

MAMA BUNNY TRACTOR  
PAPA DOGGY

## What Do Babies Know About Language?

by Fiona Cowie

DA

BA

MA



PA

GA

# WHY

Nativism is the view that there are ideas, beliefs, knowledge, or concepts that are inborn or innate. It's not just the notion that we have innate capacities to acquire knowledge from our experience; instead, it's the idea that some of what we know is already in us to start with. Some very eminent thinkers have held this view. Plato thought that ideas of the Good, the Beautiful, Virtue, and Justice were all innate; René Descartes thought our ideas of God, mathematics, and logic were innate; and Gottfried Leibniz thought that our ideas of necessity and possibility were innate.

Of course, not everyone shares this view. Prominent nonnativists include Aristotle, the Enlightenment philosophers David Hume and John Locke, and the 20th-century psychologist B. F. Skinner. I don't claim to be in the same league as these people, but I, too, am a nonnativist, or empiricist. We have in common the idea that most of what we know is empirical—that is, comes through learning. This could strike many of you as somewhat uncontroversial because, after all, learning is such a ubiquitous feature of our lives. So you might wonder why anyone, especially the eminent nativists I've named above, would deny that learning, at least in some areas, is possible.

Nativists often support their case by an argument that, in general terms, goes as follows. We know about something, X, where X could be God, the truths of mathematics, what virtue is, what goodness is, or many other things. But, it's claimed, there's too little information about X in the environment to enable us to have learned what we know about it. So if our knowledge of X couldn't have been learned, it must have been inborn. After all, there's nowhere else it could have come from! This is called the "poverty of the stimulus argument."

Which brings us to my topic, for the MIT linguist Noam Chomsky uses this argument to reason that linguistic knowledge is innate. We know facts about language, he argues, that we couldn't possibly have heard people say to us, or overheard people saying around us, at the time we were learning to speak. Nor could we have inferred these facts from what we heard around us. These facts could not have been learned, so they must be known innately—we were born knowing them. To give you an idea of how this argument goes, let's look at a particular case, the four sentences shown in the illustrations. You'll be surprised at what you know about them.

Take the simple sentence "John loved him." You know it can't mean that John loved John. It has to mean that John loves somebody else, perhaps René. What about the next sentence, "John loved himself"? You know without even thinking about it that it can't mean that John loved René; it has to mean that John loved John. What about "John thought that he loved him"? You know that it could mean a bunch of things. It could mean that John thought that he, John, loved René. Or it could mean that John thought that René loved John.

It could also mean that John thought René loved some third person, Gottfried. But you also know that it can't mean certain things too: it can't mean that John thought that he, John, loved John, and it can't mean that he was thinking about Gottfried's self-obsession. What about "John thought that he loved himself"? It could mean that John thought that John loved John, or it could mean that John was thinking about who Gottfried's object of affection was, namely Gottfried, but it can't mean a lot of other things. You know this without even thinking about it—you just automatically understand these sentences and can tell what are the possible meanings and what aren't.

The rules of grammar that govern when two terms, like "John" and "he" can refer to the same

1. John loved him.



2. John loved himself.



3. John thought that he loved him.



4. John thought that he loved himself.



**We know what the sentences on the right can and can't mean without having to open a grammar book. Are humans born with this knowledge, or is it learned? Author's son Jacob, left, can't wait to find out.**



object and when they can't are known as binding theory. Here are the principles of binding theory:

- A. Anaphors (like "himself") are bound in their binding domain.
- B. Pronominals (like "he") are free in their binding domain.
- C. R-expressions (expressions, like noun-phrases, that are used to refer to things and events in the world) are free.

(The binding domain of a noun-phrase is the smallest clause that contains the noun-phrase, its case-marker, and a subject. An expression is bound if its reference is the same as the reference of some other expression within the binding domain.)

