The 'No Negative Evidence' Problem: How Do Children Avoid Constructing an Overly General Grammar?

Melissa Bowerman

The story of language acquisition is, to a large extent, the story of how children make linguistic generalizations. Traditionally, the emphasis in describing this process has been on how learners go beyond the specific utterances they hear to draw out regularities that will enable them to produce and understand an infinite number of novel sentences. However, in the last few years the converse question has also come under intense scrutiny: how do children avoid generalizing too broadly, ending up with grammars that not only generate all the well-formed constructions of their language but a number of ungrammatical ones as well?

The problem of how children avoid constructing an overly general grammar was first posed in 1971 by Martin Braine. Braine used the problem to argue against the nativist position set forth by Chomsky (1965) and in favor of the idea that language is learned largely from scratch. It was later revived by Baker (1979), who, in an interesting turn-around, made it the cornerstone of the argument that children must be guided by innate constraints in their acquisition of language. That both nativist and empiricist theorists have been able to adapt the puzzle to their own use indicates that the problem transcends party lines. Indeed, I believe it constitutes one of the most intriguing and difficult challenges for all students of language acquisition.

By now, a range of solutions has been proposed for the problem. My goal in this chapter is not to advance still another hypothesis but rather to take stock of where we now stand in our efforts to crack the puzzle. My focus is on how children learn rules that have lexical exceptions, since these constitute a core learnability problem to which different theoretical approaches have offered very different solutions.

1 Statement of the Problem

According to Chomsky's (1965) account of language acquisition, the child's task is to construct an internal grammar. The learner does this by

using incoming language data, together with innate linguistic knowledge, to formulate hypotheses about possible grammatical rules. He then tests these hypotheses against further data, discards or revises them as necessary, and eventually applies an evaluation metric to surviving candidate grammars in order to select the best one.

Braine (1971b) observed that there is a critical discrepancy between the kind of data needed for a hypothesis testing procedure to work and the kind of data children actually receive. Hypothesis testing requires feedback about the correctness of predictions, pointed out Braine. In particular, it requires evidence not only about what *is* an instance of what is to be learned (an acceptable sentence of the language, in this case), but also about what is *not* an instance. Negative evidence is essential for the learner to revise hypothesized rules that are overly general, since these rules will generate all acceptable instances (sentences) and err only in that they generate unacceptable instances as well (sentences that fluent speakers find ill-formed).

Language is full of partial regularities that might suggest attractive but ultimately overinclusive hypotheses to children. Consider, for example, the following sentences (adapted from Baker 1979):

(1) a. Dad told a story to Sue.
    b. Dad told Sue a story.
(2) a. I gave a book to John.
    b. I gave John a book.
(3) a. Jim showed the model to Bob.
    b. Jim showed Bob the model.
(4) a. Mom baked a cake for Jack.
    b. Mom baked Jack a cake.

A child exposed to sentences like these might come to see a relationship between the (a) and (b) versions. He might conclude that, given a sentence of the (a) form, he can convert it into the (b) form (corresponding to various linguists' proposed rule(s) for an optional Dative Movement transformation). The child would accordingly produce many well-formed novel (b) sentences such as *I sent my cousin a birthday present* and *Linda knitted me a sweater*. But he would also produce sentences that fluent speakers find unacceptable, for example the (b) forms below:

(5) a. Dad said something nice to Sue.
    b. *Dad said Sue something nice.
(6) a. I reported the death to the police.
    b. *I reported the police the death.
(7) a. Mary donated a book to the library.
    b. *Mary donated the library a book.

(8) a. Jim demonstrated the model to Bob.
    b. *Jim demonstrated Bob the model.
(9) a. Mom buttoned the coat for Jack.
    b. *Mom buttoned Jack the coat.

If the child is testing a hypothesized rule of Dative Movement, how does he learn that *say, demonstrate*, etc., are exceptions to the rule? This would seem to require someone telling him that sentences like those above are not acceptable; otherwise he would have no reason to alter his hypothesis. But do children in fact get information of this type? After reviewing the available data (e.g. Brown and Hanlon 1970), Braine concluded that negative evidence is rare in the input to children; moreover, children appear to be relatively impervious to what little correction they do receive. In consequence, argued Braine, the hypothesis testing approach cannot be correct. The child must acquire language with procedures for which positive evidence alone – that is, exposure to sentences of the language – is sufficient. The procedures proposed by Braine will be considered in a later section of this chapter.

2 A Nativist Response

Although nativists did not respond immediately to Braine's challenge, the problem was eventually taken up and refocused in an important article by Baker (1979). Baker agreed that children get no reliable and systematic evidence about what is not a sentence and that they must be prepared to learn from positive evidence only. He argued, however, that this situation is not damning to the innatist program in general, but only to grammatical frameworks that allow types of rules that children could not acquire from positive evidence alone. He also observed that the 'no negative evidence' problem – as the puzzle has gradually come to be called – is a challenge not only to nativist theorizing, as argued by Braine, but also to empiricist approaches to language acquisition that invoke learning by generalization. Sentences (5a)–(9a) are 'similar' in many important respects to sentences (1a)–(4a). Why then do fluent speakers balk at (5b)–(9b), given that (1b)–(4b) are acceptable? What blocks this generalization?

2.1 *Benign versus embarrassing exceptions*

Baker pointed out that not all overgeneralizations a child might make are troublesome for a theory of language acquisition. He distinguished between rule exceptions that are 'benign', in that errors involving them can in principle be corrected without negative evidence, and those that are 'embarrassing', because their correction seems logically to depend on evidence of a type that children do not regularly receive.

Benign exceptions are irregular forms in an inflectional or other

paradigm. If children overgeneralize a rule of (say) past tense or plural marking, producing forms like *runned* or *foots*, positive evidence alone – in this case consistent exposure to the irregular adult forms *ran*, *feet*, etc. – is in principle sufficient to show them that their form is wrong.

For an embarrassing exception, however, there is *no exact, irregular counterpart*, and therefore no positive evidence that the form is an exception to the rule. Consider dative alternation, for example. The child who on a particular occasion considers saying sentence (5b) might hear sentence (5a) instead. But this cannot be taken as evidence that (5b) is not possible, any more than hearing (1a) counts as evidence against the wellformedness of (1b). Sentences of both kinds can co-exist.

In general, observed Baker, benign exceptions involve rules proposed to be *obligatory*, whereas embarrassing exceptions involve rules characterized as *optional*. He emphasized that the 'no negative evidence' problem affects a large proportion of the rules proposed in classical transformation theory.

2.2 *Baker's solution: innate constraints on the child*

Baker took the 'no negative evidence' problem as strong evidence that children must be *innately constrained* in their acquisition of language, and a number of other theorists have since followed his lead (for example, see chapters in Baker and McCarthy 1981). If certain hypotheses about language structure are incorrect, but, once made, could not be corrected except by evidence of a type that is not available in the learner's input, then there must be something that keeps children from making them in the first place. Different investigators have suggested different types of constraints, depending in part on their theoretical orientation and also on the domain of grammar they are concerned with.

Baker's proposal was to limit the grammatical rules allowed by linguistic theory to those that can be learned by children on the basis of positive evidence alone. Constraints on permissible grammatical rules can also, of course, be phrased in terms of constraints on the child – that is, children are seen as so constituted that they will not entertain, as a possible grammatical rule, any rule which, if incorrect, could not subsequently be corrected on the basis of positive evidence alone.

According to Baker, the appropriate constraints on grammatical rules were offered by the *lexicalist* approach to linguistic theory, which was just beginning to attract widespread attention at the time he wrote (see also Roeper et al. 1981). Lexicalists called for reducing the power of the syntactic component of grammar in favor of a richer lexicon. For example, they argued that regularities with lexical exceptions (and certain other properties) should not be treated as syntactic transformations. Instead, the partial regularities that transformations like Dative Movement and Passive had been designed to capture should be represented within the lexicon.

A special device for doing this is the lexical redundancy rule, proposed by Jackendoff (1975). This kind of rule does not derive one construction from another, but instead simply captures a passive generalization across lexical items that are already fully specified in the lexicon. For example, a redundancy rule representing the relationship between 'shifted' and 'non-shifted' datives would state that corresponding to the subcategorization frame $[\_ NP_1 \text{ to } NP_2]$ is the frame $[\_ NP_2 NP_1]$. Items to which the rule applies are marked accordingly, and items to which it does not apply are left unmarked.

