31–38.

Oliphant, M. (1997). Formal approaches to innate and learned communication: laying the foundation for language. Unpublished PhD thesis, Department of Cognitive Science, University of California, San Diego.

Origgi, G. & Sperber, D. (2000). Evolution, communication, and the proper function of language. In P. Carruthers & A. Chamberlain, eds., *Evolution and the Human Mind: Language, Modularity and Social Cognition*. Cambridge: Cambridge University Press, pp. 140–169.

Owren, M. J., Rendall, D. & Ryan, M. J. (2010). Redefining animal signalling: influence versus information in communication. *Biology and Philosophy*, **25**, 755–780.

Pika, S. (2008). Gestures of apes and pre-linguistic human children: similar or different? *First Language*, **28**, 116–140.

Reddy, M. J. (1979). The conduit metaphor: a case of frame conflict in our language about language. In A. Ortony, ed., *Metaphor and Thought*. Cambridge: Cambridge University Press, pp. 284–324.

Reich, W. (2010). The cooperative nature of communicative acts. *Journal of Pragmatics*, **43**, 1349–1365.

Rendall, D., Owren, M. J. & Ryan, M. J. (2009). What do animal signals mean? *Animal Behaviour*, **78**, 233–240.

Scott-Phillips, T. C. (2008). Defining biological communication. *Journal of Evolutionary Biology*, **21**, 387–395.

Scott-Phillips, T. C. (2010a). Animal communication: insights from linguistic pragmatics. *Animal Behaviour*, **79**, e1–e4.

Scott-Phillips, T. C. (2010b). The evolution of relevance. *Cognitive Science*, **34**, 583–601.

Scott-Phillips, T. C. & Kirby, S. (2010). Language evolution in the laboratory. *Trends in Cognitive Sciences*, **14**, 411–417.

Scott-Phillips, T. C., Blythe, R. A., Gardner, A. & West, S. A. (2012). How do communication systems emerge? *Proceedings of the Royal Society of London Series B*, **279**, 1943–1949.

Seyfarth, R. M., Cheney, D. L., Bergman, T. *et al.* (2010). The central importance of information in studies of animal communication. *Animal Behaviour*, **80**, 3–8.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, **27**, 379–423.

Smith, A. D. M. (2005). The inferential transmission of language. *Adaptive Behavior*, **13**, 311–324.

Smith, K. (2004). The evolution of vocabulary. *Journal of Theoretical Biology*, **228**, 127–142.

Sperber, D. & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.

Steels, L. (1999). *Words and Meanings. The Talking Heads Experiment*, Vol. 1. Laboratorium: Antwerp.

Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Science*, **7**, 308–312.

Tomasello, M. (2008). *The Origins of Human Communication*. Cambridge, MA: MIT Press.

Tomasello, M., Carpenter, M., Call, J., Behne, T. & Moll, H. (2005). Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, **28**, 675–691.

Vogt, P. (2002). The physical symbol grounding problem. *Cognitive Systems Research*, **3**, 429–457.

Werner, G. & Dyer, M. (1992). Evolution of communication in artificial organisms. In C. Langton, C. Taylor, D. Farmer & S. Rasmussen, eds., *Artificial Life II*. Redwood City, CA: Addison-Wesley, pp. 659–687.

Commentaries

We agree that what signals *do* is primary, and talk of information can serve this stance in several ways (cf. Horn, 1997). Here are three. First, signals only *do* things on average, so talk of information helps us to articulate the best case (normative/selected-for) scenario. Second, because the effect of the signal itself is only one piece of the puzzle of how it *does* things (the rest being context), talk of information helps us articulate what piece the signal contributes (that's what semantics is). Third, if information is inherent in particular manifestations of the physical world (e.g. sight of prey) that, when they are signals, are evolved behaviours of other animals, then surely the clearest shorthand (*not* metaphor) for speaking of both influences in a unified way is to talk of information.

Horn, A. G. (1997). Speech acts and animal signals. In D. W. Owings, M. D. Beecher & N. Thompson, eds., *Communication. Perspectives in Ethology*, Vol. 12. New York: Plenum Press, pp. 347–358.

Andrew G. Horn and Peter McGregor

Communication without perceiver adaptation. Scott-Phillips and Kirby adopt a pragmatic approach to the information construct, emphasising information as correlations between signals and objects or events in the world. We concur with this view, while noting that such regularity exists whether or not anyone attends to it. As elaborated in our commentary on Scarantino's contribution, a conceptual firewall is therefore needed between information as correlations in the world and observer representations of those correlations. In defining signalling and communication, Scott-Phillips and Kirby require both that signallers be specialised to influence perceiver behaviour and that perceivers be specialised to be so influenced. Thus, the authors argue that an interaction is merely coercive if an actor relies on its own physical power to influence the behaviour of another. In contrast, communication involves a signaller-actor who capitalises on the senses and muscles of a perceiver-other. Only if the latter has evolved to respond to the actor's behaviour can one conclude that true communication has occurred. As Scarantino argues, however, even in cases of clearly coercive behaviour, possible perceiver specialisation cannot be ruled out. In fact, one can argue that such cases are particularly likely to create selection pressures shaping reactor responses over evolutionary time. We must also point out that while it is difficult to rule out possible reactor specialisation in coercion, ruling it in even in archetypal cases of communication can be still more challenging. In predator-specific alarm signalling in vervet monkeys, for example, adaptive specialisation is clearly evident on the production side. Vervet infants produce recognisable, predator-appropriate alarm calls without first needing to experience how others use them or to practise making these sounds. In contrast, there is no evidence of specialisation in responding to the calls. The same infants that competently produce the vocalisations initially show no sign of knowing how to respond when others give them. Responding is acquired over the first year of life through evidently standard, generalised learning mechanisms. In other words, there is no indication of perceiver specialisation even in this quintessential instance of animal signalling.

