

Part One

Theoretical and methodological approaches

PROOF

2

Innateness and learnability

Virginia Valian

2.1 Introduction

This chapter addresses five questions. 1. What is the debate between nativism and empiricism about? 2. If there is innate linguistic content, what are good candidates for it? 3. What are the arguments for and against nativism? 4. What acquisition mechanisms are there? 5. What kind of empirical evidence do we presently have that would allow us to decide whether humans innately have some linguistic knowledge?

2.2 The nativism–empiricism debate

2.2.1 The central question

The central question about nativism is whether the child's mind has content independent of experience. The important word is 'content'. By content I mean knowledge, in the form of concepts and propositions. It is not controversial that humans are more sophisticated learners and users of information than any other species. Researchers may disagree about just how to characterize learning and memory mechanisms, but everyone agrees that all species have built-in methods of acquiring information. The nativism–empiricism debate is about content: does the mind have any content prior to experience? All learning mechanisms operate on content of some sort. It is the origin of the content that divides nativists and empiricists.

The least sophisticated content is primitive categories for classifying sense data, categories like colour and form. Those categories allow us to group together stimuli that share properties (such as redness). Perceptual categories such as lines and angles allow us to recognize a stimulus we have encountered before. Empiricists and nativists alike accept rudimentary categories that are based on physical properties. It is when we move

beyond perceptual categories to concepts that differences between empiricists and nativists arise. Strict empiricism rules out any innate knowledge in any realm, but it is possible to accept innate concepts in some domains and reject them in others. To take one example, it is possible to be a nativist with respect to nonlinguistic concepts but an empiricist with respect to language. A concept in the cognitive domain might be the notion of an agent of an action or the notion of logical (predicate-argument) structure in thinking, concepts that might be useful in the acquisition of language. A concept in the linguistic domain might be the notion of syntactic categories like noun or verb. According to content nativism in linguistics, some abstract linguistic concepts, such as syntactic categories, are necessary in order to explain the child's eventual knowledge. Empiricism denies such innate content.

Is there a middle ground between nativism and empiricism, or a way of avoiding the nativism-empiricism controversy altogether? To say, for example, that humans are 'biased' or 'predisposed' to learn language might seem to be a middle ground. But it is only while they retain their vagueness that biases or predispositions appear to be a middle ground. If, once they are fleshed out, the biases involve the absence of innate syntactic content, then they are empiricist; if they involve innate syntactic content, then they are nativist. Interactionism (Elman *et al.* 1996, Thelen & Smith 1994) is sometimes presented as an alternative to either nativism or empiricism, as is constructivism or usage-based theories (e.g. Tomasello 2003; see also Chapter 5). In both cases, the organism is seen as actively contributing to whatever knowledge is acquired. But a mind can be active without having prior linguistic content, and it is the postulation of innate content that marks the nativist. Since both interactionism and constructivism either argue against innate syntactic content or assume that it does not exist, those positions are also forms of empiricism.

2.2.2 Preliminaries and terminology

The question of what linguistic concepts are innate can be asked about every aspect of language, from phonology to pragmatics, but this chapter will focus on syntax (and morphosyntax), since that is where debate is concentrated. Although syntactic concepts are no more complex or abstract than semantic concepts, there is nevertheless less debate about semantics, perhaps because it is (incorrectly) seen as part and parcel of cognition.

In the key arguments advanced by nativists and empiricists, conceptions of the 'final state', that is, the individual's mature mental grammar, are closely related to conceptions of the 'initial state', that is, the child's mental grammar before she has had access to linguistic input. (The foetus has access to prosody (the melody of the language) but not to structured syntactic input.) Much of the dispute between nativists and empiricists

follows from their different judgments about the correctness of formal linguistic descriptions of language as an approximation of people's mental grammars. With different conceptions of the final state, different conceptions of the initial state are likely. The more abstract and complex the final state, the more likely a rich initial state is. A nativist need not adopt a complex picture of the final state, but adopting a complex picture makes it more likely that one will be a nativist, because input can only provide examples, not abstract structure itself.

In this chapter I assume that a formal grammar is an approximation of the child's final state. Formal theories cover a broad range of syntactic phenomena and aim for systematicity and coherence. Although I am assuming a 'generativist' framework, sometimes referred to as the principles-and-parameters framework, other frameworks are also possible. The key feature of formal theories is that they state what the theoretical vocabulary is – syntactic features and categories; they state the conditions under which categories can be combined to form larger constituents; they state the conditions under which word forms (such as the nominative and accusative cases) are possible, and so on. Formal grammars offer specific proposals about language universals, which in turn can be the basis for hypotheses of what is innate. Within the generativist framework are a variety of proposals, such as one or another form of minimalism (Chomsky 1995) and cartography (see, e.g., Shlonsky 2010). (See Carnie, Siddiqi & Sato 2014, Borsley & Börjars 2011 for reviews of issues in grammar and summaries of different formal grammatical theories.)

Nativism does not commit someone to a particular grammatical theory but is compatible with a wide range of theories. Nor does nativism commit someone to a particular philosophy of linguistics. Nativism is equally compatible with a view of language as a theory of people's psychological (or biological) states (Chomsky 2006) or as a theory of abstract objects (Katz 1981). Finally, the choice of a formal grammar as a model of the child's syntactic knowledge is compatible with seeing language as a vehicle for a wide range of communicative intentions. The formal grammar is one tool the child (and adult) has for expressing those intentions.

2.2.3 Examples of what is acquired: categories and word order

Two 'simple' aspects of language are acquired early by all speakers: syntactic categories and word order. (Sections 2.6.5 and 2.6.6 provide more detail.) Syntactic categories fall into two main linguistic types: **lexical** and **functional**. The lexical categories are nouns, verbs, adjectives and adverbs, and, in some cases, prepositions. Functional categories include determiners (words like *the* and *my*), inflectional elements (such as tense on a verb and auxiliaries in English) and complementizers (such as the *that* of 'I knew that she was happy'). Functional categories typically contribute less to the meaning of a sentence than lexical categories do.