Got that? I don't really need to explain what it means, do I? Because, according to Chomsky, you already know these principles! You're not conscious that you have this knowledge, of course, and you may not even be able to understand what binding theory, as formulated here, is telling you. But this is the knowledge that apparently underlies your ability to understand the meaning of those sentences about John, René, and Gottfried.

Now it's almost certain that, unless you've studied linguistics, no one ever told you these principles till now. Yet you've been using them since you were a child. How did you acquire this subtle linguistic knowledge? Surely children can't figure out such complicated principles just from listening to what people say around them? In order to do that, they'd surely need to know what sentences containing anaphors and pronominals *can't* mean—for instance, that "he loves himself" cannot mean that he loves some other person. But no one ever tells you what sentences can't mean. So it's mysterious how you could possibly have inferred these difficult principles from the information you had access to as a child. Chomsky argues that you couldn't have and that you didn't. He thus concludes that binding theory must be known without learning—it must be innate knowledge.

Chomsky and his colleagues (like Steven Pinker, a psychologist at Harvard and author of an excellent popular treatment of these issues, *The Language Instinct*) run similar kinds of poverty-of-the-stimulus arguments for the innateness of various other principles of what is called "universal grammar." These are principles of structure and organization held to apply to all natural languages, no matter what their superficial differences of vocabulary and syntax. Chomsky holds that our innate knowledge of universal grammar is embodied in a special language-specific learning device, or module, that evolved only in humans, presumably by natural selection—although he refuses to comment on how exactly this language module developed in our brains. If humans have a specialized language module that embodies their knowledge of universal grammar, then there's no need for them to learn all the deepest and darkest properties of natural language, like binding theory, for they know it already. All that children have to learn as they listen to people and try to talk to them are the superficial features of their language, such as the vocabulary and the rules governing such things as word order or past-tense formation. As a result, language learning is quick, efficient, and easy.

Having given you a review of the reasons why a nativist like Chomsky holds that a large amount of linguistic knowledge is innately known, I'd like to give you some reasons why I don't believe it is. First of all, the poverty-of-the-stimulus argument might seem convincing at first glance, but it doesn't provide any real data showing that children *don't* get adequate linguistic information. Literally none of the Chomskyans' specific claims about what children do not hear are supported by developmental studies, and many of the claims have not withstood careful scrutiny by developmental linguists. And the more general claim that children have very little access to information about language is one of those "facts" that looks plausible or not from different points of view. A child learning language hears



**With five-year-old sister Katie around, Jacob, aged three, sure doesn't suffer much from an impoverished linguistic environment. Katie and Jacob are the author's children.**

about 7,000 utterances a day over the six or seven years that language learning typically takes. Is that a *little* information, as Chomsky maintains, no doubt thinking of the infinitude of other sentences that a natural language contains? Or is it rather a lot, as it seems to me, thinking that people routinely master other infinite areas such as arithmetic, logic, or even cooking, in much less time and with much less input and practice than a child spends on learning a language? But it's not the sheer *bulk* of information coming in that's important, as anyone knows who's had a bad teacher or read a bad textbook, it's whether the information is of a kind that a person *can make use of*: a lot can be learned from very little, given the right preparation.

This brings me to my second quarrel with the argument from the poverty of the stimulus, which is that it is based on a very simplistic concept of learning. It concludes that Mother Nature has prepared us for language learning by building in most of what we end up knowing. But in stating its case for this conclusion, the argument overlooks the other kind of "preparation" with which Nature might have furnished young minds, namely, a more general, non-language-specific suite of learning capacities, abilities that allow children to take information from the environment, organize it, analyze it, and render it in forms that are more useful to them.

