

Language as an Adaptation to the Cognitive Niche*

Steven Pinker

Introduction

This chapter outlines the theory (first explicitly defended by Pinker and Bloom 1990), that the human language faculty is a complex biological adaptation that evolved by natural selection for communication in a knowledge-using, socially interdependent lifestyle. This claim might seem to be anyone's first guess about the evolutionary status of language, and the default prediction from a Darwinian perspective on human psychological abilities. But the theory has proved to be controversial, as shown by the commentaries in Pinker and Bloom (1990) and the numerous debates on language evolution since then (Fitch 2002; Hurford et al. 1998).

In the chapter I will discuss the design of the language faculty, the theory that language is an adaptation, alternatives to the theory, an examination of what language might be an adaptation for, and how the theory is being tested by new kinds of analyses and evidence.

The Design of Human Language

The starting point in an analysis of the evolution of language must be an analysis of language itself (for other overviews, see Bickerton 1990; Jackendoff 2002; Miller 1991). The most remarkable aspect of language is its *expressive power*: its ability to convey an unlimited number of ideas from one person to another via a structured stream of sound. Language can communicate

*Supported by NIH grant HD-18381. I thank Morten Christiansen and the members of his seminar on the evolution of language for helpful comments on an earlier draft.

anything from soap opera plots to theories of the origin of the universe, from lectures to threats to promises to questions. Accordingly, the most significant aspects of the language faculty are those that make such information transfer possible (Pinker 1994; 1999). The first cut in dissecting the language faculty is to separate the two principles behind this remarkable talent.

Words

The first principle underlies the mental lexicon, a finite memorized list of words. As Ferdinand de Saussure pointed out, a word is an arbitrary sign: a connection between a signal and a concept shared by the members of the community. The word *duck* does not look like a duck, walk like a duck, or quack like a duck, but I can use it to convey the idea of a duck because we all have learned the same connection between the sound and the meaning. I can therefore bring the idea to mind in a listener simply by making that noise. If instead I had to shape the signal to evoke the thought using some perceptible connection between its form and its content, every word would require the inefficient contortions of the game of charades.

The symbols underlying words are bidirectional. Generally, if I can use a word I can understand it when someone else uses it, and vice versa. When children learn words, their tongues are not moulded into the right shape by parents, and they do not need to be rewarded for successive approximations to the target sound for every word they hear. Instead, children have an ability, upon hearing somebody else use a word, to know that they in turn can use it to that person or to a third party and expect to be understood.

Grammar

Of course, we do not just learn individual words; we combine them into larger words, phrases, and sentences. This involves the second trick behind language, grammar. The principle behind grammar was articulated by Wilhelm von Humboldt as 'the infinite use of finite media'. Inside every language user's head is a finite algorithm with the ability to generate an infinite number of potential sentences, each of which corresponds to a distinct thought. For example, our knowledge of English incorporates rules that say 'A sentence may be composed of a noun phrase (subject) and a verb phrase (object)' and 'A verb phrase may be composed of a verb, a noun phrase (ob-

ject), and a sentence (complement)'. That pair of rules is *recursive*: a phrase is defined as a sequence of phrases, and one or more of those daughter phrases can be of the same kind as the mother phrase. This creates a loop that can generate sentences of any size, such as *I wonder whether she knows that I know that she knows that he thinks she is interested in him*. By means of generating an infinite number of sentences, we can convey an infinite number of distinct thoughts (see also Studdert-Kennedy and Goldstein, Chapter 13 below), since every sentence has a different meaning (most linguists believe that true synonymy is rare or nonexistent).

Grammar can express an astonishing range of thoughts because our knowledge of grammar is couched in abstract categories such as 'noun' and 'verb' rather than concrete concepts such as 'man' and 'dog' or 'eater' and 'eaten' (Pinker 1994; 1999). This gives us an ability to talk about new kinds of ideas. We can talk about a dog biting a man, or, as in the journalist's definition of 'news', a man biting a dog. We can talk about aliens landing in Roswell, or the universe beginning with a big bang, or Michael Jackson marrying Elvis's daughter. The abstractness of grammatical categories puts no restriction on the content of sentences; the recursive, combinatorial nature of grammar puts no limits on their complexity or number.

A grammar comprises many rules, which fall into subsystems. The most prominent is *syntax*, the component that combines words into phrases and sentences. One of the tools of syntax is linear order, which allows us to distinguish, say, *Man bites dog* from *Dog bites man*. Linear order is the most conspicuous property of syntax, but it is a relatively superficial one. Far more important is *constituency*. A sentence has a hierarchical structure, which allows us to convey complex propositions consisting of ideas embedded inside ideas. A simple demonstration comes from an ambiguous sentence such as *On tonight's program Dr Ruth will discuss sex with Dick Cavett*. It is composed of a single string of words in a particular order but with two different meanings, which depend on their constituent bracketings: [*discuss*] [*sex*] [*with Dick Cavett*] versus [*discuss*] [*sex with Dick Cavett*]. Of course, most sentences in context are not blatantly ambiguous, but ambiguity illustrates the essential place of constituency in interpreting meaning from sentences. As with other symbolic systems that encode logical information, such as arithmetic, logic, and computer programming, it is essential to get the parentheses right, and that's what phrase structure in grammar does.