Nativists and most empiricists agree that children's grammars – at some point – include abstract syntactic categories and represent word order in terms of abstract categories. Disagreements concern the origin of categories (and when they are acquired; see Section 2.6.5). Nativists typically start with the hypothesis that at least some syntactic categories, or the features that make up those categories, are innate; empiricists start with the hypothesis that no syntactic features or categories are innate, but rather are induced based solely on exposure to the distribution of those elements across the language (plus cognitive categories of some sort).

Does this mean that nativists leave no role for learning? No, learning still has an important role, for example, in determining what categories particular words belong to and what language-particular combinations are possible. For example, in English, the child has to learn that *the* can precede any noun, but *some* can only precede mass nouns and plural nouns. Crucially, however, what is learned is not the abstract categories themselves. Instead, learners acquire a mapping between the innate abstract categories and the particular words in the learner's target language that belong to each category. For empiricists, the hypothesis that no categories are innate means that the only way of acquiring them is by learning. Among the earliest such proposals is one by Braine (1963), proposing that children construct a pivot-open grammar in which certain words or word combinations, like *here's a*, act as pivots which the child can finish with a wide range of words (almost always nouns). More recent proposals include lexically specific formulae (Pine & Lieven 1997, Pine & Martindale 1996), lexically based learning and usage-based learning (Tomasello 2003, ch. 5). After the child has amassed a number of such cases, she creates categories for the different words.

The agreement about the child's state, at least by age 5, with respect to syntactic categories and word order makes it possible in principle to examine different learning mechanisms to see what innate content, if any, is required in order for the mechanism to arrive at those categories by that age.

2.3 Candidates for innateness: linguistic universals

The generativist framework offers linguistic universals as candidates for innateness. Linguistic universals are principles and properties that (a) are true of language and (b) define what it is to be a language. It is not enough just to say (a). Even properties that are true of all languages may hold because of irrelevant properties of speakers rather than because of properties of language.

The existing sentences in all languages are, for example, of finite length. But the finite length of any given sentence is due to speakers' limited

cognitive systems (and limited lifetimes), rather than due to speakers' language. We would not want to say that finite length is a linguistic universal. Speakers acquire a theory of their language that allows for sentences of any length whatsoever, even though people cannot physically produce sentences that would take more than a lifetime to utter.

For two reasons linguistic universals are good candidates for what could be innate syntactic content. First, universals set the defining conditions on what could be a language. Whatever is innate should not be particular to a single language but to language. Second, any child can learn any language. If anything is going to be innate, it is the abstract linguistic features that allow a child to be an omniscient language learner.

Linguistic universals are of two types: **absolute** and **relative**. Absolute universals are syntactic principles or structures that appear in every language (Chomsky 1981b). One reason to expect all absolute universals to be innate is that, by definition, they hold for every language. They are the best linguistic survival kit a child could have. Another reason for hypothesizing their innateness is that absolute universals are abstract and cannot be directly perceived from exposure to sentences. Later in this chapter I will give the case filter as an example.

What I am calling relative universals are of two types. One type is the building blocks of syntax – **syntactic features and categories**. The entire stock of features and categories may be innate, or only a subset may be innate. Not every language uses every feature and category. Some languages, for example, have a genuine future tense, but English does not. Tensed main verbs in English are either present or past tense. The 'future' in English is carried by the modal *will*, or combined forms like *be going to*; main verbs themselves do not have a future form. French main verbs, in contrast, have present, past and future tenses. Even if all features are innate, they will not all surface in any particular language, just as future tense does not surface in English. But it is also possible that only a subset of features is innate.

In addition, the members of a given category may differ from one language to another. For example, in English, possessive pronouns behave like articles and cannot be combined with them ('the my ball' is impossible in English). In Italian, however, possessive pronouns behave like adjectives and can be combined with articles. Thus, the innate specification of categories must be abstract. An innate syntactic category will not come with a list of examples, because the exact examples will vary (if only within a narrow set of boundaries). Similarly, no particular word order can be innate. In some languages, like English, function words tend to precede lexical categories within a phrase (*the ball*), but in other languages, function words tend to follow lexical categories. The dominant English word order is subject–verb–object, but in other languages other orders are possible.

The second type of relative universal is **parameters**. Parameters define dimensions of linguistically significant variation, such as whether the

subject of a verb must be overt. Another parameter concerns word order: in English the verb comes before its object, but in Japanese the object comes before its verb. Parameters are typically two-valued; each language takes one value or the other for each parameter. Parameters are an important type of linguistic universal, since they map out what syntactic variation is possible. By hypothesis, all parameters are innate, and each is independent of every other. The child's task is to choose, over the course of development, which value of each parameter characterizes her language. Parameters are relative universals because, for a given language, only one value can be correct.

An alternative to localizing crosslinguistic syntactic variation within parameters is placing them in the functional category system (the so-called Chomsky-Borer hypothesis; see Borer 1984, 2005). Variation in the features of functional categories will determine such phenomena as whether a language obligatorily represents overt subjects, as English does. A number of variations of this hypothesis exist. Wherever variation is addressed within the grammar, the grammar must both constrain possible variation and account for the actual variation that exists.

A useful heuristic for identifying candidates for innateness is that they be universal in either the absolute or relative. Within linguistics, the set of absolute and relative universals is referred to as Universal Grammar. Universal Grammar forms the upper bound of innate syntactic content. But the upper bound is not necessarily also the lower bound. A nativist could take a much more modest position and propose that only some universals are innate, while others can be inferred.

2.4 The logical argument for innateness

2.4.1 Types of linguistic evidence

The main logical argument supporting the claim of innate syntactic content is the argument from the poverty of the stimulus. This argument states that input contains too little information for children to reach the final state; the input is impoverished. Some examples of poverty of the stimulus arguments are: subject-auxiliary inversion in questions in English (see Pullum & Scholz 2002, and responses by, among others, Fodor & Crowther 2002, Legate & Yang 2002; see also Pullum 2012) and anaphoric *one* in English (see Hornstein & Lightfoot 1981, Lidz, Waxman, & Freedman 2003, and responses by Regier & Gahl 2004, Tomasello 2004; for an overall response to arguments against the poverty of the stimulus, see Berwick, Pietroski, Yankama & Chomsky 2011). I will not review these examples but, in Section 2.3.1, consider a syntactic phenomenon commonly referred to as the case filter, which I think is a good candidate for a poverty of the stimulus argument.

