As far as historical linguistics is concerned, perhaps the best current overview is Campbell (1998). The articles in Durie and Ross (1996) debate the success of the comparative method. Two important works explore the possibility of establishing links of language relatedness going back tens of thousands of years: Greenberg (1987) and Nichols (1992). However, it needs to be stressed that the former has received a far more favourable reading outside the field of linguistics than within. In addition to the references cited in the chapter itself, the following works provide a look at current approaches to grammaticalization: Traugott and Heine (1991*a*; 1991*b*), Bybee et al. (1994), and Heine and Kuteva (2002*b*).

Finally, there is a rich literature developing diverse functionalist approaches to grammar. Haiman (1985), Givón (1995), Van Valin and Lapolla (1997), and Dik (1997) are a small sample. Newmeyer (1998) is devoted to the reconciliation of formal and functional approaches. Functional explanations for the typological distribution of languages are presented in Comrie (1989) and Croft (1990). These works attempt to link typological generalizations to principles governing language change, and are therefore relevant to an understanding of language evolution.

5

Symbol and Structure: A Comprehensive Framework for Language Evolution

Derek Bickerton

Speaking as a Linguist[1]

I approach the evolution of language as a linguist. This immediately puts me in a minority, and before proceeding further I think it's worth pausing a moment to consider the sheer oddity of that fact. If a physicist found himself in a minority among those studying the evolution of matter, if a biologist found himself in a minority among those studying the evolution of sex, the world would be amazed, if not shocked and stunned. But a parallel situation in the evolution of language causes not a hair to turn.

Why is this? Several causes have contributed. Linguists long ago passed a self-denying ordinance that kept almost all of them out of the field, until quite recently. Since nature abhors a vacuum, and since the coming into existence of our most salient talent is a scientific question that should concern anyone seriously interested in why humans are as they are, other disciplines rushed to fill that vacuum. Then again, language doesn't look as if it should be all that complex, not like genetics or quantum mechanics. We all speak at least one, that one we acquired without a lick of conscious effort, and most non-linguists, in the unlikely event that they opened a copy of *Linguistic Inquiry* or *Natural Language and Linguistic Theory* only to find stuff every bit as hard going as genetics or quantum mechanics, would in many cases react by saying 'What's all this nonsense about? Why are they making such a fuss about something that's perfectly simple and straightforward?' And they

[1] Scout's honour, I hadn't read Fritz Newmeyer's contribution to the present volume when I wrote this chapter, nor had he read mine when he wrote his. The impressive similarity of our introductions was quite spontaneous, a highly natural reaction to the circumstances, and should serve as a wake-up call to linguists and non-linguists alike.

would probably go on to say, 'What do *I* need this stuff for? I'm a systematic biologist/palaeo-anthropologist/ evolutionary psychologist/computational mathematician [strike out whichever do not apply]—I do not need this.'

Well, the reason they do need this is simple and straightforward. Language is a means of communication (among many other functions) that differs radically from the means of communication of any other species. At the same time, we are a species that differs radically from other species in our creativity and variability of behaviour (anyone who confuses this statement of plain fact with the claim that we are the pinnacle of evolution or divinely created is herself seriously confused). There is a good chance that these two uniquenesses are not coincidental; in other words, we most likely are the way we are because we have language and no one else does. If that is so, then it must be because there are specific properties of language which, if other species had had them first, would have produced similar results. What we have to do is determine what these properties are, which of them are essential and which accidental. We have to determine how we came to have, not just 'language' in the abstract—whatever that might mean—but the precise set of linguistic properties that happens to correlate with, and most probably causes, our unique nature. We have to do this if we are ever going to explain how humans evolved. But we can't explain why language has the set of properties that it does have, and no other set, if we do not know what those properties are.

The Interdisciplinary Stance

This is not, yet, a particularly widespread view. A lot of writers believe one can treat language as a given, a black box, in effect, and account for its evolution simply by selecting the selective pressure that gave rise to it. Was it a grooming substitute (Dunbar 1993; 1996)? Or maybe setting up a menstruation ritual for female bonding (Power 1998; 2000)? Or letting men know if their women had cheated on them (Ridley 1993: 229)? Or initiating marriage, so men would know who they weren't supposed to cheat with (Deacon 1997)? The fact that these and similar explanations flourish side by side tells one immediately that not enough constraints are being used to limit possible explanations.