Proponents of the argument talk of how little children can "learn" from what they hear, but they don't take account of the fact that learning is not just a matter of what the philosopher Karl Popper referred to as "bold conjecture" and refutation. For instance, their idea that children's hypotheses about language may be constrained by their ability to perform sophisticated inductive and statistical inferences is not followed up. The argument simply *assumes* that children are not good at analyzing large amounts of data, nor at making accurate generalizations going beyond the data they have access to. Yet this is now known to be false: an impressive body of experimental work by psycholo-

gist Jenny Saffran of the University of Wisconsin, Madison, and colleagues has shown that even very young babies take extraordinarily little time to extract high-level regularities from their analyses of the statistical properties of rule-generated inputs, linguistic and otherwise. For example, Saffran showed that after a mere two minutes' exposure to a stream of artificial speech, eight-month-old infants are able to recognize what is and is not a "word" of the artificial language, based solely on the probabilities of certain sounds going together.

The argument also ignores the fact that children are able to use myriad kinds of information to evaluate their hypotheses about how language works: the "linguistic data" that kids have access to is not just a set of sentences, a list of what other people say. Instead, it includes information about meaning and context, information about which things are sometimes said differently (and what these differences imply), and information about what is not said and when and why. It includes information about how children's own linguistic sorties and those of others are received, and whether their or others' demands, requests, and questions are understood and effective. It includes, in other words, information about what chunks of language are *for*, and about what they can *do*—language being *for* communication and able to *do*, well, just about anything, from getting someone to buy you a toy to starting a war.

Here's the point: the poverty of the stimulus argument in effect contends that if you took a child who lacked innate knowledge of universal grammar, locked her in a room for seven years, and made her listen to recordings of around 18 million sentences of, say, English, then she would come out of the room unable to speak or understand that language. Well, maybe so. But what the argument does *not* show is that if you took a child who lacked innate knowledge of grammar, put her in the world, and gave her the vast amounts of information about language and its works and workings

that actual children have access to, she would fail after seven years to have learned her language. By putting an impoverished conception of the data together with an impoverished conception of children's remarkable abilities to learn about their world, Chomsky's argument looks overwhelming. But once you enrich your conception of the child, and of the linguistic data, the argument seems a lot less compelling.

Although nativists about language may have given poor arguments for their view, we can nonetheless test the hypothesis of innate linguistic knowledge in another way—on its merits. How well does linguistic nativism fare? Not well at all, or so I will try to convince you. In order for a scientific theory to be fully validated, it needs to do two things: it needs to explain or account for the data within its area, and it needs to be consistent with other things we know. Linguistic nativism fails on both fronts.

What we want from a theory of how language is learned is—a theory about how language is learned! That is, we want a theory about the psychological mechanisms used in language acquisition, and about the data used by children, that accords with what we know about children's psychology and the data they have access to, and that predicts the actual course of language acquisition.

Nativists say that our innate knowledge of universal grammar, together with a theory of parameter setting (a process in which the—very few—variables in universal grammar are nailed down to a particular value, as when the basic word order within phrases is determined), explains how language is learned, given the paucity of linguistic information and the stupidity of young children. However, nativists have failed almost completely to provide any detailed, testable theories about how actual children go about the task of learning their language. (A notable exception here is Steven Pinker, who developed a theory about verb-learning in the 1980s that initially looked promising but is now widely held to be inadequate.) So nativists' theory of the innateness of a language organ embodying universal grammar has not delivered on its promise of productivity: while it explains how language acquisition might in principle work, it has not even attempted to tell us in any kind of detail how this story is supposed to explain the actual course of language learning. It's as if Newton had rested content with "There's this weird force out there that explains how planets and other things move. Let's call it gravity." But scientific validity, not to mention God, is in the details. So nativism fails the first test: as it stands (indeed, as it's stood for almost 50 years), nativism is not clearly enough articulated to provide an adequate scientific explanation of language acquisition.