To say that the input is impoverished is different from claiming that it is noisy or degenerate. The former claim is that input to children lacks information that would allow children to acquire certain syntactic principles or regularities. The latter claim is that input to children includes run-on or incomplete sentences, false starts and perhaps some outright ungrammaticalities. Speech to children tends to be short, free of hesitations and generally free of outright errors, though it does contain a reasonable number of fragments as well as (in English) sentences without subjects about 5 per cent of the time. The language acquisition mechanism is obviously built to withstand a certain amount of noise in the input.

The important question is how the mechanism copes with impoverished input. Input, in the form of speech to the child (or speech that the child hears), is called positive evidence. That speech illustrates sentences of the language. It is evidence that certain words and phrases occur. Two other possible types of evidence are negative evidence and indirect negative evidence. Negative evidence is responses from the child's interlocutor either that a certain way that the child has just spoken is ungrammatical or that the child should replace her formulation with the one the interlocutor has just produced. If, for example, the child says 'I knowed it' and the parent says 'Oh, you knew it', the use of *knew* for *knowed* could constitute negative evidence (sometimes also called implicit correction, negative feedback, a recast, or a reformulation). Similarly, if a child says 'That the last one' and the parent says 'That's the last one' the use of *that's* for *that* could constitute negative evidence.

Indirect negative evidence is the absence of a structure that the child would expect to see, given a starting hypothesis. If, for example, an Italian child thought that subjects might be required, their consistent absence in sentences like *Piove* 'It's raining' might be sufficient for the child to revise that hypothesis.

All three sources of evidence are imperfect and require inferences on the child's part. Although adults' grammatical violations in talking to children are few, they might temporarily mislead the child. Negative evidence is also imperfect, both because it does not occur every time the child makes a mistake and because the child might not recognize it as a correction. Data from my laboratory, based on 21 child-mother pairs, suggest that parents provide 'implicit' corrections for about 25 per cent of children's ungrammatical utterances, leaving most ungrammatical utterances uncorrected. More to the point is that the child might not recognize the use of *that's* for *that* as a correction. Finally, indirect negative evidence requires the child both to have a specific hypothesis and to determine whether the absence of an expected sentence type is due to syntactic or nonsyntactic reasons. People never produce triply-embedded sentences to children, for example, but children should not take that absence as evidence that triple embeddings are ungrammatical.

2.4.2 An example of a poverty-of-the-stimulus argument: the case filter

Consider examples (1)–(5); only (1) is grammatical. (The * indicates ungrammaticality.) What distinguishes the examples is that (2)–(5) all have the incorrect case for one or both pronouns. *Case* (specifically, abstract or structural case) refers to the syntactic function that a noun or pronoun has in a sentence. It is not the same as the semantic role, as is apparent by the contrast in (1) and (1'). The first person is the person doing the greeting in both sentences, but in (1) the pronoun has nominative case (*I*) and in (1') it has objective (or accusative) case (*me*). Similarly, the third person is the one being greeted in both sentences, but in (1) the correct form is *him* and in (1') it is *he*.

- (1) I greeted him yesterday; (1') He was greeted yesterday by me.
- (2) *Me greeted him yesterday; (2') *Him was greeted yesterday by I
- (3) *My greeted him yesterday; (3') *Him was greeted yesterday by my
- (4) *I greeted he yesterday; (4') *He was greeted yesterday by I
- (5) *I greeted his yesterday; (5') *His was greeted yesterday by me

Case is a syntactic property that noun phrases (NPs) have as a function of their relation to another category, such as a verb, a preposition, an inflectional element like tense, or another noun phrase. English has three cases: nominative, objective (or accusative) and possessive (or genitive; see Carnie 2006, for an introduction to case and other syntactic properties and relations). Although case is only visible on pronouns in English, the case filter claims that it is *invisibly* present on all overt nouns in English. If we replace *I* with *the girl* in (1), *the girl* has nominative case even though the case is not overt. In some languages, such as Hungarian, most cases are overt on all overt noun phrases, both pronouns and nouns. And some languages, again like Hungarian, have other kinds of cases, sometimes called morphological case, upwards of ten.

The case filter is an example of an *absolute* universal. It is the requirement that all overt nouns and pronouns in every language have case; different cases may have distinct morphological forms, as with first-person pronouns in English, or may be abstract and have no external form, but only a positional relation to another grammatical element that can assign case to the noun or pronoun in question, as with full lexical noun phrases in English and all nouns and pronouns in Thai. (The word 'filter' is used because structures containing an overt NP that is not cased are filtered out. See Rezacs 2013 for discussion of case.)

Even though case is largely morphologically absent in English, there are examples that show it is grammatically present and, in the example of objective/accusative case, assigned by verb (or preposition) just to its left. Without the concept of case, the ungrammaticality of certain

sequences is otherwise inexplicable. In (6), the verb *consider* assigns objective case to *Jane*.

- (6) Lee considered Jane to be happy
- (7) *Lee considered she to be happy
- (8) Lee considered her to be happy

The ungrammaticality of (7) and grammaticality of (8) show that the position right after the verb, if filled by a noun or pronoun, is one that receives objective case; otherwise *she* would be an acceptable substitution for *Jane*. *She* would be acceptable if the following verb, instead of being an infinitive, were tensed, as in *Lee considered she would be happy [to receive the package]*. In that case, the tensed verb assigns nominative case to the pronoun.

If an element intervenes between the verb and the following noun, case cannot be assigned and the resulting string of words is ungrammatical.