Simply taking into account what we know about language should form an adequate constraint, since all these proposals run up against some language feature quite incompatible with them. Take the grooming proposal. It is far from clear why, if language simply substituted for grooming when group size became too large, language should invariably convey factual information, indeed be incapable of *not* giving factual information, even in flattering someone (*That outfit really suits you—matches your eyes*). Surely a similar result could have been achieved simply by using pleasant but meaningless noises. Lovers often do just that, even now, with all of language at their disposal. Or take the proposal that the driving force behind the emergence of language was gossip and/or some sort of Machiavellian manipulation. Since there undoubtedly was a time when the vocabulary was zero, there must also have been a time when the vocabulary was vanishingly small, no more than three or four units/signs/words, whatever those may have been. The question is simply whether the gossip one could transmit with such a vocabulary would be of any interest whatsoever. It seems unlikely. Similarly, it is highly implausible that, with a small initial stock of symbols, one could do much in the way of social manipulation.

Both gossip and manipulation require a vocabulary of some size, but such a size could hardly have been achieved unless earlier and smaller vocabularies had already served some useful purpose. However, the issue of the minimal vocabulary size required for implementing functions such as these is simply not addressed by those who claim a social-intelligence source for language (see Bickerton 2002*a* for a fuller treatment of these and related questions).

Ignorance of both language and linguistic theory seemed to me for a long time to be the most serious deficiency among writers on the evolution of language. Then I reviewed two books by linguists (Loritz 1999; Jenkins 2000; see Bickerton 2001) and I'm no longer so sure. Such ignorance now appears as simply a special case (perhaps the most serious, though by no means the only serious one) of a much more widespread tendency. There are at least six fields–linguistics, palaeo-anthropology, evolutionary biology, neurology, psychology and primatology–that cannot possibly be ignored in any study of language evolution, and a number of others, such as genetics or palaeo-climatology, that bear on it perhaps somewhat less directly. All too often, a writer whose home is in one or other of these disciplines will make a proposal that is unacceptable in terms of one or more of the other relevant disciplines. This is not inevitable. It certainly does not result from the impossibility of acquiring the necessary knowledge, since anyone of average intelligence should, given goodwill and a little effort, be able to mas-

ter enough of the literature in all relevant disciplines to avoid making gross errors.[2]

Having mastered linguistics and all the other relevant disciplines, are we now ready to make sense of language evolution? I don't think so. There's something that is implicit in much of my own writing which I fully realized only quite recently. It is that the biggest obstacle to understanding the evolution of language is thinking about it as 'the evolution of language'.

Divide and Rule

Language as we know it today involves the coming together of three things: modality, symbolism, and structure. I can see no reason for supposing that all three evolved as a package deal, and good reasons for supposing that they evolved separately. Let's look at each of these three things in turn.

First modality: that includes speech and sign. For many, 'speech' and 'language' are interchangeable; how depressingly often one turns to the index of a book on human evolution to find the damning entry: 'Language: see Speech'. Most recently Mithen (2000: 216) has claimed that Neanderthals had probably acquired 'a degree of vocalization that is most appropriately described as language'. And even some linguists take a similar approach: Lieberman (1984; 1991; 1996), for example, assumes that once speech was there, the rest followed.

However, I take the arguments in Burling (2000) and Sperber and Origgi (2001) to be quite unanswerable. Before any of the three components of language could exist, let alone come together, there had to be comprehension of some kind, however primitive; pre-humans at some stage had to start trying to figure out one another's intentions. This largely solves the problem of what I once called the 'magic moment' (Bickerton 1990): how did the first hearer of a meaningful signal know that it was a meaningful signal (as opposed to a cough, a grunt of pain, or whatever)? Answer: if our ancestors were already trying to interpret the behaviour of their conspecifics, even

[2] I speak here with all the zeal of the converted, having myself violated biological probabilities with the 'macro-mutation' scenario of Bickerton (1990) and neurological probabilities with the 'different bits of the brain getting linked' scenario of Bickerton (1995; 1998). Although I know there are neurologists who do not buy them, I know of no evidence, neurological or other, that rules out the proposals of Calvin and Bickerton (2000), which have now replaced those earlier ones, and to which I return later.

perhaps to the extent of reading meanings where none was intended, they surely wouldn't take long to recognize intentional meanings. It also neatly solves the problem of whether language began as sign or speech. The answer is that it probably began as both—a mixture of anything that might serve to convey meaning. The original mixture of isolated grunts and gestures may have eventually settled on the vocal mode merely through the exigencies of communicating at night, over distance. or in dense vegetation. For a small initial vocabulary, no vocal improvement would have been needed.

