1 Constraining our theory of language

1.1 Language and evolution

Language is a defining characteristic of what it is to be human. One of the most fundamental questions that the field of linguistics must therefore try to answer is why language is an ability specific to our species. Attempts to provide such an answer have focused on investigating a number of peculiarities of language, best thought of as posing the following highly interconnected questions: (1) How do human children acquire language in a relatively rapid and effortless manner based on typically sparse input? (2) What are the idiosyncratic linguistic features that are found in all human languages? (3) What knowledge must the language user have? (4) How is language dealt with in the human brain?

Question (1) asks about the nature of coming to be a language user. What does the child bring to the task of acquiring his first language? Must he have prior knowledge that is specifically linguistic? Is such prior knowledge innately encoded? The second question deals with language universals. The many different languages spoken across the world all appear to conform to certain common rules or guidelines. What are these commonalities, and more importantly, why are these features, and not some others, universal? Question (3) asks about the internal state of a mature language user's mind. What does the language user have to know about both his specific language, and language generally, in order to have successful communicative interactions? The fourth question inquires into the nature of the human brain. Are there specific areas in the brain dedicated to linguistic tasks? How is the processing of language carried out by these areas? What is unique in the human mind, and are these differences directly responsible for language? Can language itself change the way we think?

None of these issues can stand alone; answers to one set of questions impact on, and are impacted by, the others. Universal characteristics of language may be so because they are aspects that make it more easily acquirable. The universals of language, in turn, will form the basis of the knowledge that the language user must have in his mind in order to successfully put language to use. The

manner in which the brain processes language may shed light on the prior knowledge brought to the acquisition task; if the brain deals with language in a very different way to other cognitive functions, this may suggest that domain-specific knowledge is required by the child. Knowing what must be innate will help to answer the question of how the human mind differs from that of other species. Characterising what the language user must know about his specific language, and what he must know about language in general, can aid answering the question of what the child's prior knowledge must consist of. Many other associations hold between these central questions.

A further question has recently received significant attention: (5) How has language arisen in the evolution of Homo sapiens? This question asks what happened in the evolution of humans, from the last common ancestor we share with our closest ape relatives (some five to ten million years ago) through the successive hominid species, that gave us language. In other words, what events in our evolutionary history resulted in our ability to communicate in a way very different to that of other species? Question (5) must thus, in its broadest sense, attend to issues in genetics, palaeontology, and archaeology. A problem, however, arises through the lack of evidence in these fields in answering this question. While the field of genetics has made huge leaps in understanding in the last decades, the empirical data in the case of language is still relatively sparse. In the case of the archaeological and palaeontological record, the story is even worse. Language leaves no fossils; even the anatomical indicators of the ability to vocalise do not all fossilise (and where they do, the human case is not as unique as one might imagine).[1] Consequently, the language evolutionist appears to face an uphill struggle.

One way to ease the language evolutionist's seemingly impossible task is to appeal to evidence that is available; that is, to call upon the knowledge gained in tackling the issues of (1) through (4) above. Other sub-disciplines within the field of linguistics offer to the evolutionary linguist a set of data which can direct his particular inquiry. For example, neurolinguistics offers evidence of the neural patterns relevant to language; comparative evidence can then be used to identify those patterns which are uniquely human. Psycholinguistics offers evidence of the behaviour of language users in both production and perception; such behaviours may signal underlying evolutionary pressures to which language responded. Studies in linguistic universals taken from the sub-disciplines of phonetics and phonology, morphology, syntax, and semantics offer evidence of the types of language that are both possible and actual; this distribution in the space of possibilities must have an evolutionary basis. Studies both in developmental linguistics, and on pidgins and creoles, offer

[1] And, of course, we must recognise that vocalisation does not necessarily entail human language abilities.

evidence of more underdeveloped forms of language; such data may inform our theories of what earlier, more primitive stages of language looked like. The field of language evolution thus clearly highlights the value of interdisciplinary research.