Syntax also involves *predicate–argument* structure, the component of language that encodes the relationship among a set of participants (Pinker 1989). To understand a sentence one cannot merely pay attention to the order of words, or even the way they are grouped; one has to look up information associated with the predicate (usually the verb) which specifies how its arguments are placed in the sentence. For example, in the sentences *The man feared the dog* and *The man frightened the dog*, the word *man* is the subject in both cases, but its semantic role differs: in the first sentence the man experiences the fear; in the second he causes it. In understanding a sentence, one has to look up information stored with the mental dictionary entry of the verb and see whether it says (for instance) ‘my subject is the one experiencing the fear’ or ‘my subject is the one causing the fear’.

A fourth trick of syntax is known as *transformations*, *movement*, or *binding traces*. Once one has specified a hierarchical tree structure into which the words of a sentence are plugged, a further set of operations can alter it in precise ways. For example, the sentence *Dog is bitten by man* contains the verb *bite*, which ordinarily requires a direct object. But here the object is missing from its customary location; it has been ‘moved’ to the front of the sentence. This gives us a way of shifting the emphasis and quantification of a given set of participants in an event or state. The sentences *Man bites dog* and *Dog is bitten by man* both express the same information about who did what to whom, but one of them is a comment about the man and the other is a comment about the dog. Similarly, sentences in which a phrase is replaced by a *wh*-word and moved to the front of a sentence, such as *Who did the dog bite?*, allow the speaker to seek the identity of one of the participants in a specified event or relationship. Transformations thus provide a layer of meaning beyond who did what to whom; that layer emphasizes or seeks information about one of the participants, while keeping constant the actual event being talked about.

Syntax, for all that complexity, is only one component of grammar. All languages have a second combinatorial system, *morphology*, in which simple words or parts of words (such as prefixes and suffixes) are assembled to produce complex words. The noun *duck*, for example, comes in two forms—*duck* and *ducks*—and the verb *quack* in four—*quack*, *quacks*, *quacked*, and *quacking*. In languages other than English morphology can play a much greater role. In Latin, for example, case suffixes on nouns convey information about who did what to whom, allowing one to scramble the left-to-

right order of the words for emphasis or style. For example, *Canis hominem mordet* and *Hominem canis mordet* (different orders, same cases) have the same non-newsworthy meaning, and *Homo canem mordet* and *Canem homo mordet* have the same newsworthy meaning.

Language also embraces a third combinatorial system called *phonology*, which governs the sound pattern of a language. In no language do people form words by associating them directly with articulatory gestures like a movement of the tongue or lips. Instead, an inventory of gestures is combined into sequences, each defining a word. The combinations are governed by phonological rules and constraints that work in similar ways in all languages but whose specific content people have to acquire. English speakers, for example, sense that *bluck* is not a word but could be one, whereas *nguck* is not a word and could not be one (though it could be a word in other languages). All languages define templates for how words may be built out of hierarchically nested units such as feet, syllables, vowels and consonants, and features (articulatory gestures). Interestingly, whereas syntax and morphology are semantically compositional—one can predict the meaning of the whole by the meanings of the elements and the way they are combined—this is not true of phonology. One cannot predict the meaning of *duck* from the meaning of /d/, the meaning of /ʌ/, and the meaning of /k/. Phonology is a combinatorial system that allows us to have large vocabularies (e.g. 100,000 words is not atypical for an English speaker) without having to pair each word with a distinct noise. The presence of these two kinds of discrete combinatorial systems in language is sometimes called *duality of patterning*.

Phonology also contains a set of adjustment rules which, after the words are defined and combined into phrases, smooth out the sequence of articulatory gestures to make them easier to pronounce and comprehend. For instance, one set of rules in English causes us to pronounce the past-tense morpheme *-ed* in three different ways, depending on whether it is attached to *jogged*, *walked*, or *patted*. The adjustment for *walked* keeps the consonants at the end of a word either all voiced or all unvoiced, and the adjustment for *patted* inserts a vowel to separate two *d*-like sounds. These adjustments often function to make articulation easier or speech clearer in a way that is consistent across the language, but they are not merely products of a desire to be lazy or clear. These two goals are at cross purposes, and the rules of phonology impose shared conventions on the speakers of a language as to exactly when one is allowed to be lazy in which way.

Interfaces of Language With Other Parts of the Mind

Grammar is only one component of language, and it has to interface with at least four other systems of the mind: perception, articulation, conceptual knowledge (which provides the meanings of words and their relationships), and social knowledge (how language can be used and interpreted in a social context). While these systems also serve non-linguistic functions, and may have been carried over from earlier primate designs, at least some aspects of them may have evolved specifically to mesh with language. A likely example is the vocal tract: Darwin pointed to the fact that in humans every mouthful of food has to pass over the trachea, with some chance of getting lodged in it and causing death by choking. The human vocal tract has a low larynx compared to those of most other mammals, an arrangement that compromises a number of physiological functions but allows us to articulate a large range of vowel sounds. Lieberman (1984) has plausibly argued that physiological costs such as the risk of death by choking were outweighed in human evolution by the benefit of rapid, expressive communication.

Is Language an Adaptation?

In the biologist's sense of the word, an 'adaptation' is a trait whose genetic basis was shaped by natural selection (as opposed to the everyday sense of a trait that is useful to the individual). What are the alternatives to the theory that language is an adaptation? And what are the reasons for believing it might be one?