If there are no general syntactic rules relating shifted and non-shifted datives, actives and passives, etc., then children do not have to learn exceptions to the rules. Instead, they proceed conservatively. Predicate by predicate, they wait for positive evidence that the form can appear in a given syntactic context. They will notice that *give*, for example, appears in both the subcategorization frame $[\_ NP_1 \text{ to } NP_2]$ and the frame $[\_ NP_2 NP_1]$. But they will observe *say* only in the frame $[\_ NP_1 \text{ to } NP_2]$. If learners never encounter a form in certain contexts, they will not make the corresponding entry in their mental lexicons. Their grammars will therefore never contain overly general rules with embarrassing lexical exceptions.

3 A Second Nativist Proposal: the Subset Principle

Before evaluating Baker's proposal, let us take a quick look at a second, qualitatively different approach to constraining the child from the outset. Recall that Baker called for restricting the types of rules children can entertain. Other investigators have appealed instead to a built-in order for generating hypotheses: the focus is not on what rules a child can conceive of, but on which rule she considers first when more than one rule is conceivable.

This approach, proposed initially by Dell (1981) and formalized by Berwick (1985; Berwick and Weinberg 1984) under the name 'the Subset Principle', exploits the fact that the grammars generated or sanctioned by candidate grammatical rules (or constraints, or parameters) may stand in a subset–superset relation to each other; that is, the grammar generated by one rule may be a proper subset of the grammar generated by the other rule. The proposal is that since learning must proceed from positive instances only, children must first hypothesize the *narrowest possible grammar* compatible with the evidence observed so far. If the grammar is too narrow, the learner will eventually discover this through positive evidence (that is, sentences in the input that the grammar does not account for). She will then posit the next larger grammar compatible with the data as she now perceives it. The learner must never hypothesize a grammar that is too general, since without negative evidence there is no way to cut back to the correct, narrower grammar.

Berwick has applied the Subset Principle to a wide variety of problems in grammar construction, including rules with lexical exceptions like Dative and Passive. See also Wexler and Manzini (1987) and Smith (1981) for related applications to parameter-setting and word formation, respectively.

4 Overproductivity: a Problem for both 'Innate Constraints' Approaches

Several questions can be raised on purely theoretical grounds about the adequacy of both the lexical redundancy rule approach and the Subset Principle as solutions to the 'no negative evidence' problem.[1] There is, however, a compelling empirical reason for questioning them as well: children *do* overgeneralize rules with 'embarrassing' exceptions, producing errors that should not occur if these approaches were correct.

At the time Baker wrote, evidence concerning such errors was sparse, so the hypothesis that children simply do not make them was perhaps not unreasonable. However, as linguists have increasingly based their theorizing on the assumption that children do not formulate overly general grammars, and so must be innately constrained, child language scholars have begun to emphasize that the critical errors do occur. New experimental work has also confirmed that children generalize more broadly than Baker, Berwick and others have supposed.

The major systematic body of evidence about the relevant errors in children's spontaneous speech comes from my diary records of my two English-speaking daughters, which I have supplemented with comparable examples from other children. The errors include overgeneralizations of dative alternation, causative verb formation, passivization, locative alternation (*spray/load*-type verbs) and *un*-prefixation (Bowerman 1974, 1982a, 1982b, 1983). Some examples are shown in table 4.1.

Spontaneous speech data are backed up by experimental evidence that children often judge ungrammatical 'shifted' datives as acceptable (Mazurkewich and White 1984; these investigators also list a few dative errors from children's spontaneous speech). Similar results for unacceptable causatives are presented by Hochberg (1986). There is also experimental evidence that children are able to create never-heard passives, shifted datives and causative forms for novel verbs (Pinker et al. 1987; Pinker 1987; Maratsos et al. 1987), which goes counter to Baker's proposal that children are conservative, waiting for positive evidence before listing a new subcategorization frame for a verb in their mental lexicon.

How damaging is such evidence to the program of solving the 'no negative evidence' problem by innately constraining children in such a way that their grammars never become overly general? Before drawing conclusions, let us consider two possible lines of counterargument.

Table 4.1 Some overgeneralizations of rules with lexical exceptions[a]

A			Dative Alternation
	1	C	3; 1	I *said her* no.
	2	C	2; 6	Don't *say me* that or you'll make me cry.
	3	L	7; 8	Shall I *whisper you* something?
	4	C	2; 6	I want Daddy *choose me* what to have. (Re: what kind of juice to have at breakfast.)
	5	M	5+	*Choose me* the ones that I can have.
	6	C	3; 4	*Button me* the rest. (Request to have remaining snaps on her pyjamas fastened.)
	7	—	6; 0	Mommy, *open Hadwen* the door. (Mazurkewich and White 1984)
	8	—	2; 3	I'll *brush him* his hair. (Mazurkewich and White 1984)

B Lexical Causatives (The regularity to which the verbs below are exceptions is exemplified by intransitive/transitive pairs like *The stick broke / I broke the stick*.)

	9	J	6+	Do you want to see us *disappear* our heads? (Then, with a friend, she ducks down behind couch.)
	10	—	2; 8+	I don't want any more grapes; they just *cough* me. (Braine 1981a)
	11	R	5; 9	I want to *comfortable* you. (R lying on sofa with mother, cuddling her.)
	12	E	3; 0	Don't *giggle* me. (As father tickles her.)
	13	E	3; 2	Will you *climb* me up there and hold me? (Wants mother to help her climb a pole.)
	14	E	3; 7	I'm gonna put the washrag in and *disappear* something under the washrag. (Playing in tub with small toys and a container into which she puts washrag.)
	15	C	3; 6	Did she *bleed* it? (After her sister falls and hits head on edge of table.)
	16	C	4; 3	It always *sweats* me. (Refusing sweater.)
	17	M	5; 8	*M*: These are nice beds. *Mother*: Yes, they are. *M*: Enough to *wish* me that I had one of those beds.

C Passive

	18	E	3; 8	Both are going to be *go-ened* in! (= gone in. Watching one child sit down on potty and another on toilet.)
	19	C	3; 6	Until I'm four I don't have to be *gone*. (= be taken to the dentist.)
	20	C	3; 6	If you don't put them in for a very long time they won't get *staled*. (Reference to crackers in a bread box. Passive of novel causative; see B above.)
	21	C	4; 3	Why is the laundry place *stayed* open all night? (= kept.)

Table 4.1 (*cont.*)

22	H	4; +	He's gonna die you, David. (Turns to mother.) The tiger will come and eat David and then he will be *died* and I won't have a brother any more.
23	E	5; 2	Mommy will get *lightninged*. (= struck by lightning)
24	C	5; 6	I don't want to be *dogeared* today. (Asking for her hair not to be arranged in 'dogears'.)
25	C	8; 9	A child wanted her doll to be *mummied*. (= made into a mummy; mummified)

D   Locative Alternation (The regularity here is exemplified by sentence pairs like these:

                a   b
. . . spray paint on the wall / spray the wall with paint.
. . . load hay into the wagon / load the wagon with hay.
. . . spread butter on the bread / spread the bread with butter.

Hall (1965) proposed deriving the *b* forms from the *a* forms with an optional transformation for verbs of a certain semantic class. Certain verbs in this class are exceptions, however, in that they are 'fixed' in either the a-pattern (e.g. *pour, spill, steal*) or the b-pattern (e.g. *fill, cover, rob*).

b-verb in a-pattern:

26	E	4; 5	I'm gonna *cover* a screen over me.
27	E	5; 0	Can I *fill* some salt into the bear? ('the bear' = a bear-shaped salt shaker.)
28	C	4; 9	She's gonna *pinch* it on my foot. (Protesting as E approaches with a toy.)

a-verb in b-pattern:

29	E	2; 11	(Waving empty container near mother.) *E*: Pour, pour, pour. Mommy, I *poured* you. *M*: You poured me? *E*: Yeah, with water.
30	E	4; 11	(Mother asks if E is going to finish toast.) I don't want it because I *spilled* it of orange juice.
31	C	6; 5	(Telling of TV episode.) *C*: Once the Partridge Family got *stolen*. *M*: The whole family? *C*: No, all their stuff.

E   Reversative *un*-prefixation.

32	E	3; 11	How do you *unsqueeze* it? (Coming to mother with clip earring dangling from ear; wants it off.)
33	E	3; 10	*Mother*: I have to capture you. (Grabbing E in a game.) *E*: *Uncapture* me!