The issues to be addressed in this book fall squarely within question (5), yet the approach taken here will not attempt to answer question (5) by using evidence from other linguistic disciplines. Instead, it will use evidence from studies in evolution to cast a critical light on linguistic theory. As the preceding paragraph noted, cross-fertilisation of evidence and ideas is crucial. In linguistics, as in any scientific endeavour, it is important to consider the effect that findings from outside the discipline have on one's theory. While a grand unified theory of language is still a very long way off, by combining knowledge from different fields we can at least begin to narrow down the possibilities. Just as neurolinguistics, psycholinguistics, or semantic theory can inform our theories of how language evolved, so too can evolutionary biology inform our theories of how language is.

1.2 Studies in syntactic theory

Linguistics is the study of human language – how it is used, how it is structured, how it is acquired by children, how it is decoded by its users, how it interacts with other mental faculties. Linguists look for commonalities across languages, and uniquenesses between languages. Language can be studied from a functional perspective as a social, cultural, or psychological phenomenon. Syntactic theory, on the other hand, studies language from a structural perspective, asking about the architecture of the unconscious linguistic knowledge that language users possess.

Syntactic theory, as a study of the intrinsic structural aspects of language, is a sub-discipline of linguistics which is prone to isolation from other fields, and hence lacks the interdisciplinarity that is appropriate in the study of language. Typically, syntactic frameworks rely to a large extent on theory-internal argumentation, predicated on a legacy of assumptions accepted without verification against new evidence garnered in other domains. In order to narrow down the range of possible theories of syntax, we must look at them through a different lens. That is, we must scrutinise syntactic theory not simply as a theory which accounts for the syntactic data, but as a theory which is, at once, consistent with the syntactic data and with the evidence from other fields (whether that be psychology, biology, or sociology). In other words, does the particular theory of syntax allow for language to be acquired, allow for language to be processed, allow for language to have evolved, allow for language to be used?

This is a considerable enterprise, given (i) the number of differing syntactic frameworks that are currently available for analysis, and (ii) the number of

possible lenses that we might view the theory through. There are thus two decisions that will determine the direction the current investigation will take. The first of these has already been stated – evolutionary biology will provide the alternative viewpoint on syntactic theory, the evidence against which syntactic theory will be judged to be viable or not. The second decision concerns the particular syntactic theory to be judged. The choice taken is the Minimalist Program (henceforth MP). This choice can be justified both by the MP's standing as a dominant framework under the rubric of *modern syntactic theory*, and also by the fact that the MP makes a number of strong predictions which necessarily invite evolutionary assessment. Before we can even begin to address these predictions, a brief survey of the backdrop against which the MP stands is in order.

1.2.1 Nativism versus empiricism

The question of what the child brings to the task of acquiring his native language can be given a nativist or an empiricist answer. This question is essentially just one realisation of the nature versus nurture question that bears on very many aspects of human cognition. The nativist reply would assert that the language learner is born endowed with a genetically predetermined body of linguistic knowledge which assists him in the not straightforward task ahead of him. The empiricist reply, in contrast, would place the burden of explanation on experience; the language learner is exposed to an environment filled with stimuli which provide enough information to allow him to determine the structure of his native language.

Linguistic nativism, first clearly articulated by Chomsky in the 1960s (Chomsky 1966), owes its debt to the early philosophy of Plato's *Meno*, where it is suggested that the paradox of inquiry[2] can be solved if learning is reconstrued as simply recollecting that which we knew in a former existence.[3] What has come to be known as the logical problem of language acquisition – the question of how a complex language can be acquired based on input which is limited and degenerate – forms the basis of the case for linguistic nativism. The proposal is that the human brain is pre-wired to include an innate language faculty – a language-specific domain. By making use of a Universal Grammar (henceforth UG) – a genetically pre-specified body of knowledge about human language – the human language faculty permits the child to have the capability of acquiring a human language despite insufficient input.

Linguistic empiricism, on the other hand, owes its debt to John Locke (Locke 1689), who argued that there are no innate ideas or knowledge, only

[2] The paradox of inquiry asks how one can inquire into something one knows nothing about *(Meno:80d)*.