Is Language a Distinct Part of the Human Phenotype?

One alternative is that language is not an adaptation itself, but a manifestation of more general cognitive abilities, such as 'general intelligence', 'a symbolic capacity', 'cultural learning', 'mimesis', or 'hierarchically organized behaviour' (see e.g. Bates et al. 1991; Deacon 1997; Tomasello 1999). If so, these more general cognitive capacities would be the adaptation.

These alternatives are difficult to evaluate, because no one has spelled out a mechanistic theory of 'general intelligence' or 'cultural learning' that is capable of acquiring human language. Intelligence, learning, symbol comprehension, and so on do not happen by magic but need particular mech-

anisms, and it is likely that different mechanisms are needed in different domains such as vision, motor control, understanding the physical and social worlds, and so on (Pinker 1997). The ability to acquire and use the cultural symbols called 'language' may require learning mechanisms adapted to that job. Attempts to model the acquisition of language using general-purpose algorithms such as those in traditional artificial intelligence or connectionist neural networks have failed to duplicate the complexity of human language (Pinker 1979; Pinker 1999; Pinker and Prince 1988).

Though it is hard to know exactly what is meant by terms like 'cultural learning' or 'general intelligence', one can see whether mastery of language in the human species resembles abilities that are unambiguously culturally acquired, like agricultural techniques, chess skill, knowledge of government, and mathematical expertise, or whether it looks more like a part of the standard human phenotype, like fear, humor, or sexual desire. Some very general properties of the natural history of language suggests that the latter is more accurate (see Jackendoff 2002; Lightfoot and Anderson 2002; Pinker 1994).

First, language is universal across societies and across neurological normal people within a society, unlike far simpler skills like farming techniques or chess. There may be technologically primitive peoples, but there are no primitive languages: the anthropologists who first documented the languages of technologically primitive societies a century ago were repeatedly astonished by their complexity and abstractness (Voegelin and Voegelin 1977). And despite stereotypes to the contrary, the language of uneducated, working-class, and rural speakers has been found to be systematic and rule-governed, though the rules may belong to dialects that differ from the standard one (Labov 1969; McWhorter 2002).

Second, languages conform to a universal design. A language is not just any conceivable code that maps efficiently from sound to meaning. The design specifications listed in the preceding section—and, indeed, far more subtle and complex properties of grammar—can be found in all human languages (Baker 2001; Comrie 1981; Greenberg et al. 1978; Hockett 1960).

A third kind of evidence is the ontogenetic development of language. Children the world over pass through a universal series of stages in acquiring a language (Brown 1973; Ingram 1989; Pinker 1994). That sequence culminates in mastery of the local tongue, despite the fact that learning a language requires solving the daunting problem of taking in a finite sample of sentences (speech from parents) and inducing a grammar capable

of generating the infinite language from which they were drawn (Pinker 1979; 1984). Moreover, children's speech patterns, including their errors, are highly systematic, and can often be shown to conform to linguistic universals for which there was no direct evidence in parents' speech (Crain 1992; Gordon 1985; Kim et al. 1994).

A fourth kind of evidence also comes from the study of language acquisition. If children are thrown together without a pre-existing language that can be 'culturally transmitted' to them, they will develop one of their own. One example, studied by Bickerton, comes from the polyglot slave and servant plantations in which the only lingua franca among adults was a pidgin, a makeshift communicative system with little in the way of grammar. The children in those plantations did not passively have the pidgin culturally transmitted to them, but quickly developed creole languages, which differ substantially from the pidgins and which have all the basic features of established human languages (Bickerton 1981). Another example comes from deaf communities, where complex sign languages emerge quickly and spontaneously. A recent study in Nicaragua has tracked the emergence of a complex sign language in little more than a decade, and has shown that the most fluent and creative users of the language were the children (Senghas and Coppola 2001).

A fifth kind of evidence is that language and general intelligence, to the extent we can make sense of that term, seem to be doubly dissociable in neurological and genetic disorders. In aphasia and in the genetically caused developmental syndrome called Specific Language Impairment, intelligent people can have extreme difficulties speaking and understanding (Leonard 1998; Siegal et al. 2001; van der Lely and Christian 1998). Conversely, in a number of retardation syndromes, such as Williams syndrome and the sequelae of hydrocephalus, substantially retarded children may speak fluently and grammatically and do well on tests of grammatical comprehension and judgement (Clahsen and Almazan 1998; Curtiss 1989; Rossen et al. 1996). Few of these dissociations are absolute, with language or non-linguistic cognition completely spared or completely impaired. But the fact that the two kinds of abilities can dissociate quantitatively and along multiple dimensions shows that they are not manifestations of a single underlying ability.

Did Language Evolve by Means Other Than Natural Selection?

A different alternative to the hypothesis that language is an adaptation is the

possibility that it evolved by mechanisms other than natural selection, a hypothesis associated with Stephen Jay Gould and Noam Chomsky (Chomsky 1988; Gould 1997; see Piatelli-Palmarini 1989 and Pinker and Bloom 1990 for discussion). On this view, language may have evolved all at once as the product of a macromutation. Or the genes promoting language may have become fixed by random genetic drift or by genetic hitchhiking (i.e. genes that were near other genes that were the real target of selection). Or it may have arisen as a by-product of some other evolutionary development such as a large brain, perhaps because of physical constraints on how neurons can be packed into the skull.