Table 4.1 *(cont.)*

34	C	4; 7	*C*:	I hate you! And I'm never going to *unhate* you or nothing! (Angry after request is denied.)
			Mother:	You're never going to unhate me?
			C:	I'll never like you.
35	C	4; 5		(C has asked mother why pliers are on table.)
			Mother:	I've been using them for straightening the wire.
			C:	And *unstraighting* it?
36	C	5; 1		He tippitoed to the graveyard and *unburied* her. (Telling ghost story.)
37	C	7; 11		I'm gonna *unhang* it. (Taking stocking down from fireplace.)

[a] Child's age is given in years; months. Sources as indicated, plus Bowerman 1983, for dative alternation and passives (I am grateful to Eric Kellerman for example 3); 1982a, 1982b, 1983 for lexical causatives; Bowerman 1982b for locative alternation and reversative *un-* prefixation.

4.1 *Limited productivity for lexical redundancy rules?*

Lexicalists who have not been specifically concerned with the 'no negative evidence' problem have suggested that although lexical redundancy rules are typically passive (that is, not called on in sentence production or comprehension), speakers might occasionally use them productively (e.g. Jackendoff 1975). Tolerance for productivity would appear to be necessary to accommodate the evidence that *adult* speakers sometimes produce passives, shifted datives, etc., with novel verbs. Perhaps, then, examples like those in table 4.1 should be dismissed under an escape clause that allows speakers an occasional burst of creativity with what is basically a passive generalization over existing lexical entries.

This defense is untenable. First, only a small portion of children's errors such as those shown in table 4.1 involves novel verbs whose properties are in principle open to creative negotiation. For the most part they are everyday verbs whose syntactic and morphological privileges are well known to fluent speakers. We are still left, then, with the problem of how children identify them as exceptions to a particular rule. More generally, the very existence of a 'creativity' escape clause destroys the advantage claimed in the first place for the lexicalist approach as a solution to the 'no negative evidence' problem. Once the door is opened a crack to productivity, how does the child know where to stop? (See also Wasow 1981 on this point.)

4.2 *Errors as analogies, not rule-governed productions*

A second way to deal with errors like those in table 4.1 within the innate constraints framework is to argue that they do not actually reflect rule use. Instead, they are on-the-spot analogies in which the child momentarily 'borrows' the legitimate syntactic or morphological privileges of a verb that is semantically similar to the one in question (see Maratsos 1979: 335 for suggestions along this line). If this is correct, there is no problem for either the lexicalist solution or the Subset Principle. The child really *does* know the syntactic or morphological privileges of the form in question; his grammar is not overinclusive.

It seems possible that some of children's errors are passing analogies rather than the output of established rules. But is this true in all cases of errors involving 'embarrassing' lexical exceptions? The 'analogy' interpretation for errors of a given genre seems most plausible when the errors are very rare, and there are semantically closely related forms to serve as models for the analogy. Errors of this type in my data include incorrect choices of verb complement, for example *Christy insisted me to make a house* (= insisted that I should make . . .; E age 7; 3; cf. *forced, persuaded, told*). But the 'passing analogy' interpretation seems strained for errors that occur more frequently and involve many different lexical items, many of which lack close semantic neighbors that undergo the rule. In my data these include novel causatives, locatives and passives. Errors with shifted datives are not frequent in my data, but ill-formed shifted datives were ofted judged grammatical by Mazurkewich and White's (1984) subjects, which is difficult to explain by reference to passing analogies.

If, as the data strongly indicate, at least some genres of children's overgeneralizations involving 'embarrassing' lexical exceptions are rule-governed and not mere analogies, then theoretical proposals for constraints that will prevent such overgeneralizations from ever occurring are on the wrong track. Instead, we must shift our attention to the procedures children use for cutting back on overgeneral rules.

5 The Criteria Approach: Conditions on Rules

An interesting hypothesis about these procedures has recently been proposed by Mazurkewich and White (1984) for dative alternations, and expanded by Pinker (1984: ch. 8; 1987) to cover causatives, passives, locative alternations and *un-* prefixation as well. This approach is based on the assumption that rules with lexical exceptions do not apply to an arbitrary set of verbs, but rather to verbs of a coherent class, characterized by shared semantic, morphological and/or phonological criteria (hence the label 'Criteria' approach; Pinker 1987). Children may initially over-generalize, but they eventually identify the criteria that define the lexical

class appropriate to the rule. When they do, they limit productive use of the rule to lexical items of the right class, and errors cease.

For dative alternation, for example, the child learns that 'shiftable' indirect objects must be 'prospective possessors' of the entity named by the direct object and that the verb itself must be of 'native [Germanic] stock', not Latinate (Mazurkewich and White 1984; Pinker 1984; these authors draw on Goldsmith 1980, Green 1974, Oehrle 1976, and Stowell 1981). For causativization of an intransitive verb or adjective, the causation must be direct (for example, physical), and the agent, manner and goal of causation must be stereotypic or conventional for the act in question (Pinker 1984, drawing on Shibatani 1976, Gergely and Bever 1986 and others). And for reversative *un*-prefixation, a 'covert semantic class' identified by Whorf (see Bowerman 1982b) is relevant: verbs that can be *un*-ed share 'a covering, enclosing, and surface-attaching meaning ... hence we say "uncover, uncoil, undress, unfasten, unlock, unroll, untangle, untie, unwind", but not "unbreak, undry, unhang, unheat, unlift, unmelt, unopen, unpress, unspill"' (Whorf 1956: 72).

This approach has some similarity to the Subset Principle, since both strategies appeal to subset–superset relations among possible rules of grammar. However, it differs importantly in that children are not expected to identify or automatically prefer, from the beginning, the grammatical rule that generates the smallest grammar; instead, they must work to identify the criteria that define the target subsets. Although the Criteria approach clearly invokes learning in a way that the Subset Principle does not, the approach is nevertheless compatible with – and may well require – some strong nativist corollary assumptions about how children identify the conditions on rules.[2]

The Criteria approach has the clear advantage over both the 'lexical redundancy rule' strategy and the Subset Principle that it can accommodate the overgeneralizations that children actually make. And there is evidence that children are indeed capable of restricting a rule that is initially overly general to verbs of the 'right' semantic class (see Bowerman 1982b on *un*- prefixation). Nevertheless, I am doubtful that it is the right answer to the 'no negative evidence' problem.

One difficulty is that it is hard to understand why children would go to so much work to identify conditions on their rules (Fodor and Crain 1987; Randall, 1987). Why should they bother to hypothesize possible constraints and check whether these constraints correctly predict the adult words that undergo a rule, especially given that, with their overly general rule, they can already understand all adult utterances of the relevant form? And how would they know which conditions to consider in connection with which rules?

A second difficulty, discussed immediately below, is that the subsets proposed for rules with lexical exceptions are themselves dotted with gaps: items that fully conform to the semantic/morphological conditions on the rule, but that still do not undergo the rule.

5.1 *Gaps in the subsets*

For dative alternation, such an item is *choose* (see examples 4–5 in table 4.1). Many speakers find something distinctively odd about 'shifted' indirect objects with *choose*, as in *I chose you a book at the library sale*.[3] Yet *choose* satisfies the putative semantic/morphological restrictions on verbs that allow dative alternation, since (i) it is of native stock and (ii) the beneficiary of an act of choosing is the 'prospective possessor' of the object named by the direct object. (Many speakers who regard 'shifted' indirect objects with *choose* as unacceptable find them perfectly normal with *pick out*, which is semantically almost identical to *choose*.)

For lexical causatives there are gaps in English like 'to **cough/*laugh/ *comfortable/*vomit* someone'. There seems to be no principled reason why a too-big bite can *choke* or *gag* us but not *cough* us (see example 10 in table 4.1), why we can *cheer* someone *up* but not *laugh* or *giggle* her (for example, with tickling or a joke) (example 12) and why we can *quiet* or *burp* a baby but neither *comfortable* (or *comfy*) her when she is uncomfortable (example 11) nor *vomit* her when she is nauseated (for example, with a finger). The acts of causation specified by the latter verbs do not seem any less 'direct' or 'conventional' than those specified by the former.

