Published version: *Synthese* Vol. 111, pp. 17–51, 1997.

FIONA COWIE

THE LOGICAL PROBLEM OF LANGUAGE ACQUISITION*

ABSTRACT. Arguments from the 'Logical Problem of Language Acquisition' suggest that since linguistic experience provides few negative data that would falsify overgeneral grammatical hypotheses, innate knowledge of the principles of Universal Grammar must constrain learners' hypothesis formulation. Although this argument indicates a need for domain-specific constraints, it does not support their innateness. Learning from mostly positive data proceeds unproblematically in virtually all domains. Since not every domain can plausibly be accorded its own special faculty, the probative value of the argument in the linguistic case is dubious. In ignoring the holistic and probablistic nature of theory construction, the argument underestimates the extent to which positive data can supply negative evidence and hence overestimates the intractability of language learning in the absence of a dedicated faculty. While nativism about language remains compelling, the alleged 'Logical Problem' contributes nothing to its plausibility and the emphasis on the Problem in the recent acquisition literature has been a mistake.

1. NATIVISM, EMPIRICISM AND THE LOGICAL PROBLEM

Nativism in the domain of language is the conjunction of the following theses:

(DS) Learning a language requires that the learner antecedently possess task- or domain-specific information about language.

(I) The information required by (DS) is innate.

(DS) asserts that in order to learn her native tongue, the child must already know something about the properties of natural languages. (I) contends that this knowledge is not learned, being encoded rather in the original fabric of the mind. Together, (DS) and (I) amount to the claim, made famous by Chomsky, that there exists a specially-built psychological mechanism for language-learning: since the information necessary for that task is both domain-specific and inborn, the human mind may be regarded as incorporating a distinct 'language faculty'.

While (DS) and (I) are rarely formulated in the linguistics literature as baldly as I have put them here, both are omnipresent. (DS) is typically voiced as a claim about the contents of the 'Universal Grammar' (UG).

UG describes the structural properties common to all natural languages. According to the nativist, this kind of information is precisely what children must have access to if they are to succeed in learning language. (I) is spelled out in linguists' speculations as to how UG is implemented in the 'language acquisition device'. Formerly (e.g., Chomsky 1975) held to be represented implicitly, as inborn limitations on and orderings of the range of grammatical hypotheses open to the learner, UG is now viewed by many as being represented explicitly in the biologically-determined organisation of the mind, experience serving only to 'fine-tune' or 'parametrize' an already highly-articulated mental organ.[1]

Not only are (DS) and (I) in fact linked in the nativist literature, both are necessary to demarcate the nativist's position from that of an empiricist. (DS) by itself is consistent with empiricism, for an empiricist may agree that knowledge (or at least informed speculation) about the nature of language is involved in language learning, yet claim that the requisite information is acquired piecemeal through experience. And (I) alone is also consistent with an empiricist approach. If the information needed for language learning were of a piece with the general principles of association or induction held by empiricists to underlie all learning, then, since no empiricist denies that we know innately how to learn, none will deny that we know innately how to learn language.[2] It is thus in *conjoining* (I) and (DS) that the nativist distinguishes himself from the empiricist. Hence it is this conjunction that needs to be defended.

In this paper, I examine an argument for nativism that has recently gained considerable currency in the literature. It derives from what is called variously the 'Projection Problem' or the 'Negative Evidence Problem' or the 'Logical Problem of Language Acquisition'. The Logical Problem is alleged to arise because there is an almost total dearth of 'negative evidence' – evidence as to which strings of words are *not* sentences of the language being learned – in the 'primary linguistic data', that is, in the sample of sentences to which children have access during learning. As a consequence, certain sorts of plausible but false grammatical hypotheses – in particular overgeneral hypotheses – would appear to be experientially incorrigible. So, for example, suppose that on hearing the sentences *It is likely that John will leave, John is likely to leave* and *It is possible that John will leave*, the learner predicted that *John is possible to leave* is a sentence of English. Then, absent some concrete reason to think that this last is *not* a sentence – absent negative evidence – she should continue in her erroneous belief that the string is grammatical. The mystery is why, given that few if any speakers have been told or otherwise informed that

John is possible to leave is not a sentence, competent speakers of English universally judge that construction to be ungrammatical.

The challenge posed by the Logical Problem is thus to explain why people do not make these sorts of errors, given that there is nothing in their experience to prevent their doing so. How could language acquisition occur successfully in the face of such massive evidential deprivation? The nativist's advocacy of (DS) and (I) is a response to this conundrum. By supposing that much of the information about language we possess is supplied by our genes, and not by the world, one can see how acquisition could occur under such conditions. Insofar as it "takes the burden of explaining learning out of the environment and puts it back in the child" (Pinker 1989, 15), the nativist's conjoining of (DS) and (I) shows how in principle to account for acquisition in the near-total absence of relevant linguistic inputs: it is learners' inborn knowledge of language that prevents them from making hypotheses the truth or falsity of which cannot be ascertained from their experience.

In its most extreme form, the nativist view would seek to rule out *all* errors of overgeneralization: the learning mechanism is so tightly constrained by the principles of UG that learners literally cannot make hypotheses that are too broad. As Pinker (1989, 18–30) has argued, however, this draconian position, which he calls the 'conservatism hypothesis', is untenable: children do, after all, overgeneralize during language learning.[3] Thus, the nativist's aim must, strictly speaking, be to develop a theory of UG such that its principles, incorporated in the language faculty, are strong enough to prevent most errors, yet not so strong as to rule out those that we know to occur. Nonetheless, the extreme position remains for nativists a kind of ideal in responding to the Logical Problem and since nothing I will say in the following will turn on this issue, I too will sometimes speak as if the nativist sought to prevent overgeneration altogether. The reader should, however, keep the qualification just mentioned in mind.

In this paper I will argue that, contrary to what appears to be the received opinion, reflection on the Logical Problem does little to motivate nativism. On the view I shall defend, this argument does support (DS): it does suggest that a learner's theorizing about language must be constrained by certain prior views about the task domain. However, I contend, it provides no support whatsoever for (I), the claim that those prior views are innately encoded. In Section 2, I give the argument as it is typically presented in the literature. In Section 3, I present some obvious empiricist challenges to the argument as formulated in Section 2, and reconstruct it so as to circumvent these objections. This version of the argument, which I take to be the strongest, does make (DS) seem plausible. As I argue in Section 4, however,

it does not legitimate (I). Hence, it gives no special support to nativism over empiricism about language. In Section 5, I briefly discuss the rhetorical role of the Logical Problem in discussions of language acquisition and explore some of the implications my arguments have for research in this area.

2. THE LOGICAL PROBLEM

The argument from the Logical Problem employs the nativist's favourite weapon against empiricism: the poverty of the stimulus. In general, poverty of the stimulus arguments point to a 'gap' between the rich stock of knowledge about language that competent speakers possess and the meagre supply of linguistic information that their experience during learning provides. This gap is so wide, the nativist asserts, that no child endowed only with an empiricist-style, general-purpose learning mechanism could hope to bridge it. What fills the chasm are (DS) and (I): the acquisition of linguistic competence requires that knowledge of linguistic universals be embodied in the innate configuration of the learning mechanism.

The argument from the Logical Problem seeks to locate the above-mentioned gap in the logical structure of the learning task.[4] Take some stage in the learning process at which the child has acquired a theory as to the nature of her language – a grammar, in other words. Call the language generated by her current grammatical hypothesis **H**. Call the language generated by the correct grammar – the target language – **L**. Then the learner is in one of five possible situations (Figure 1).

Situation (v) represents the endpoint of the learning process: the child's grammar generates all and only the sentences of the target language. The question is how, supposing she is in any of the other situations, she can use the primary linguistic data to get there. In cases (i), (ii) and (iii), the learner will be forced to revise her hypothesis if she hears a sentence in **L** that is not a sentence of **H** – one that is not generated by her grammar. That is, she can move towards the correct grammar by exposure to 'positive evidence', evidence (as is provided by someone's uttering a sentence) that a given string is a sentence of the target language. Crucially, however, if she is in the situation represented by (iv) – if her grammar generates all of **L** and also some strings that are not in **L** – she will never be able to discover that she errs just by exposure to further sentences of the target language. For, every sentence of **L** is equally a sentence of **H**. What she needs is the information that a string that is generated by her grammar – a string in **H** – is *not* a sentence of **L**: only then can she 'shrink' her hypothesis and bring it back in line with that of the linguistic community.

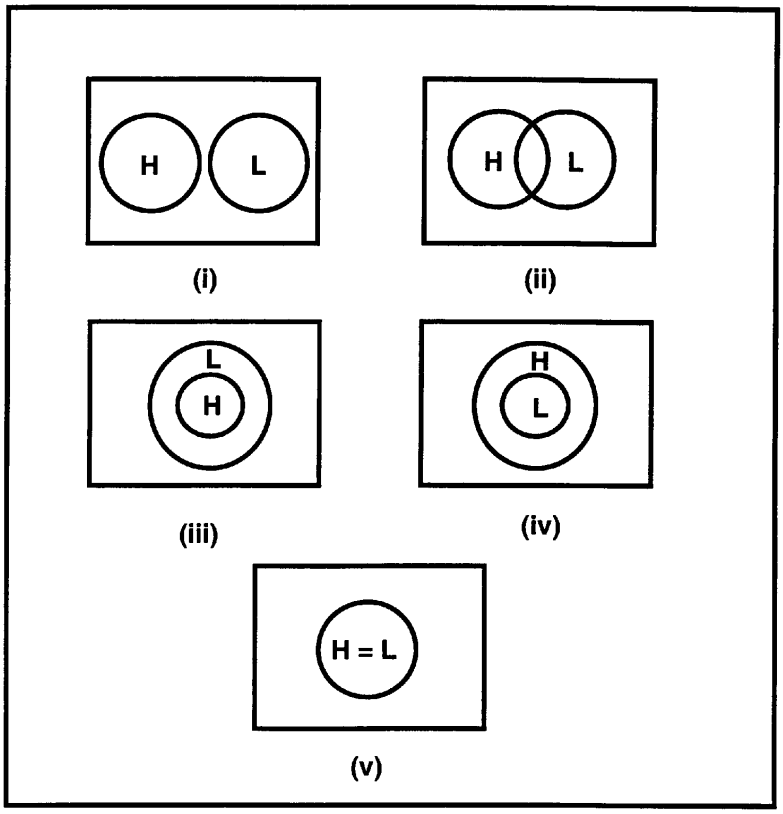

Figure 1.