It is hard to evaluate this theory (though, as we shall see, not impossible), because there have been no specific proposals fleshing out the theory (e.g. specifying the physical constraint that makes language a neurobiological necessity). So what is the appeal of the non-selectionist theories?

One is a general misconception, spread by Gould, that natural selection has become an obsolete or minor concept in evolutionary biology, and that explanations in terms of by-products (what he called 'spandrels') or physical constraints are to be preferred in principle (e.g. Piatelli-Palmarini 1989). This is a misconception because natural selection remains the only evolutionary force capable of generating complex adaptive design, in which a feature of an organism (such as the eye or heart) has a non-random organization that enables it to attain an improbable goal that fosters survival and reproduction (Dawkins 1986; Williams 1966). Moreover, natural selection is a rigorous concept which can be modelled mathematically or in computer simulations, measured in natural environments, and detected by statistical analyses of organisms' genomes (Kreitman 2000; Maynard Smith 1988; Przeworski et al. 2000; Weiner 1994).

A second appeal of non-selectionist theories comes from a scepticism that language could have provided enough reproductive benefits to have been selected for. According to one objection, popular among linguists, language has arbitrary features that do not obviously contribute to communication. However, *all* communication systems have arbitrary features (such as the particular sequences of dots and dashes making up Morse code), because arbitrary ways of linking messages to signals are useful as long as they are shared by sender and recipient. Moreover, since a feature that eases the task of the speaker (by omitting information or reducing the complexity of the signal) will complicate the task of the listener (by making the message more ambiguous or vulnerable to noise), a shared code must legislate arbi-

trary conventions that do not consistently favour any single desideratum (Pinker and Bloom 1990).

Another argument for non-selectionist theories is that grammar is more complicated than it needs to be to fulfil the communicative needs of a hunter-gatherer lifestyle. As one sceptic put it, 'How does recursion help in the hunt for mastodons?' But as Bloom and I pointed out, complex grammar is anything but a useless luxury: 'It makes a big difference whether a far-off region is reached by taking the trail that is in front of the large tree or the trail that the large tree is in front of. It makes a difference whether that region has animals that you can eat or animals that can eat you.' Since selection can proceed even with small reproductive advantages (say, one per cent), the evolution of complex grammar presents no paradox.

A third misconception is that if language is absent from chimpanzees, it must have evolved by a single macromutation. This is seen as an argument for a macromutational theory by those who believe that human language is qualitatively distinct from the communicative abilities of chimpanzees, and as an argument that human language cannot be qualitatively distinct from the communicative abilities of chimpanzees by those who believe that macromutations are improbable. But both arguments are based on a misunderstanding of how evolution works. Chimpanzees and bonobos are our closest living relatives, but that does not mean that we evolved from them. Rather, humans evolved from an extinct common ancestor that lived six to eight million years ago. There were many other (now-extinct) species in the lineage from the common ancestor to modern humans (*australopithecines*, *habilis*, *ergaster*, *archaic sapiens*, etc.) and, more important, many individuals making up the lineages that we group into species for convenience. Language could well have evolved gradually *after* the chimp/human split, in the 200,000–300,000 generations that make up the lineage leading to modern humans. Language, that is, could be an autapomorphy: a trait that evolved in one lineage but not its sister lineages.

The final appeal of the non-selectionist hypothesis is that language could only have been useful once it was completely in place: a language is useless if you are the only one to have evolved the ability to speak it. But this objection could be raised about the evolution of any communicative system, and we know that communication has evolved many times in the animal kingdom. The solution is that comprehension does not have to be in perfect synchrony with production. In the case of language, it is often possible to decode parts of an utterance in a language one has not completely mastered.

When some individuals are making important distinctions that can be decoded by listeners only with cognitive effort, a pressure could thereby develop for the evolution of neural mechanisms that would make this decoding process become increasingly automatic and effortlessly learned (Pinker and Bloom 1990). The process whereby environmentally induced responses set up selection pressures for such responses to become innate, triggering conventional Darwinian evolution that superficially mimics a Lamarckian sequence, is known as the Baldwin Effect (Hinton and Nowlan 1987).

Opposing these spurious arguments for the non-selectionist hypothesis is a strong *prima facie* reason to favour the selectionist one: the standard argument in evolutionary biology that only natural selection can explain the evolution of complex adaptive design (Dawkins 1986; Williams 1966). The information-processing circuitry necessary to produce, comprehend, and learn language requires considerable organization. Randomly organized neural networks, or randomly selected subroutines from an artificial intelligence library, do not give rise to a system that can learn and use a human language. As we saw, language is not just a set of symbolic labels for concepts, not just the use of linear order, not just the use of hierarchical structure, and not just a blurting out of a sequence of sounds. It is an integrated system containing a lexicon, several components of grammar, and interfaces to input-output systems, possibly with language-specific modifications of their own. And this complexity is not just there for show, but makes possible a remarkable ability: language's vast expressive power, rapid acquisition by children, and efficient use by adults.