For *un-* prefixation, the verb *squeeze* is an inexplicable exception (example 32). You can squeeze somebody's hand but you can't **unsqueeze* it, even though *squeeze* falls into Whorf's covert class of 'centripetal' verbs. (Note, for example, that you can both *clench* and *unclench* your teeth or fists; *squeeze* is similar to *clench* in specifying a continuous pressure toward a center point.)

Gaps in a target subset – let's call them 'negative exceptions' – constitute a grave challenge to the Criteria approach. The approach can tolerate exceptions, but only if they are positive – that is, items to which a rule applies even though they do NOT belong to the target class. These latter exceptions are generally regarded as tractable because in principle they could be learned on the basis of exposure to positive evidence. For example, children could learn that *assign* allows dative alternation – even though it is Latinate – through hearing sentences like *The teacher assigned John a desk in the back row*. However, there is no comparable evidence to mark *squeeze*, *choose*, *cough*, etc., as exceptions that *cannot* undergo rules whose conditions they otherwise satisfy. They are invisible holes within their subsets. How do children identify such holes?

5.2 *Partitioning subsets more finely*

In very recent work, Pinker (1987) argues that where there have appeared to be negative exceptions to target subsets (as argued above and in Bowerman 1987), the semantic or other constraints on the items that can undergo a particular operation have not yet been adequately pinned

down. Since it is often difficult or impossible to identify a single subclass per rule that satisfactorily includes all and only the desired lexical items, Pinker now proposes that for each rule there may be *sets* of coherent subclasses of items that do and do not undergo the specified operation. In other words, if we slice up the lexicon more finely, alleged 'negative exceptions' will disappear.

Drawing on work by Levin (1985) and Rappoport and Levin (1985), Pinker outlines a number of candidate subclasses. For example, he proposes that dative alternation applies to 'verbs of giving' (*give, pass, send* ...), 'verbs of future having' (*offer, promise, bequeath* ...), 'verbs of successful communication' (*tell, show* ...) and a couple of other subclasses. However, it does *not* apply to 'verbs of choosing' (*choose, select, pick* ...), 'verbs of manner of speaking ' (*shout, whisper* ...) and a few other subclasses that were encompassed by the class previously specified as relevant for dative alternation in Pinker (1984).

Similarly, causativization applies only when the resulting transitive verb specifies a 'cause of a change of physical state' (*open, melt, shatter* ...), 'cause of motion in a particular manner' (*float, roll, bounce* ...), 'coerced or encouraged locomotion in a particular manner' (*walk, gallop, run, jump* [your horse] ...) and 'enabling and accompanying of willful transportation in some manner' (*fly, boat, motor* [someone to New York] ...). It does *not* apply to verbs specifying 'motion in a direction' (cf. **I went my son to school; *I rose the flag*) and verbs of 'internally caused acts' (**Bill vomited Jill, *Fred laughed the baby, *John died Harry*), although there may be suppletive forms with the same meaning (*I took my son to school, I raised the flag, John killed Harry*), and some verbs are 'positive exceptions', in that they causativize even though they belong to a class whose members in general do not (*Mom burped the baby*, 'internally caused act'; *Mary dropped the ball*, 'motion in a direction').

Do these changes strengthen the viability of the Criteria solution to the 'no negative evidence' problem for rules with lexical exceptions? Although the fit of the subclasses to the data appears at first glance to be better, I believe that this is somewhat illusory: repartitioning has simply redistributed problems to places where they seem superficially to do less damage; it has not eliminated them.

Recall that the motivation behind partitioning lexical items into smaller semantic classes was to get rid of alleged negative exceptions like (**'choose* + indirect object' and causative **cough*, since the original version of the Criteria approach was not able to account for how a child could identify them. These items have now been excised from the original target subclasses in a principled way, by being declared members of semantically coherent subclasses of their own, none of whose members undergo the rule (except as positive exceptions). These 'negative' subclasses are in turn used in defining the boundaries of the new, multiple 'positive' subclasses to which the rule *does* apply.

For example, 'verbs of locomotion in a particular manner' can causativize, unless they also happen to be members of the non-causativizing class of 'verbs of motion in a direction'; this accounts for why *climb*, apparently a perfectly good verb of locomotion, does not causativize. *The riders walked/galloped/*climbed their horses up the hill* (Pinker, personal communication). Similarly, 'verbs of future having' can undergo dative alternation, but this class must be construed so as to omit 'verbs of choosing', even though the one for whom something is chosen could reasonably be viewed as a 'future haver'.

Now, the problem for a plausible theory of language acquisition is to explain how children could learn the critical negative subclasses that serve to define and delimit the positive subclasses. I will argue that the cost of eliminating negative exceptions in the way proposed is to so increase the number of new *positive exceptions*, relative to the size of the now smaller subclasses, that many of the necessary generalizations about what classes do not undergo a rule are likely to be unlearnable.

An example will illustrate the problem. As noted above, in Pinker's (1987) version of the Criteria approach the child must learn that 'verbs of motion in a direction' do not causativize. This will account for the cessation of causative errors with inherently directional verbs like *climb*, *go* and *rise*. But could a child in fact learn this? There are many positive exceptions to the generalization – that is, many 'verbs of motion in a direction' that DO causativize: for example, *drop*, *lower*, *sink*, *topple*, *recline*, *dip*, *tilt* (direction of motion 'down'); drain ('out and down'), *exude* ('out'); *lift*, *levitate* ('up'); *transfer* ('across' [Latin *trans*-]); *withdraw* ('away'); *filter*, *percolate* ('through'), *retract* ('back'). Most of the other critical negative subsets outlined by Pinker also have positive exceptions.

In principle, of course, positive exceptions to negative subsets are learnable through positive evidence. The problem, though, is that there must not be so many positive exceptions that the boundary between items that can and can't undergo the rule is hopelessly blurred. The threat posed by positive exceptions is a function of their number relative to the number of items correctly predicted by the criteria for distinguishing positive and negative classes. If they are few, and the number of correctly predicted items is large, a robust learning system can still make the generalization. But as the number of exceptions rises or the size of the classes falls, the generalization becomes steadily weaker.

In the earlier version of the Criteria model, in which there was a single critical subclass for rule application and a residual negative class of items to which the rule did not apply, the classes were large enough that some positive exceptions could be tolerated. But in the present version, the single split has been replaced by *multiple subclasses*, both positive and negative; this means that the size of the classes against which any positive exceptions must be weighed has dropped sharply. Consequently, even a few positive exceptions can so severely pollute the generalization that it is unlikely a child would ever consider it.

In sum, efforts to make the Criteria approach work as a solution to the 'no negative evidence' problem have veered from the Scylla of unlearnable negative exceptions, as argued in the last subsection, to the Charybdis of too many positive exceptions, as just discussed. Is it possible that the approach can be strengthened by still further attempts at partitioning aimed at reducing the number of positive exceptions to tolerable levels? I think this is unlikely.

I do not question the claim that the sets of words that undergo rules like dative alternation or causativization are not arbitrary; there are clearly correlations, often very strong, between certain meanings and eligibility for certain rules. But is the nature of language such that there *necessarily* are systematic principles governing which words do and do not undergo a particular rule? The Criteria approach insists that this is the case. Wasow (1977), however, presents a different view:

> I assume that if a tree satisfying the structural conditions of a transformation is prohibited from undergoing the structural change, then some explanation is called for, but *a lexical item that does not undergo a lexical rule whose conditions it satisfies is perfectly normal*. Transformations are crucial to the generation of all and only the sentences of the language (and hence have infinite domains); in contrast, lexical rules express subregularities within a finite lexicon ... Hence, I assume (following Jackendoff, 1975) that *lexical rules ... will typically have unsystematic exceptions*. (p. 331, emphasis added)

The Criteria approach to the 'no negative evidence' problem recognizes and can tolerate only *one* kind of unsystematicity in lexical rules: positive exceptions. But there seems to be no principled reason to rule out the opposite kind of unsystematicity: negative exceptions. If Wasow's view is correct, as I believe it is, then drawing and redrawing subset boundaries to get rid of negative exceptions is a futile exercise. Negative exceptions may be just as characteristic of the lexicon as positive exceptions. If so, children must have techniques for identifying them. And if they do have such techniques, they could presumably apply them directly to *all* lexical items to which a rule has been overgeneralized, rather than first eliminating some items through the discovery of appropriate subclasses before bringing in stronger methods to detect any remaining stragglers.