This latter kind of evidence, that a string is not in the language, is called 'negative evidence'. And the Logical Problem arises because children have no systematic access to it: information as to what strings are not sentences of the language they are learning is in general unavailable.

First, they are certainly not given lists of ungrammatical strings! And they are rarely corrected when they make syntactic errors.[5] Worse, the primary linguistic data contain a proportion of ungrammatical utterances that are not flagged as such.[6] Worst of all, however, is the fact that there are infinitely many strings of (say) English words that the child will *never* encounter in the data. Some of them she will not hear because they are not, in fact, English sentences; others, however, are perfectly good sentences that are absent from the data for the simple reason that no-one has gotten around to uttering them. What this means is that the mere non-occurrence of a string in the data cannot by itself constitute negative evidence. So even if she were always corrected when she made a mistake; and even if

her interlocutors invariably spoke impeccably, there will always remain an infinity of strings, some of them English and some of them gibberish, that the child has no information one way or the other about.

Hence, if a child's hypothesized grammar were to overgenerate, there would be no data that would force her back towards the correct theory. So, since children do eventually converge on the right theory, they must never be in situation (iv): they must never hypothesize a grammar that is 'too large'. Indeed, since cases (i) and (ii) will eventually collapse into case (iv),[7] children must never project *any* hypothesis implying statements of the form 'String S is a sentence of **L**' where S is not in fact a sentence of **L**. In other words, their theorizing must be constrained such that they are only ever in situation (iii). The Logical Problem of Language Acquisition is the problem of explaining how this desideratum is achieved.

Nativists argue that (DS) and (I) are the solution to the Logical Problem. First, let us take the hypothesis that domain-specific information is required for language-learning. The argument is that it is only by invoking her foreknowledge about language that the learner is prevented from making wild and incorrigible overgeneralizations. The empiricist's requirement that the learner prefer simpler, more general, more elegant (etc.) hypotheses provides little in the way of guidance: too little, surely, to prevent a learner's falling irretrievably into the myriad possible pitfalls revealed by the Logical Problem. Language is so strange, its workings so abstruse, that a learner utterly uninformed as to its character surely *must* go wrong. The Logical Problem thus illuminates, in a particularly vivid and intuitively compelling way, the need for a task-specific helping-hand in the linguistic domain.

What of (I), the contention that the requisite information about language is not only known, but known innately? Pinker (1984, 33–37) canvasses an argument to the effect that (I) should be accepted on grounds of parsimony. While recognizing that considerations of simplicity could in principle go either way – depending on the details, a single, general-purpose learning mechanism might well turn out to be simpler than a bevy of task-specific ones – Pinker claims that "a conclusion of task-specificity of the learning mechanisms quickly leads one to a tentative conclusion of nativism because of certain parsimony considerations" (1984, 33). He argues (1984, 34–35) that although the nativist hypothesis that the principles specified in (DS) are known innately is not more parsimonious that *any possible* rival hypothesis, it *is* more parsimonious than any *plausible* rivals, and is therefore to be preferred, pending further empirical findings (1984, 36 and note 10). It seems to me, however, that it is a mistake to propose, even tentatively as does Pinker, that "the innateness debate in the study of language acquisition be recast as a parsimony debate concerning rival acquisition theories for

language and other cognitive domains" (Pinker 1984, 36). First, neither the nativist's theory nor its rivals would seem to be sufficiently well-articulated at this time to be reliably evaluated in terms of parsimony. Second, it is far from obvious that simplicity is in fact the correct criterion of theory-choice in the cognitive domain. For it is by no means clear that the brain, an evolved organ cobbled together over millenia by Mother Nature, operates according to principles that we find maximally parsimonious.

There is, however, a second argument for (I) available. Suppose that domain-specific information, such as is specified in UG, is required for language learning to occur. Then, we have two choices. Either the principles of UG are learned; or they are innate. If they are learned, then the Logical Problem simply arises all over again: the child attempting to infer the Universal Grammar from the data faces exactly the same difficulties with regard to the paucity of available evidence and the danger of overgeneralizing as does the child attempting to infer the grammar of a particular language. Claiming that UG is learned simply pushes the Logical Problem back a step. At the risk of an infinite regress, then, we must conclude that UG is innate.

So, it seems, the argument from the Logical Problem mandates a nativist approach to language acquisition. Further, it mandates nativism while making comparatively minimal empirical assumptions: (i) that speakers learn by some form of hypothesis-testing; (ii) that what they learn is some kind of generative grammar; and (iii) that the primary data consist overwhelmingly of positive data. As a consequence, the present argument is not hostage to the vagaries of fact and theory to quite the same extent as are other kinds of poverty of the stimulus argument.[8] And, while its assumptions may be – indeed have been – challenged,[9] that is not my strategy here. Although I will raise some questions of detail in Section 4.3, I will argue that even if we *accept* the basic picture of learning and competence that the argument from the Logical Problem assumes, it *still* fails to support nativism about language learning.

As a preliminary to my argument, it is necessary to clarify and reformulate the nativist's case so that it resists some obvious empiricist objections. That is the task of the next section. Armed with what I take to be the strongest statement of the Logical Problem of Language Acquisition, I shall then turn, in Section 4, to a consideration of its merits.

3. GUARANTEEING LEARNABILITY

An empiricist will immediately counter the argument presented above with the following. First, in defending (DS), the nativist asserts that no general-

purpose learning device could sufficiently constrain the child in her choice of hypotheses. He has not, however, established that there is *no possible* empiricist model that could explain how a child could learn a language merely through exposure to a sample of sentences.[10] Second, the empiricist will note, the arguments for both (DS) and (I) proceed as if there were *no* negative data available to the learner, when, in fact, there are sources of both direct and indirect negative evidence in the primary linguistic data. Direct negative evidence – explicit information to the effect that string S is not a sentence of the target language – may be comparatively rare, but it nonetheless exists. And indirect negative evidence is surely even more plentiful, being provided, perhaps, by a parent's failure to understand an utterance of the child,[11] or by the repeated failure of certain forms predicted by her grammar to appear in the data. Thus, the nativist's motivation for claiming that overgeneration must be prevented by the incorporation of domain-specific information into the learning device is inadequate. And so too is his attempt to iterate the argument from the Logical Problem in establishing that that incorporation is achieved by nature, rather than by learning. The argument from the Logical Problem is inconclusive. It fails to establish either (DS) or (I).

The argument, however, can be restated so that it represents a formidable challenge to the empiricist. The reformulated argument turns on the need of a learning theory to *guarantee language learnability*. Normal children are virtually certain to master a language given even minimal exposure to it. An acquisition theory must account for this fact. It must explain how learners are equipped such that, *if* they go wrong, they *only* go wrong in ways that turn out to be corrigible on the basis of the evidence available to them. What the Logical Problem underscores is that explaining this is no easy matter. The linguistic evidence available to learners is radically impoverished. Negative data rarely occur. When they do, they occur haphazardly. How, then, can the child be guaranteed only to make errors that are corrigible given precisely the data she happens to stumble upon? And not only do we need to account for the individual case. All children in (say) an English-speaking community learn English despite their being exposed during learning to arguably very different samples of English sentences. How do we guarantee that everyone gets exactly the evidence he or she needs to arrive at the same grammatical hypothesis as everyone else?[12]

Clearly, a learner who *never* overgenerates would be unfazed by the scarcity and sketchiness of negative evidence in the primary linguistic data. The theorist, therefore, should seek to construct a model that approximates, as closely as is consistent with the developmental facts, this conservative ideal. In this endeavor, the nativist is at liberty to impose domain-specific

constraints on the hypothesis-generating mechanism. By adopting (DS) – by supposing, as Demopoulos puts it, "that the range of the learning function is constrained by the learner's knowledge of what constitutes a possible language" (1989, 80) – he may hope to ensure that learners do not project hypotheses requiring negative evidence for their disconfirmation, thus rendering the learnability problem in the individual case tractable. And, by adopting (I) as well – by supposing that the task-specific information required to guarantee learnability is available as a consequence of human biology – the nativist can explain also how it is that everyone learns the grammar of their language.

The empiricist, by contrast, appears at a loss. His view is that the mind, ab initio, is equipped only with certain domain-neutral injunctions concerning learning from experience: "Look for similarities", "Prefer simpler hypotheses to more complex ones", perhaps "Seek out underlying regularities" and so on. But, the nativist argues first, there is no reason to think that these sorts of principles, even if coupled with some 'mechanism of inductive inference' (whatever that is), could constrain the learner so that she *never* needed a negative datum to evaluate a grammatical hypothesis. On the contrary there is every reason to believe that a child employing only general-purpose principles of theory construction will frequently overgeneralize. Second, and supposing that the empiricist's learner does overgeneralize, there is no reason to expect that she will chance upon just the negative evidence she needs to discover her mistake. Unlike the linguist, who already *knows* the language she is theorizing about and who can, therefore, *construct* negative data relevant to her hypotheses, the child learning language must take her negative data when and as she finds them. Since the empiricist can provide no guarantee that the child will find just the data she needs, when she needs them, he can provide no guarantee of language learnability.

Thus the point of the argument from the Logical Problem, on this reconstruction, is to underscore the difficulty of the empiricist's task: he must explain how it is that learners manage only to err in ways that turn out to be rectifiable on the basis of the paltry and haphazard data that they have access to. Certainly no empiricist model proffered to date has managed to do this; and there are no grounds, the empiricist's protestations notwithstanding, to expect any future attempt to succeed either. Of course one cannot rule out the *logical possibility* that empiricism might be true. But, as Hume has put it:

> 'Tis impossible to refute a system, which has never yet been explain'd. In such a manner of fighting in the dark, a man loses his blows in the air, and often places them where the enemy is not present.[13]

So unless and until the empiricist delivers a theory worth fighting against, the argument from the Logical Problem indicates that we should pursue the nativist paradigm.