As with other complex organs that accomplish improbable feats, the necessary circuitry for language is unlikely to have evolved by a process that is insensitive to the functionality of the end product, such as a single mutation, genetic drift, or arbitrary physical constraints. Natural selection is the most plausible explanation of the evolution of language, because it is the only physical process in which how well something works can explain how it came into existence.

What Did Language Evolve For?

If language is an adaptation, what is it an adaptation for? Note that this is different from the question of what language is typically *used* for, especially

what it is used for at present. It is a question about the 'engineering design' of language and the extent to which it informs us about the selective pressures that shaped it.

What is the machinery of language trying to accomplish? The system appears to have been put together to encode propositional information—who did what to whom, what is true of what, when, where and why—into a signal that can be conveyed from one person to another. It is not hard to see why it might have been adaptive for a species with the rest of our characteristics to evolve such an ability. The structures of grammar are well suited to conveying information about technology, such as which two things can be put together to produce a third thing; about the local environment, such as where things are; about the social environment, such as who did what to whom, when where and why; and about one's own intentions, such as *If you do this, I will do that*, allowing people to convey the promises and threats that undergird relations of exchange and dominance.

The Cognitive Niche

Gathering and exchanging information is, in turn, integral to the larger niche that modern *Homo sapiens* has filled, which John Tooby and Irven DeVore (1987) have called 'the cognitive niche' (it may also be called the 'informavore' niche, following a coinage by George Miller). Tooby and DeVore developed a unified explanation of the many human traits that are unusual in the rest of the living world. They include our extensive manufacture of and dependence on complex tools, our wide range of habitats and diets, our extended childhoods and long lives, our hypersociality, our complex patterns of mating and sexuality, and our division into groups or cultures with distinctive patterns of behaviour. Tooby and DeVore proposed that the human lifestyle is a consequence of a specialization for overcoming the evolutionary fixed defences of plants and animals (poisons, coverings, stealth, speed, and so on) by cause-and-effect reasoning. Such reasoning enables humans to invent and use new technologies (such as weapons, traps, coordinated driving of game, and ways of detoxifying plants) that exploit other living things before they can develop defensive countermeasures in evolutionary time. This cause-and-effect reasoning depends on intuitive theories about various domains of the world, such as objects, forces, paths, places, manners, states, substances, hidden biochemical essences, and other people's beliefs and desires.

The information captured in these intuitive theories is reminiscent of the information that the machinery of grammar is designed to convert into strings of sounds. It cannot be a coincidence that humans are special in their ability to outsmart other animals and plants by cause-and-effect reasoning, and that language is a way of converting information about cause-and-effect and action into perceptible signals.

A distinctive and important feature of information is that it can be duplicated without loss. If I give you a fish, I do not have the fish, as we know from sayings like *You can't have your cake and eat it*. But if I tell you how to fish, it is not the case that I now lack the knowledge how to fish. Information is what economists call a non-rival good, a concept recently made famous by debates about intellectual property (such as musical recordings that can be shared without cost on the internet).

Tooby and DeVore have pointed out that a species that has evolved to rely on information should thus also evolve a means to *exchange* that information. Language multiplies the benefit of knowledge, because a bit of know-how is useful not only for its practical benefits to oneself but as a trade good with others. Using language, I can exchange knowledge with somebody else at a low cost to myself and hope to get something in return. It can also lower the original acquisition cost—I can learn about how to catch a rabbit from someone else's trial and error, without having to go through it myself.

A possible objection to this theory is that organisms are competitors, so that sharing information is costly because of the advantages it gives to one's competitors. If I teach someone to fish, I may still know how to fish, but they may now overfish the local lake, leaving no fish for me. But this is just the standard problem of the evolution of any form of cooperation or altruism, and the solution in the case of language is the same. By sharing information with our kin, we help copies of our genes inside those kin, including genes that make language come naturally. As for non-kin, if we inform only those people who are likely to return the favour, both of us can gain the benefits of trade. It seems clear that we do use our faculties of social cognition to ration our conversation to those with whom we have established a non-exploitative relationship; hence the expression 'to be on speaking terms'.

Language, therefore, meshes neatly with the other features of the cognitive niche. The zoologically unusual features of *Homo sapiens* can be explained parsimoniously by the idea that humans have evolved an ability to encode information about the causal structure of the world and to share it among themselves. Our hypersociality comes about because information

is a particularly good commodity of exchange that makes it worth people's while to hang out together. Our long childhood and extensive biparental investment are the ingredients of an apprenticeship: before we go out in the world, we spend a lot of time learning what the people around us have figured out. And because of the greater pay-off for investment in children, fathers, and not just mothers, have an incentive to invest in their children. This leads to changes in sexuality and to social arrangements (such as marriage and families) that connect men to their children and to the mothers of those children.

Humans depend on culture, and culture can be seen in part as a pool of local expertise. Many traditions are endemic to a people in an area because know-how and social conventions have spread via a local network of information sharing. Humans have evolved to have a long lifespan (one end of the evolutionarily ubiquitous trade-off between longevity and fecundity) because once you have had an expensive education you might as well make the most out of it by having a long period in which the expertise can be put to use. Finally, the reason that humans can inhabit such a wide range of habitats is that our minds are not adapted to a narrow, specialized domain of knowledge, such as how to catch a rabbit. Our knowledge is more abstract, such as how living things work and how objects collide with and stick to each other. That mindset for construing the world can be applied to many kinds of environment rather than confining us to a single ecosystem.