Afterwards, as more (and more complex) information gradually came to be exchanged, attempts to convey it would have strongly selected for improved vocal capacities. I think there can be no doubt that the capacity to transmit information was what selected for improved speech capacity, rather than vice versa. Being able to speak more clearly does not, in and of itself, give you more to say. There is thus good reason to believe that the speech modality, far from being the driving force behind language, was entirely contingent on the two other components, the symbolic and the structural, and was developed and refined in response to their development.

Those components, the symbolic and the structural, are also distinct and can also be dissociated from one another. They are actually dissociated in several forms of development than can still be observed in the world around us: in early-stage pidgins, in early stage second-language learning, and in the productions of trained apes and other animals.[3] Of course you can't get a double dissociation: syntactic structure without symbolic content, the kind of thing you saw in old-fashioned phrase-structure rules, cannot be used by animate beings to communicate with one another. But symbolic representation does not require any kind of structure—telegrams and headlines are immediately comprehensible with relatively little grammatical structure, and we can perfectly well understand utterances that have no syntactic structure at all, like Pinker's *skid–crash–hospital* (Pinker 1994). Indeed a variety of factors I have discussed at length in earlier work suggest that, in the evolution of our species, symbolism may have preceded syntax by as much as two million years.

[3] I have previously claimed early-stage first-language learning as an example (Bickerton 1990). This needs caveatting, if you'll pardon the Haigism. Children learning inflected languages can and do acquire morphological affixes (they would have a hard time not doing this, in languages where bare stems are virtually or completely non-appearing) and at least sometimes use them correctly. They also acquire some basic facts about word order in the target language. But these are not syntax (see 'Fear of syntax' for what is).

Perhaps the clearest evidence for phylogenetic dissociation lies in the fact that while symbolic representation (at some level; and under instruction, at a near-human level, see Savage-Rumbaugh 1986) is within the reach of a number of non-human animals, syntax, regardless of the quantity or quality of instruction, remains beyond the reach of any other species ('putting symbols in a regular order' does not, of course, come anywhere close to being syntax). This inevitably suggests that the genetic and neural substrates for the two are quite distinct, and that they therefore must have distinct evolutionary histories. Consequently, explaining how language evolved requires us to answer two separate and quite distinct scientific questions. The first is: how and why one particular primate species, or one primate line of descent, developed a system of communication involving symbolic representation that allowed the transfer of (potentially) unlimited factual information, and the basic principles of which differed from those of all previous systems of communication. The second is how such a system acquired the very specific structural characteristics that the syntax of modern human languages exhibits. If one abbreviates these questions to 'How did meaningful units (words or signs) evolve?' and 'How did syntax evolve?', little is lost. But if only one question is answered, or if the two issues are mixed in together and confused with one another, we will continue to get the conflicting and unsatisfying accounts of language evolution that have predominated to this date. Let us therefore keep them clearly separate and deal with each in turn.

Symbols

The most crucial thing to grasp about the emergence of symbolic representation is that it must have been primarily a cultural rather than a biological event. This idea, again implicit in some of my earlier work, has not previously been stated in quite these terms. However, it follows inevitably from the fact that a neural substrate adequate for some level of symbolic representation exists not simply in other great apes but among creatures as phylogenetically distant from us as African Grey parrots (Pepperberg 2000). This widespread nature of potential for symbolic representation suggests an analogous rather than a homologous development, the kind of development that produced fins in sharks, ichthyosaurs, and cetaceans. Probably the potential for symbolism exists in any animal with a brain of sufficient complexity, and this would hardly be surprising, given the still wider spread

of iconic and indexical precursors of symbolism (for Pavlov's dogs, the ringing of a bell was an indexical representation of food, for example).