4. THE CASE FOR (DS) – AND THE CASE *AGAINST* (I)

I think that we should concede that the argument of the previous section establishes (DS). Although the nativist's reasoning is not apodeictic, the need to guarantee learnability of language from the primary linguistic data strongly suggests that domain-specific information is implicated in the formulation of linguistic hypotheses. Without the guidance it provides, there are just too many ways to go wrong about the grammar of one's language and too few data around to guarantee that one's mistakes will be rectifiable. The fact that the adoption of (DS) offers at least the hope that learnability can be ensured, where no such hope existed before, is a powerful pragmatic argument in its favor. Thus while some influential empiricists (Putnam 1971, for example) have expended considerable energy in contesting (DS), those energies have, in my view, been misplaced. Not only is there the present pragmatic argument to consider, there are also broader reasons to think that an empiricist should be neither dismayed nor surprised by the inclusion of (DS) in our theory of language learning. Dismay is unwarranted since, as I argued above (Section 1), (DS) is perfectly consistent with an empiricist approach to acquisition. So long as the domain-specific information needed for the acquisition of grammatical knowledge is itself acquired through experience, the empiricist commitment to the primacy of our experience over our genes is unassailed. And surprise at the likely truth of (DS) is misplaced too. For it is something of a commonplace in the philosophy of science that induction cannot work in the absence of a prior theoretical framework.[14] That the nativist's Logical Problem extends this 'problem of induction' beyond the boundaries of scientific inference and into the realm of individual psychology is interesting, but hardly world-shattering.

Thus, the anti-empiricist (or pro-nativist) force of the argument from the Logical Problem must come, if it comes at all, from its power to establish (I), the contention that the information specified in (DS) is innate. In Sections 2.2 and 3 above, I outlined two arguments from the Logical Problem to this conclusion. In this section I show that neither is successful. The first fails on empirical grounds; the second for more conceptual or 'philosophical' reasons.

4.1. *Why the Learnability of UG is a Red Herring*

The first of the nativist's arguments for (I) was as follows. The information specified in (DS) could not be gleaned from experience as the empiricist claims. For learning the highly abstract and unintuitive principles of UG from the data is an even more difficult and error-prone enterprise than that of learning the grammar for a particular language. Thus, the contention that the child learns UG simply raises the Logical Problem all over again. UG must be innate. This argument contains a suppressed premiss, namely, that the information posited in (DS) as being necessary to language learning *is the information specified in the Universal Grammar*. On the view many linguists subscribe to, this premiss is something of a conceptual truth: to theorize about UG *just is* to theorize about the information needed for language-learning.[15] But on the view I shall defend in this section, the claim that UG is implicated in learning is an empirical claim in need of empirical defense. Furthermore, I shall suggest, it's an empirical claim that in fact *lacks* any sound empirical basis. There is no good reason to believe that the knowledge posited in (DS) is knowledge of UG. Hence, I will contend, the learnability or otherwise of UG is inconclusive with respect to establishing (I).

UG describes syntactic properties shared by all natural languages. As currently understood, these include such features as the following: sentences conform to the structural constraints of X′ theory; transformations are limited in their domain of application by the principle of Subjacency; Case is always assigned to nounphrases; anaphors are bound, and pronominals are free, in their Governing Category, and so on. The similarities among languages that linguists have discovered and systematized under the umbrella of Government-Binding theory are deep and important, accounting in a remarkable way for a large variety of superficially unrelated linguistic facts.[16] But, and without in any way disparaging this achievement, it is one thing to claim that the principles of modern syntax describe, as one might say, the essential properties of languages; and it is quite another to claim that these principles are what speakers must know in order to learn a language.

First, it is in general false that theorizing about a thing's essential properties is the same enterprise as theorizing about what we need to know in order successfully to learn about that thing. Biologists may worry about what makes cats cats, but a child's grip on cathood predates her excursions into zoology. Philosophers attempt to uncover the essential properties of persons, but it is not your knowledge of that essence that is helping you deal with your boss. Mechanics, perhaps, theorize about the properties shared by all internal combustion engines, but no cognisance of those universals,

thankfully, is necessary for our learning to drive. And linguists, finally, theorize about UG, the essence of natural languages. But just as was the case with respect to cats and persons and cars, it is a very substantial additional claim that UG describes not just languages' essential properties, but also what *we need to know* in order for language acquisition to occur.[17]

A second line of argument for my claim that UG's relevance to explaining language learning should not simply be taken for granted may be extracted from Soames' (1984) insistence that linguistics and psycholinguistics are 'conceptually distinct' disciplines: they concern different domains of fact (natural languages vs. the psychology of those who speak them) and are supported by different kinds of evidence (facts about languages vs. facts about speakers). So while it is of course unlikely that theories in the two domains will prove to be totally unrelated, and while it is of course possible that linguists' theories (such as UG) might fortuitously turn out also exactly to describe the psychology of learners, this convergence cannot simply be assumed: it requires further empirical support.[18]

In sum, UG on the face of it embodies a claim about the defining characteristics of the class of natural languages. Learning theories, on the other hand, embody claims about the psychological mechanisms of language acquisition. And it's at least questionable whether the two kinds of claim will turn out to have all that much to do with each other: certainly it does not seem to be true in general that learning is driven by knowledge of essential properties. Hence it cannot simply be taken for granted that the learnability of UG bears on the question that concerns us in this section, namely, whether the information postulated in (DS) is innately specified.

But suppose the linguist were to grant my point that it is not a conceptual truth that learning requires knowledge of UG. He can nonetheless construct a strong prima facie case for this hypothesis. First, it constitutes the only even remotely well-articulated acquisition-theory around. In contrast to the empiricist's hand-waving appeals to a 'something-we-know-not-what' as a mechanism for language-learning, nativists have produced a variety of proposals as to how, given knowledge of UG, children could learn language from mostly positive evidence. Second, and more compellingly, some of these theories have been formalized, using the techniques of formal learning theory pioneered by Gold (1967), and provably guarantee the learnability of languages without access to negative evidence. Wexler, Culicover and Hamburger, for example, have demonstrated that the transformational component of a grammar is learnable from positive data if two UG-like constraints, which they call the 'Binary Principle' and the 'Freezing Principle', are imposed on the learner's possible hypotheses.[19] The Binary Principle, which is similar to UG's Subjacency Principle, states that no

transformation may have a structural description that refers to symbols in more than two adjacent levels in the deep structure. The Freezing Principle states that no transformation may apply to a configuration of symbols that could only have been created by the previous application of another transformation. These principles serve to delimit the class of grammatical hypotheses the learner considers during learning: any grammar that does not satisfy them is simply not entertained. Wexler, Culicover and Hamburger prove that given these constraints, natural languages are learnable without access to negative evidence. And since Wexler et al.'s constraints are similar to those specified by UG, their work provides support for the nativist's identification of the principles of UG with the domain-specific information implicated in learning.[20] Hence, the nativist may assert, the argument given at the beginning of this section *does* support the second conjunct, (I), of his position.

It is unclear, though, how much support results like these really provide for the hypothesis that it is our innate knowledge of UG that is making language acquisition possible. For while some mathematical learning models (including Wexler et al.'s) are indeed constructed with at least an eye towards their psychological cogency, their primary aim is nonetheless to prove learnability, and not to provide an accurate account of how real children learn languages. But learnability, unfortunately, appears often to be purchased at the expense of psychological plausibility.

One immediate locus of dubiety concerning work in mathematical linguistics generally is the concept of learnability itself. To prove a class of languages learnable is to prove that, given data from any one of the languages in that class, the learner will in some finite amount of time hit upon the correct grammar for that language and thereafter not be forced by the data to change her mind.[21] It is not, however, to prove the class learnable within some developmentally credible amount of time. Hence a standard learnability proof is at most a demonstration that it is *not impossible* that language could be learned from the available data given certain assumptions about the class of natural languages (that is, given certain assumptions about the innate UG). But in general, proofs that some phenomenon is not impossible given so and so assumptions lend those assumptions little credence: as any child fed the cabbage-patch-and-storks account of baby-making instinctually recognizes, a 'how possible' story is not always a 'how actually' account.[22]

Another problem with viewing language learning through the lens of formal learning theory derives from its idealisations of the learning strategy, the learner's method for formulating and testing hypotheses. Early models (e.g., Gold 1967) were criticized (e.g., Pinker 1979) because of

their requirement that in order to test his new hypotheses, the learner must remember all the data to which he has hitherto been exposed. This is plainly an untenable requirement to impose on actual children. Furthermore, such models conceive of the learner as adopting and testing whole grammars, rather than individual rules; and see him also as adopting whatever new hypothesis happens to be next in the enumeration rather than modifying his beliefs based on what he already knows. But there is no reason to think that this is how children actually proceed. A child whose grammar fails to generate some string in the data (say, *John left and Mary wept*) will surely not wipe the slate clean and adopt some wholly new (and likely unrelated) grammar for testing. Rather, she will simply add a rule for conjunction to the grammar she already possesses. The actual child, in contrast to the mathematical linguist's learner, approaches learning in a piecemeal and conservative manner.

Wexler et al. (1975, 225) claim that the learning function employed by their learner is much more plausible than the one just described, for their learner does not need to store all the data with which he has been presented, instead making his theoretical decisions (that is, whether to retain or modify his hypothesis) on the basis of his present state of grammatical knowledge, together with the current input. Further, he does not project and test whole grammars, working instead at the level of the retention or rejection and/or replacement of individual rules. But while these are clearly advances over the earlier Gold-style learner, Wexler et al.'s creation remains a very poor model of the learning child. First, he modifies his hypothesis only when it is explicitly falsified in the data (that is, when it fails to generate an input sentence). Yet, as I argue below (Section 4.3), other kinds of evidence (such as evidence that certain forms do not appear in the linguistic data) can also bear on learners' hypotheses. Second, it is assumed that the learner can always *tell* when his grammar fails to generate a datum: his derivational powers are assumed to be unlimited. Yet it is far from clear that actual children may be assumed to be logically omniscient.[23] Thirdly, when confronted with an unruly datum and deciding what action to take in response to it, Wexler et al.'s learner has unreasonably limited options: he may either add a single rule to his hypothesized grammar so as to generate the datum in question; or he may delete one of the rules used in the unsuccessful derivation. But while restricting the learner's options in this way enables Wexler and Co. to avoid the charge (aired above with respect to Gold's learner) that the learning function is implausibly holistic, the fact that children do seem at times to make major revisions in their state of linguistic knowledge suggests that the gradualness forced upon their learner perhaps errs too much in the other direction. Finally, and most

importantly, the revisions Wexler et al.'s learner actually makes in the light of a disconfirming datum are unrealistically insensitive to the character of the evidence that provokes them. Wexler et al. claim that when their learner decides whether to add or drop a rule (and if the latter, which rule to drop) "[t]his, of course, is done in a reasonable, not arbitrary, manner" (1975, 225). As revealed in Wexler and Culicover (1980, 99ff), however, 'reasonable' here means 'reasonable from the mathematical linguist's point of view', not 'reasonable from the psychologist's'. For, on their model, all possible changes in response to a disconfirming datum are held to be equiprobable, the learner selecting one of them at random.[24] Yet as even Quine in his most holistic moods recognized, some revisions in the light of recalcitrant experience are always more likely than others.[25]

A third problem with Wexler et al.'s model concerns its assumptions about what gets learned during language acquisition. In early versions of the theory (described in Wexler, Culicover and Hamburger 1975, 221–222), it was surmised that the learner learns only the transformational component of a generative grammar, the base being held to be innate. The latter assumption, however, is disputable. One difficulty, later acknowledged by the authors, is that since different languages order words and phrases differently, the base rules governing word-order, at least, must be fixed by experience.[26] Another, unacknowledged difficulty is that in order correctly to apply base rules, the learner must know the subcategorisation properties of lexical items. Yet it is implausible to hold that the lexicon too is known innately.