Nor is nativism consistent with other things we know. A group of psychologists, including Jeff Elman and the late Elizabeth Bates, both of UC San Diego, Michael Tomasello of the Max Planck

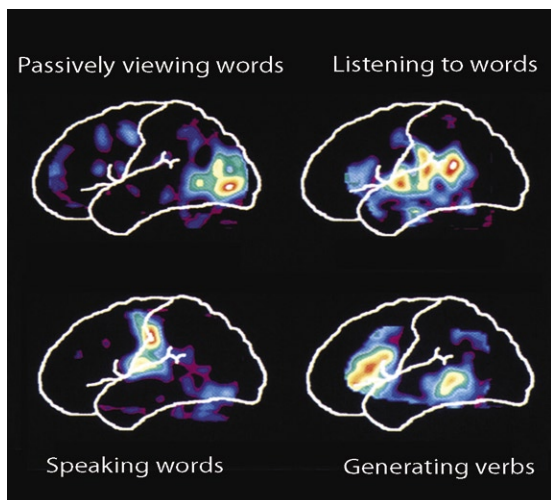

Institute for Evolutionary Anthropology in Leipzig, and Annette Karmiloff-Smith of University College London, have been developing an alternative theory of how language acquisition works, and my current research is aimed at bolstering their case. My aim is to show how this alternate theory coheres better with other areas of the mind sciences, especially developmental psychology and neuroscience, and to bring home the implications of this diverse body of research for the orthodox nativist position.

In this alternative, "constructivist" view of language learning, there is no evolved, specialized language module. Instead, numerous faculties with different evolved functions cooperate to make language learning possible. For instance, children have the capacity to focus their attention on the same thing that somebody else is attending to, and this underlies their earliest attempts at word learning. They can perform extremely sophisticated statistical inferences from data, as we have already seen in the Wisconsin studies, and this accounts for their initial ability to extract words from the incoming stream of "noise" and their progressive understanding of ever-more-general rules about how language works. They also have the capacity to understand other peoples' intentions, particularly their communicative intentions. This again is a critical skill for a language learner: language is, after all, primarily a vehicle for communication. And they have the ability to learn by imitation, an ability that is exhibited in the virtually ceaseless stream of "practice language" that is both the pleasure and despair of the parents of young children. This last may be a peculiarly human ability—it's not clear whether any other animals can learn by imitation, though many researchers believe that if they can, they find it very, very difficult—but it's not an ability that's specific to the task of learning language. On the contrary, it plays a role in many other kinds of learning as well. Finally, children also have the ability to perform what's called "categorical perception," which I'll elaborate on later. The key feature of this alternative view is that all the capacities that are used in language acquisition are also useful for other tasks. There may be innate knowledge and innate capacities that enable us to learn a language, but none of them are *specific* to the task of learning language.


One of the nice things about being at Caltech is you don't have to stay in your own disciplinary pigeonhole, which is just as well, because in order to defend this view of language acquisition I've had to become familiar with, or at least know people who are familiar with, a lot of different things outside philosophy: neuroscience, genetics, psychiatry, developmental psychology and psycholinguistics, historical and comparative linguistics, anthropology, and evolutionary biology.

Let's look at the evidence from neuroscience. If linguistic knowledge were innate, you'd expect it to be expressed somewhere in the brain. Indeed, until

Right: Until recently, it was thought that two areas of the brain's left hemisphere shared language duties, with Broca's area responsible for syntax, the production of sentences, and Wernicke's area handling semantics, the meaning of what's being said. It's true that the two areas, connected by a thick bundle of nerve fibers (pink), work closely together, but they're not the only areas involved. In the PET scans below left, the brain of someone looking at words, listening to words being spoken, speaking words, and turning nouns into verbs, lit up all over the place. And Broca's area also analyzes things other than language syntax. In the experiment below right, Broca's area and its right-hemisphere equivalent lit up (as measured by magnetic field strength) when listeners heard music that unexpectedly hit a dud chord. The brain was probably trying to work out what had gone wrong with the expected harmonic syntax.



From *Images of Mind* by M. J. Posner & M. E. Raichle. ©1994, 1997, Scientific American Library, reprinted by permission of Henry Holt & Co., LLC.