This view may be disturbing to some who, with Deacon, see the Rubicon between us and other animals as being symbolism rather than syntax. Part of the reason may be that when people think of symbolism, they think of the sophisticated version we enjoy today—a vast branching network of symbols each of which is interpretable in terms of other symbols—and not of the primitive version, compounded mainly of indexical and iconic associations, that may have come into existence two million years ago. Another part may be disbelief that, given the benefits of communication—so visible to us, with the twenty-twenty vision of hindsight—any animal capable of communication would fail to use it under natural conditions. Such a view ignores the essential unreliability of language. Words require little energy to produce; they are 'cheap tokens' and can be used with little or no risk or cost to deceive, just as easily as to inform. Body language is much more reliable for most animal purposes.[4]

I suspect that the only things missing for any relatively large-brained species were the two Ms–modality and motivation. Modality was perhaps the lesser problem, given the low requirement of an initial minimal vocabulary. Any modality capable of differentiating a half-dozen or so symbols would do for a start. Motivation was another matter. To us, able to appreciate the myriad potential uses of language, its possession seems too obvious a boon. But since no species has the gift of foresight, we should ask what benefits this minimal vocabulary would have bestowed on any other species. The answer is simple: none. Solitary species do not need to communicate. Other social species get along fine with non-linguistic methods. Since only one social species has even begun to develop the language mode, the logical place to look for motivation is in circumstances unique to that species. Elsewhere (Bickerton 2002*a*) I have argued in detail that the initial protolanguage arose through the exigencies of extractive foraging in mainly dry savannah-type environments. I will not repeat those arguments here; suffice it to say that hunger and a high risk from predation would have engendered social

[4] An actual illustration may be relevant here. Karl Muller, a part-Hawaiian who runs a shelter for the homeless in Honolulu, had a confrontation with the shelter's cook. Karl removed his false teeth. The cook ran away. He knew that the removal of the teeth (an expensive set) meant that Karl, an impressive street fighter, was ready to accept serious damage in going the limit with him. No mere verbal threat would have deterred the cook, a skilled amateur boxer ('He could have been a contender', according to Karl).

systems in which individuals were more interdependent than they are in most primate societies, and where, accordingly, a degree of trust sufficient to overcome the 'cheap tokens' problem would necessarily be engendered.

Symbolism arose culturally, then, because the minimal necessary biological equipment was already in place and the exploitation of symbolism directly benefited both individuals and groups (groups by optimizing foraging under the fission–fusion constraints that a wide day-range plus vulnerability to predation imposed, individuals by enhancing the status of those who located and led the group to the best food sources). The only question that remains is what the earliest symbols were like.

Holistic Versus Synthetic

Until quite recently, it was generally assumed that ontogeny and phylogeny, though far from indissolubly wedded, were at least alike to this extent: the earliest units of pre-human utterances were pretty similar to the earliest units of contemporary infants. That is to say, they were basically single units with ostensively definable referents, perhaps somewhat broader in meaning than the units of an adult vocabulary. Recently, however, this notion has been challenged from a variety of perspectives, all of which converge on the assumption that holistic utterances, semantically equivalent to one-clause sentences in modern languages, formed the earliest linguistic utterances.

Wray (2000) sees this proposal as solving the 'continuity paradox' (Bickerton 1990: 7). Calls can be interpreted holistically, so there could be a seamless transition from a non-linguistic communication system to some form of protolanguage. Carstairs-McCarthy (1999; 2000), assuming syntax to be modelled on the syllable (but see the next section), has to have a holistic protolanguage for the syntactic equivalents of onset, nucleus, and coda to be factored out of. Computational linguists (Batali 1998; Kirby 2000; Briscoe 2002*a*)[5] likewise assume initial holistic utterances, so that, by a process described as 'self-organization', vast quantities of variable and random

[5] Hurford (2000*b*: 225–6) claims that his approach to computational evolution is synthetic, rather than analytic. This is somewhat disingenuous, in light of his own statement that 'speakers were prompted to express atomic meanings (e.g. BERTIE, SAY, or GIVE) 50% of the time, and simple or complex whole propositions (e.g. HAPPY (CHESTER), LIKE (JO, PRUDENCE) or SAY (BERTIE (HAPPY (JO))) 50% of the time' (Hurford 2000: 334). The other three authors cited here appear to make their 'speakers' produce holistic propositions 100 per cent of the time.

utterances could gradually converge on fixed forms with fixed meanings. This holophrastic approach to initial symbols can even claim some kind of history, since holistic beginnings are implicit in the 'singing ape' conjectures of Darwin, Jespersen, and others.