Later models attempted to mitigate the implausibility of this 'Universal Base Hypothesis' by surmizing that, antecedently to the learning of transformations, the base is learned from information concerning sentences' deep-structures.[27] But the increase in acceptability gained with this move highlights a third kind of psychologically implausible hypothesis made by Wexler et al., this time concerning the inputs to learning. Learning the base in the above model requires information about deep-structures. But deep-structures are not plausibly viewed as part of the primary linguistic data. They are theoretical entities postulated by syntacticians which, while perhaps playing a computational role in the production and/or comprehension of sentences, are surely not apparent to the casual ear of even competent speakers of a language. A fortiori, they are not apparent to the ear of the child who lacks even a grip on the most rudimentary phrase-structure – not to mention the transformational – rules of his language.[28] This problem, moreover, cannot be evaded by simply reinstating the Universal Base Hypothesis. For on all Wexler et al.'s models, learning transformations requires access to sentences' deep-structures too: transformations are

learned from inputs consisting of ordered pairs (*b, s*) of deep-structures and surface strings. It is true, as Wexler et al. (1975, 222) point out, that their model gains in plausibility by using the more accessible surface *sentences*, not surface-*structures*, as the second element, s, of the input-pairs.[29] However the problem of explaining how the learner extracts the first, deep-structure element, b, remains. Wexler and Culicover (1974) and Wexler et al. (1975, 237ff) argue that deep structures may be extracted from data concerning sentences' meanings.[30] But the notion that meanings can be in any straightforward way 'read off contextual clues' is itself highly questionable. So, while one of the virtues of Wexler et al.'s learner is that he does not need *negative* data to learn a language, he nonetheless requires positive data of kinds that are, arguably, equally unavailable. For what actual children have access to during learning is streams of noise set against a blooming, buzzing background of people and things. Hence a model assuming their access to (*b, s*) pairs during learning is at most only potentially adequate.[31]

I should emphasize that my object in this section has not been to denigrate the work of Wexler, Culicover and Hamburger. Rather, my purpose has been to stress that achieving psychological verisimilitude is not the primary aim of formal learning theory. The nativist, recall, identifies the principles specified in (DS) as being needed for language learning with the principles of the Universal Grammar, pointing then to the difficulty of learning UG from the data as a motivation for claiming that those principles are known innately. I argued above that the identification of UG with the knowledge demanded by (DS) cannot be made a priori and urged the need to supply empirical justification for this move. Here, I have argued that because its primary goal is not the provision of psychologically accurate models, work in formal learning theory offers no such validation. Since its intention is instead to expand our understanding of the computational demands of various types of learning task by showing how in principle learnability can be guaranteed, psychological plausibility frequently takes a back seat in the construction of the learning theorist's long and intricate proofs.[32] Precisely because "[m]any of the learning assumptions are motivated because of their simplicity or because of technical reasons involving proofs of learnability", as Wexler and Culicover (1980, 100–101) put it, the appearance of a learner like Wexler et al.'s on the scene gives little support to the idea that UG is implicated in language acquisition. Although making this hypothesis does indeed enable us to envisage learners who may be sure eventually to discover the correct transformations without negative data, this certainty is bought only at the expense of making other undesirable assumptions about the learner and the context in which he learns.

Pinker has eloquently pressed the point that guaranteeing learnability is only a part of the task of explaining language acquisition. He articulates (1979, 218) a number of conditions that a learning theory must satisfy in order to be a plausible psychological model. Weakest are the 'Learnability' and 'Equipotentiality' conditions, which jointly require that a model must entail the learnability of all natural languages. Stronger are the 'Time' and 'Input' conditions, which stipulate, respectively, that the model must allow for learning in a reasonable amount of time and under reasonable assumptions about the available evidence. And strongest are Pinker's 'Cognitive' and 'Developmental' conditions. These require that the model respect what is known about the general cognitive capacities of the learning child (by not assuming, for example, that she can remember every sentence she has heard); and that it reflect the actual course of language-acquisition (by predicting, for instance, that children will erroneously regularize the past-tenses of irregular verbs).

In his (1979), Pinker argued that "no current theory of language learning satisfies, or even addresses itself to, all six conditions" (1979, 218). The situation is little better today as regards the provision of empirical validation for the notion that an antecedent knowledge of UG is what drives language learning. Although much energy has been spent, particularly by proponents of the 'parameter-setting' approach to learning, in showing how prior knowledge of UG would enable a child to learn language from positive data, virtually no attention has been paid to establishing that theories incorporating UG satisfy Pinker's criteria.[33] As a consequence, the linguist's identification of the principles of UG with the information specified in (DS) remains unwarranted. Hence the question of whether or not those principles could be learned – the question of whether or not the nativist is entitled to (I) as well as (DS) – remains academic.

In this section, I have argued that since there are no convincing empirical grounds for thinking that knowledge of UG is what is required for language learning, there is no argument from the fact UG would be unlearnable from the primary linguistic data to the claim that UG is innate. The most that the nativist is entitled to is the conditional claim: *if* some plausible learning theory – one that (at least) satisfies all six of Pinker's criteria – ascribes to learners prior knowledge of the Universal Grammar, *then* the Logical Problem provides us with grounds for thinking that that knowledge is innate. Pending the establishment of the antecedent of this conditional, the argument discussed in this section provides no support for the nativist's (I).

4.2. *A Curry is Itself and Not Another Thing*

The nativist's second argument for (I) turned on the need for guaranteeing not just the individual's ability to learn language, but that of the entire species. An individual might by chance encounter just the information she needs to keep her grammatical theorizing on track. Yet although everyone exposed to English will learn English, we cannot all be assured such luck: not everyone can be guaranteed access to those vital data. The only way our collective success in language learning can be accounted for is to suppose that the 'luck' is built in. No matter what the information needed for learning language turns out to be, that information must be biologically encoded. The argument from the Logical Problem entails not only (DS), but (I) as well.

The nativist's case seems compelling. Nonetheless, I contend, it is flawed. For the Logical Problem is a completely general problem arising for *all* learning involving projection beyond our experience. There is, that is to say, a dearth of negative evidence not just in the domain of language, but in every domain in which people learn. So, if the need to guarantee learnability in the absence of negative evidence is indeed a reason to adopt (I) as well as (DS), we should 'go nativist' about everything. But this is untenable. The nativist's argument, therefore, must be invalid.

Consider, to take a mundane example, learning about food. Given some minimal gastronomic experience, virtually all normal humans are guaranteed to acquire 'culinary competence', the ability to recognize and distinguish a variety of foods from each other and from non-foods. Yet, when learning about curries, say, no-one ever systematically informs us that Irish stews, tacos and quiches – let alone boats and babies and bison – are not curries. There are infinitely many things that *might* be curries, but which are *not* curries, and which we have no information one way or the other about. Yet, and despite the dearth of negative evidence, we all manage to converge on the view that a curry is itself and not another thing. We all, that is to say, manage more or less to converge in our intuitive judgments as to what is and is not an instance of the kind: curry.[34]

What follows from this, for reasons identical to those given above in our discussion of language-acquisition, is (DS). General directions to the effect that we should choose simple, elegant, powerful etc. hypotheses allow too much room for overgeneralisation: "All the world's a curry". Since the data needed to correct such errors are unavailable, domain-specific information – perhaps some idea to the effect that curries are a kind of food, or that they have a characteristically spicy taste – must be constraining our choices of curry hypotheses.

But what surely does *not* follow from consideration of this case is (I). Although it is not the case that everyone can be guaranteed access to the same curry samples, all normal people exposed to a curry or two arrive at more or less the correct view about what curries are. In the case of language, the nativist urged us to accept (I) as an explanation of this kind of convergence: we all learn the same language because the information we need is built in. Are we then to accept the same inference here? Surely not! It is just *absurd* to suppose that the domain-specific principles required for learning about curries are innate, biologically encoded in a special 'culinary faculty'! But, and this is the point of this example, if we are reluctant to posit a 'culinary faculty' as guarantor of our ability to acquire curry-competence, we should be equally reluctant to accept (I) as the conclusion of the precisely parallel argument offered us by the nativist in the linguistic case.

My aim here is emphatically not to suggest that nativism about language is implausible in the way that nativism about curries is: my argument is not an attempted *reductio* of linguistic nativism. Instead, my purpose is to show that the truth of (I) in the linguistic domain cannot be established by the argument set out at the beginning of this section, just as the truth of (I) in the culinary domain cannot be established by the argument in its latter part. Humans learn an awful lot, about a bewildering variety of topics, from sketchy and largely positive data. That they can do so – that they are virtually *guaranteed* to do so given some minimal degree of information and motivation – is miraculous and mysterious. It is not, however, a reason to accept a nativist explanation of the miracle or solution to the mystery. To put the point another way, the curry case shows that, contrary to the argument given above, (I) *could not be* the only explanation of how we are all be guaranteed to master a certain domain from the (largely positive) data available. That the 'Logical Problem' in *this* instance is so clearly a sham suggests that the Logical Problem of Language Acquisition may be something of a pseudo-problem too. I shall argue that this is indeed the case. Arguments from the Logical Problem are based on conceptions of the learning task that are seriously inadequate. Once those inadequacies are recognized, the nativist's case loses much of its force.