Michael J. Owren and Drew Rendall

The distinction between the 'code model' and the 'ostensive-inferential model' of communication depends primarily on whether a signal uniquely specifies the signaller's meaning or just provides evidence for it: in other words, whether the correlation between signal and meaning is 1.0 or less than 1.0 (but not equal to 0). In both models the meaning of a signal is something (perhaps a state of mind) other than the correlation. This approach is explicit in those simulations of communication that score a successful interaction when the receiver interprets a signal in the same way the signaller does. The authors' new model of communication makes an important improvement by incorporating 'functional effects' (the consequences of producing or responding to a signal). It is then easy to show that communication must have benefits for both signallers and receivers. It is important to emphasise, however, that signallers and receivers must receive benefits only on average. In this chapter, the 'on average' is missing from the accounts of 'functionality'.

R. Haven Wiley

Response

I thank the commentators for their thoughtful and engaged responses to the chapter I co-authored with Simon Kirby. I have divided my responses into three main areas: (i) the difference between the code and ostensive-inferential models of communication; (ii) the role of information in communication; and (iii) various issues around the functionality of signals and responses.

First, I wish to clarify the difference between the code and ostensive-inferential models of communication. **Wiley** interprets the distinction between the two as whether or not a signal uniquely specifies the correct meaning. This is not quite right. In some computational models agents have different meaning–form mappings, and their responses are determined probabilistically (e.g. Smith, 2005). In these models signals do not uniquely specify meaning, but these are still code models – because the basic paradigm remains one that involves 'looking up' signal–meaning mappings, both in production and reception. The ostensive-inferential model, in contrast, involves the expression and recognition of intentions (Sperber & Wilson, 1986). It is this difference in the way that communication is achieved that distinguishes the two models. One view that we wished to express in our chapter was that a potentially profitable avenue for future research is the development of agent-based or mathematical models of the evolution of language that are ostensive and inferential in this way.

Second, I should respond to the comments of both **Horn and McGregor** (H&M) and **Owren and Rendall** (O&R) on information. H&M appear to suggest, with a parenthetical aside, that we think that use of the term *information* is metaphorical, and so I should clarify that this is not the case. What is metaphorical is the code model of how communication works; this is not the same thing as information. Otherwise, H&M, and also O&R, agree with us that although information transfer is, on its own, an insufficient characterisation of communication, there is information in whatever correlations exist between organismic behaviour and features of the world. For H&M, this is a reason to use information as shorthand for the different ways in which external influences can affect organisms. In contrast, for O&R, it is a reason to be very careful about how we use the term information. Both are correct: if we simply want to study how the inputs received by one organism affect its behaviour, talk of information is good shorthand for both signals and other sources – but at the same time, if we are focused on communication, we should not, as our chapter emphasised, equate communication with the use of information to inform behaviour.

Finally, the commentaries raise a variety of issues around the functionality of signals (and responses). The first is the need for the clause 'on average' in our definition of communication. Following Maynard Smith and Harper (2003), we defined communication in functional terms: as any pair of behaviours (or structures) that have symbiotic functionality. This implies that signals should on average be beneficial for both parties. Both Wiley and H&M point out that we neglected to mention this, and they are correct: 'on average' is a necessary clause in nearly all population-level analysis of social behaviour (Davies, Krebs & West, 2012). Our definition of communication is no exception.

As Wiley comments, our definition makes it explicit that communication must (on average) be beneficial to *both* parties. O&R, however, resist this conclusion, and argue that communication can be (but need not necessarily be) maladaptive for the receiver. As I argued in my commentary on their chapter, I do not see how this can be the case. Communication is a symbiotic interaction between two organisms, and if it were not beneficial for both parties then it would collapse (Scott-Phillips *et al.*, 2012). In the specific case of receivers, this means that if responses were maladaptive (on average) then they would be selected against, and the interaction would no longer be communicative.

O&R argue against this conclusion by pointing to some uncontroversial instances of communication in which there is no receiver specialisation (an argument that is reflected in the title of their commentary). They illustrate this with the example of vervet calls, the response to which appears to be acquired through general learning mechanisms. I do not see the relevance of this. Our definition of communication is based upon functionality, not on the specifics of any mechanism, specialised or otherwise. The mechanisms involved in responses produce adaptive outcomes (hiding from predators), and it is at least in part because of this that they are successfully passed on from one generation to the next. This makes them functional as response mechanisms, and that is both necessary and sufficient for our definition (see Millikan, 1984, for extensive discussion of what it means for a trait to have a function). If instead the consequences of the interaction are neutral for the 'receiver', or at least insufficiently strong to undergo selection, then that would be a case of coercion.

Davies, N. B., Krebs, J. R. & West, S. A. (2012). *An Introduction to Behavioural Ecology*, 4th edn. Oxford: Wiley-Blackwell.

Millikan, R G. (1984). *Language, Thought, and Other Biological Categories*. Harvard, MA: MIT Press.

Scott-Phillips, T. C., Blythe, R. A., Gardner, A. & West, S. A. (2012). How do communication systems emerge? *Proceedings of the Royal Society of London Series B*. 279, 1943–1949.

Smith, A. D. M. (2005). The inferential transmission of language. *Adaptive Behavior*, **13**, 311–324.

Sperber, D. & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.

Thomas C. Scott-Phillips