The approach is, however, beset by a variety of problems. Initially, there is the problem of comprehension: while one can deduce many things about others' intentions from their behaviour, anyone who has visited a country with a strange language knows that such understanding does not extend to linguistic behaviour. Since Quine's discussion (1960; see also Premack 1986) the problem of the speaker who says *Gavagai!* to the researcher when a rabbit runs past has been well known. We may believe, unlike Quine, that the word is rather unlikely to mean 'undissociated rabbit parts'. But that may be largely because we benefit from a history, both ontogenetic and phylogenetic, of learning what words are likely to mean. It is a task of considerable difficulty, although well worth attempting, to try to imagine oneself as not knowing any words at all, not even knowing what words were or could do. Certainly the passage of a rabbit at the moment of utterance is no guarantee that *Gavagai!* has anything to do with the rabbit: events do not occur in a vacuum; something else that might be being referred to is always going on. Even if we could somehow know that *Gavagai!* and the rabbit were connnected, they could merely be connected in the way that *God bless you!* is connected with a sneeze.

Now these are problems that would affect understanding even if words or other symbols referred to isolated objects or events and were initially used in the presence of those objects and events (an unlikely proviso—there is little point in talking about what people can see for themselves, and indeed the whole point of language as opposed to animal communication systems in general is that the former, but not the latter, can be used to inform about things that are not physically present). Those problems of understanding are compounded infinitely if the initial utterances of a language do not correspond to anything tangible or easily identifiable, but refer to some set of circumstances that may or may not be apparent from the surrounding context. If the intended meaning is apparent from that context—if, say, initial utterances were things like *Give that to me!* or *Stay away from her!*—you wouldn't need language to express them. Such things are much more unambiguously expressed by behaviour already in an animal's communicative repertoire, such as begging gestures or threat gestures. If the intended meaning is not apparent from that context, the receiver would never be able

to select, from a potentially infinite range of possible meanings, the one that the sender meant to express.

It is no accident that in most, if not all, computer simulations of language evolution, the self-organizing 'agents' already *know what their interlocutor means to say*. If the problem space were not limited in this way, the simulations simply wouldn't work—the agents would never converge on a workable system. But such unrealistic initial conditions are unlikely to have applied to our remote ancestors.

Let us suppose, counter to probability, that our ancestors somehow developed a holistic language. They would then have confronted the problem of how to go from such a language to languages that are built up from discrete units with single meanings (as all languages are today). Even today, modern children, equipped with all the bells and whistles of full human language, have a hard time segmenting adult utterances which, though they may sound holistic to the child, consist already of ready-made word units. How much more difficult for creatures with no experience of language to segment strings that were genuinely holistic!

There are two logical possibilities, one of which must be fulfilled by any such holistic utterances. One is that the units that would eventually dissolve into discrete words already contained regularities within the holistic utterance—a phonetic sequence like *-meg-*, for instance, might occur in any holophrase that made reference to 'meat'. This would remove the problem at the same time as it removed any possible justification for supposing that language began in this way. For if such utterances could be straightforwardly decomposed into the equivalent of words, then words as we know them already existed and there would have been no point in starting out with holophrastic units.

The other possibility is that the sequences were truly holistic, in other words that their sound structures bore no relation to one another: 'Give the meat!' might then be *megalup* and 'Take the meat!' might be *kokubar*. From these, or any similar examples, it would simply be impossible to factor out a single symbol for 'meat'.

A holophrastic account has yet more difficulties to face. It may seem easy to translate a holophrase (given that we understand it correctly) into an equivalent sentence of discrete words. But suppose some large male hominid, in the presence of a female, aggressively utters **&*&x@**! We may reasonably take this to mean, 'Stay away from her!' But it could just as easily mean 'Do not go near her!' Or, 'Stay right where you are!' Or even 'You——,

you, get out of my sight!' How smaller units can be factored out from holophrases when even their global meanings are so potentially ambiguous remains unclear.

But perhaps the biggest problem with the holistic approach is that it doesn't explain anything worth explaining. All the substantive problems in language evolution—how symbolism got started and fixed, how, when, and why structure emerged, where and how and to what extent any of this got instantiated in neural tissue—remain to be solved, whether one accepts a holistic account or not.