The first thing that the curry case makes clear is that there is much more negative evidence around than a proponent of the Logical Problem would allow. What makes learning here possible is the vast quantity of *indirect* or *implicit* negative evidence about curries available. For example, the fact that we call hamburgers 'hamburgers', not 'curries', is surely evidence – not explicit evidence, but evidence nonetheless – that hamburgers are not curries. A second and related point raised by this example is that learning is

a gradual and piecemeal process. Our theories about curries are constantly being updated and modified in the light of the evidence, both positive and negative. This is important because it means that the child can use the hypotheses she has already formulated to find new sources of negative evidence. Once she has decided, for example, that curries are a kind of food, then the fact that we do not typically eat babies, whereas we do eat curries, is negative evidence: if we don't eat them, babies likely aren't curries. The relation between theories and their evidence, in other words, is holistic: what can count as evidence for or against an hypothesis depends on what else we believe. That relation is also, to draw a third moral from the curry case, probabilistic. Assessing the bearing of incoming data on our theories is a complex (and poorly-understood) matter of mutual accommodation and adjustment. A single piece of evidence may not suffice either to verify or to falsify a theory: cherished hypotheses may be retained in the light of apparently falsifying data and many pieces of evidence, each in themselves inconclusive, may have to combine their individual weights to overthrow them. Thus, whereas one experience of a bland and tasteless curry may not be enough to make me give up my hypothesis that all curries are spicy, repeated experiences of that kind may lead me to change my view.

A final point raised by our consideration of the curry connoisseur concerns the nature of 'culinary competence', the end-point of the learning process. First, there probably *is* no such final point: we *never* stop learning about curries. Second, it is likely that none of us has *exactly* the same theory about curries as anyone else. You and I may disagree, for example, about whether Mulligatawny Soup is a curry. On the one hand, it's a delicious concoction of meats, vegetables and spices with that characteristically spicy taste; on the other, though, it's a soup, not a stew, and is not a part of indigenous Indian cuisine, having originated in the kitchens of the British Raj in the 19th century. The fact that such disagreements can occur indicates that people's curry-theories are subtly different – as is to be expected given their different experiences. But, and this is what is crucial, this fact is completely consistent with their agreeing, as surely they do, on the vast majority of cases. What needs to be explained, then, is not how we all manage to converge on the exact same theory; but rather how we converge exactly enough to make both communication *and disagreement* about curries possible.

Thus the Logical Problem, as it pertains to curries, makes three serious errors about the learning task. First, it fails to recognize the many sources of indirect or implicit negative evidence in the data. Second, it fails to recognize the piecemeal nature of theory construction and the holistic and probabilistic relations of evidence to theories. Third, it assumes that what

needs explaining is how we all arrive at the same theory when, in fact, the existence of disagreements about curries shows that we don't. Once these misconceptions about learning are made explicit, the 'Logical Problem' in the curry case simply dissolves. Of course the problem of explaining *learning* does not dissolve. But the task of building a bridge between what we learn and the data we learn it from *without* the help of a biologically-specified culinary faculty looks, as it should, much more tractable.

I contend that the Logical Problem of Language Acquisition is flawed in exactly these ways. In misconceiving the nature of the learning task, it makes the nativist position look far more inescapable, and the empiricist approach far more implausible, than either of them is. I take it for granted that the general points made about theory-construction and confirmation made above apply also in the case of language-learning. So too, I think, does my contention that the 'final states' of the learning process may differ subtly from person to person. Although we all agree about the grammaticality or not of overwhelmingly most sentences, there are nonetheless problem cases. Is *Colorless green ideas sleep furiously* ungrammatical (as some current theories would suggest) or merely semantically anomalous? Is *Who you believe the rumours about*? ungrammatical (as Lightfoot (1982, 114)) maintains? Or is it, as I think, grammatical (though rather inelegant)? For the purposes of constructing a grammar of English, the Chomskian invocation of an 'ideal speaker hearer' as arbiter of these disputes is surely acceptable. But when the issue is language learnability, this idealisation leads us to see a problem – that of explaining complete convergence among different speakers – where no such problem exists.

Most important of the points raised above, however, is the issue of negative evidence, for it is here, I think, that the crux of the alleged learnability problem for languages resides. This topic is the focus of the next section of this paper.

4.3. *Substitutes for Negative Data*

The argument from the Logical Problem asserts that a learner who believes that all Fs are Gs, when in fact only some (or, indeed, no) Fs are Gs, will need negative data – the information that some Fs are *not* G, in order to shrink her hypothesis and arrive at the correct theory. But, as the argument goes on to point out, explicit information to the effect that Fs are not Gs is typically unavailable. Hence the nativist's insistence that the learner's hypotheses be so strictly constrained by domain-specific innate information that she never (or very rarely) entertains an hypothesis that is too large.

Now linguists do pay lip service to the idea that there are indirect sources of negative evidence in the linguistic data: the possibility is typi-

cally mentioned, only to be set aside. The arguments in the previous section, however, show that this possibility should not be ignored. Just as there are many sources of negative evidence in the data concerning curries, there must be substantial sources of negative evidence in the data concerning language.

As a preliminary to exploring some of these sources, let me introduce some terminology. Hitherto, I have been using 'data' and 'evidence' interchangeably. Let me now stipulate that by 'data' I shall mean the facts as they are presented to experience; and by 'evidence' I shall mean those facts as they bear on the (dis)confirmation of some theory. The point of this – admittedly vague but, I think, serviceable enough – distinction is to underscore the fact that positive data can constitute *both* positive and negative evidence for a theory – and similarly for negative data. The positive datum that a is G can be positive evidence for (i.e., confirm) a theory T if, for instance, T predicts that a is in fact G; conversely, if T predicts that a is *not* G, then that same positive datum will be negative evidence for (i.e., disconfirm) T. Similarly, the negative datum that b is not G can be positive evidence for T (if T predicts that fact) or it can be negative evidence for T (if T predicts the contrary).

What the Logical Problem points out is that, in the case of language-learning, there are very few *negative data* available: the child has access to very little explicit information to the effect that such and such string of words is not a sentence of the language he is learning. And what the generalized argument (introduced with the curry example in Section 4.2) maintains is that the failure of experience to supply negative data is universal. But what poses a problem for a learner whose hypothesis is overgeneral is not a lack of negative data per se; it is a lack of *negative evidence* – of the means to disconfirm that hypothesis. So, what we need to find, assuming that we cannot rule out altogether the possibility that her learner will project hypotheses that are too large, are sources of negative evidence other than those provided by explicitly negative data.

4.4. *Non-Occurrence as Negative Evidence*

In discussions of the Logical Problem of Language Acquisition, it is often emphasized that the non-occurrence of a particular string in the linguistic data cannot function as a negative datum: from the mere fact that she has not heard S to date, the child cannot conclude that S is not a sentence of L. For, most sentences of a natural language never have and never will be spoken. Thus the child has no way of telling whether S's non-occurrence is due to its not being a sentence; or whether it's simply due to the fact that

no one has had occasion to utter it yet. This being so, the child cannot use non-occurrence as negative evidence.[35]

I want to suggest, however, that in many cases, the non-appearance of a string in the primary data can legitimately be taken as constituting negative evidence. The first, fairly obvious, way that a learner could use non-occurrence as a source of negative evidence is by making predictions on the basis of her current state of knowledge, and having them disconfirmed. Think of Edna, in the grip of the idea, widespread among pre-school children, that all transitive verbs can be used as causatives.[36] Her grammar will generate sentences like **I giggled her*, meaning 'I caused her to giggle', in addition to the perfectly acceptable *I melted it*, meaning 'I caused it to melt'. Consider the scene witnessed by Edna: Father knocks over his coffee and utters the sentence: *I caused the cup to fall from the table*. Edna's grammar, however, had generated **I falled the cup off the table* as appropriate to those circumstances. Her prediction was disconfirmed by her father's utterance. The non-occurrence, in that context, of the string predicted by Edna's grammar constitutes evidence that her grammar is inadequate. Note that in this case, the *datum* is positive: Edna's datum is that *I caused the cup to fall from the table* is a sentence. But the *evidence* it supplies is negative: the occurrence of something other than what Edna's theory led her to expect constitutes evidence against her hypothesis. This evidence is not, to be sure, *decisive*. Other explanations for the failure of match are available. Perhaps Edna's father misspoke; perhaps she did not hear him aright; perhaps both of *I falled it* and *I caused it to fall* are grammatical in that context, etc. But, as I argued above (Section 4.2), the fact that a theory may always be saved in the face of recalcitrant experience is hardly news; and it's not news either that more than a single failure of prediction may be needed to overturn a cherished hypothesis. But these general facts about the holistic and probabilistic nature of the relation between theories and their evidence should not lead us to ignore the fact that non-occurrence *can* constitute negative evidence against a theory: *indecisive* negative evidence, but negative evidence nonetheless.

There is a second, more interesting, way in which non-occurrence can serve as negative evidence: it can serve as evidence that whole *classes* of objects are not instances of the kind whose extension is being learned. Consider the following strings of English words, neither of which, I'll assume, have ever been uttered:

(1) Steve enjoyed the curry

(2) Enjoyed curry Steve the

I want to suggest that, while the non-occurrence of (1) is *not* evidence that *that* string is not a sentence, the non-occurrence of (2) *is* in fact evidence that *that* string is not a sentence of English. The reason for the asymmetry is that in the case of (1), but not in the case of (2), we have independent evidence that strings of that kind *do* occur in English. First, distinguish a string *qua string* from a string *qua instance of a particular syntactic structure*, where by 'particular syntactic structure', I mean a phrasemarker with the lexical items left out, or a tree of a certain shape with the terminal nodes left empty. So, (1a) and (2a) are rudimentary assignments of syntactic structures to strings (1) and (2) respectively.