[3] As the title of Chomsky's 1966 book suggests, Descartes too held that knowledge is innate, his argument being that it is inscribed in the mind by God.

minds that are capable of undertaking a variety of tasks, of which learning, and more specifically, language learning, is one. The strongest version of empiricism states that the mind is at birth a *tabula rasa*, a blank slate onto which experience writes itself. Language acquisition is then simply a case of statistical learning, where the child employs domain-general learning processes used in the acquisition of many other cognitive abilities. A recent defence of empiricism (Sampson 2005) cites creativity as the reason for rejecting nativism. That is, if we are entirely pre-programmed, with all knowledge and ideas lying dormant, just waiting to be aroused, humans cannot be the creative beings that history has shown us to be.

The problem with Sampson's suggestion is that it is based on the strongest nativist claim possible. Just as the furthest extreme at the empiricist end of the spectrum – that we are born with nothing in our minds – is untenable, so too is the furthest extreme at the nativist end – that we are born with all knowledge that we might ever need (and more) already intact. Nativism and empiricism are not two components of a strict dichotomy, but rather two ends of a scale of possibilities. Empiricism can be made more amenable to the nativist by, for example, assuming that the mechanisms used in learning language from experience are not domain-general but domain-specific; in other words, although knowledge of language is not innate, language-specific learning processes are. Similarly, nativism can be made more amenable to the empiricist by reducing the contents of UG; allowing certain linguistic knowledge to be derived from general cognitive intelligence leaves us with less innate machinery for genetics to account for.[4]

The MP is a nativist theory in a strong sense. It assumes a genetically predetermined language faculty, and a UG. Yet the MP differs from its predecessor nativist theories in reducing the amount of innate knowledge.[5] As will be detailed later in this chapter, the MP envisions a very different system of language, in which the innate machinery required to account for the linguistic evidence is smaller, more atomic, and more economical than previously assumed.

1.2.2 What is a generative grammar?

Generative grammar dates from the end of the 1950s, and is associated primarily with two of Chomsky's early publications (Chomsky 1957, 1975b). A *generative* theory began as one which assumes that the possible sentences of

[4] In fact, we should not be making a choice between nature and nurture; but recognise that both are relevant. Caspi *et al.*'s (2007) recent paper on the relation between breastfeeding and IQ makes this point especially clearly.

[5] At least this is how the situation appears; chapter 5 will show the truth to be somewhat different.

a language are generated from a set of grammatical rules (a generative grammar). Viewing language in this way promotes the notion of creativity; language is creative because the underlying grammar consists of rules which permit the language user to produce (and understand) infinitely many novel combinations of lexical items. Generative theories of language thus rely heavily on syntax as the driving force of the innovation; they are what have come to be known in some parts as *syntactocentric* theories (Jackendoff 1998). The nativist doctrine is an intrinsic part of a generative theory; generativism assumes that the framework for grammar forms part of the newborn's genetic endowment.

The MP is a generative theory of language because (i) it places syntax at the core, relegating non-syntactic aspects of the system to peripheral components, (ii) it understands the possible utterances of a language to be generated from a grammar (although in more recent times, the underlying set of rules has been replaced by three underlying operations applied to lexical items), and (iii) it judges the underlying machinery to be innately encoded.

1.2.3 What syntactic theory tells us about the language faculty

The question then is what we know about the language faculty from what syntactic theory says. In other words, what constraints on a unified theory of the system of language do theories of syntax in general, and the MP specifically, impose?

Language is typically thought of as a two-way relation holding between a signal (be it spoken, signed, or written) and a meaning. The first thing that syntactic theory tells us is that the grammatical structure of language is the mediator between signal and meaning. The MP, as a generative, syntactocentric theory, imposes the more specific condition that the syntax actually creates this relation; the grammatical architecture works to construct a mapping between phonology and semantics. Syntactic theory furthermore tells us that the language faculty must incorporate a means for permitting an unlimited repertoire of utterances – a productive system of some sort. A generative theory, and hence the MP, on top of this tells us that the productive system in question is coded in our genes, and pre-wired into our brains from birth. Thus, a generative theory bids us to focus on I(nternal)-language – the competence or knowledge of language that is in the mind of the individual – rather than E(xternal)-language – the performance of the individual, or language as a property of a community (Chomsky 1986a). The MP additionally dictates that the genetically endowed linguistic component should be the most minimal machinery required to allow its possessor to acquire and use his native language.

