

Communication and language have always been key topics for research at the interface of cognition and culture. Rightly so, given the central role that linguistic communication plays in human social and cultural life. In fact, communication and language are doubly important, since they occupy both sides of the cognition and culture coin. On the one side is a cognitive ability, to engage in linguistic communication in the first place. On the other side are cultural objects, namely languages themselves, which are collections of communicative conventions shared within a population. This double importance is reflected in the ambiguity of the phrase “language evolution”, which is used both to describe the study of how humans evolved to communicate in the ways that they do, and the study of how languages themselves evolve, culturally, to take the various forms that they do.

In 1866 the Société de Linguistique de Paris declared that it would no longer consider correspondence on the topic of language origins (“La Société n’admet aucune communication concernant, soit l’origine du langage – soit la création d’une langue universelle”). This “ban” is often said to have curtailed language evolution research, only for it to be reawakened in the 1990s, but this clichéd history has little truth to it. A great deal of language evolution research occurred between 1866 and 1990. Darwin himself speculated on the origins of language just a few years after the Parisian edict, and several 20th-century research agendas directly address language origins. The clearest and most well-known example is the many attempts to teach human language to non-human apes. Contrary to common assertion, 1990 was not year zero for language evolution research.

It is however fair to say that it was during the 1990s that something like an academic field of language evolution began to emerge. Researchers from different disciplines began to share ideas, the first conferences focused on the topic took place, and some pioneering individuals began to consider it their main research interest. This growth has continued since. We should not exaggerate the extent of this, but we can reasonably say that there is now a field of language evolution, it is here to stay, there is some outstanding science done in its name (some less than outstanding science too), and also that the various disciplines concerned with language and communication are increasingly open to evolutionary perspectives. This is most clear within linguistics itself. For many years mainstream linguistics was at best uninterested in evolutionary approaches, and at worst actively hostile to them. This is no longer the case, as a browse of the major journals of the discipline will quickly show. It is no longer unusual to find work that is informed by an evolutionary perspective in one way or another.

This influence of language evolution on mainstream linguistics has been mirrored by a corresponding influence in the other direction, although the impact of this is more subtle, and less easy to observe. For most of its history as an academic discipline, linguistics’ principal intellectual concern has been with the form and structure of languages. The classic divisions between different levels of language structure are nicely illustrated in this figure:

[[Major levels of linguistic structure.svg](#)] 

All but the outermost circle of this figure describe items – sounds, phonemes, words, phrases, and literal meaning – that are either discrete in character, or can be treated as such. The sixth, outermost circle is where the linguistic rubber hits the communicative road, and this interface with the outside world brings with it a raft of issues that complicate linguistic analysis. Depending on one’s specific questions, it can be productive to abstract away from these issues, but the unintended consequence of doing so is that, despite good intentions and well-meaning lip-service, pragmatics is kept at the margins of linguistics. The fact that it is concerned with the boundary between language and the outside world is, ironically, a major reason why pragmatics is kept on the periphery of the discipline itself.

This marginalisation has affected the direction of language evolution research, and continues to do so. If you doubt this, turn to the index of the [Oxford Handbook of Language Evolution](#), a book billed

(accurately) as “a wide-ranging summation of work in all the disciplines involved [in language evolution]”. There you will find 213 pages listed under ‘syntax’ and related terms, 100 pages listed under ‘phonetics’, ‘phonology’ and related terms, but just 8 pages under ‘pragmatics’. Or, alternatively, take a look at the lists of plenary speakers from the ten [Evolang](#) conferences that have taken place to date. You will find a relative dearth of pragmaticists. Two especially conspicuous omissions, particularly in comparison with the relevance of some who have delivered Evolang plenaries, are Stephen Levinson and Dan Sperber, both high-profile pragmaticists who have written extensively about the origins and evolution of human communication and language.