From B. Maess et al., *Nature Neuroscience*, 2001, Vol. 4, 540-545, with permission.

reading, which is a tactile task. This suggests that functional specialization in the cortex is determined less by genetic than by experiential factors. So if there's localization for language in the cortex, it's an open question where that specialization came from.

In any case, it's beginning to appear that there really isn't much localization of language in the brain. New imaging techniques developed in the last few years have revealed that language processing is much more widely distributed than the earlier picture supposed. You can see this in PET scans

late last century, studies of brain lesions and aphasias (a condition where people lose the power to use or understand language) were thought to show that language was relatively localized to Broca's and Wernicke's areas of the left hemisphere. Broca's area was thought to be responsible for syntax (the form of the utterance) and Wernicke's area for semantics (what the utterance means). This apparent localization of language function in the brain appeared to support linguistic nativism: Broca's and Wernicke's areas were plausible candidates for the repositories of our innate linguistic knowledge.

However, it's not actually clear that functional localization tells us very much about whether or not the function is *innate*. As Elman and Bates, among others, have argued, functional localization can occur from virtually any developmental trajectory—learning, genetic determination, and everything in between. We know, for example, that the brain has a lot of plasticity, and can adapt to changed circumstances. The congenitally deaf use their auditory areas for the processing of sign language, which is a visual task, and the congenitally blind use their visual cortex for Braille


of someone passively viewing words, listening to words, speaking words, and generating verbs from nouns. When I look at this kind of scan, I think to myself, where *is* the language module? It seems to be everywhere! Some nativists have responded to these kinds of brain imaging data by saying it doesn't matter if there are lots of language areas in the brain; the important thing is that *some* areas of the brain are destined to encode the specialized linguistic knowledge that our genes represent. The genes can put linguistic knowledge in the brain wherever they like, so long as they *do*.

The idea that it is *language-specific* information that the genes encode in the brain is brought into question by the fact that areas once thought to be specialized for linguistic tasks, such as Broca's area, can also perform tasks other than the processing of linguistic syntax, as shown by a recent study in which this area lit up on an MEG (magnetoencephalography) scan while people were listening to harmonious and disharmonious music. It even lit up in one place when harmonious music was played, and in another place when disharmonious music was played. What Broca's area seems to be



**Top: Part of the first “language” gene to be identified, transcription factor *FOXP2*, with the mutation that causes a severe speech and language disorder colored red. *FOXP2* is a gene orchestrating the development of brain circuitry for the precise coordination of movement in mammals.**

**When it’s faulty, humans lose the ability to accurately control the muscles used in speech. The image is from Dr. Simon Fisher of the University of Oxford, who was part of a team that identified this gene with the help of three generations of the KE family, whose pedigree diagram is shown below. Family members with the inherited disorder are shown by the red squares (males) and circles (females).**



doing is processing not just linguistic form, but also musical form. This kind of functional overlap, like the fact that language processing seems to be “smeared out” over much of the brain, suggests that the processes responsible for language are not *specific* to language. The evidence from neuroscience seems to support an empiricist rather than a nativist view.

What about the nativist counterargument that it doesn’t matter *how* language is implemented in the brain; what matters is that linguistically specific knowledge encoded in the genes is expressed during language acquisition? First, it’s not clear how knowledge of universal grammar could actually be “encoded” in the genes. For one thing, as Bates pointed out, half facetiously, there may not be enough of them! Recent estimates give us around 20–25,000 genes, which have a lot more to do beside encode for universal grammar. In addition, many noted biologists and philosophers, including Richard Dawkins of Oxford University and Peter Godfrey-Smith, of both the Australian National University and Harvard, argue that although genes can be said to code for proteins and transcription factors, they do not in any real sense “encode” higher-level traits like knowledge of universal grammar at all (though they are certainly involved in producing them). More damagingly, nativists have never given even the barest hint as to *how* linguistic knowledge (or any other knowledge, for that matter) might be genetically encoded. What, exactly, are the processes, genetic and otherwise, by which this genetically coded information gets expressed?