- (9) Lee considered quickly whether to go
- (10) *Lee considered quickly the matter
- (11) Lee quickly considered the matter
- (12) Lee considered the matter quickly

In (9) it is possible to put the adverb *quickly* directly after the main verb *considered*, although it is a bit awkward. In (10) the sequence is worse than awkward; it is not grammatical. The important difference between (9) and (10) is that in (9) there is no overt object NP, whereas in (10) there is (*the matter*). Since, in (10), an adverb intervenes between the verb and the NP to which it would otherwise assign accusative case, the sentence is ungrammatical. If the adverb is moved so that it does not intervene between the verb and its object, as in (11) or (12), the sentences are grammatical. In English, then, if an element intervenes between the verb and its noun, objective case cannot be assigned.

A sequence like (13), which is easily understood, and is very similar in surface form to sentences like (6), (8) and (9), is nevertheless ungrammatical. The NP *Jane* is uncased: *whether* intervenes between *considered* and *Jane*, preventing the verb from assigning case verb to the NP.

- (13) *Lee considered whether Jane to go

Neither *she* nor *her* can substitute for *Jane*, also showing that the position is one which cannot receive case. If it could, at least one cased form of the pronoun would be legitimate. (Again, the sentence can be saved by changing the infinitive to a tensed verb, as in 'would go'.) Without the case filter, the ungrammaticality of (13) is inexplicable. (13) violates the case filter, and is thereby ungrammatical.

The concept of the case filter presupposes the concept of grammatical case, the category of NP and a syntactic mechanism for assigning case. That mechanism in turn involves reference to syntactic categories like verb and preposition. The claim that all NPs in every language must have case is thus embedded in a linguistic system. Only within that system does the claim have meaning. If the case filter is innate, so are the concepts that comprise it.

The case filter is a good example of a poverty of the stimulus argument. Native speakers of English show, by their acceptance or rejection of the sequences in (1)–(13), that a concept like the case filter is part of their mental grammar. It is part of their grammar even though they have no explicit knowledge of case or the case filter. Crucially, there is no evidence in the input that could lead speakers to put case in their grammar. Case does not correspond to concepts that might be more easily inferred from context, such as ‘agent of an action’ or ‘object of an action’. Case is purely syntactic (and, in languages with overt case, morphosyntactic).

There is no way to acquire the case filter from positive evidence. Unlike examples with subject–auxiliary inversion, where there is disagreement about how many possibly informative examples might exist in speech to children, in this example there are no examples. There is also no way to acquire the case filter from negative evidence. Even if children spontaneously produced sequences like (10) and (13) (of which there are no known examples), and received reformulations by their caregivers, nothing in the reformulation could allow the child to infer the case filter or the concepts that make it up. Indirect negative evidence could lead children to wonder why no sequences like (13) are in their input. They might expect to hear combinations of sequences like (6) and (9). But there is no path that could take children from the absence of such combinations to the syntactic components of the case filter.

2.5 Arguments against nativism

Arguments against nativism generally take the form of parsimony arguments. If acquisition can be explained without recourse to innate content, then no innate content should be proposed. The fewer entities – mental or otherwise, innate or acquired – the better. Nativism seems to posit more entities than empiricism and thus to be less preferable. But parsimony is a comparative notion that demands (a) two theories for (b) the same body of facts. Parsimony chooses between two specific competing explanations of the same set of phenomena. If one theory accounts for more data than another, the fact that it uses more entities than another theory is not a violation of parsimony. Parsimony never comes into play.

The need for a comparable set of data is one reason that the conception of the final state is so important in language acquisition theories. If very

little knowledge of an abstract character is acquired, very few mental concepts – innate or otherwise – will be required to explain that knowledge. If a great deal of abstract knowledge is acquired, many more concepts will be encompassed. The example of the case filter is a case in point. Empiricist theories have not addressed its acquisition. Since nativists and empiricists tend to disagree about the nature of the final state, parsimony is usually an irrelevant principle: the two positions are not explaining the same set of phenomena and thus cannot be evaluated with respect to parsimony.

An alternate approach is to stay closer to the data. Some investigators have analysed corpora from early child speech and concluded that the child does not – during the specific time period when the observations are made – have one or another abstract syntactic category, such as determiners (e.g. Pine & Martindale 1996). Instead, the child has local and limited knowledge about particular words that function as verbs or determiners in the adult system. If syntactic categories like determiners play no role in the young child's performance, they appear otiose. One can achieve a simpler and more parsimonious account of the child's behaviour by omitting such innate categories and postulating that they develop later, after the child has abandoned narrow, lexically specific generalizations.

But if the child does eventually acquire knowledge of an abstract category, as almost everyone agrees is the case, she must – within this empiricist approach – shift at some point from a set of unrelated small-scale word patterns to an organized category. Such qualitative differences must be accounted for in some fashion, either by invoking additional concepts or additional mechanisms. Something may have been saved by ruling out innate categories, but something will be spent by postulating as yet undetermined mechanisms. The extent to which the initial parsimony yields a net saving is thus unknown.

One important goal of language acquisition theories is an explanation of how the child arrives at her final state. It is not enough to describe one or another point in development. If the child shows no clear knowledge of a concept at one time, but does show knowledge of it at a later time, the theory of development must state how that change takes place.

Nativists solve part of the problem of syntactic development by postulating a continuous process in which the child learns how to map innate categories and structures onto input. The initial learning mechanism continues until learning is complete. Thus, the nativist does not postulate unknown learning mechanisms of unknown complexity. Rather, the complexity of the system is known, in principle, at the outset: this innate content, this learning mechanism. The contrast between the two approaches to development demonstrates their incomparability. They are not explaining the same phenomena and thus neither can be rated as more or less parsimonious than the other.

Whether development actually *is* continuous in the nativist's sense is irrelevant to the logic of the continuity argument. What is important is

that development *could* be continuous in the way the nativist postulates. The possibility of continuity, coupled with a final state of knowledge of abstract categories, means that we cannot assess theories with respect to parsimony or simplicity until we have competing theories of how knowledge develops to an agreed-upon final state.