In summary, it is not clear that children's identification of lexical subclasses plays any necessary role in their retreat from an overly general grammar, although the process of discovering correlations between lexical rules and the meanings of the lexical items that typically undergo them is an interesting aspect of language acquisition in its own right (see Bowerman 1982b).

6 Principles of Universal Grammar

One question about the Criteria approach to the 'no negative evidence' problem, as mentioned earlier, is whether it is plausible that children would go to all the effort of identifying conditions on a rule, given that their overly general version of the rule accounts perfectly well for all the utterances they hear and allows them to create new ones. Some theorists have argued that we need stronger principles than the Criteria approach provides to explain what would cause a child to reject a successful but overly general rule – in particular, principles of Universal Grammar (Grimshaw 1987; Randall 1987).

Principles of Universal Grammar are assumed, by theorists who work within this framework, to be built into the human capacity for language. Thus, if a child's rule should be incompatible with such a principle, this would trigger automatic rejection. But, we may ask, if the child had a principle that would eventually lead to rejection of a rule, why would this principle not operate from the beginning to prevent the rule from ever being constructed? The answer is that the child may not at first 'realize' that the principle applies to the forms on which he errs. When he discovers that it does, errors will cease.

There is a danger to this logic, as noted by Matthews (1983): it can be used to reconcile the 'innate constraints' approach with any data, no matter how uncooperative. To avoid being totally *ad hoc* in appealing to a particular principle to explain the cessation of certain errors, an investigator must at least propose a plausible account of why the child does not at first realize that the principle applies, and of what triggers the realization that it does apply. An intriguing proposal that attempts to satisfy these basic requirements has been made by Randall (1985, 1987) to explain how children eventually restrict dative alternation to verbs of the right class.

The proposal comprises two key ingredients. First, Randall argues that verbs that do and do not undergo dative alternation differ in their basic argument structure: verbs that take the double-object form (the *give*-class) have three mandatory arguments, whereas those that do not take double objects have only two (*donate*-class), although they can optionally take a third, the dative object. This difference in argument structure can be seen from differences in the way verbs of the two types behave in a variety of syntactic contexts; most saliently, it is perfectly normal for *donate*-class verbs to appear without a dative argument, whereas this is more marked for *give*-class verbs, requiring special contexts where the missing dative argument can be inferred.

The second critical ingredient of Randall's proposal is the claim that children are innately equipped with a principle central to X-bar theory (Jackendoff 1977), which dictates that obligatory elements are attached inside of optional elements, relative to the phrasal head and its projec-

tions. This means that optional arguments cannot precede obligatory arguments in the predicate of an English sentence: *Pablo invited (*to the art opening) Doris*; *Dylan spent (*on drink) a lot of money*. Notice that this principle rules out shifted datives with *donate*-class verbs, according to Randall, since the dative argument of these verbs, unlike that of *give*-class verbs, is optional, and so must appear *after* the obligatory argument.

What children are lacking at first, argues Randall, is an understanding of the basic argument structure of *donate*-class verbs. They assume that *both* the direct object *and* the dative object arguments are obligatory; hence, the double-object construction that reverses the order of these two arguments is permissible. Once they discover that the dative object is optional for these verbs – possibly through observing that the verbs can occur without dative objects in unmarked discourse contexts – the principle that obligatory arguments attach inside optional elements serves to eliminate errors.

This principle is so strong that it will also rule out certain double-object constructions that English in fact allows, in particular those with 'beneficiary' arguments such as *Mom baked Jack a cake*. *Bake*, *knit*, etc., do not subcategorize for a beneficiary. Why then do these verbs allow the double-object construction? Randall (1987 and personal communication) proposes that there is a lexical rule of English that converts two-argument verbs like *bake* into three-argument verbs like *give*, subject to certain semantic conditions (the beneficiary must be the 'prospective possessor' of the thing baked, knitted, etc.). These new three-argument verbs are free to appear with 'shifted' datives, since all their arguments are obligatory.

Children would construct this lexical rule as follows. When they discover that the beneficiary of verbs like *bake* is optional, they at first simply rule out the double-object construction for these verbs entirely, since it would violate the universal principle of attachment. Then they would learn, one by one from positive evidence, that *bake*, *knit*, etc., in fact can take the double-object construction. After they have assembled a collection of such 'exceptions', all of which share the 'prospective possessor' property, they recognize that there is a pattern to the exceptions and formulate a lexical rule to account for it. This will allow them to create double-object constructions for novel verbs of cooking, creation, etc., as long as the semantic condition on the rule is met. Although this account makes reference to a semantic criterion, it does so in a much more restricted way than Mazurkewich and White's (1984) and Pinker's (1984, 1987) approach. Most of the work of eliminating undesired double object constructions is done by a universal constraint on the syntactic realization of verb argument structure; semantics is brought in only after the grand cutback has taken place, in order to capture what all the positive exceptions to the cutback have in common.

Randall's proposal raises a number of immediate questions. For example, why should children at first assume that *donate*-class verbs have three obligatory arguments, especially since (1) the critical evidence that

they have only two arguments – sentences without a dative object like *John said something nice*; *Mary chose a dress* – is surely available in the input all along, and (2) to judge from their spontaneous speech, children seem more disposed to think that *obligatory* constituents are *optional* than that *optional* ones are *obligatory* (Brown 1973). (Notice that if children did not assume that *donate*-class verbs have three obligatory arguments, but only two, the universal principle of attachment would apply immediately and no overgeneralizations of dative alternation would occur – a wrong prediction.)

More generally, it is important to ask whether a solution of the same general type as Randall proposes for datives can be found for all rules with lexical exceptions. For example, is there a universal principle that accounts for why some verbs of English causativize and others do not? One particular difficulty may be to formulate principles that are strong enough to block undesirable generalizations in one language without at the same time ruling out, or gratuitously rendering as 'marked' or unusual, what appear to be the wholly normal constructions in other languages (Bowerman 1983).

If the proposal that the 'no negative evidence' problem for rules with lexical exceptions can be solved through appeal to principles of Universal Grammar is to be viable, it is essential that comparable explanations for other rules with lexical exceptions can be found. As long as there are rules with lexical exceptions for which such an explanation cannot be constructed, children must have some other techniques for identifying lexical exceptions. And if they have such techniques, children can presumably apply them to any rule with lexical exceptions; it thus becomes superfluous to call upon principles of Universal Grammar.

Two possible candidates for such very general techniques are 'preemption' and 'discovery procedures'. Let us consider their strengths and weaknesses.

7 'Benign Exceptions' Reconsidered: the Role of Preemption

Almost every investigator who worries about the 'no negative evidence' problem has assumed that at least part of the answer lies in preemption: when children formulate overly general rules, they eventually give up overgeneralized forms if they are consistently faced with positive evidence for other forms expressing the same meanings.[4]

A precondition for preemption is that there must be a conventional adult counterpart for a child's overgeneralized word or construction: that is, a form that means what the child's form means and that occurs consistently in the same contexts. This precondition is met most clearly in cases of inflectional overregularization (*runned/ran*; *foots/feet*, etc.). Some overgeneralized causatives also have plausible preempting counterparts (for example, causative **die/kill*; **come/bring*), as do a very few

locative verbs (*rob/steal jewels from the bank; *steal/rob the bank of its jewels).

But what about causative *disappear or *climb, or *choose + indirect object? For these there are no obvious candidate preempting forms: they are 'embarrassing exceptions' by Baker's (1979) definition. Nevertheless, some researchers have suggested that the notion of preemption might be stretched to cover overgeneralizations of these types as well. For example, Clark (1987) proposes that children will give up their overgeneralized causatives like *'to *disappear* (something)' in favor of 'to *make* (something) *disappear*', since in every context where they would say the former, they hear adults say the latter.

A stretch is involved here because *make disappear* is not a perfect semantic match to causative *disappear* (as *kill* is, for example, to causative *die*). In general, lexical causatives and their periphrastic counterparts differ with respect to the directness and conventionality of the act of causation specified (compare, for example, *John stood the baby up* [direct physical causation] with *John made the baby stand up* [indirect causation, e.g. through giving an order]). The weight of this meaning distinction, pervasive throughout English, ought to work against children's willingness to let periphrastic causatives like *make disappear* preempt nonexistent lexical causatives like *disappear*. Still, we might be willing to accept this solution, if only for lack of a better idea.[5]

But, even if it is correct, this extended view of preemption buys us only a little more help with the 'no negative evidence' problem, not a cure. For the approach to work, there must at least be a *consistent relationship* between the child's overgeneralized form and an adult counterpart, even if that counterpart is not identical in meaning to the overgeneralization. This condition is indeed met with causative *disappear* and its periphrastic counterpart.