Worse still, recent attempts to locate genes specialized for language have resulted in the discovery of genes whose functions are *non-linguistic*. Let me give you an example of this. In England there’s a family called the KE family who have an inherited language disorder. As you can see in the pedigree, about half the people in the family have what’s called Specific Language Impairment in quite a severe form. They have deficits in the production of various grammatical morphemes like the “s” at the end of a plural, the “ed” at the end of a past tense, the “ing”—all those niggly little bits of language that carry certain kinds of grammatical and semantic information. In 1991, to great fanfare, the Canadian linguist Myrna Gopnik suggested that the gene responsible for the family’s language problems was a “grammar gene” encoding grammatical morphology, based on the argument that the disorder showed a Mendelian inheritance pattern corresponding to a single dominant gene, that a fault in this gene resulted in grammatical deficits, and that it must therefore be a gene for grammar. The faulty gene was recently identified as *FOXP2* on chromosome 7. But it’s not obvious that *FOXP2* can be called a gene for grammar. For one thing, other animals also have this gene, yet we’re the only species (as far as we know) that uses language. Other species communicate symbolically, to be sure, but it’s generally thought that

because their symbols cannot be recombined to express different thoughts—compare “The dog bit the man” with “The man bit the dog”—their communication systems are not languages proper. For another, this gene seems to play a role in motor development, rather than linguistic development *per se*. In the rat, it encodes a transcription factor implicated in the normal development of the corpus striatum, a part of the brain involved in the planning and sequencing of motor behaviors. This supports an alternative view of what is wrong with the KE family. It’s not that they lack a grammar gene, it’s rather that they have an articulatory problem in moving the mouth, lips, and tongue so as to form certain language sounds. On this alternative view, it is this articulatory problem that in the first instance hinders the affected family members from learning some of the relevant grammatical rules (lack of practice makes imperfect) and in the second instance prevents them from expressing what linguistic knowledge they do have. The gene isn’t language-specific, and it isn’t even species-specific, so there’s no support for the innateness of language here. On the contrary, the fact that a gene concerned with motor development is closely implicated in a disorder of language supports the empiricist view that lots of different abilities have come together to enable language learning.

To defend the alternative argument that children learn language from the information they hear around them rather than having large chunks of it built in, we also have to explain why human languages across the world are so similar to one another. Proponents of the innateness hypothesis have argued that all languages—described at a suitable level of abstraction, anyway—are the same. They reason that this is because all people have a universal grammar embedded in their heads and, of course, all languages conform to this universal grammar. The features that are common to all the world’s languages, the linguistic universals, are somewhat controversial, and if you had five linguists in a room, they wouldn’t reach any agreement about what they are. But one relatively uncontroversial feature common to all, or nearly all, languages is the syntactic distinction between nounlike words and verblike words. Most languages treat nounlike words—words that refer to things—differently from the way they treat verblike words—words that refer to actions, processes, or states. Nativists claim that these similarities across languages arise because universal grammar is known innately. But there are other explanations.

For example, some broad similarities among languages are almost certainly due to universal features of the communication situation. We use language to communicate, so precious necessities for communication are going to shape everybody’s language. Indeed, in 1921, the linguist and anthropologist Edward Sapir proposed that the noun-verb distinction arose because to communicate, you need a way of picking out something as



**Was language invented only once, in Africa, about 125,000 years ago by a group of early *Homo sapiens* that eventually populated much of Asia and Europe? The 90,000-year-old skull shown here was found in the Skhul cave near Mount Carmel, Israel.**

the topic of your utterance (e.g., the bee), and then a way of saying something about it (stings). So, of course, all languages are going to develop ways to do those things.

Other common features of language are probably due to nonlinguistic features of human cognition, such as processing or attentional constraints. Most people don’t use sentences that are 15,000 words long, and it’s not because of anything deep, it’s because peoples’ memories and attention spans just don’t last that long.