Accordingly, it is more parsimonious to assume that language began as it was to go on—that discrete symbols, whether oral or manual, were there from the beginning. I do not know of a single coherent argument why they shouldn't have been. If they were, it is most likely that, once these symbols exceeded the merest handful, they began to be strung together in some ad hoc fashion. One hears frequently of 'proto-syntax', which seems to mean one-clause sentences with fixed word order, and there is a widespread but wholly erroneous belief that this does not merely constitute a step in the direction of real syntax, but that once one has achieved such a level of structure, real syntax follows automatically. In other words, they account for *The cat sat on the mat* and then cross their fingers,[6] confident that 'self-organization' will take care of the rest.

Fear of Syntax

Perhaps the most depressing aspect of language evolution studies is fear of syntax, which, the present collection suggests, is as widespread as ever. I know of no other field of study in which the work of a large body of highly intelligent specialists is so systematically misinterpreted, ignored, or even trashed. As a matter of plain fact, we have learned more about syntax in the last forty years than in the preceding 4,000, but you'd never guess that from reading most books on language evolution, including, alas, this one. Syntax forms a crucial part, arguably the most crucial part—since no other species is capable of it—of human language. If we are going to explain how language evolved, we have to explain how syntax evolved. If we are going

[6] I am not the originator of this sentence, but I have sought in vain for years to find its true begetter. I believed it to be Lila Gleitman, but she (p.c.) has denied authorship; if its real author contacts me, I will be happy to make full acknowledgment.

to explain how syntax evolved, we have to explain how it came to have the peculiar properties it has, and no others. It just will not do to dismiss it as due to self-organization, or the grammaticization of discourse, or analogies with motor systems or syllabic structure, or any of the other one-paragraph explanations that writers on language evolution typically hope to get away with.

The trouble is that most non-syntacticians think syntax is just a matter of regular word order (I wonder what they think syntacticians do all day!) plus perhaps a few prefixes and suffixes. As a corrective to this view, I offer the following brief test on Real Syntax (an asterisk before a word or sentence means that it is ungrammatical):

(1) (a) Bill wants someone to work for.
(b) Bill wants someone to work for him.

Why does a pronoun at the end of the sentence change the understood subject of *work*?

(2) (a) Who was it you said you didn't wanna/want to see?
(b) Who was it you said you didn't *wanna/ want to see you?

Why does a pronoun at the end of the sentence stop you from contracting *want to* to *wanna*?.

(3) (a) Which letter did you throw away without opening?
(b) That letter you threw away without opening contained anthrax.
(c) *You were wise to throw away that letter without opening.

Why is it okay to leave out an *it* after *opening* in the first and second sentence but not the third?

(4) (a) We wanted the chance to vote for each other.
(b) *We wanted the champ to vote for each other.

Why, given that the second sentence is perfectly logical and comprehensible—I wanted the champ to vote for you and you wanted him to vote for me—is it ungrammatical?

Two things need to be emphasized here. First, these sentences do not exhibit weird quirks peculiar to English or other western European languages (it is worth noting that the first full-length generative grammar of any language dealt with Hidatsa (a native-American language of North America)—Matthews (1961)—and that there is a vast generative literature on Australian, Austronesian, Native American and countless other non-western European languages that unfortunately doesn't seem to have had a wide readership). To the contrary, the phenomena these sentences illustrate arise from broad general principles familiar to anyone who is up to speed on generative syntax (for those who aren't, I can only refer them to said literature). Second, it should be apparent that phenomena of the type illustrated in these sentences are vanishingly unlikely to have come about through social factors, or self-organization, or the streamlining of discourse, or any of the many alternative explanations currently on offer.

This does not, however, mean that they must remain mysterious. It is a good bet that they are as they are because that is the way the brain works—that when syntax is finally and fully understood, it will become apparent that the algorithms the brain actually uses to produce sentences will necessarily produce (as epiphenomena, one may suppose) the features of (1)–(4). It should therefore be the task of anyone seriously interested in the evolution of language to work at either one end or both ends of the mystery: finding out the most parsimonious description of syntax that will satisfy the syntactic facts, or trying to determine (through neuro-imaging or any other available means) how the brain actually puts sentences together. Once we know exactly what evolved, we may begin to approach a final answer as to how it evolved.