On this view, then, three key features of the distinctively human lifestyle—know-how, sociality, and language—co-evolved, each constituting a selection pressure for the others.

Alternatives to the Cognitive Niche Theory

Several alternative hypotheses acknowledge that language is an adaptation but disagree on what it is an adaptation for. One possibility, inspired by an influential theory of the evolution of communication by Dawkins and Krebs (Dawkins 1982), is that language evolved not to inform others but to manipulate and deceive them. The problem with this theory is that, unlike signals with the physiological power to manipulate another organism directly, such as loud noises or chemicals, the signals of language are impotent unless the recipient actively applies complicated computations to decode them. It is impossible to use language to manipulate someone who does not understand the language, so hominids in the presence of the first linguistic

manipulators would have done best by refusing to allow their nascent language systems to evolve further, and language evolution would have been over before it began.

Another possibility is that language evolved to allow us to think rather than to communicate. According to one argument, it is impossible to think at human levels of complexity without a representational medium for propositions, and language is that medium (Bickerton 1990). According to another argument, we spend more time talking to ourselves than talking to other people, so if language has any function at all, it must be thought rather than communication (Chomsky 2002). These theories have two problems. One is that they assume the strongest possible form of the Whorfian hypothesis—that thought depends entirely on language—which is unlikely for a number of reasons (see Pinker 1994; 2002; Siegal et al. 2001; Weiskrantz 1988). The other is that if language evolved to represent information internally, much of the apparatus of grammar, which converts logical relationships into perceptible signals, would be superfluous. Language would not need rules for defining word orders, case markers, phonological strings, adjustment rules, and so on, because the brain could more efficiently code the information to itself silently, using networks of variables and pointers.

Considerations of language design rule out other putative selectional pressures. Language is unlikely to have evolved as a direct substitute for grooming (Dunbar 1998), or as a courtship device to advertise the fitness of our brains (Miller 2000), because such pressures would not have led to an ability to code complex abstract propositions into signals. A fixed set of greetings would suffice for the former; meaningless displays of virtuosity, as in scat singing, would suffice for the latter.

New Tests of the Theory That Language is an Adaptation

Contrary to the common accusation that evolutionary hypotheses, especially ones about language, are post hoc ‘just so’ stories, the hypothesis that language is an evolutionary adaptation can be made rigorous and put to empirical test. I will conclude by reviewing two new areas of research on the evolution of language that have blossomed since my 1990 paper with Bloom and which are beginning to support its major predictions.

Language and Evolutionary Game Theory

Good theories of adaptation can be distinguished from bad ones (Williams 1966). The bad ones try to explain one bit of our psychology (say, humor or music) by appealing to some other, equally mysterious bit (laughing makes you feel better; people like to make music with other people). The good ones use some *independently established* finding of engineering or mathematics to show that some mechanism can efficiently attain some goal in some environment. These engineering benchmarks can serve as predictions for how Darwinian organisms ought to work: the more uncannily the engineering specifications match the facts of the organism, the more confidently one infers that the organism was selected to carry out that function.

Evolutionary game theory has allowed biologists to predict how organisms ought to interact with other organisms co-evolving their own strategies (Maynard Smith 1982). Language, like sex, aggression, and cooperation, is a game it takes two to play, and game theory can provide the external criteria for utility enjoyed by the rest of evolutionary biology. Modellers assume only that the transmission of information between partners provides them with an advantage (say, by exchanging information or coordinating their behaviour), and that the advantage translates into more offspring, with similar communicative skills. The question then is how a stable communication system might evolve from repeated pairwise interactions and, crucially, whether such systems have the major design features of human language.

The first such attempt was a set of simulations by Hurford (1989) showing that one of the defining properties of human language, the arbitrary, bi-directional sign, will drive out other schemes over evolutionary time (Hurford 1989). More recently, Nowak and his collaborators have now done the same for two of the other central design features of language (Nowak and Krakauer 1999; Nowak et al. 1999a; Nowak 2000).

Nowak and his colleagues pointed out that in all communication systems, errors in signalling or perception are inevitable, especially when signals are physically similar. Imagine organisms that use a different sound (say, a vowel) for every concept they wish to communicate. As they communicate more concepts, they will need additional sounds, which will be physically closer and hence harder to discriminate. At some point adding new signals just makes the whole repertoire more confusable and fails to increase its

net communicative power. Nowak and colleagues showed that this limitation can be overcome by capping the number of signals and stringing them together into sequences, one sequence per concept. The sequences are what we call words, and as I mentioned earlier, the combination of meaningless vowels and consonants into meaningful words by rules of phonology is a universal property of language, half of the trait called 'duality of patterning'. Nowak and his colleagues have shown how its evolution is likely among communicators with a large number of messages to convey, a precondition that plausibly characterizes occupants of the cognitive niche.

Nowak and his colleagues have recently motivated another hallmark of language. Imagine a language in which each message was conveyed by a single word. For any word to survive in a community, it must be used frequently enough to be heard and remembered by all the learners. As new words are added to the vocabularies of speakers, old words must be used less often, and they are liable to fade, leaving the language no more expressive than before. Nowak et al. point out that this limitation can be overcome by communicators who use compositional syntax: rather than pairing each word with an entire event, they pair each word with a *component* of an event (a participant, an action, a relationship), and string the words together in an order that reflects their roles (e.g. *Dog bites man*). Such communicators need not memorize a word for every event, reducing the word-learning burden and allowing them to talk about events that lack words. Syntax and semantics, the other half of the duality of patterning, will evolve.