This vision of the language faculty raises numerous questions for evolutionary linguistics. Firstly, why would a process for mapping between signal and meaning have evolved? In other words, assuming prior signals and meanings,

what evolutionary pressures might have given rise to a systematic means for creating correspondences between them? Secondly, how could the language faculty have come to be innate? That is, what has happened in the evolution of our species that has led to not simply the emergence of linguistic abilities, but the emergence of the genetic machinery for these abilities? Thirdly, how could a complex capacity come to be represented internally to the mind in a minimal fashion? In other words, what sort of evolutionary processes lead to low structural realisation for high functional effect, and could the language faculty have evolved in this manner?

1.3 Studies in evolution

Evolution can be understood as a change in the proportions of units in populations of organisms. Those units can be genetic, cultural, linguistic, etc. It is evolutionary theory in the typical biological sense that is relevant to the arguments of this book.[6] In this sense, evolution is the study of changes in the genepool of a population.

Evolutionary theory posits a number of requirements for evolution, and a number of mechanisms of evolution. In order for evolution to take place, there must be genetic variation in a population of individuals. Every individual in a population possesses genes, and many of these genes will be common for individuals of the same species. However, in order for evolution to occur, some genes must differ between individuals. Differing chemical variants of genes (alleles) provide this variation in the population. Genes are replicated across generations of individuals. In the case of asexual species, the chromosome (the body on which genes are carried) is copied indentically to the offspring. In sexual species, each parent passes genes to the offspring, and the process of recombination merges together parts of each parent's chromosome in the offspring, giving rise to genetic variation. The mechanism of mutation can also produce variation in both asexual and sexual populations. Mistakes can occur in the copying of genes, and external environmental factors can impact on a gene, in both cases altering the DNA sequence of a gene, and resulting in biological innovations in the population.

With variation in a population, two mechanisms are at work to decide which individuals will survive and produce offspring. Selection is normally understood as natural selection, but selection at the level of kin is also known to operate (Haldane 1955; Hamilton 1963).[7] Genes produce phenotypes – the physical characteristics of an individual. Those phenotypes which are more

[6] See Skelton (1993) for a good introduction to biological evolution.
[7] Somewhat more controversial is the claim that selection operates at the level of the group (Wynne-Edwards 1986).

adaptive in some way will survive above the others; this is natural selection. In other words, variants in a population must exhibit differences in fitness. Fitness is a measure of the ability of an individual to survive and leave offspring that themselves survive and leave offspring (Skelton 1993). If the fitness of variants does not differ significantly, selection may not be enough to ensure the new (slightly fitter) variant becomes established in the population (Zuidema 2004). The individuals that survive produce offspring whose genotypes (the full allellic make-up of an individual) are closer to theirs than to those individuals who were not selected, and so there will be a larger proportion of these gene variants in future generations of the population. Selection may also be interpreted as sexual selection. In this case, traits will be selected for as a result of attractiveness to potential mates.[8]

The second mechanism at work here is known as drift. This occurs when the proportion of a particular gene variant in a population changes statistically due to chance events. For example, binomial sampling errors occur when parents have a small number of offspring, the frequency of a particular gene not being reproduced exactly in the next generation (Suzuki *et al*. 1989: 704). Gene drift also proceeds in such a fashion that, over time, a particular gene variant will become more and more common in the population. Thus, both selection and drift decrease rather than increase the variation in a population of individuals. Numerous instances of selection and drift lead to the modification of a species over time. The core of modern evolutionary theory is what became known as the Modern Synthesis – the unified theory of Darwinian natural selection and Mendelian genetics. The core teaching of evolutionary theory from a population biology perspective is often summed up in the following three sentences: Genes mutate. Individuals are selected. Populations evolve.

This somewhat simplified view is made more complex by a number of considerations. Firstly, genes are not entirely independent units, but are associated with each other in numerous ways. Genes that occur together on a chromosome are physically linked, but such linkages can be upset by recombination. This is the process whereby strands of DNA are broken up and combined with other such pieces in different ways, usually occurring as chromosomal crossover during meiosis (the cell division occuring in the reproductive cells in order to reduce the number of chromosomes in the offspring by half). Associations between genes can also be created as the result of epistasis. The joint effect of two or more genes may hugely outweigh their separate effects, and so selection will ensure that these genes appear in the individual in combination more frequently than would be expected by chance.