As a linguist, I can only attempt the first course. One promising avenue of inquiry, briefly sketched in the appendix to Calvin and Bickerton (2000), further developed in yet to be published work (Bickerton 2003; in preparation) and now described as 'surface minimalism', would reduce syntax to only three components:

(5) (a) Conditions on the attachment of words to one another.
(b) Cycles of attachment yielding domains that consist of heads and their modifiers (phrases and clauses).
(c) Principles derived from the order in which constituents are attached to one another.

If language can run on these resources and these only, nothing like the massive amount of task-specific innate equipment many researchers have very reasonably feared would be required. (5a), or a great deal of it, can be derived directly from semantics. (5b) can be derived via a shift in function of a kind of social score-keeping device such as may have developed in several primate lines: the mapping of every event in episodic memory into a simple

schema incorporating who did what to whom (see further Bickerton 2000; Calvin and Bickerton 2000). (5c) can be derived from the way the brain processes any kind of material. The brain is adept at merging series of discrete inputs into coherent wholes (it does this every time you look at anything), and it can keep track of the sequence of its own operations through the gradualness with which neuronal activity decays (Pulvermuller 2002). All that is needed to run such a system is a far higher number of neurons and more of both cortico-cortical and cortico-cerebellar connections than we find in the brains of other primates.

Hurford (2000*a*: 223) has, very reasonably, expressed doubts as to whether such a stripped-down system could handle 'many of the examples given by Lightfoot [2000]'. Consider Lightfoot's *pièce de résistance*, to which he devotes almost half his paper: the asymmetry between subjects and objects that allows much freer extraction of the former than the latter (Lightfoot's (11)–(13)):

(6) (a) Who do you think that Ray saw?
(b) *Who do you think that saw Ray?

(7) (a) Which problem do you wonder how John solved?
(b) *Who do you wonder how solved which problem?

(8) (a) This is the sweater which I wonder who bought.
(b) *This is the student who I wonder what bought.

According to Lightfoot, this asymmetry represents a serious dysfunction in language (it 'conflicts with the desire/need to ask questions about the subjects of tensed clauses': Lightfoot 2000: 240), so cannot in itself be adaptive, but must result from some more general condition 'that presumably facilitates parsing' (p. 244). However, he has no explanation for such a condition beyond the suggestion that 'complex, dynamical systems can sometimes go spontaneously from randomness to order' (p. 245).

In surface minimalism, order of attachment (5c) yields two crucial principles, priority and finality (see Bickerton 2002*b*), which are involved in many different syntactic relations. Only the second of these need concern us here. A constituent X is final in a domain Y if there is no constituent Z such that X could be attached within Y before Z is attached. In other words, final attachments (in Lightfoot's examples, final referential attachments, or final arguments) mark the boundaries of domains (phrases or clauses). Non-final arguments (as in the (a) sentences) can be moved freely, since no

information about domain boundaries is lost if they are moved. But if final arguments are moved, as they are in the (b) sentences above, those boundary markers are removed and the sentences consequently become harder to parse, because it is less clear where one clause ends and another begins, and therefore more difficult to assign (unambiguously *and automatically*, as syntax must do) an argument to the domain to which it belongs. Accordingly, restrictions are placed on the mobility of final arguments: the *that* which introduces the Theme argument of a verb can't be attached unless a final argument is in place (6b), and question words (*who*, *which*, etc.) that are final arguments cannot be moved at all if any other question word has been moved (7b, 8b).

Timing Syntactic Emergence

One question that remains is when syntax emerged. If it emerged gradually, as many (see especially Pinker and Bloom 1990) think it did, there is no problem. A gradually enlarging brain, by providing, not greater intelligence *per se*, but more available neurons and more specialized connections between neurons, could have gradually provided more and more syntax—sentences just got longer and longer, as I once naively supposed (Bickerton 1981: ch. 5).

At least two things are seriously wrong with this. First, the principles involved are across-the-board principles: they apply everywhere, to all structures. At any given time, either they were in place or they weren't. Once they were in place, what was to stop syntax becoming immediately like it is today?

The other involves cognitive development. If we can measure cognitive development by the artefacts our ancestors produced (and what other way do we have?), there was something close to cognitive stagnation over the two million or so years that preceded the appearance of our species (if you doubt this, check out Iain Davidson's deconstruction of the palaeo-anthropological progress myth, chapter 8 below). Then, suddenly, creativity blossomed.