This neglect of pragmatics is, I believe, a profound mistake. For some fields (theoretical syntax, say) and for some questions, abstraction away from the messy realities of human social interaction, and towards the more abstract, idealised world of discrete linguistic items, can be a reasonable and profitable agenda. However, I do not believe that language evolution is such a field, at least not in general, and this is true whether one is concerned with its cognitive or its cultural dimensions. How languages are used in communication is, instead, critical to any evolutionary explanation of why humans are able to communicate in this way, and of why languages themselves take the forms that they do. If I am right about this, then rather than being peripheral, pragmatics should be a key concern of language evolution – and indeed of cognition and culture studies in general. One way to read *Speaking Our Minds* is as a demonstration of just how much we can learn about language evolution by taking pragmatics seriously. Put simply, it is my belief that with a pragmatic perspective, it is possible to develop cogent, compelling answers to all the major questions asked about language evolution (and I think this is true even if one does not agree with the specifics of my own answers).

Let me expand a little, and in doing so provide a little background for my own claims. In *Relevance Theory*, not to mention many publications since, Dan Sperber and Deirdre Wilson argued that linguistic communication exists on a continuum with other, non-linguistic forms of communication, canonical examples of which include points, shrugs, and other non-verbal gestures. The overall category here is that of ostensive communication: communication that involves the expression and recognition of intentions. These intentions are, specifically, communicative intentions (which we can roughly gloss as an intention to make apparent to the audience that one is trying to communicate) and informative intentions (which we can roughly gloss as an intention to make apparent to the audience what one is trying to communicate). Ostensive communication can be used for a great many communicative ends (there is much you can do just with non-linguistic grunts and gestures), but its expressivity is hugely increased by the addition of words, grammar, and the other communicative conventions that collectively comprise a language. I can make a request of others by ostensively pushing unchopped vegetables in their direction, but with specific conventions I can make requests about things remote in time and space. Linguistic communication is, then, a special case of ostensive communication, namely one in which expressivity is hugely increased by the existence of shared communicative conventions.

If all this is correct, then already a pragmatic perspective has earned its keep, since it makes clear what the two most general questions for language evolution should be. First, how and why did humans evolve ostensive communication (and do any other species communicate ostensively?). Second, how do collections of communicative conventions develop, and how and why do they transition towards the forms that they do? The short versions of my answers to these questions are that ostensive communication is uniquely human, it evolved as a by-product of enhanced social intelligence, conventions emerge through interaction and use, and they gravitate towards particular forms through a process of cultural attraction, in order to most closely fit the goals and dispositions of language users. These answers collectively motivate my subtitle: *Why Human Communication Is Different, & How Language Evolved To Make It Special*.

Chapter 1: Two Approaches to Communication

In the first chapter of *Speaking Our Minds*, I distinguish ostensive communication from code model communication. Code model communication is achieved through pairs of association: one between a state of the world and a signal, and the other between the signal and a response. Following arguments developed by Sperber and Wilson in *Relevance Theory*, I defend the claim that ostensive communication cannot be reduced to pairs of associations, and hence that the difference between ostensive communication and code model communication is not one of degree, but one of kind. More precisely, the difference between these two types of communication is at bottom a difference in the sort of mechanism that makes communication possible in the first place. Ostensive communication is made possible by mechanisms of social cognition, whereas code model communication is made possible by mechanisms of association. (Whether or not any species other than humans communicates ostensively, or whether all animal communication is code model communication, is an empirical question that I defer until chapter four.)

Where does this distinction between ostensive and code model communication leave language and linguistic communication? Clearly associations are involved in some way: there is a linguistic code. However, these associations do not function to make linguistic communication possible in the first place (and hence linguistic communication does not operate according to the code model). As I discussed above, linguistic associations (words, grammar, etc) instead function to make ostensive communication more precise and more expressive than it otherwise would be. I adopt the labels natural codes and conventional codes to distinguish between, respectively, codes that function to make communication possible in the first place, and codes that function to make a different type of communication more expressive. Languages are not natural codes, but instead (sets of) conventional codes.

Chapter 2: The Emergence of Communication Systems

How do different types of communication system emerge? I argue in chapter two that the process is quite different for the two different types of communication distinguished in chapter one.