2.6 Mechanisms of acquisition and learnability

Any theory of acquisition has to show that the knowledge postulated for the prior state, plus a particular learning mechanism, plus the input, will yield the knowledge postulated at the subsequent state. Learnability theories seek to lay out those elements: what combination of the learner's initial stock of concepts, mechanism of acquisition and input will yield a particular intermediate or final state (see, for example, Berwick & Niyogi 1996, Fodor 1998, Gibson & Wexler 1994, Lightfoot 1989, Sakas & Fodor 2012, Wexler & Culicover 1980, Yang 2002, 2012). When learnability researchers try to model acquisition of an entire language, they discover enormous difficulties even when they provide the model with a great deal of innate content.

Some learnability models ~~often~~ propose a form of acquisition called triggering (e.g. Sakas & Fodor 2012). A trigger is a minimal input – perhaps only a single sentence – which is sufficient to set the correct value of a binary-valued parameter. On such a model, parameter values are not learned. Rather, a parameter is like a switch, set in one position or the other by positive evidence. There are a number of difficulties with the model of triggering, but for our purposes the important point is that triggering is not psychologically plausible. It idealizes acquisition as instantaneous once the appropriate datum arrives (to a mind prepared to receive it). But since children do not appear to make instantaneous decisions, the idealization appears to misstate the actual acquisition process. Alternatives to triggering include variationist models that include constraints but also allow for a period of data-gathering (e.g. Legate & Yang 2007, Yang 2012).

One possible model of acquisition is hypothesis-testing (e.g. Valian 1990), which can be constrained or unconstrained. In nativist theories, hypothesis-testing is constrained by absolute and relative universals. The analogy is to theory confirmation in science, although there is no implication that the child consciously tests hypotheses. For parameters, or functional categories, the hypotheses are constrained by the possible values, which incoming data are used to choose between. In the case of syntactic rules, such as subject-auxiliary inversion, the hypotheses will be constrained by innate knowledge of possible syntactic structures – the fact that linguistic rules are structure-dependent.

Thus, the child will never entertain the structure-independent hypothesis that the first auxiliary in a sentence with an embedding (*The girl who is*

happy is singing) is the one which is inverted (yielding the incorrect *Is the girl happy is singing?* instead of *Is the girl who is happy singing?*; see Crain & Nakayama 1987 for relevant data, and Ambridge, Rowland & Pine 2008 for contrary data). Rather, the child will only entertain the hypothesis that the auxiliary of the matrix clause can be inverted.

In the case of syntactic categories, hypotheses will be directed to which specific categories are instantiated in the learner's language. Hypothesis-testing need not be nativist. It can be unconstrained by any innate syntactic content, though it might be constrained by cognition. Nativist hypothesis-testing differs from triggering not in whether linguistic content is assumed to be innate – in both sets of theories, there is innate linguistic content – but in what mechanism is proposed. In hypothesis-testing, learning takes time.

Bayesian approaches to language acquisition are a form of hypothesis testing, but of a very different sort than the frequentist hypothesis testing I have just described. In the Bayesian framework the individual is faced with data and wants to draw an inference about the process that gave rise to the data, based in part on (an estimate of) the prior probability of a hypothesis plus the likelihood that one would observe the incoming data if that hypothesis were true. The combination of the prior probability and the likelihood yields a posterior probability. (See Griffiths & Yuille 2006, Perfors, Tenenbaum, Griffiths & Xu 2011, for introductions.) They may or may not be nativist, depending on how one interprets some of the assumptions. One Bayesian analysis of question formation claims that the hierarchical structure that is required to allow the child to front the correct auxiliary can be established by a learner that chooses, via Bayesian methods, between a linear and a hierarchical grammar that are initially set to be equally probable (Perfors, Tenenbaum & Regier 2011). A nativist interpretation of such an analysis is that it is indeed nativist, since the grammars that are being evaluated are all part of the prior set of hypotheses. (For a critique, see Marcus & Davis 2013.)

Any form of hypothesis-testing uses one or another form of distributional analysis to evaluate the incoming data. Distributional analysis is essentially a form of pattern analysis in which learners observe what elements of a sequence go where, what elements can substitute for other elements, and what elements tend to occur together. Many different instantiations of such models have been proposed for different aspects of language acquisition (Cartwright & Brent 1997, Freudenthal, Pine & Gobet 2006, Mintz 2003, Redington, Chater & Finch 1998; see Chapter 3, for a discussion of statistical learning). Models differ in what units they presuppose. For example, most models aimed at acquisition of syntactic categories assume that individual words (and sometimes morphemes) are available to the child; the bracketing of speech into words is assumed already to have taken place. Non-nativist theories try to eliminate any morphosyntactic information, such as information about what categories to aim for. Models of isolated pockets of syntax at particular points in the acquisition sequence can achieve at least limited

success with relatively little by way of innate content, although even models limited to acquisition of syntactic categories have had only partial success (e.g. good accuracy but low completeness, Mintz 2003, or the reverse). Given the failure of taxonomic linguistics, it seems unlikely that a purely taxonomic approach to language acquisition could be successful. There are no non-nativist theories that have tackled acquisition of the entire grammar.

I am omitting here a range of curve-fitting models like dynamical change models and connectionist models. In these models, learning mechanisms are seen as continuous and what is learned is seen as discontinuous. What a given network learns appears to change qualitatively over the course of development, even as the mechanisms remain constant. In some cases, proponents of such models see knowledge acquisition as a mirage: knowledge does not genuinely take place but only appears to (Thelen & Smith 1994; see Spelke & Newport 1998 for a reinterpretation). For such models, no comparison is possible with models of knowledge acquisition, since they are explaining different things. Proponents of connectionism also sometimes propose the models as knowledge acquisition devices that lack innate concepts (Elman *et al.* 1996). In that case, the issues are whether the models presuppose some of the concepts that are supposedly learned and whether they succeed in modelling acquisition. Critiques of these models vary (for a summary of critiques of connectionism and replies, see Bechtel & Abrahamsen 2002, also see Marcus 2003, Valian 1999).