Somehow, there is a threshold effect. Somehow, it has to be explained. The advent of fully syntacticized language is the best candidate explanation so far. If anyone can think of a better alternative, or can explain (instead of merely explaining away) the suddenness of the transition, I'll be delighted

to hear it. Until then, with all its problems (e.g. why Neanderthals, with bigger brains than ours, didn't win out against us), the best explanation is still that syntax as we know it developed in our species but in no other.

A second timing problem associated with the origins of syntax involves the connection between fully syntacticized language and what has been called the 'Great Leap Forward' (the explosion of human culture that allegedly took place some 30,000–40,000 years ago). If syntax emerged with the human species (say,120,000 years or more ago), what accounts for the long delay before any tangible consequences appeared?

The answer is, of course, that syntacticized language enables but does not compel. Even today, in the Amazon and Congo jungles, there exist (barely—we are killing them off as quickly as we can) human societies whose toolkits, in an age of spacecraft and supercomputers, show relatively little advance over those of Cro-Magnons. What human language confers is not a technological straitjacket, but freedom—freedom to develop any way you can or think you want to (into societies where you work ever-lengthening hours at some servile and soul-destroyingly repetitive job so that you can afford to buy labour-saving devices, or into societies where you gather all you need for subsistence in fifteen to twenty hours a week and hang out the rest of the time drinking and gossiping). Instead of wondering why culture didn't explode the moment language emerged, maybe we should be wondering why, having acquired language, we chose the path that led to mass poverty, exploitation, perennial warfare, and perennial injustice.

Conclusion

I have tried in this chapter to present a framework adequate to include all the processes that uniquely produced language in the human species. This framework may be summarized as follows. Driven by climate changes into habitats where predators were fierce and common but food was scarce, at least one primate species began to exchange basic information about the environment in order to survive. But proto-language is not bee language. Once invented symbols begin to be used, they can be used to describe anything, consequently language can be adapted for social or any other purposes. So symbols multiplied but structure probably did not, prevented from developing by inadequate numbers of neurons and the right kind of connectivity. Once both of these had developed, things the brain could already do enabled protolanguage to develop quite rapidly into language as we know it today. The first group to cross some threshold that allowed unlimited combinations of words *and ideas* happened to be ours. And the rest, as they say, is history.

FURTHER READING

The attitude of many non-linguists to linguists who concern themselves with language evolution is expressed by Ingold (1993). The collection that contains his article (Gibson and Ingold 1993) is itself fairly representative of a variety of non-linguistic approaches to the subject. Nowak et al. (2002) illustrates the problems that may arise when a high level of non-linguistic sophistication (in this case, in computer science) mixes with a lower level of linguistic sophistication; linguists (and some biologists too) may boggle at the assertion that 'during primate evolution there was a succession of U(niversal) G(rammar)s that finally led to the UGs of human beings' (p. 615). However, some non-linguists do show a more sophisticated level of understanding, among them Maynard Smith and Sazthmary (1995) and Szathmary (2001).

The latter source develops the idea that although the language faculty must be biologically determined, it does not rely on hard-wired modules as sensory and motor faculties do. This position is related closely to the position on brain size held by Calvin (1996*a*; 1996*b*), who takes as critical the development of 'excess' neurons, which are not committed to any specific function but can be recruited for a number of tasks (including linguistic tasks), depending on what the brain is concerned with at any given moment. This view is consistent with results derived from both brain imaging and lesion studies (Damasio et al. 1996; Crosson 1993; Indefrey et al. 2001). It is also consistent with a view of brain activity held by Dennett (1991; 1997), in which there is no central homunculus or 'executive suite' in the brain. Rather than an individual thinking a thought and expressing it in words (the conventional view), sentences are constantly forming and reforming in the mind, but only the ones that can recruit enough neurons get to be consciously thought or spoken.

Certain developments within the Minimalist Program suggest possibilities of reconciling generative syntax with Darwinian evolution. Berwick (1998) is among the few writers who tackle this explicitly, but although they make no reference to evolution, protagonists of the derivational approach to minimalism (Epstein et al. 1998) are producing analyses that are easier to reconcile with biological and neurological constraints than alternative theories.