Most work on the origins of communication systems focuses, quite reasonably, on code model communication. Here, there is a chicken-and-egg problem: which comes first, the signal or the response? After all, one without the other is pointless, and natural selection does not act with foresight. This problem may be especially serious for the emergence of combinatorial codes, in which two existing signals are combined to achieve an effect that is different to the sum of the effects of the two component parts. The obvious solution to the chicken-and-egg problem is that one of the two behaviours first evolves for reasons independent of communication, and this then provides the right selective environment for the other. The processes of emergence described in the animal communication literature, both in theory and in the empirical data, match this prediction.

Ostensive communication is quite different. The ability to express and recognise intentions means that any behaviour can, in principle, be used communicatively, so long as it is produced in an ostensive way. This means that there is no chicken-and-egg problem to contend with, and new signals can be created as and when required. Some of my own experimental work shows this process in action, but it can also be observed in natural data, such as in the emergence of homesign (novel systems created by deaf children of hearing parents).

Chapter 3: Cognition and Communication

I said above that ostensive communication is made possible by mechanisms of social cognition (by definition it involves the expression and recognition of intentions). Chapter three discusses these

mechanisms in more detail. I begin by coining the term pragmatic competence: the ability to use ostensive communication in a competent way. As signallers, this is the ability to produce the right sort of behaviour to express one's intended meaning; as listeners, it is the ability to make the right sort of inferences about that behaviour. The question is: what cognitive mechanisms make pragmatic competence possible? In answering this question, I take as my starting point Relevance Theory, which I believe provides a bone fide scientific paradigm for pragmatics. I argue in particular that recursive mindreading – the ability not only to infer what others believe (think, desire, etc), but also to infer what others believe about what further others believe – is a necessary component of any computational description of what is involved in ostensive communication.

There is a tension between, on the one hand, the theoretical cogency of these arguments, and on the other, the intuition that recursive mindreading is a cognitively demanding activity, certainly too demanding to be involved in everyday communicative interaction. To date, most researchers who have grappled with this problem have chosen to accept the intuition, or some version of it, and argue against the theory. That is to say, they have tried to develop arguments that ostensive communication in fact requires (supposedly) less demanding cognitive resources. I adopt the other approach: I accept the theory, and argue against the intuition that recursive mindreading is cognitively demanding. In particular, I discuss how these arguments can be squared with the data on the mindreading abilities of children, who demonstrate pragmatic competence from a young age.

Chapter 4: The Evolution of Ostensive Communication

Does any other species communicate ostensively? Do human children? There is a sizeable literature on intentional communication in other species, but intentional communication is not, I believe, the same thing as ostensive communication. Intentional communication involves signals being used in a goal-directed way (this is clear from looking at the criteria used to test for this), whereas ostensive communication involves the expression and recognition of intentions. Consequently, even if a species or group communicates intentionally, this does not mean that what is expressed are themselves intentions (which is what ostensive communication consists of, by definition). How, then, can we identify ostensive communication?

Since by definition it consists of the (1) expression and (2) recognition of both (a) communicative and (b) informative intentions, what we must look for is evidence of each of these four behaviours (i.e. 1a, 1b, 2a, and 2b). We have good experimental evidence for three of the four in human infants (the exception is the recognition of informative intentions), but the corresponding experiments in non-human species have not, to my knowledge, been conducted. It is therefore possible that other species do communicate ostensively, but my best guess, which I defend in this chapter, is that this is unlikely. It is more likely that ostensive communication is uniquely human. Why did it evolve in humans and only humans? My suggestion is that, since it critically depends upon recursive mindreading abilities, ostensive communication is a byproduct of increased skills of social cognition, which were selected for as an adaptation to the extremely social nature of human life. Once ostensive communication had emerged, it likely led to the selection of traits which allow ostensive communication to occur more easily than it otherwise would. An example is white sclera (the whites of the eyes), which allow us to more easily infer the direction of eye gaze, which is an important means of ostensive expression. Natural pedagogy may be another such adaptation.