Secondly, the association of genes to traits, or the genotype–phenotype map, is not at all transparent. Some of the reasons for this will be discussed in

[8] Sexual selection is often invoked in accounts of the evolution of language; some further details will be introduced in chapter 4.

more detail in section 3.8, but some hints can be given at this stage. Genes can be pleiotropic. In this case, a single gene influences multiple phenotypic traits simultaneously. A mutation on this gene will then affect all of the relevant phenotypic traits. An example of pleiotropy is seen in the disease phenylketonuria (PKU), where the phenotype arising out of a mutation in a particular gene exhibits mental retardation as well as reduced pigmentation in skin and hair colour (Nash 1999). On the other hand, phenotypic traits can be degenerately coded for. In this case, multiple genes influence one and the same phenotypic trait. This has been shown by the lack of phenotypic effect exhibited when a particular gene is deleted in what are known as *knockout* experiments (Tononi *et al*. 1999).

1.3.1 Misconceptions about evolution

A number of misunderstandings of evolution are prevalent in domains outside biology. Before leaving this brief introduction to evolutionary theory, this section will examine and clarify some such misconceptions, in order that the topics in the chapters which follow will not be misinterpreted. Johansson (2005) offers a clear overview of some of the most frequent misunderstandings about evolutionary theory, and much of what follows can be seen in more detail in his book.

(i) *Evolution works in a vacuum* – In section 2.5.3, some constraints on reaching perfection through adaptation will be considered. The general point is that there are numerous types of environmental factors that influence and constrain in what direction an organism can evolve.
(ii) *Evolution has foresight* – An organism can evolve only towards what is adaptive at that point in time; there is no way for the evolutionary process to know what will be useful, or indeed maladaptive, when the environment changes in the future.
(iii) *Evolution works for the good of the species* – A trait is selected in an individual organism because it is beneficial to that individual, not because it benefits the species to which that individual belongs. Evolution is concerned with transmitting the genes of that individual.
(iv) *Features can evolve multiple times* – Although we find similar complex features in multiple organisms, it is unreasonable to assume that the feature has evolved multiple times independently. Instead, it is more likely that the organisms have a common ancestor, and the feature is homologous.
(v) *Evolution reaches global optima* – Section 2.4 will be concerned with showing that this statement is a misunderstanding. Although evolution

can reach local optima, it frequently will become stuck there unable to climb down in order to reach the global optimum.

(vi) *Evolution climbs a ladder* – The evolutionary record shows that there is not a clear path from lower species to higher species, the lower one dying out as the higher one invades. Instead, the typical family tree is bush-like, with multiple branches taking different species in different directions.

(vii) *A trait's current function is the same as its original function* – Looking at the function that some trait fulfils in the present tells us nothing about whether it evolved *for* that function. It may have been later on exapted (Gould and Vrba 1982) to fulfil a new function.

(viii) *Gradual evolution results in useless intermediate stages* – It is often argued that half an eye is of no use to its possessor, and thus evolution cannot be gradual. Yet, the visual capacity that half an eye might afford is likely to be considerably better than no vision at all. Or in the case of language, having a rudimentary system akin to protolanguage (see section 4.5.2) would be considerably better than having no language at all.

1.4 Studies in language evolution

Dobzhansky's famous quote (actually the title of a 1973 essay) – 'nothing in biology makes sense except in the light of evolution' – will be seen in the chapters which follow to be relevant also to language. The field of language evolution uses what knowledge we have from studies in evolution more generally to posit hypotheses about how humans have come to have language. Although this question is not a new one, it is only in recent decades that the knowledge we have amassed from other scientific fields has permitted the type of careful and credible argument whose lack the 1866 ban of the Société Linguistique de Paris mitigated against.

1.4.1 A taxonomy of theories

The following sub-sections will characterise evolutionary accounts of human language as being classifiable along multiple dimensions.[9] I will outline the central tenets and principal motivations of each class of theory, supplementing each with examples of the particular types of argument advanced. Importantly,

[9] West-Eberhard's (2003) book is useful for an understanding of such different classes of theory which hold not just in evolution, but in biology more generally.