2.7 Empirical evidence that language is special

Several characteristics of language acquisition show that language is *special*. (i) Only humans acquire a full language. (ii) Language appears to be independent of other cognitive abilities: even profoundly cognitively impaired individuals have close to normal syntax; syntactic deficits occur in individuals with no cognitive impairment. (iii) Acquisition occurs most easily and fully during early childhood. (iv) Some linguistic impairments appear due to certain genetic mutations. Let us consider these characteristics of acquisition in turn.

2.7.1 Animals and language

That only humans acquire a full language is clear. Some species have communication systems that encode a limited amount of information, but no species encodes remotely as many concepts as those encoded by the languages humans acquire, and no other animals' communication systems have the form of the languages that humans acquire.

Take the dance of the honeybees, for example, which encodes the distance and direction of a source of food or possible new site for a hive. The dance does not encode the altitude of the site, despite the possible relevance

of that information (von Frisch 1967). Nor does the dance differentiate between food or a new hive. In addition, the nature of the encoding is very different from that of languages humans acquire: direction is encoded by the angle of the dance and distance by the number of waggles in the dance. This system is thus a continuous rather than discrete system of the sort used in human language (Janda 1978). There is nothing akin to grammatical categories and nothing akin to a phenomenon like word order. Vervets have alarm calls that appear to differ depending on the identity of the predator, but, again, there is nothing akin to syntactic categories or word order (Seyfarth & Cheney 2003).

Thus, naturally occurring animal communication systems differ from the languages humans acquire in two ways. First, they are not effable (Katz 1978): they do not contain the machinery that would allow communication of more than a tiny number of concepts and there is no evidence that any of the communications are propositional in nature. Second, they bear no syntactic similarity to the languages that humans acquire. Although the lack of language among animals shows that animals differ from humans, it does not entail that humans have innate syntactic concepts and animals lack them. Humans might differ from animals in their computational power alone, or in the extra-syntactic concepts they have.

Studies that attempt to expose animals to language or to teach them language might provide a better comparison. Animals that have been studied include chimpanzees, bonobos, dolphins and grey parrots. The results suggest that animals can use symbols (at least occasionally) in connection with the objects they refer to, can make limited requests using symbols, and can follow limited commands made by humans (see Kako 1999 for discussion and summary). None of these animals, however, shows evidence of syntactic categories.

If no special innate endowment were required to acquire language, then any two species with identical abilities to learn and remember information and with identical repertoires of cognitive concepts should be able to acquire language on the basis of the input provided. If one of the two species is nevertheless unable to learn language, that provides an argument for innate content. The problem, however, is that it is impossible to be certain that we have creatures who are cognitively identical. Bonobos (one of two species of chimpanzee, sometimes called a pygmy chimpanzee) and humans, for example, have highly similar learning abilities and similar cognition; they also share about 98 per cent of their DNA. But the small differences between bonobos and humans might be just those that are relevant to language. Because arguments for innate content based on cross-species differences crucially rely on the assumption of cross-species similarity of the nonlinguistic systems and of learning mechanisms, the arguments can only be suggestive.

With those caveats in mind, consider a particular bonobo, Kanzi. Kanzi's experimenters spoke English to him, attempting as much as possible to

duplicate conditions in which a human child acquires language (Savage-Rumbaugh *et al.* 1993). The experimenters also accompanied their speech by points to lexigrams on a keyboard for major words, including 'nouns' and 'verbs'. Lexigrams did not include function morphemes, so the system did not fully duplicate the auditory system. Since Kanzi could not produce speech, he had to use a combination of points to objects, gestures and lexigrams, a clear handicap compared to a normal child, a handicap that precluded Kanzi's using function words like *a* and *the*. Kanzi began learning the lexigrams for single words when just a few months old. By the age of 5 years, his sequences were 1.15 items long (only 10 per cent of his utterances were more than 1 item long; Greenfield & Savage-Rumbaugh 1990). They remained at that length for the next three years. In his short utterance length and failure to develop more complex utterances Kanzi was very different from a human child.

In comprehension tests at age 8, Kanzi appeared to understand a wide range of sentences, such as *Take the snake outdoors*, *The surprise is hiding in the dishwasher*, *Get Rose with the snake* (Savage-Rumbaugh *et al.* 1993). After hearing such sentences, Kanzi carried out the correct action almost 75 per cent of the time on average. (For *Get Rose with the snake*, Kanzi tossed the snake to Rose.) Indeed, Kanzi at age 8 was correct more often than a child aged 1½ to 2 years who was tested on similar materials and averaged 65 per cent. (Kanzi did not understand more than individual words or fixed expressions for his first 3–4 years.) Impressive though Kanzi's achievements were, he may typically have answered correctly on the basis of his knowledge of the individual items, the most plausible combination of those items, or an order of mention strategy. For example, when told to *Pour the milk in the bowl*, Kanzi performed the correct action. In this particular case, other than by eliminating one of the items, it is hard to see how Kanzi could get the command wrong. The correct action is the most plausible combination of the individual words and follows order of mention. Few minimal pairs (cases where the sentence could reverse the agent and object with full grammaticality) were included.

In production at age 5, Kanzi did not demonstrate the agent–action word order that dominates English. Instead, he systematically used action–agent order (120 times) instead of agent–action order (13 times), despite the input (Savage-Rumbaugh & Rumbaugh 1993). In English, objects typically follow actions; for Kanzi, action–object was used 39 times and object–action was used 15 times. If input determined what rules a learner would form, then 'smart' animals like bonobos would acquire a regularity as simple, obvious and robust as the agent–action–object order. Kanzi seemed to have the concepts of agent and action, he was a good learner generally and he had an enriched environment. But he did not learn the agent–action order. Kanzi does not seem to bring the same syntactic concepts to the task that children do. (See also Terrace 1987 for discussions of earlier failures with chimpanzees.)