Nowak et al. note that syntax has a cost: the requirement to attend to the order of words. Its benefits exceed the costs only when the number of events worth communicating exceeds a threshold. This 'syntax threshold' is most likely to be crossed when the environment, as conceptualized by the communicators, has a combinatorial structure: for example, when any of a number of actors (dogs, cats, men, women, children) can engage in any of a number of actions (walking, running, sleeping, biting). In such a world, the number of words that have to be learned by a syntactic communicator equals the sum of the number of actors, actions, places, and so on, whereas the number that must be learned by a nonsyntactic communicator equals their *product*, a potentially unlearnable number. Nowak et al. thus proved the theoretical soundness of the conjecture of Pinker and Bloom (1990) that syntax is invaluable to an analytical mind in a combinatorial world.

Language and Molecular Evolution

Mathematical models and computer simulations can show that the advantages claimed for some features of language really can evolve by known mechanisms of natural selection. These models cannot, of course, show that language *in fact* evolved according to the proposed scenario. But recent advances in molecular and population genetics may provide ways of testing whether selection in fact occurred.

Evolution is a change in gene frequencies, and the first prediction of the theory that language is an evolutionary adaptation is that there should be genes that have as one of their distinctive effects the development of normal human language abilities. Such a gene would be identifiable as an allelic alternative to a gene that leads to an impairment in language. Since pleiotropy is ubiquitous, one need not expect that such a gene would affect *only* language; but its effects on language should not be consequences of some more general deficit such as a hearing disorder, dysarthria, or retardation.

Clinical psycholinguists have long known of the collection of syndromes called Specific Language Impairment (SLI), in which a child fails to develop language on schedule and struggles with it throughout life (Bishop et al. 1995; Leonard 1998; van der Lely et al. 1998). By definition SLI is not a consequence of autism, deafness, retardation, or other non-linguistic problems, though it may co-occur with them. In one form of the syndrome, sometimes called 'Grammatical SLI', the children are normal in intelligence, auditory perception, and the use of language in a social context, but their speech is filled with grammatical errors and they are selectively deficient in detecting ungrammaticality and in discriminating meaning based on a sentence's grammar (van der Lely et al. 1998; van der Lely and Stollwerck 1996). Though it was once thought that SLI comes from a deficit in processing rapidly changing sounds, that theory has been disproven (Bishop et al. 1999; Bishop et al. 2001; van der Lely et al. 1998).

SLI runs in families and is more concordant in monozygotic than in dizygotic twins, suggesting it has a heritable component (Bishop et al. 1995; Stromswold 2001; van der Lely and Stollwerck 1996). But the inheritance patterns are usually complex, and until recently little could be said about its genetic basis. In 1990 investigators described a large multi-generational family, the KEs, in which half the members suffered from a disorder of speech and language, distributed within the family in the manner of an

autosomal dominant gene (Hurst et al. 1990). Extensive testing by psycholinguists showed a complex phenotype (Bishop 2002). The affected family members on average have lower intelligence test scores (perhaps because verbal coding helps performance in a variety of tasks), but their language impairment cannot be a simple consequence of low intelligence, because some of the affected members score in the normal range, and some score higher than their unaffected relatives (Bishop 2002; Lai et al. 2001). And though the affected members have problems in speech articulation (especially as children) and in fine movements of the mouth and tongue (such as sticking out their tongue or blowing on command), their language disorder cannot be reduced to a motor problem, because they also have trouble with identifying phonemes, understanding sentences, judging grammaticality, and other language skills (Bishop 2002).

In 2001, geneticists identified a gene on Chromosome 7, FOXP2, that is perfectly associated with the syndrome within the KE family and in an unrelated individual (Lai et al. 2001). They also argued on a number of grounds that the normal allele plays a causal role in the development of the brain circuitry underlying language and speech, rather than merely disrupting that circuitry when mutated.

A second crucial prediction of the language-as-adaptation theory is that there should be *many* genes for language. If human language can be installed by a single gene, there would be no need to invoke natural selection, because it is not staggeringly improbable that a single gene could have reached fixation by genetic drift or hitchhiking. But if a large set of co-evolved genes is necessary, probability considerations would militate against such explanations. The more genes are required for normal language, the lower the odds that our species could have accumulated them all by chance.

It seems increasingly likely that in fact many genes are required. In no known case of SLI is language wiped out completely, as would happen if language was controlled by a single gene which occasionally is found in mutated form. Moreover, SLI is an umbrella term for many distinct syndromes (Leonard 1998; Stromswold 2001; SLI Consortium 2002). Grammatical SLI, for example, is distinct from the syndrome affecting the KE family, which in turn is distinct from other cases of SLI known to clinicians (van der Lely and Christian 1998). In yet another syndrome, language delay, children are late in developing language but soon catch up, and can grow up without problems (Sowell 1997). Language delay is highly heritable (Stromswold 2001), and its statistical distribution in the population suggests that it is a

distinct genetic syndrome rather than one end of a continuum of developmental timetables (Dale et al. 1998). There are yet other heritable disorders involving language (Stromswold 2001), such as stuttering and dyslexia (a problem in learning to read which may often be a consequence of more general problems with language). Both have been associated with specific sets of chromosomal regions (Stromswold 2001).