Some features of language may just be historical accidents, like driving on the right. There’s nothing inherently correct about driving on the right side of the road as opposed to the left, but someone, somewhere, just decided that was how we were going to do it, and we all conformed (except the British, some of their former colonies, and the Japanese, who are still holding out), because it was easier to do so than to effect a change to the other side of the road.

The same could be true of some language universals, particularly those that seem inexplicable in terms of communicative necessities or general features of our brains and minds. What matters for communication is not so much *what* rules we all follow, but that we all follow the *same* rules. Seemingly arcane or strange rules might thus be adopted, and might subsequently persist, because changing our linguistic conventions would lead to communicative breakdown. If certain rules became fixed in a common ancestor language, and if changing those rules was more bother than it was worth, it could explain why all languages spoken today share certain features. Is there evidence for such a common ancestor language? It used to be thought not, but recent developments in historical linguistics, archaeology, and genetics suggest that all human languages are descendants of the language spoken by a group of people coming out of Africa about 125,000 years ago. Arbitrary convention plus common descent, rather than constraints

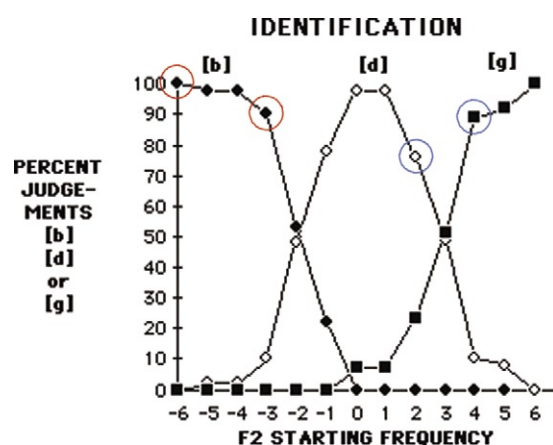


imposed by an innately known universal grammar, can explain linguistic universals.

The crux of the issue between nativists and their opponents is this: Are the processes by which we learn language specific only to learning language, or not? The nativist says Yes: after all, innate knowledge of universal grammar would be useful for learning language, but not for much else. The empiricist or constructivist says No: there's innate stuff involved in language acquisition, of course, but that stuff is used for other learning tasks as well.

As a kind of test case, let's look at phonological learning, which has for many years been touted as a convincing defense for nativism. Phonemes are the smallest linguistic units relevant to meaning.

I. G. Mattingly, *American Scientist*, 1972, 60, 327-337, with permission.



The nativist position is undermined, however, when you look at the mechanism by which phoneme perception occurs. Then you find that it's initially inborn but shaped by learning, that it's not language specific, and that it's not even specific to our species. Our brains distinguish phonemes by a mechanism called categorical perception. The brain takes the continuous speech stream, which is a continuously varying acoustical signal—a bunch of noise, basically—and segments it into chunks that map onto the phonemes of our language. It's a very complicated process, as you can see by looking at the graph of an artificially engineered acoustical signal that varies continuously along one dimension, and what people's response to that sound is. When the starting frequency is -6,

Listening to a series of computer-generated sounds in which the frequency at the start of each sound changed over a continuum, the subject in the experiment at left heard only three syllables: ba, da, or ga. And although the difference in frequency between the red-circled markers was more than that between the blue-circled ones, both reds were heard as ba, while one blue was judged to be da, and the other ga. The brain "chunks" these sounds into familiar categories and doesn't hear the nuances in between. This is also true of newborn babies, like 10-day-old Ella, below, and even of the chinchilla on the facing page.