But what about errors with reversative *un-*, for example? Here the child meets with no consistent alternatives in the adult input. For instance, in contexts where *unsqueeze* would be appropriate, if it existed, adults might say *loosen*, *ease up*, *release*, *let go*, *remove*, and so on. None of these is in direct semantic competition with *unsqueeze*, since none of them specifies or requires that the event referred to is the reversal of an act of 'squeezing'. Nor should the child take the existence of such forms as having any bearing on the possibility of *unsqueeze*: reversative *un-* forms coexist harmoniously with various related constructions, for example *unwrap* and *take the wrapper off*, *unzip* and *pull the zipper down*, *unload* and *empty*.

For overgeneralizations of the type shown in table 4.2 (for example, *untie it off*; see Bowerman 1982b), the problem is even more severe. These sentences are based on a highly productive construction pattern of English that specifies the combination of a causing event and a resulting change of state or location: for example *shoot your enemy dead*, *pat the baby dry*, *wipe the table clean*, *eat oneself into a stupor*, *pull your socks up*,

Table 4.2 Resultative overgeneralizations

1	C	3;8	I *pulled* it *unstapled*. (After pulling stapled booklet apart.)
2	C	3;10	*Untie* it *off*. (Wants mother to untie piece of yarn and take it off tricycle handle.)
3	C	4;0	I'm *patting* her *wet*. (Patting sister's arm after dipping her own hand into a glass of water.)
4	C	6;2	It's hard not to knock them down 'cause whenever I breathe I *breathe* them *down*. (Having trouble setting up a paper village; when she exhales, the houses fall down.)
5	E	3;11	I took my bunny out . . . I *pinched* him *out* with [= by] his fur. (Telling how she reached into bedroom through door, and extracted her toy rabbit by pulling on a bit of fur.)
6	E	6;3	His doggie *bited* him *untied*. (Telling how tied-up man in a TV show was freed.)
7	M	5;6	Are you *washing* me *blind*? (As mother wipes corners of her eyes.)
6	M	5;10	Feels like you're *combing* me *baldheaded*. (As mother combs her hair.)
7	A	4;3	When you get to her, you *catch* her *off*. (A is on park merry-go-round with doll next to her. Wants a friend, standing nearby, to remove doll when it comes around to her.)
8	R	4;9	I'll *jump* that *down*. (About to jump on bath mat M has just put on top of water in tub.)

cut your hair off, *chop a tree down*. However, the pattern is subject to constraints that are still incompletely understood (Green 1972; Mc-Cawley 1971; Randall 1983; Simpson 1983). How do children come to appreciate that there are any restrictions at all, much less what these restrictions are?

The special difficulty is that novel utterances of this type – both those that are acceptable to adult ears and those that sound odd – are usually 'one time only' constructions, designed to fit a certain passing configuration of cause and effect such as pulling on a book and the book's becoming unstapled, or combing the hair and becoming bald. This means that learners do not have repeated opportunities to observe 'the way other people express this particular meaning.' Even if a particular configuration of cause and effect should arise quite frequently (say, 'untying a rope' so that it 'comes off' of something, as in example 2), the alternatives a child might hear from an adult – such as *untie the rope and take it off*, *take the rope off by untying it*, or just plain *take the rope off* – have no bearing on the grammaticality of the child's version, just as the verb *empty* has no bearing on the well-formedness of *unload*. As Fodor and Crain (1987) point out, a learner cannot take every sentence he hears as precluding all sentences that express somewhat related messages; natural languages are too rich for this.

In summary, children make a number of overgeneralizations for which preemption, even if interpreted liberally, fails to provide a correction.

8 Braine's Discovery Procedures Proposal

Braine's original (1971b) empiricist proposal for how to solve the 'no negative evidence' problem has received relatively little attention. According to this approach, the language learner is equipped with 'discovery procedures' that register the properties of incoming sentences. As properties are repeatedly registered, they are passed along a series of intermediate memory stores until they reach permanent memory. Since general properties are characteristic of more sentences than are specific properties, they will be encountered the most often and will reach permanent memory first. Further acquisition involves registering more detailed information – for instance, about the syntactic contexts in which specific lexical items appear. Specific knowledge comes to stand before general knowledge. This means that if a general rule that has arrived in permanent memory has lexical exceptions, it will be overridden, where necessary, by information about the syntactic and morphological behavior of individual words.

This solution has certain advantages. One is that although a rule that is initially overgeneralized will eventually no longer cause problems for familiar lexical items, it is still available. It can therefore be applied to novel words whose properties have never been registered, which is necessary if we are to account for productivity with the passive, etc.

Another advantage is that the learning system has a 'sieve-like' property: only properties that are encountered repeatedly in the input get retained and passed along to permanent storage, and any pattern that is not repeatedly instantiated eventually decays and drops out of the intermediate stores, never making it into permanent memory. This means that occasional errors by the child's conversational partners do no harm. (In contrast, in Baker's (1979) lexicalist approach, a single exposure to a faulty sentence like *Don't say me that* could cause the child to add a wrong subcategorization frame that could never be expunged; see note 1.)

Finally, Braine's model provides a way to formalize native speakers' strong intuition that they do not use familiar lexical items in certain ways simply because they have never heard them used that way. Baker's 'lexical redundancy rule' account also captures this intuition, but, unlike Braine's model, it does not allow for children's overgeneralizations or for productivity with novel lexical items.

Despite these advantages, Braine's account makes the wrong prediction about the *timing* of children's overgeneralizations. If overgeneralizations of a particular pattern are made at all, they should be made early, before word-specific information has been registered. Once such information *has* been registered, the errors should drop out. In fact, however, most of

children's overgeneralizations follow the opposite course (see Bowerman 1982b for discussion). The child first treats individual lexical items, whether exceptional or unexceptional with respect to some rule, with remarkable syntactic and morphological accuracy. This lexically specific stage of learning is followed by the extraction of general rules. Once a generalization is made, errors begin. The errors thus do not reflect initial ignorance of the handling of specific words, which could be corrected simply by further learning of the relevant details, but rather the *overriding* of lexically specific knowledge.

To deal with this problem, Braine (in press) has recently suggested that there may indeed be special circumstances that cause a specific pattern to be learned before a general one. Sometimes, for example, 'the categories of the general pattern may not be available to the [child] at the time the specific pattern is being learned, so that the specific pattern has a long lead through the filter system.' Whenever a specific pattern is learned first, for whatever reasons, 'it may not take precedence over a later-acquired general pattern.' This means that early specific knowledge does not block the later onset of overregularizations involving the same forms.

This modification is too sketchy at present for proper evaluation. One obvious question is how the learning system 'knows' whether information about a specific pattern has preceded or followed information about a general pattern: where and how is information about the relative progress of different patterns stored and evaluated? A second critical question is how overgeneralizations that have been preceded by lexically specific knowledge are eliminated. Must the same lexically-specific information be relearned *after* the generalization, this time coming to override it? A third question is how Braine's learning system would handle the learnability puzzles posed by rules which, if overly general, could not be corrected by learning about the contexts in which specific lexical items can occur.

For example, consider again overgeneralizations of the resultative pattern, shown in table 4.2. One of the constraints on sentences of this type, as Green (1972) has pointed out, is that the 'result' must not be expressed with a past participle (compare, for example, *She combed her hair SMOOTH/*UNTANGLED*; *She cooked the roast DRY/*BURNED/ *OVERDONE*; *he smashed the box FLAT/*BROKEN*).[6] (This accounts for the ungrammaticality of examples 1 and 6.) Children following Braine's learning strategy might conclude that *untangled*, *burned* and *broken* cannot appear as result complements, because they have never observed them in these positions. But since this information is specific to particular words, there is nothing to stop the learner from creating resultatives with the past participles of novel or newly learned verbs. How do children form a general block against ALL past participles as result complements, regardless of past opportunities to observe a particular participle's behavior? (See Bowerman 1987: 461 for discussion of the problem of identifying properties shared by items that do *not* undergo a rule.)