Kanzi's data illustrate the argument that the speech data to which children are exposed underdetermines what they will acquire. Bonobos' failure to absorb the regularities in their input demonstrates that no matter how 'transparent' and input-dictated a regularity appears to those of us who acquire it, it is opaque to a learner who cannot represent that regularity in its hypothesis space. We cannot know why Kanzi did not represent word order as human children do, even after massive exposure. Although it seems likely that bonobos lack the innate syntactic ideas that humans have, it is also possible that they have different learning mechanisms or different cognition. Kanzi's data, however suggestive, cannot prove that humans have innate syntactic concepts. His data are primarily useful to us in showing that rich input does not by itself yield learning. Even under conditions of great enrichment, animals do not develop anything like a full language, while humans, even under conditions of great impoverishment, do (see below). Infra-humans' lack of natural language demonstrates that language is special, but it does not demonstrate what makes language special.

2.7.2 Dissociation between language and cognition

When we turn to individuals with various forms of cognitive impairments, we find some conditions where syntax is close to normal, as with individuals with Williams syndrome, whose cognition is extremely limited (Grant, Valian & Karmiloff-Smith 2002, Thomas & Karmiloff-Smith 2005; see Chapters 27 and 28). And there are forms of linguistic impairment, such as Specific Language Impairment, that leave cognition relatively intact (Leonard 2014; also see Chapters 24–6). Such examples suggest that language is special and at least in part distinct from other cognitive systems. But they do not tell us in what way language is special.

2.7.3 Sensitive period

Language acquisition is most likely to be complete if acquired in childhood, though there are exceptional examples of individuals acquiring native-like fluency in a new language as adults. This argues that language is different from other aspects of cognition which people typically improve at with age, until reaching a plateau. It demonstrates that language is special, but that special quality may be due to learning mechanisms that are specific to humans, rather than to innate content. (For review see Pallier 2013. Note also that birdsong development is also typically time-limited.)

2.7.4 Genetic involvement

A family known as KE has been studied for years because of the language difficulties of some of its members, difficulties which are now known to

be due to a mutation in just one gene, *FOXP2*, involving one nucleotide change (see Graham & Fisher 2013, Marcus & Fisher 2003, Chapter 24). Even though only a single change on a single gene is involved, that gene has multiple effects, perhaps by influencing the actions of other genes (Marcus & Fisher 2003). Tests of syntax comprehension and production are not the only deficits that individuals with the mutation show. Affected individuals also have difficulties telling apart words and non-words; indeed, that difference alone can distinguish family members with and without the mutation (Watkins, Dronkers & Vargha-Khadem 2002); affected individuals have some cognitive and motor difficulties as well. Further, *FOXP2* is found in a number of species and, even in humans, is related to lung and other organ functions as well as cognitive function. Finally, other forms of language delay and impairment show no mutation on *FOXP2*. Like the considerations we have examined in Sections 2.6.1 through 2.6.3, the genetic data strongly suggest that humans are wired to learn language. But the data leave unanswered the question of whether the wiring involves syntactic content or a special linguistic tuning of some sort that does not involve content.

2.8 Empirical evidence that aspects of syntax are innate

The poverty of the stimulus argument reviewed in Section 2.4 is a logical argument in favour of nativism. The empirical evidence reviewed in Section 2.7 shows that language is unlike other cognitive achievements. This section considers what an empirical argument for innate syntactic content would look like and what the best candidates are for that content.

A nativist view of category acquisition places an abstract specification of categories in the child's grammar as part of the child's initial state. For a nativist, the child's task is then to find out what words fall into each category and how that category behaves in the child's target language; input plays the role of providing specific information. On an empiricist view, the child creates the categories on the basis of regularities in the input and context.

2.8.1 'Creation' of language in the absence of linguistic input

Deaf children born to hearing parents who do not want their children to learn sign language create a limited gesture system that uses some of the devices, such as word order and inflection, that natural languages use (Goldin-Meadow 2003b, 2005; see Chapter 9). One deaf child who was studied intensively appeared to use hierarchical structure to represent noun phrases (Hunsicker & Goldin-Meadow 2012). Such children are a challenge to empiricist theories.

Another challenge comes from the evolving sign language of deaf individuals in Nicaragua. Before 1977, Nicaraguan deaf individuals had no access to other deaf individuals or to schooling. After the revolution, in 1977, twenty-five deaf individuals were brought together to a school and others joined them in successive years. By 1983 there were 400 individuals of various ages receiving education together (Senghas 2003). The first group developed a common, albeit limited, gestural system. Young individuals who entered the school later, and who were exposed to the limited sign system of the first group, developed the system further, so that it now encoded properties that were not initially present, such as a syntactic means for representing the positions of objects (Senghas 2003, 2011, Senghas & Coppola 2001, Senghas, Kita & Özyürek 2004).

The examples of children with greatly impoverished or no input contrast strikingly with the example of Kanzi described earlier. Children with minimal input develop basic syntactic properties. Such children are evidence in favour of innate syntactic content.

2.8.2 Syntactic categories in typically developing children

What about typically developing children? Children appear to have knowledge of categories, including functional categories, very early. Consider, for example, the class of determiners: articles like *a* and *the*, demonstratives like *this* and *that*, possessive pronouns like *my*, and quantifiers (see summaries in Valian 2009, 2013). Spontaneous speech data demonstrate that children use determiners appropriately as soon as they start putting words together – between the ages of 18 and 28 months (Abu-Akel, Bailey & Thum 2004, in a longitudinal investigation of seventeen 18-month-olds; Ihns & Leonard 1988, in a longitudinal investigation of a 2-year-old; Valian 1986, in a cross-sectional study of six 2-year-olds; Valian, Solt & Stewart 2009, in a cross-sectional study of twenty-one 2-year-olds).