In English, they are sounds like be, ke, pe, te, and ah (which are often written as /b/, /k/, /p/, /t/, and /a/). According to nativists, all phonemes for all possible languages are represented in our brains at birth, and all that our experience does during language learning is to prune away the phonemes we don't need for the particular language we're learning. If we were learning Japanese, for example, the distinction between the English /l/ and /r/ sounds would be pruned away. There is some support for this account. Phoneme perception begins in the womb, and newborns prefer the sound of their mother's voice and the sound of their parents' language minutes after birth. Infants aged between one and six months can reliably discriminate many different natural-language phonemes, even ones not occurring in the language being spoken around them, but after 12 months they have lost that ability. So although a Japanese six-month-old can discriminate /l/ and /r/ sounds, a one-year-old can't. It does look as if we're all born with innate representations of these sounds and they wither away if we're not using them.



-5, -4, and -3, people judge it's the sound of the letter /b/, then when it gets up to 0, they judge it to be /d/, and at +4, +5, and +6, they think they're hearing /g/. Moreover, the two sounds that have red circles differ from one another physically much more than the two markers encircled in blue, yet both red sounds are judged to be the sound /b/, whereas one blue sound is judged to be /d/ and the other is judged to be /g/. Such responses are characteristic of categorical perception, in which some things that are physically or acoustically different are counted as being the same, while other things that differ physically by exactly that same amount are counted as being different.



**The brain also chunks faces into familiar categories, as when George W. morphs into Arnie.**

The ability to perform categorical perception is inborn, but it's not language specific. The "chunking" of continuously varying stimuli into discrete categories is a general feature of human perception, and we do it with meaningless sounds such as cheeps, chirps, and bleats, we do it with musical sounds such as those from a violin, and we do it with faces. For example, if a digitized picture of George W. Bush's face is "morphed" gradually into one of Arnold Schwarzenegger, there will come an abrupt point when people change their response from "It's George" to "It's Arnold." No in-betweens.

It's also not specific to our species; crickets, birds, chinchillas, and other animals all "chunk" their acoustical input. Indeed, chinchillas respond to human speech by chunking it into /b/, /t/, and /d/ in exactly the same way newborn babies do. So even the case of phonological knowledge, in which innate abilities do figure largely, does not support a nativist picture of language acquisition. Instead, it supports an alternative picture whereby our linguistic abilities are cobbled together out of preexisting and nonlinguistically specific mechanisms.

The same is very likely true of our other linguistic capabilities. It's unlikely that there's a highly specialized language-acquisition mechanism and much more likely, I think, that language acquisition draws on mechanisms of far more ancient lineage such as the ability to "chunk" incoming perceptual signals into larger units, the ability to recognize statistical regularities among these signals and generalize from them, the ability to deploy attention to important tasks, the ability to share attention with others of the same species, and (of more recent origin) the ability to figure out what other people are thinking, to learn by imitation, and to use tools (like language) as a means of manipulating the world.

To be sure, these abilities would have been honed by the positive selection pressure that came into play as soon as language got up and running,

because language is so useful that any trait that enhanced the ability to learn it would have been massively selected for. But it's unlikely that natural selection created a radically new language organ embodying knowledge of universal grammar. Which is just as well, since, as I've argued here, there's not much reason to think we'd need one. □

*A native of Sydney, Australia, associate professor of philosophy Fiona Cowie gained a BA at the University of Sydney in 1988, followed by an MA and PhD in philosophy at Princeton in 1992 and 1994, respectively. She came to Caltech as an instructor in 1992, became an assistant professor in 1993 and an associate professor in 1998. Her book What's Within? Nativism Reconsidered gained her the 1999 Gustave O. Arlt Award in the Humanities from the Council for Graduate Studies. She is presently working on another book with James Woodward, the Koepfli Professor of the Humanities, to be entitled Naturalizing Human Nature: Beyond Evolutionary Psychology. This article is adapted from a talk given on Seminar Day in May.*



PICTURE CREDITS:  
12, 13, 21 – Doug Cummings; 12 – Fiona Cowie; 13 – Caltech Archives; 14, 15 – Herb Shoebriidge; 19 – Smithsonian; 20 – Kate Quirk; 21 – Vienna Veterinary Dept.