Conversely, why does the child who follows Braine's strategy eventually come to regard sentences like *Whenever I breathe I breathe them down* and *I pinched him out with his fur* (examples 4 and 5 in table 4.2) as peculiar? The problem here is not with the specific lexical items, but with their interaction. For example, both *down* and *out* occur frequently as result complements (for example, *blow/chop/push NP DOWN*; *pull/yank/press NP OUT*), and *breathe* and *pinch* both occur as main verbs in resultative sentences (*BREATHE NP in/out*; *PINCH NP black and blue/to death*). Why then do they resist occurring together? We cannot solve this puzzle within Braine's framework by arguing that these words have never been heard together in this construction (i.e., by appealing to information specific to the contingencies between TWO lexical items). To do so would be to lose sight of the fact that the child whose learning system we are trying to characterize must end up with a productive rule that allows him to create *novel combinations* of main verb and result complement.

It is not clear whether these problems can be solved within Braine's general framework. Clearly the model will require more work before its potential as a solution to the 'no negative evidence' problem can be properly evaluated.

9 Does the 'No Negative Evidence Problem' Really Exist?

In view of all the difficulties I have raised for various approaches to the 'no negative evidence' problem, the reader might have concluded that children must get negative feedback after all. Perhaps the 'no negative evidence' problem is just a myth.

Several investigators have argued for a resolution to the problem along these lines. They concede that children do not get many *explicit corrections* of their grammar, but they argue that learners do get negative feedback in the form of misunderstandings, requests for clarification, repetitions and recastings (e.g. Demetras et al. 1986; Hirsch-Pasek et al. 1984). I do not think the answer lies in this direction, however, for the following reasons.

First, researchers who argue that children do get negative evidence from their speaking partners have not distinguished in the necessary way among feedback to different categories of ungrammatical child utterances. Whenever negative feedback is observed for child utterances that are imperfect by adult standards, it is indiscriminately taken to count against the seriousness of the 'no negative evidence' problem. But most of this feedback is simply irrelevant. Many utterances, especially among younger children, are ungrammatical not because the speaker's rules are overly general but because the speaker hasn't yet constructed the necessary rules at all. (Omissions of grammatical morphemes are a case in point.) Adult misunderstandings, recastings and the like might or might

not hasten rule construction in these cases, but they do not bear on the problem of how children cut back on overly general rules.

Even where such feedback does follow errors resulting from overly general rules, these errors often involve 'benign' rule exceptions, which in principle could be corrected by preemption and so do not require negative feedback. I suspect that the amount of negative evidence that remains after we eliminate these two types of irrelevant feedback is not very great, especially since overgeneralizations of the types discussed in this chapter are produced by relatively old children and rarely cause misunderstanding (as noted also by Mazurkewich and White 1984).

Second, listener misunderstandings, requests for clarification, repetitions and recasts are not reliably diagnostic of ungrammaticality on the speaker's part: they follow well-formed utterances as well as those that are ungrammatical (Hirsch-Pasek et al. 1984). If a child's first impulse on hearing such responses is to question the adequacy of his grammar, he would continually be trying to revise perfectly acceptable rules. It seems unlikely that children are so readily led astray. Even if a child does on occasion question his grammar, only recasts give information about where the problem lies—misunderstandings, repetitions and 'what' questions are silent about what is wrong.

In sum, the 'no negative evidence' problem is not a myth, but a very real and serious challenge for the construction of an adequate theory of language acquisition.

10 Conclusions

Many linguists have assumed that the lack of negative feedback to children creates a logical problem for language acquisition that can be solved only by reference to innate constraints that prevent children from ever formulating overly general grammars. The data discussed in this chapter show that this assumption is untenable. Children do construct overly general grammars, yet somehow they are able to recover. How this recovery is accomplished is still uncertain.

It is possible, as some proponents of Universal Grammar have argued, that recovery is effected through innate linguistic constraints that come into play *after* rather than before the critical overgeneralizations have been made. But evidence for this is as yet slight. Existing hypotheses make questionable assumptions about why the proposed innate knowledge does not block the child's overgeneralizations from the beginning. In addition, no proposals have been developed to deal with most of children's error genres, and it is not clear that this will be possible, especially in view of cross-linguistic variation, even among closely related languages, in what lexical items constitute exceptions to a particular rule.

At present, I think it is most likely that children cut back on rules that are too broad with the help of relatively general learning mechanisms that

apply across a wide range of error types. However, developing a plausible account of such mechanisms has proved remarkably difficult, and I have raised a number of problems for the three approaches of this type considered in this chapter: Criteria, Preemption, and Discovery Procedures.

Despite these criticisms, and also those I have leveled at Baker's verb-frame-by-verb-frame approach, the Subset Principle, and the Universal Grammar hypothesis, I believe we have made significant progress in the last few years in understanding the 'no negative evidence' problem. I hope that this chapter may contribute to its eventual resolution by highlighting some of the subtle and complex difficulties that future research on this puzzle will have to address.

Notes

1 One problem is that these approaches do not seem sufficiently robust against misleading input (Bowerman 1983). Children are often exposed to errors in the speech they hear, especially if they have a parent who is not a native speaker of the language they are learning. For example, Dutch adults often make shifted dative errors with *say* and *suggest* when speaking English (parallel structures with the Dutch cognate verbs are grammatical). Once having registered an incorrect syntactic frame for a verb, how could a child ever get rid of it? In addition, as Fodor and Crain (1987) point out, the Subset Principle solution to the 'no negative evidence' problem has an unattractive corollary. When the output of one rule is a proper subset of the output of another rule, then the narrower rule is typically more complex than the broader one since it is annotated for one or more constraints on application that the broader rule does not respect. This means that if children always start out by hypothesizing the narrowest rule possible, their first rules are routinely more complex than their later ones. But it clashes with our intuitions to imagine that children move consistently from more complex to less complex rules – that language acquisition proceeds by the successive removal of constraints on rule application.

2 For example, to make the Criteria approach work, Pinker (1984, pp. 333–341) credits children with foreknowledge of (1) a finite, universal set of features, possibly hierarchically organized, that are potentially relevant to constraining the class of lexical items to which a given rule applies, 2) the 'Unique Entry' principle (see footnote 4), violations of which will lead children to sample from among the potential constraints mentioned in (1) if they have not already done so spontaneously, and (3) a universal set of notions like 'passive', 'causative', and 'inchoative' that define the levels of the word-formation paradigms across which the 'Unique Entry' principle operates.

3 I base this claim on an informal survey of about twenty native speakers of English, most of whom rejected examples like these. Failure to find 100 percent agreement is not surprising: for every rule with lexical exceptions there are items about which speakers disagree or are uncertain. This is not important for my argument. As long as there are mature speakers who find such examples odd, we must explain how in their grammars the lexical item involved came to be excepted from the rule under examination.

4 There are both nativist and empiricist accounts of how preemption works. Some theorists argue that there is some property 'in the child' that rejects the idea that two forms should have exactly the same meaning. For example, Pinker (1984) postulates a 'Unique Entry' principle, according to which children resist having more than one entry in an inflectional or derivational paradigm unless they are faced with strong positive evidence that both forms exist. Clark's (1987) 'Principle of Contrast' is even stronger, stating that children assume that every two forms differ in meaning, and will resist acquiring or retaining two forms that seem synonymous. A contrasting, 'empiricist' hypothesis explains preemption as the outcome of *competition* among alternative forms for expressing the same meaning (MacWhinney 1987; Rumelhart and McClelland 1987): over time, the conventional adult form gains in activation strength and the child's overregularized form weakens until it dies out entirely. Grounds for deciding between these alternative views are at present unclear (see Bowerman 1987 for discussion).

5 See also Pinker (1981), Maratsos and Chalkley (1981) and Bowerman (1983) for the related proposal that children may identify items that are exceptions to their overly general rules through continually failing to encounter those items in discourse contexts where they 'expect' them; this is sometimes called 'indirect negative evidence', following Chomsky (1981: 9).

6 *Closed* and *shut* seem to be the only exceptions to this constraint (cf. Green 1972).