With recent advances in genomics, the polygenic nature of language is likely to become more firmly established. In 2002, an 'SLI Consortium' discovered two novel loci (distinct from FOXP2) that are highly associated with SLI but not associated with low non-linguistic intelligence (SLI Consortium 2002). Moreover, the two loci were associated with different aspects of language impairment, one with the ability to repeat non-words, the other with expressive language, further underscoring the genetic complexity of language.

The most important prediction of the adaptation theory is that language should show evidence of a history of selection. The general complaint that evolutionary hypotheses are untestable has been decisively refuted by the recent explosion of quantitative techniques that can detect a history of selection in patterns of statistical variation among genes (Kreitman 2000; Przeworski et al. 2000). The tests depend on the existence of neutral evolution: random substitutions of nucleotides in non-coding regions of the genome, or substitutions in coding regions that lead to synonymous codons. These changes have no effect on the organism's phenotype, and hence are invisible to natural selection. The genetic noise caused by neutral evolution can thus serve as a baseline or null hypothesis against which the effects of selection (which by definition reduces variability in the phenotype) can be measured.

For example, if a gene has undergone more nucleotide replacements that alter its protein product than replacements that do not, the gene must have been subject to selection based on the function of the protein, rather than having accumulated mutations at random, which should have left equal numbers of synonymous and amino-acid-replacing changes. Alternatively, one can compare the variability of a gene among the members of a given species with the variability of that gene across species; a gene that has been subjected to selection should vary more between species than within species. Still other techniques compare the variability of a given gene to estimates of the variability expected by chance, or check whether a marker for an allele is found in a region of the chromosome that shows reduced vari-

ation in the population because of a selective sweep. About a dozen such techniques have been devised so far. The calculations are complicated by the fact that recombination rate differences, migrations, population expansions, and population subdivisions can also cause deviations from the expectations of neutral evolution, and therefore can be confused with signs of selection. But techniques to deal with these problems have been developed as well.

It is now obvious how one can test the language-as-adaptation hypothesis (or indeed, any hypothesis about a psychological adaptation). If a gene associated with a trait has been identified, one can measure its variation in the population and apply the tests for selection. The day that I wrote this paragraph, the first of such tests has been reported in *Nature* (Enard et al. 2002). A team of geneticists examined the FOXP2 protein (the cause of the KE family's speech and language disorder) in the mouse, several primate species, and several human populations. They found that the protein is highly conserved among mammals: the chimpanzee, gorilla, and monkey versions of the protein are identical to each other and differ in only one amino acid from the mouse version and two from the human version. But two of the three differences between humans and mice occurred in the human lineage after its separation from the common ancestor with the chimpanzee. And though the variations in the gene sequence among all the non-human animals produce few if any functional differences, at least one of the changes in the human lineage significantly altered the function of the protein. Moreover, the changes that occurred in the human lineage have become fixed in the species: the team found essentially no variation among forty-four chromosomes originating in all the major continents, or in an additional 182 chromosomes of European descent. The statistical tests showed that these distributions are extremely unlikely to have occurred under a scenario of neutral evolution, and therefore that the FOXP2 gene has been a target of selection in human evolution. The authors further showed that the selection probably occurred during the last 200,000 years, the period in which anatomically modern humans evolved, and that the gene was selected for directly, rather than hitchhiking on an adjacent selected gene. Alternative explanations that rely on demographic factors were tested and at least tentatively rejected.

This stunning discovery does not *prove* that language is an adaptation, because it is possible that FOXP2 was selected only for its effects on orofacial movements, and that its effects on speech and language came along for the

ride. But this is implausible given the obvious social and communicative advantages that language brings, and the fact that the deficient language in SLI is known to saddle the sufferers with educational and social problems (Beitchman et al. 1994; Snowling et al. 2001).

The studies I reviewed in this section are, I believe, just a beginning. I predict that evolutionary game theory will assess the selective rationale for an increasing number of universal properties of human language, and that new genes for language disorders and individual variation in language will be discovered and submitted to tests for a history of selection in the human lineage. In this way, the theory that language is an adaptation, motivated originally by the design features and natural history of language, will become increasingly rigorous and testable.

FURTHER READING

For general introductions to the structure and function of language, see Baker (2001); Bickerton (1990); Jackendoff (1994; 2002); Lightfoot and Anderson (2002); Miller (1991); Pinker (1994).

Good overviews of natural selection and adaptation include Dawkins (1986); Dawkins (1996); Maynard Smith (1986; 1989); Ridley (1986); Weiner (1994); Williams (1966). The debate over whether language is a product of natural selection may be found in the target article, commentaries, and reply in Pinker and Bloom (1990).

Specific Language Impairment is explained in Leonard (1998); van der Lely et al. (1998). An overview of the genetics of language can be found in Stromswold (2001). Evolutionary game theory is explained by its founder in Maynard Smith (1982). Methods for detecting natural selection in molecular genetic data are reviewed in Aquadro (1999); Kreitman (2000); Przeworski et al. (2000).