Chapter 5: Building a Language

As I noted in the opening remarks above, ostensive communication is a powerful tool, but its expressive potential is increased hugely by the creation of communicative conventions. How, then, do conventions emerge and become linguistic? Over the past ten or so years a literature has emerged in which these questions are studied in the laboratory. A variety of different experimental

paradigms have been developed. Individuals motivated to communicate typically develop conventions quickly, although not without hiccups, and it seems likely that the same was true for the first “language” users. Over time, as these conventions are used repeatedly by the same individuals, or as new individuals acquire them, the sorts of structural properties that makes a system recognisably linguistic begin to emerge. A similar process has been observed in naturalistic settings, such as in the recent emergence of new sign languages in Nicaragua and elsewhere. These findings, and others, have been used to motivate a partial alternative to historically influential explanations of language structure. According to the traditional view, language structure is the consequence of a specialised faculty of language that constrains language form, and which allows humans to acquire and use languages. The alternative hypothesis is that language structure is a consequence of repeated instances of use and acquisition. This turns the previous view on its head: languages fit language users, not the other way around.

I believe this finding is best understood within the broad framework of Cultural Attraction Theory, which argues that, in the process of propagation, mental representations (e.g. beliefs, knowledge, intentions) and their public expressions (e.g. words, behaviours, artefacts) are non-randomly modified, in the direction of a better fit with the goals and dispositions of the human mind. They thus gravitate towards particular forms, called attractors, and away from others. This is attraction in its technical sense, taken from dynamical systems theory. Identification of the factors that influence this process can provide genuinely casual explanations of why cultural items take the forms that they do. The field in which this perspective has been most fruitfully applied to date is the cognitive science of religion, but this is certainly not the only example, and there is much potential to apply it more broadly. In the case of languages, models and experiments, such as those mentioned above, together suggest that two especially important factors are learnability (languages must be learnable by new users) and expressivity (languages must be able to express the meanings their speakers want them to). Speaking Our Minds presents these findings in terms of cultural attraction, but I probably could and should have stressed more clearly that by framing language evolution in this way, we can link explanations of structure in language with explanations of structure in other cultural domains.

Chapter 6: Evolutionary Adaptation

Chapter six aims to provide some answers to the following question: How should an adaptationist think about human communication and language?

To answer this question, it is important to keep in mind what the natural objects of study are here. Linguistic communication is not such a natural object of study, since there is no natural dividing line between it and other forms of ostensive communication. What might be a natural object of study is an innate cognitive mechanism – sometimes called a Universal Grammar – without which we would not be able to acquire and use languages. I say that this only “might be” a natural object of study simply because whether such a mechanism actually exists is a disputed and much vexed issue, on which I am personally agnostic. However, if there is such a mechanism, which shows signs of design for language acquisition and use, then the Darwinian conclusion that it is a biological adaptation seems inescapable. Such a mechanism could have been selected for once the use of communicative conventions, themselves built to enhance the efficacy of ostensive communication, has become widespread within a population.

Although linguistic communication is not a natural object of study, ostensive communication is such an object, and its ultimate evolutionary function is social navigation. Put simply, speakers speak in order to mentally manipulate their audience, and audiences listen in order to gain access to the minds of speakers. These direct functions can of course be subdivided into numerous derived functions, such as gossip, courtship, hunting, and all the other ends we use language and communication for. (It is important not to conflate direct and derived functions, which are two

distinct levels of analysis.)

Finally, an important adaptationist question is what keeps human ostensive communication evolutionarily stable. My answer is the banal one that there are social costs to dishonesty that often outweigh the potential benefits. If we lie too often or in a too cynical way, our reputation will lead to rejection in future collaborations - a potentially severe cost for a species as social as humans. Some more elaborate explanations have been proposed, but these are typically based on misunderstandings of the underlying evolutionary theory.

Epilogue: The Big Questions Answered

The final chapter is very short. It lists the nine questions that leading researchers have proposed as the most important for language evolution. These questions include, for instance, "Why do only humans have language?", and "Where does language structure come from?". I then summarise how, unlike any previous work, *Speaking Our Minds* has answered all of them.