(1a) $[_{\text{S}}[_{\text{NP}}][_{\text{VP}}[_{\text{V}}][_{\text{NP}}[det][_{\text{NP}}]]]]$

(2a) $[_{\text{V}}][_{\text{N}}][_{\text{N}}][det]$

The relation of syntactic structures to strings is one-many: for any given structure, there will be many strings that are instances of it and that can be obtained by substituting morphemes of an appropriate syntactic category into the vacancies left in the phrasemarker. Intuitively, then, there are many strings that would serve as evidence that a particular structure is allowed by the language. The occurrence of any of the following strings, for example, would be evidence that the structure (1a) is a structure of English:

(3) George loves a fight

(4) Cats hate the rain

(5) Curries are an art form

Further, because these sentences are all evidence that (1a) is a structure of English, the occurrence of any one of them is evidence that all the other strings instantiating that structure are sentences of the language too, even if some of those other sentences never in fact appear in the primary linguistic data. So in cases where a given structure has been encountered before in the data, the non-occurrence of a particular *instance* of that structure – like (1) – *is not* evidence that that string is not a sentence. For, we already have positive evidence that structures of which (1) is an instance *do* occur in the language: structures, so to say, trump strings. In cases where a particular structure has *never* been instanced in the data, by contrast, further non-occurrences of strings of that kind constitutes *further* evidence that that structure is not allowed in the language. Hence, the non-occurrence of a sentence like (2), or indeed of any sentence having the

structure (2a), provides evidence that the structure (2a) is not allowed by the language. And this, in turn, is reason to believe that (2), which exhibits that structure, is not a sentence of English. The nonoccurrence of a string like (2) constitutes negative evidence. But what makes it negative evidence is not merely its non-occurrence *qua string*: what makes it negative evidence is its non-occurrence *qua* instance of a particular syntactic structure. Again, not decisive evidence, but evidence nonetheless.

For similar reasons, non-occurrence can be evidence also against the existence of certain grammatical *rules*. The non-occurrence of **On the table is tough to dance*, for example, is evidence that there is no rule of 'Prepositional Phrase Raising' in English; whereas the nonoccurrence of *The table is tough to dance on* is *not* evidence against the existence of a rule for NP Raising. For in the latter case, unlike the former, we have plenty of positive evidence that NP raising is permissible.[37]

4.5. *Positive Data as Negative Evidence*

The Logical Problem asserts that the primary data can provide a learner who overgeneralizes with no motive to 'shrink' her hypothesis. We have just seen how, on the contrary, the non-appearance in those data of forms predicted by her grammar may supply the learner with the means to disconfirm an overgeneral hypothesis. Pinker describes a method, which he calls 'constraint sampling', whereby positive data too can lead the learner towards a less inclusive grammar.[38] In constraint sampling, the learner uses her analyses of the primary linguistic data as a basis for restricting the application of an over-general rule. Given an input sentence, the learner randomly selects one feature of that sentence and applies it as a constraint on the rule. Each possible constraint has a non-zero probability of being hypothesized at any given time and every constraint adopted is retained until a sentence violating it is encountered in the input, at which point it is dropped from the constraint-set.[39] Eventually, the child attains the correct set of constraints on the rule, and her grammar no longer overgenerates in that respect.

A feel for how the procedure works can be got by considering how a child might learn when to affix the suffix '−s' to the verb stem in a declarative sentence.[40] In order to do this, the child must learn that '−s' encodes present tense, imperfective aspect and third person singular subject. Suppose she heard the sentence *The boy wants a curry*. The most general rule she could propose consistent with this input would be something like 'Always add '-s' to the verb stem'. This rule, however, is over-general and would lead the child to make errors like **I likes it, *We wants them* etc. On Pinker's theory, a child who projected such a rule would immediately

start hypothesizing constraints on it. Hearing, say, *Dad wants a beer*, she might propose that the '−s' encodes for an animate subject. Hence, her rule would be 'Add an '−s' only when the subject is animate'. This rule would generate errors like **This pen write well* and would be straightforwardly falsified by an input like *The curry tastes good*, at which point the learner will project another constraint – say that the '−s' encodes for present tense. No inputs violating this constraint would be encountered, so it would be retained and further constraints added to the rule. Eventually, the learner will acquire the correct set of constraints on the rule and never be forced by the input to abandon them.

Although his argument is too complex to rehearse here, Pinker (1984) shows how the hypothesis that children employ constraint sampling predicts quite specific facts to do with the kinds of errors they make and the overall course of acquisition. He points out that this approach involves "abandoning the assumption that the child is a completely rational hypothesis-tester, one who never abandons an hypothesis unless it is contradicted by some input datum" (1986, 69). The constraint-sampling approach, that is to say, *builds in* the holistic and probabilistic view of confirmation that I urged in Section 4.2 above. It is perhaps for this reason that Pinker is able to see sources of negative evidence in the linguistic data that other theorists in the grip of the Logical Problem have missed.

In this section, I have attempted to elucidate some of the ways that a learner might obtain negative evidence in the absence of negative data. The aim has been to buttress my contention, made in Section 4.2, that the Logical Problem of Language Acquisition is misconceived. The generality of that problem, I argued, was reason to be suspicious of its claim to establish (I) in the linguistic case. The fact that it is based on a number of erroneous assumptions about the nature of learning indicates why those suspicions are justified. The argument from the Logical Problem, in sum, fails to support nativism.

5. CONCLUSIONS AND IMPLICATIONS

In the foregoing sections of this paper, I have contended that the argument from the Logical Problem fails to support nativism about language learning. In response to this, it may be objected that the Logical Problem is not in fact *intended* by nativists to be an argument for their position. According to Pinker (in conversation), for example, linguists' discussion of the Logical Problem is not meant to incline us toward any one type of learning theory over another: it is meant only to delineate the constraints that any acceptable such theory – of whatever stamp – must satisfy. But

while I am prepared to believe that Pinker himself means to invoke the Logical Problem only to make vivid the difficulties involved in providing an acquisition theory, I do not think that his view of the Problem's overall dialectical role is entirely accurate. It seems to me that consideration of the Logical Problem, especially when supplemented with additional considerations like those adduced in Sections 2 and 3 above *is* an argument for nativism over empiricism. (Though not, as I have argued, a particularly good one.) And it seems to me further that when linguists *present* the Logical Problem of Language Acquisition, something very much like this argument is what they have in mind. Take Lightfoot, for instance. Immediately after presenting the Problem and discussing an example where children learn a rule in the absence of negative data, he writes – in what certainly appears to be an argumentative spirit:

> If the child's linguistic experience does not provide the basis for establishing some particular aspect of linguistic knowledge, some other source for that knowledge must exist. That aspect must be known a priori in the sense that it is available independently of linguistic experience. It may be available genetically or arise of a consequence of some other non linguistic experience; I shall gloss over this distinction and speak only of genetic determination . . . . (Lightfoot, 1982, 18)

Or consider Lasnik. Although he does not actually make the argument, Lasnik also seems to view the Problem as the basis of an argument for nativism, beginning his (1989) paper with the following summation of the status quo:

> Much of the recent discussion of language learnability has centered around the absence for the learner of negative evidence and the implications of that absence. The basic argument has been reiterated many times: If the child does not have access to negative evidence . . . then Universal Grammar presumably does not make available choices that can only be resolved by such evidence. (Lasnik 1989, 89)

Even Pinker himself, in what may be a slip of the pen, seems at times to invoke the Logical Problem in an argumentative vein. In his (1990), for instance, he appears to be saying that the Problem provides us with reason to abandon (a version of) empiricism and adopt (a version of) the nativist's (DS):

> . . . without negative evidence, if a child guesses too large a language, the world can never tell her she's wrong. This has several consequences. For one thing, the most general learning algorithm one might conceive of – one that is capable of hypothesizing any grammar . . . – is in trouble without negative evidence. Without negative evidence . . . there is no general-purpose, all-powerful learning machine; a machine must in some sense "know" something about the constraints in the domain in which it is learning. (1999, 204)

In fact, the argument from the Logical Problem has now more or less usurped other arguments for nativism (such as arguments from the ubiquity

of linguistic universals, arguments from species specificity, arguments from structure dependence of grammatical rules etc.) in the introductory chapters and sections of books and papers on language acquisition. So while Pinker may be right to maintain that the Logical Problem is not *always* intended to be an argument motivating nativism, I think it must be conceded that the Problem's dialectical role is at least highly ambiguous. If for no other reason than to resolve this ambiguity, then, the argument deserves examination.

I have attempted to provide the requisite examination in this paper. My overall conclusion has been that although the Logical Problem does support (DS), the claim that domain-specific information is necessary for language learning, it does nothing to support (I), the claim that that information is innate. Thus, since nativism is the conjunction of those two theses, and since the argument from the Logical Problem supports only one of the conjuncts, the motivation it provides for the nativist position is at best very partial.

In Sections 2 and 3, I argued that although that argument does not constitute an absolutely watertight case for (DS), the hypothesis that learners are constrained in their theorizing about language by task-specific principles does represent a plausible suggestion as to how they might avoid making extravagant overgeneralizations – and hence from needing access to amounts and kinds of negative evidence that may not exist. In Section 4, however, I argued that the Logical Problem provides no support whatsoever for (I). First, it is irrelevant to the truth of (I) that the principles of UG are unlearnable from the primary data. For, there is no good reason to accept the linguist's assumption that those principles are implicated in language-learning. Second, the very generality of the Logical Problem undercuts its power to support the hypothesis of a biologically-specified language faculty: if it proves anything, it proves too much. This latter point is made vivid by examining a version of the Logical Problem argument for another position, "culinary nativism", which I take to be wildly implausible. I argue that since the Logical Problem argument doesn't do much to incline us to accept culinary nativism, it can't do much to incline us towards linguistic nativism either. There are many domains in which learning unproblematically takes place in the absence of negative data. So unless we are prepared to accept 'culinary faculties' for learning about curries, 'automotive faculties' for learning about cars, and 'botanical faculties' for learning about trees and flowers, and so on, we should resist the inference to faculties in the linguistic case.

As is by now no doubt clear, my view is that we should indeed resist the global nativism that is entailed by the generalized form of the argument from the Logical Problem – and hence resist also its nativist implications in

the domain of language. In Sections 4.2 and 4.3 I attempted to show how the Logical Problem's apparent force derives from a conception of the learning task that is misleading in a number of ways, thus providing something more than mere intuition in support of my contention that an inference to nativism based on it is invalid. Among the problematic features of this conception are its failure to recognize the probabilistic and holistic nature of confirmation and its related conflation of negative evidence (which may have many sources) with explicitly negative data (which are few and far between). Given these shortcomings, I conclude, the recent obsession in the literature on language learning with the argument from the Logical Problem is just a mistake. To the extent that linguistic nativism is plausible (and in this paper, I take no stand on this matter), its plausibility does not derive from the Logical Problem argument: it must have other sources.