References

Baker, C. L. (1979) 'Syntactic theory and the projection problem'. *Linguistic Inquiry*, 10, 533–581.
    and McCarthy, J. J. (1981) *The Logical Problem of Language Acquisition*. Cambridge, Mass.: MIT Press.
Berwick, R. (1985) *The Acquisition of Syntactic Knowledge*. Cambridge Mass.: MIT Press.
    and A. Weinberg (1984) *The Grammatical Basis of Linguistic Performance*. Cambridge, Mass.: MIT Press.
Bowerman, M. (1974) 'Learning the structure of causative verbs: a study in the relationship of cognitive, syntactic, and semantic development'. *Papers and Reports on Child Language Development* (Stanford University Department of Linguistics), 8, 142–78.
    (1982a) 'Evaluating competing linguistic models with language acquisition data: implications of developmental errors with causative verbs'. *Quaderni di Semantica*, 3, 5–66.
    (1982b) 'Reorganizational processes in lexical and syntactic development'. In E. Wanner and L. R. Gleitman (eds), *Language Acquisition: The State of the Art*. Cambridge: Cambridge University Press.

(1983) 'How do children avoid constructing an overly general grammar in the absence of feedback about what is not a sentence?' *Papers and Reports on Child Language Development* (Stanford University Department of Linguistics), 22, 23–35.

(1987) 'Commentary: mechanisms of language acquisition'. In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum.

Braine, M. D. S. (1971a) 'The acquisition of language in infant and child'. In C. Reed (ed.), *The Learning of Language*. New York: Scribners.

(1971b) 'On two types of models of the internalization of grammars'. In D. I. Slobin (ed.), *The Ontogenesis of Grammar*. New York: Academic Press.

(in press) 'Modeling the acquisition of linguistic structure'. In Y. Levy, I. M. Schlesinger and M. D. S. Braine (eds) Categories and Processes in Language Acquisition. Hillsdale, NJ: Erlbaum.

Brown, R. (1973) *A First Language*. Cambridge, Mass.: Harvard University Press.

and C. Hanlon (1970) Derivational complexity and order of acquisition in child speech. In J. R. Hayes (ed.) *Cognition and the Development of Language*. NY: Wiley.

Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.

(1981) *Lectures on Government and Binding*. Dordrecht: Foris.

Clark, E. (1987) 'The principle of contrast: a constraint on language acquisition'. In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum.

Dell, F. (1981) 'On the learnability of optional phonological rules'. *Linguistic Inquiry*, 12, 31–7.

Demetras, M. J., K. N. Post and C. E. Snow (1986) 'Feedback to first language learners: the role of repetitions and clarification questions'. *Journal of Child Language*, 13, 275–92.

Fodor, J. D. and S. Crain (1987) 'Simplicity and generality of rules in language acquisition'. In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum.

Gergely, G. and T. Bever (1986) 'Related intuitions and the mental representation of causative verbs in adults and children'. *Cognition*, 23, 211–77.

Goldsmith, J. (1980) 'Meaning and mechanism in language'. In S. Kuno (ed.), *Harvard Studies in Syntax and Semantics* III, Department of Linguistics, Cambridge, Mass.: Harvard University Press.

Green, G. (1972) 'Some observations on the syntax and semantics of instrumental verbs'. *Proceedings of the Chicago Linguistic Society*, 8, 83–95.

(1974) *Semantics and Syntactic Regularity*. Bloomington, Ind.: Indiana University Press.

Grimshaw, J. (1987) 'Discussion: Session on "Acquisition of verbs and argument structure"'. Boston Child Language Conference, October.

Hall, B. (1965) 'Subject and object in modern English'. Doctoral diss., MIT (Published under B. Partee, [same title], New York: Garland, 1979).

Hirsch-Pasek, K., R. Treiman and M. Schneiderman (1984) 'Brown & Hanlon revisited: mothers' sensitivity to ungrammatical forms'. *Journal of Child Language*, 11, 81–8.

Hochberg, J. G. (1986) 'Children's judgements of transitivity errors'. *Journal of Child Language*, 13, 317–34.

Jackendoff, R. (1975) 'Morphological and semantic regularities in the lexicon'. *Language*, 51, 639–71.

(1977) *X̄ Syntax: a Study of Phrase Structure*. Cambridge, Mass.: MIT Press.

Levin, B. (1985) 'Lexical semantics in review: an introduction'. In B. Levin (ed.), *Lexical Semantics in Review*. Lexicon Project Working Papers 1. Cambridge, Mass.: MIT Center for Cognitive Science.

McCawley, J. D. (1971) 'Prelexical syntax'. In R. J. O'Brien (ed.), *Georgetown University Round Table on Language and Linguistics*. Washington, DC: Georgetown University Press.

Maratsos, M. (1979) 'How to get from words to sentences'. In D. Aaronson and R. Reiber (eds), *Perspectives in Psycholinguistics*. Hillsdale, NJ: Erlbaum.

and M. Chalkley (1981) 'The internal language of children's syntax'. In K. E. Nelson (ed.), *Children's Language*, vol. 2. New York: Gardner Press.

R. Gudeman, P. Gerard-Ngo and G. DeHart (1987) 'A study in novel word learning: the productivity of the causative'. In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum.

Matthews, P. H. (1983) Review of Hornstein and Lightfoot (eds), *Explanation in Linguistics: The Logical Problem of Language Acquisition* (1981). *Journal of Child Language*, 10, 491–3.

Mazurkewich, I. and L. White (1984) 'The acquisition of the dative alternation: unlearning overgeneralizations'. *Cognition*, 16, 261–83.

Oehrle, R. T. (1976) 'The grammatical status of the English dative alternation'. Doctoral dissertation, MIT.

Pinker, S. (1981) Comments on the paper by Wexler. In C. L. Baker and J. J. McCarthy (eds), *The Logical Problem of Language Acquisition*. Cambridge, Mass.: MIT Press.

(1984) *Language Learnability and Language Development*. Cambridge, Mass.: Harvard University Press.

(1987) 'Resolving a learnability paradox in the acquisition of the verb lexicon'. *Lexicon Project Working Papers 17*, Center for Cognitive Science, MIT. (To appear in R. L. Schiefelbusch (ed.), *The Teachability of Language*.)

D. S. Lebeaux and L. A. Frost (1987) 'Productivity and constraints in the acquisition of the passive'. *Cognition*, 24.

Randall, J. (1983) 'A lexical approach to causatives'. *Journal of Linguistic Research*, 2, 77–105.

(1985) 'Negative evidence from positive'. In P. Fletcher and M. Garman (eds), *Child Language Seminar Papers* (Proceedings of the 1985 Child Language Seminar, University of Reading).

(1987) 'Indirect positive evidence: overturning generalizations in language acquisition'. Indiana University Linguistics Club.

Rappoport, M. and B. Levin (1985) 'A case study in lexical analysis: the locative alternation'. Unpub. MS, Center for Cognitive Science, MIT.

Roeper, T., S. Lapointe, J. Bing and S. Tavakolian (1981) 'A lexical approach to language acquisition'. In S. L. Tavakolian (ed.), *Language Acquisition and Linguistic Theory*. Cambridge, Mass.: MIT Press.

Rumelhart, D. E. and J. L. McClelland (1987) 'Learning the past tenses of English verbs: implicit rules or parallel distributed processing'. In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum.

Shibatani, M. (1976) 'The grammar of causative constructions: a conspectus'. In M. Shibatani (ed.), *The Grammar of Causative Constructions*. New York: Academic Press.

Simpson, J. (1983) 'Resultative attributes'. In A. Zaenen, M. Rappoport, and

THE 'NO NEGATIVE EVIDENCE' PROBLEM
B. Levin (eds), *Papers in Lexical-Functional Grammar*. Indiana University Linguistics Club.

Smith, C. (1981) Comments on the paper by Roeper. In C. L. Baker and J. J. McCarthy (eds), *The Logical Problem of Language Acquisition*. Cambridge, Mass.: MIT Press.

Stowell, T. (1981) 'Origins of phrase structure'. Unpub. Ph.D. diss., MIT.

Wasow, T. (1977) 'Transformations and the lexicon'. In P. W. Culicover, T. Wasow and A. Akmajian (eds), *Formal Syntax*. New York: Academic Press.

Wasow, T. (1981) 'Comments' on C. L. Baker 'Learnability and the English auxiliary system'. In C. L. Baker and J. J. McCarthy *The Logical Problem of Language Acquisition*. Cambridge, Mass.: MIT Press.

Wexler, K. and M. R. Manzini (1987) 'Parameters and learnability in binding theory'. In T. Roeper and E. Williams (eds), *Parameter Setting*. Dordrecht: Reidel.

Whorf, B. (1956) *Language, Thought and Reality* (ed. J. Carroll). Cambridge, Mass.: MIT Press.