Experimental data show that very young children attend to and understand determiners, using them to aid in noun repetition (Gerken, Landau & Remez 1990, with 2-year-olds) or to pick out a stuffed animal or block (Gelman & Taylor 1984, with 2-year-olds). Eighteen-month-olds and older infants parse a speech stream better if they hear a genuine determiner than a nonsense form or function word from a different class (such as *and*), and, often, better than if they hear no determiner (Gerken & McIntosh 1993, Kedar, Casasola & Lust 2006, Zangl & Fernald 2007). Even though children at 18 months seldom produce determiners, their comprehension is improved when they hear real determiners, indicating that they have a determiner slot which they expect to be filled appropriately. Eleven-month-olds prefer to look at monosyllabic nouns that are preceded by real, rather than nonce, determiners (Hallé, Durand & de Boysson-Bardies 2008); 14- to 16-month-olds listen longer to test passages where a nonsense noun is in a verb context rather than a noun context (Höhle,

Weissenborn, Kiefer, Schulz & Schmitz 2004); 18-month-olds look longer at a visual target if it is described by a sentence with a determiner before the noun than if a different short word precedes the noun (Kedar *et al.* 2006). Thus, there is strong evidence that even infants have the category determiner.

Sceptics have questioned whether 2-year-olds actually have a category determiner, proposing instead that children have lexically specific formulae (Pine, Freudenthal, Krajewski & Gobet 2013, Pine & Lieven 1997, Pine & Martindale 1996), but other work suggests that children are not bound by frames in their use of determiners (Valian *et al.* 2009, Yang 2013). One difficulty in examining the corpora of 2-year-olds is having a sample size that is large enough to fully reveal children's knowledge. With small samples, children can appear much less abstract than they actually are (Valian *et al.* 2009). Children's only error with respect to determiners is their failure to use them in all the contexts where they are required. The reason for those omissions may be prosodic rather than syntactic (Chapter 11, Gerken 1994): if unstressed syllables do not fit a prosodic template for a language, they will tend to be omitted.

Twelve-month-olds exposed to a miniature artificial language are able to use the combination of high-frequency markers yoked with either one- or two-syllable words to form categories (Gómez & Lakusta 2004). Even though the items in the language have no meaning, infants form the categories quickly. Since these categories are not natural language categories, they do not argue for innateness, but they demonstrate that children do not form item-specific representations as their first hypothesis, but more abstract representations.

Precursors to a full understanding of determiners are revealed by experiments with very young infants: 8-month-olds use *the* to segment speech using nonce nouns, but find the nonsense syllable *kuh* equally useful (Shi, Cutler, Werker & Cruickshank 2006). Young infants thus appear initially to have an underspecified representation, accepting a high-frequency vowel whether it appears in *the* or *kuh*. Fourteen-month-olds learning French treat different determiners similarly, suggesting that they have a category and not just individual words (Shi & Melançon 2010; see also Chapters 7 and 8 for more discussion of early uses of segmentation cues).

2.8.3 Abstract representations

Word order and categories are intimately entwined. To get word order right, the child either has to have memorized a very large number of sequences or to have coded those sequences in terms of categories. Children do get word order right, both within a phrase (for example, for English, placing determiners in front of adjectives, and placing determiners and adjectives in front of nouns) and within a sentence (correctly ordering the major elements of a sentence, such as the subject, verb and

object). As with categories, children's spontaneous speech is ordered appropriately as soon as children put words together. A language-specific word order cannot be innate. Rather, the issue is whether children's ordering of word sequences is in terms of abstract syntactic categories, or only as individual words or only in terms of semantic roles.

Sceptics have proposed that 2-year-olds do not understand that English word order is subject-verb-object (Akhtar 1999, Akhtar & Tomasello 1997), based on studies with nonce verbs, in which 2-year-olds do not correct wrong word orders that experimenters use with nonce verbs, although older children do. One difficulty with experiments with 2-year-olds is ensuring that the child can cope with the cognitive demands of the task (Naigles 2002). Even 15 minutes may be too long a time period for a 2-year-old to demonstrate the abstract knowledge that she can demonstrate in 10 minutes (Dittmar, Abbot-Smith, Lieven & Tomasello 2011). Fifteen minutes of concentration on a task apparently can tax the cognitive abilities of 2-year-olds to such an extent that the knowledge that they previously demonstrated is no longer available. Even children younger than 2 are sensitive to word order (Gertner, Fisher & Eisengart 2006). Sixteen-month-olds, for example, listen longer to sequences displaying correct word order than to those with incorrect word order (Shady 1996).

Precursors to word order sensitivity are apparent in infants ranging from 7 to 12 months of age. Seven-month-olds exposed to artificial language sequences quickly acquire order-dependent patterns (Marcus, Vijayan, Bandi Rao & Vishton 1999), and work with miniature artificial languages demonstrates sensitivity to order among 12-month-olds (Gómez & Gerken 1999). Eight-month-olds are sensitive to whether high-frequency items like determiners occur first or last in a phrase: Japanese 8-month-olds preferred to hear highly frequent nonce syllable after low-frequency syllables, while Italian children preferred the reverse pattern (Gervain, Nespor, Mazuka, Horie & Mehler 2008). Such experiments demonstrate that children's first hypotheses are abstract, rather than item-based.

2.8.4 What is innate – inference to the best explanation

The observational and experimental data on two-year-olds' behaviour suggest that, at least as soon as children can string words together, they are operating with abstract syntactic categories and use those categories to establish the basic word order of their language. Experiments with children younger than 2 demonstrate that infants under the age of 1 year form abstract categories and rules rather than lexically specific patterns. Taken together, the data provide specific empirical evidence about innate syntactic content. Language is not just special, it is special in a particular way. The data of course do not prove that there are innate syntactic ideas, since no empirical study can prove such a conclusion. But the data do support the

inference of innate content as the best explanation. That is as good as it gets in the empirical world.

Suggestions for further reading

- Berwick, R. C., Pietroski, P., Yankama, B. & Chomsky, N. (2011). Poverty of the stimulus revisited. *Cognitive Science*, 35, 1207–42.
- Jackendoff, R. (2011). What is the human language faculty?: Two views. *Language*, 87, 586–624.
- Valian, V. (in press). *Input and Innateness: Controversies in Language Acquisition*. Cambridge, MA: MIT Press.