Let me end now by emphasizing that my aim in this paper has *not* been to argue either against nativism or for empiricism in the domain of language. It has rather been to show that one highly-influential argument for the nativist position is ill-founded. My own view is that it remains an open empirical question whether or not language acquisition is subserved by a task-specific faculty. It is very likely, I think, that people possess *both* general-purpose *and* task-specific abilities to learn, hence that nativism is likely to be true in some domains; and empiricism in others. Learning how to recognize the faces of people one knows, for instance, seems a likely candidate for a nativist explanation. Minutes-old infants respond selectively to faces and, despite the computational complexity of the face-recognition task, two-day old babies can reliably discriminate their mother's faces from those of strangers (Ellis 1992a, b): the extremely subtle cues used for this purpose do not have to be learned. Furthermore, face recognition can be selectively impaired due to brain damage: prosopagnosic patients not only lose their ability to recognize faces but are also unable to re-learn the techniques needed to perform this task. Finally, since it is so clearly to an organism's advantage to be able to distinguish members of the group not only from each other but also from (possibly hostile) strangers, a satisfying evolutionary explanation is available for why a special-purpose mechanism for face recognition might have proliferated in the species.

But there are other abilities we have that do not display these sorts of characteristics and that seem to be more plausibly viewed as products of a *general* capacity we have to learn from experience. Knowing what curries are; being able to recognize a Mozart symphony; being able to distinguish an electron-track in a cloud chamber from that of an alpha-particle; being able to write a philosophy paper or drive a car or sing the National Anthem

– these are abilities for which it seems, to say the least, *farfetched* to postulate domain-specific learning devices.

So, first, *whatever* is the truth about language-learning, we are *still* going to need an empiricist-style general learning mechanism to account for some acquisition phenomena. Maybe it will turn out to be also the mechanism for language-learning, maybe not. But nativism is unlikely to be true in *every* domain. Hence it is a mistake to allow the argument from the Logical Problem to forestall – as it has forestalled – the development of non-nativist theories of learning. And, second, the truth in a particular domain – whether it be language-learning or the acquisition of other competencies – must be discovered the hard way, that is, by finding the best overall account of the facts of acquisition. The correct account may be nativist – and there are other arguments, which I have not addressed in this paper, suggesting that this might well be so. But it may also be empiricist – as is surely the case when we learn about curries and cars. There is no way of settling the issue in advance. Nativism cannot be ruled in, nor empiricism ruled out, by an argument from the Logical Problem.

NOTES

[*] I am indebted to Alex Byrne, Michael Devitt, Gilbert Harman, David Hilbert, Scott Soames, Kim Sterelny and Jim Woodward for their comments on various earlier versions of this paper. Many thanks also to William Demopoulos and Stephen Pinker for their very helpful referees' reports. I am grateful also to a third, anonymous, *Synthese* referee for his or her suggestions.

[1] See, e.g., Chomsky (1990), Lightfoot (1982, 1989, 1991), Pinker (1984, 34–37).

[2] No empiricist will deny it: but none would assert it either. For an empiricist, it would be highly misleading to single out our language-learning ability in this way and imply that it is somehow special or distinctive in being innate.

[3] An oft-cited example being their tendency to apply the rule for constructing past tenses of regular verbs to irregular verbs, generating, e.g., the incorrect **goed*, **swimmed* and **runned* in addition to the correct *walked*, *hurried* etc. Analysis of the errors children do and do not make may influence theorizing about the exact contents of UG.

[4] What follows is a synthesis of the many statements of the argument to be found in the literature. See, e.g., Braine (1971, 157); Baker (1979); Hornstein and Lightfoot's 'Introduction' to their (1981); Matthews (1984); Lasnik (1989, 89–90); Pinker (1986, 55–56, 1989, 5–6). My presentation is particularly indebted to those of Pinker.

[5] Chomsky (1959), Brown, Cazden and Bellugi (1969), Brown and Hanlon (1970).

[6] For example, incomplete sentences, sentences containing 'Umms', 'Uhhs', etc. Note, however, that according to one estimate (Newport, Gleitman and Gleitman 1977), an astonishing 99.7% of sentences spoken by mothers to their children are grammatically impeccable.

[7] Take case (ii) in Figure 1 above, for example. Positive evidence will force the child to expand the boundary of **H** rightwards (so to speak) so that it converges with the rightmost boundary of **L**. But negative evidence is still required to deal with the leftmost, overgenerating, portion **H**.

[8] I have elsewhere distinguished the Logical Problem argument from another form of nativist argument which I call 'De facto' or 'A posteriori' poverty of the stimulus arguments. Whereas the Logical Problem argument is a very *general* way of making a case for linguistic nativism, abstracting from many specific empirical details, De facto arguments assert the impossibility of acquiring *particular* grammatical rules (e.g., 'The Principle of Structure Dependence' (Chomsky 1975, 33ff) or Principle B of Binding Theory, or the argument structures of particular verbs (Pinker 1989)), given *particular* assumptions about the contents of the primary linguistic data. They thus seem more vulnerable to empirical attack than are Logical Problem arguments. For example, a De facto argument which appealed to the paucity of evidence in the primary linguistic data regarding deep structures as evidence for the innateness of the deep-structure/surface-structure distinction would be on shaky ground these days. For if, as some linguists (e.g., Chomsky 1995, 390ff.) argue, the deep-structure/surface-structure distinction is misdrawn, then the question of how we come to know that distinction dissolves. Nothing I say in this paper is intended to speak to the adequacy of De facto arguments for nativism. See Cowie Section 2 (forthcoming) and Chapter 6 (in preparation) for further discussion.

[9] Proponents of connectionist-style language learning models (e.g., Rumelhart and McClelland 1986), for instance, have challenged the view that learning is a matter of hypothesis-testing. Soames (1984) and Devitt and Sterelny (1989) have questioned whether there is any interesting sense in which speakers know grammars. And linguists themselves have speculated that there might be unrecognized sources of *indirect* negative evidence in the primary linguistic data – see especially Lasnik (1989). See also Section 4.2.

[10] The results of Gold in his seminal (1967) paper may, I think, have been misunderstood in this connection. Pinker (1979, 232), for example, claims that Gold's proof to the effect that the learner's hypothesis set must be constrained means that we are "committed to at least a weak form of nativism, according to which 'the child approaches the data with the presumption that they are drawn from a language of an antecedently well-defined type' (Chomsky 1965, 27)". That is, we are committed to nativism's (DS): the child must be constrained by prior views about language. However what Gold in fact proved was that a learner who is utterly unconstrained as to possible hypotheses will fail to learn a language from positive data. He did not show that no learner constrained *merely by general-purpose constraints* (such as 'Prefer simpler hypotheses' etc.) must fail in that task. Since nativist and empiricist agree that there must be *some* constraints on the hypothesis space – since their dispute concerns not the need for constraints per se, but rather the nature of the constraints needed – Gold's results are wholly irrelevant to their dispute. Demopoulos reaches a similar conclusion in his (1989).

[11] Note, however, that one study (Brown and Hanlon 1970) has argued that failures of understanding play little role in children's acquisition of syntax.

[12] The necessity that our acquisition theory guarantee learnability is emphasized in, e.g., Pinker (1989, 11); Wexler (1982, 291–292); Hornstein and Lightfoot (1981, viiff). What is new in my presentation is the drawing of a clear distinction between this point and the much more general point (discussed in Section 2) that a dearth of negative evidence in some domain makes it difficult for a learner to recover from overgeneralizations. The distinction matters because the two points suggest quite different arguments for (I) which deserve to be treated separately – see Sections 4.1 and 4.2.

[13] *Treatise* III.1.i (1978, 464). Wexler and Culicover (1980, 9) trenchantly put the point thus: ". . . it is perfectly possible logically that the purported linguistic constraint is a special case of, or follows from, a more general cognitive constraint, that is, a constraint that applies to all cognitive systems . . . but the cognitive domains that are referred to . . . are not for the

most part subject to any kind of reasonably well specified theory, so it doesn't even make sense to assert the existence of a constraint that applies to these domains. One might almost as well suggest that the linguistic constraint applies to black holes, which would make the constraint even more general".

[14] Thanks to Jim Woodward for this point.

[15] Cf. Chomsky (1988, 6): "We may ask whether the linguist's constructions are correct . . . . But there are few meaningful questions about the "reality" of these constructions – the "psychological reality", to use the common but highly misleading term – just as there are few meaningful questions about the "physical reality" of the chemist's constructions".

[16] The 'government-binding' or 'principles and parameters' approach to syntax was introduced in Chomsky (1981). See Sells (1985, Chap. 1) and Friedin (1994) for overviews of the theory.

[17] I would not want to claim that beliefs about the essences of Fs *never* play a role in learning about Fs. In some domains, e.g., mathematics, learning does seem to proceed at least partly through knowledge of essential properties.

[18] The discussion in this paragraph does not do justice to this hotly-debated topic in the philosophy of linguistics. The reader may consult e.g., Antony (1991) and Chomsky (1988) for defenses of the linguist's view that to do linguistics is to do psychology; and Devitt and Sterelny (1989) and Soames (1984) for defenses of the opposing position.

[19] See Hamburger and Wexler (1975), Wexler, Culicover and Hamburger (1975), Wexler and Culicover (1980).

[20] This is clearly the moral Lightfoot draws from the work of Wexler et al. Their model, he writes (1982, 35n), "specifies the boundary conditions for a theory of language-learning . . . . In fact it is likely that grammars are more restricted than they suppose, but at this stage it is useful to have a mathematical proof that at least some kind of transformational grammar is learnable under reasonable assumptions about the language learner".

[21] Gold (1967). Wexler and Culicover (1980, 40–41, 99) express the criterion of learnability in terms of there being some finite *data set* which enables the learner to select the correct grammar for her language. These formulations are equivalent for our purposes, for in these models the data are assumed to be presented sequentially, one datum per unit time. Hence finite data implies finite time, and vice versa.

[22] Thanks to Kim Sterelny for the catch-phrase.

[23] Thanks to Kim Sterelny for this point.

[24] See Wexler et al. (1980, 102 and note 30). For example, suppose there were 10 rules used in the unsuccessful derivation (hence 10 rules any one which could be dropped from the hypothesis set); and 20 new rules any one of which could be added to the set. Then, the learner has 30 options, each with a probability of 1/30 of actually occurring.

[25] Cf. Quine's (1953) insistence that, despite the fact that any experience can be taken to bear on any hypothesis, some hypotheses are more "germane" to sense experience than others, hence more likely to be revised.

[26] See Wexler et al. (1975, 237).

[27] See Wexler and Culicover (1974), Wexler et al. (1975, 237ff), Wexler and Culicover (1980, 99ff).

[28] For example, even if one *knew* all the transformational rules (which this learner doesn't, since they haven't been learned yet) it is still no simple matter to extract a deep-structure from a surface string.

[29] It's nonetheless worth stressing here that the difficulty of extracting syntactic information, even something as mundane as a string of words, from the acoustic signal is one of

the reasons nativists have argued that an empiricist-style general-purpose learner could not 'bootstrap' his way up from the data to grammatical competence.

[30] By assuming (i) that the ordering of elements in sentences' semantic representations are invariant across languages and (ii) that the hierarchical relations in the semantic representations are retained in the deep structure (the 'Invariance Hypothesis'), they attempt to explain how learners extract deep structures from the linguistic input: once the rules governing the linear ordering of base constituents have been figured out, deep structure representations may be read off sentence meanings, which in turn are extracted from contextual information.

[31] Another potential source of controversy concerning the inputs to learning derives from Wexler et al.'s supposition that the primary linguistic data contain surface-structures of degree-two complexity. Some theorists, however – most notably David Lightfoot (1989, 1991) – have argued that the data cannot be assumed to include such complex constructions. Lightfoot (1989) argues, indeed, that language must be learnable from data of degree-zero complexity. (A sentence of degree two complexity contains two embedded sentences – e.g., *The crook* [$_S$ *who stole the gun* [$_S$ *that was used in the robbery*]] *was arrested.* A sentence of degree-zero complexity contains no embeddings.)

[32] Wexler and Culicover, e.g., are quite explicit about this. Cf. their (1980, 97–104).

[33] Pinker's own work, as might be expected, is an exception. I discuss some of his ideas, and question his invocation of UG, in Section 4.3.

[34] It's perhaps worth stressing here that the task at issue is not that of learning the meaning of the word 'curry'. It's rather that of learning *what curries are*. Cf. language learning: the problem is not learning how to apply the words 'grammatical' and 'ungrammatical' to strings; it's that of learning *what the sentences of one's language are.*

[35] Baker (1979, 536), Lightfoot (1982, 17) and Pinker (1989, 14).

[36] See Pinker (1986, 61).

[37] Note that my argument does not depend on the assumption, made here for ease of exposition, that the various types of rules and structures hypothesized by the child are stated in terms of familiar grammatical categories like NP, VP etc. The argument works equally well on the perhaps more realistic assumption that children define syntactic types differently at different stages of language acquisition. So, for example, a child might reason that since there have been plenty of sentences with the structure ACTOR-ACTION in the data, the non-occurrence of a particular sentence (say, *Barney sucks*) is not evidence against that structural hypothesis; whereas the fact that there have been no sentences of the form ACTION-ACTION-ACTOR-THING is evidence against that structure, and hence against instances of it like *Hit kick Barney dinosaur*.

[38] See his (1986, 67–70) for an overview; the theory is presented in detail in his (1984).

[39] I should note that although Pinker's constraint-sampling heuristic is quite general, he is a nativist, holding that the space of possible constraints for a given rule-type is specified innately by UG. It seems to me, though, that if we allow (as all models of language-learning in fact do) that a learner can use her current state of grammatical knowledge to perform some preliminary syntactic analysis of the input, possible constraints might be suggested by that analysis, rather than (as in Pinker's model) being selected from an innately-specified set.

[40] This example is mine, not Pinker's, and is intended only to provide a simple illustration of the constraint-sampling procedure: it is not meant to be an account of the actual process by which children learn the rule in question.

REFERENCES

Antony, L.: 1991, 'Rabbit Pots and Supernovas', unpublished manuscript.

Braine, M. D. S.: 1971, 'On Two Types of Models of the Internalisation of Grammars', in D. I. Slobin (ed.), *The Ontogenesis of Grammar: A Theoretical Symposium*, Academic Press, New York.

Brown, R., C. Cazden, and U. Bellugi: 1969, 'The Child's Grammar from I to III', in J. Hill (ed.), *Minnesota Symposium on Child Psychology*, Vol. II, University of Minnesota Press, Minneapolis.

Brown, R. and C. Hanlon: 1970, 'Derivational Complexity and Order of Acquisition in Child Speech', in J. R. Hayes (ed.), *Cognition and the Development of Language*, John Wiley, New York.

Chomsky, N.: 1959, 'Review of Skinner's *Verbal Behavior*', *Language* **35**, XX-YY

Chomsky, N.: 1975, *Reflections on Language*, Fontana, London.

Chomsky, N.: 1981, *Lectures on Government and Binding*, Foris, Dordrecht.

Chomsky, N.: 1988, *Language and Problems of Knowledge, The Managua Lectures*, MIT Press, Cambridge, Massachusetts.

Chomsky, N.: 1990, 'On the Nature, Use and Acquisition of Language', in W. G. Lycan (ed.), *Mind and Cognition: A Reader*, Basil Blackwell, Oxford, pp. 627–645.

Chomsky, N.: 1995, 'Bare Phrase Structure', in G. Webelhuth (ed.), *Government and Binding Theory and the Minimalist Program*, Basil Blackwell, Oxford, pp. 383–439.

Cowie, F.: (forthcoming), 'Innateness of Language', *Routledge Encyclopedia of Philosophy*, Routledge, London.

Cowie, F.: (in preparation), *Innate Ideas and Ideas of Innateness.*

Culicover, P. and K. Wexler: 1977, 'Some Syntactic Implications of a Theory of Language Learnability', in P. Culicover, T. Wasow and A. Akmaijan (eds.), *Formal Syntax*, Academic Press, New York.

Demopoulos, W.: 1989, 'On Applying Learnability Theory to the Rationalism-Empiricism Controversy', in R. J. Matthews and W. Demopoulos (eds.), *Learnability and Linguistic Theory*, Kluwer, Amsterdam, pp. XX-YY.

Demopoulos, W. and Marras, A.: 1986, *Language Learning and Concept Acquisition: Foundational Issues*, Ablex, Norwood, New Jersey.

Devitt,M. and K. Sterelny: 1989, 'Linguistics: What's Wrong with "The Right View" ', in J. E. Tomblin (ed.), *Philosophical Perspectives 3: Philosophy of Mind and Action Theory*, Ridgeway, Atascaterola, CA, pp. 495–531.

Ellis, H. D.: 1992a, 'The Development of Face Processing Skills', *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)* **335**, 47–54.

Ellis, H. D.: 1992b: 'A Wise Child: Face Perception by Human Neonates', *Behavioral and Brain Sciences* **15**, 514–15.

Friedin, R.: 1994, 'The Principles and Parameters Framework of Generative Grammar', in R. E. Asher (ed.), *The Encyclopedia of Language and Linguistics*, Pergamon Press, Oxford.

Gold, E. M.: 1967, 'Language Identification in the Limit', *Information and Control* **10**, ??.

Hamburger, H. and K. Wexler: 1975, 'A Mathematical Theory of Learning Transformational Grammar', *Journal of Mathematical Psychology* **12**, 137–177.

Hornstein, N. and D. Lightfoot (eds.): 1981, *Explanation in Linguistics: The Logical Problem of Language Acquisition*, Longman, London.

Hume, D.: 1978, *A Treatise of Human Nature*, 2nd edn., L. A. Selby-Bigge and P. H. Nidditch (eds.), Clarendon Press, Oxford.
Lasnik, H.: 1989, 'On Certain Substitutes for Negative Data', in Demopoulos and Marras, pp. 89–105.
Lightfoot, D.: 1982, *The Language Lottery: Towards a Biology of Grammars*, MIT Press, Cambridge, Massachusetts.
Lightfoot, D.: 1989, 'The Child's Trigger-Experience: Degree-0 Learnability', in *Behavioral and Brain Sciences* **12**, 321–334.
Lightfoot, D.: 1991, *How to Set Parameters: Arguments from Language Change*, Bradford/MIT Press, Cambridge, Massachusetts.
McClelland, J. L. and D. E. Rumelhart.: 1986, 'On Learning the Past Tenses of English Verbs', in J. L. McClelland, D. E. Rumelhart et al. (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 2, MIT Press/Bradford Books, Cambridge, Massachusetts, pp. 216–270.
Newport, E., Gleitman, H. and Gleitman, L.: 1977, 'Mother, Please, I'd Rather Do It Myself: Some Effects and Non-Effects of Maternal Speech Style', in C. Snow and C. Ferguson (eds.), *Talking to Children: Language Input and Acquisition*, Cambridge University Press, New York, pp.109–150.
Pinker, S.: 1979, 'Formal Models of Language Learning', *Cognition* **7**, 217–282.
Pinker, S.: 1984, *Language Learnability and Language Development*, Harvard University Press, Cambridge, Massachusetts.
Pinker, S.: 1986, 'Productivity and Conservatism in Language Acquisition', in Demopoulos and Marras, pp. 54–79.
Pinker, S.: 1989, *Learnability and Cognition: The Acquisition of Argument Structure*, Bradford/MIT Press, Cambridge, Massachusetts.
Pinker, S.: 1991, 'Language Acquisition', in D. N. Osherson and H. Lasnik (eds.), *Language, An Invitation to Cognitive Science*, Vol. 1, Bradford Books/MIT Press, Cambridge, Massachusetts.
Putnam, H.: 1971, 'The "Innateness Hypothesis" and Explanatory Models in Linguistics', in J. Searle (ed.), *The Philosophy of Language*, Oxford University Press, London.
Quine, W. V. O.: 1953, 'Two Dogmas of Empiricism', in *From a Logical Point of View*, Harvard University Press, Cambridge Massachusetts.
Sampson, G.: 1953, 'Language Acquisition: Growth or Learning?', *Philosophical Papers* **XLVII**, 203–240.
Sells, P.: 1985, *Lectures on Contemporary Syntactic Theories: An Introduction to Government-Binding Theory, Generalized Phrase Structure Grammar and Lexical-Functional Grammar*, CSLI, Stanford University, Stanford, California.
Soames, S.: 1984, 'Linguistics and Psychology', in *Linguistics and Philosophy* **7**, 155–179.
Wexler, K., P. W. Culicover and K. Hamburger: 1975, 'Learning-Theoretic Foundations of Linguistic Universals', in *Theoretical Linguistics* **2**, 215–253.
Wexler, K. and P. W. Culicover: 1980, *Formal Principles of Language Acquisition*, MIT Press, Cambridge, Massachusetts.

California Institute of Technology
Division of the Humanities and Social Sciences
Pasadena, CA 91125
U.S.A.