4

How argument structure constructions are learned[1]

Surely it is premature to give up hope that humans, with our rich cognitive abilities, complex social skills, predilection to imitate, and 100-billion-neuron brains, can learn language from the available input. Children have ample motivation to learn language, hear thousands of utterances every day, and receive constant indirect statistical feedback (see Chapter 5 for discussion of the notion of feedback).

Recent empirical work has demonstrated that even honeybees, with brain volumes of approximately 1 mm^3, are capable of quite advanced learning (Collett et al. 1997; Giurfa 1996; Menzel and Giurfa 2001). In one study, honeybees were shown to be able to learn the abstract concepts of "sameness" and "difference" in the following way. Bees were trained to find food in a Y-shaped maze; they individually entered the maze through the bottom of the Y where they encountered a stimulus, A. The entrance-way led to the fork of the Y, where the bee would have to choose a path. At the entrance to one fork was placed matching stimulus A, or a different stimulus B. If the bee matched the stimulus, it was rewarded in that there was food in the fork beyond the matching stimuli but not beyond the non-matching stimuli. Bees were found to learn to match stimulus A to A, B to B at a rate well above chance within four to seven trials. After learning to match the training stimuli, the bee was tested with entirely new stimuli C and D. The bees successfully transferred their new knowledge to the new stimuli with an average of 75 per cent accuracy. In fact they were also able to generalize "different" in that they could be trained to learn "if A is at the entrance then take the fork marked by non-A"; and they then could extend that learning to new stimuli C and D as well (Giurfa et al. 2001).

[1] The corpus work described in this chapter was done in collaboration with Devin Casenhiser and Nitya Sethuraman, and the experimental work was done in collaboration with Devin Casenhiser. Please see Goldberg, Casenhiser, and Sethurman (2004, 2005) and Casenhiser and Goldberg (forthcoming) for full results and analysis.

Adding to the impressiveness of these findings is the fact that A and B, and C and D represented distinct modalities. For example in one experiment, A and B were different colors (yellow and blue), whereas the novel testing stimuli C and D were lines of different orientations (horizontal or vertical); in another experiment, A and B were different smells (lemon and mango) and C and D were different colors. That is, in a "difference" experiment, a honeybee could learn that if a yellow patch was at the entryway, then the food would be behind the non-yellow patch; and once that idea was learned, the bee would also know that if there were horizontal lines at the entryway, then the food would be behind the non-horizontal lined (i.e. vertical lined) entrance. This learning goes well beyond simple associative learning in which a specific stimulus triggers a conditioned response. The authors conclude that "bees can, not only learn specific objects and their physical parameters, but also master abstract interrelationships, such as 'sameness' and 'difference' "(Giurfa et al. 2001).

These sorts of advances in our understanding of what even insects are capable of learning could not be envisioned in the 1950s and 1960s when Chomsky asserted that critical aspects of syntax were "unlearnable" by human beings and therefore must be innate; yet the assertion became dogma in our field and led to the continuing, widespread belief in the necessity of a biological endowment that contains knowledge representations that are specific to language: i.e. "universal grammar."

4.1 Learning word segmentation, phrase boundaries, grammatical categories

In the past decade, we have witnessed major discoveries concerning children's ability to extract statistical regularities in the input. Children are able to extract word forms from continuous speech based on transitional probabilities between syllables (Saffran, Aslin, and Newport 1996). For example, the phrase *bananas with milk*, contains four transitional probabilities across syllables (*ba* to *na*; *na* to *nas*; *nas* to *with*; and *with* to *milk*). The probability that *ba* will be followed by *na*, and the probability that *(ba)na* will be followed by *nas* is higher than the probability that *nas* will be followed by *with*. That is, transitional probabilities are generally higher within words than across words. Eight-month-old infants are sensitive to these statistical cues (Saffran, Aslin, and Newport 1996) and treat these newly acquired words as part of their lexical inventory (Saffran 2001b).

These basic learning abilities are neither unique to humans nor specific to language. Cotton-top tamarin monkeys also track transitional probabilities,

and show evidence of discovering word boundaries (Hauser, Newport, and Aslin 2001). Moreover, children can discover regularities in the boundaries between sequences of tones (Saffran et al. 1999), and visual patterns (Fiser and Aslin 2002), using the same type of statistical cues.

Statistical cues provide a powerful means by which initial language learning can begin. Children are able to discover syntactic regularities between categories of words as well as the statistical regularities in sound patterns (Marcus et al. 1999). For example, the presence of an article (*the* or *a*) predicts a noun somewhere downstream, and learners can use this type of cue to learn syntactic phrase boundaries (Saffran 2001a). Elman has also demonstrated that grammatical categories can be distinguished on the basis of distribution (Elman 1990). Moreover, children are able to combine word level and syntactic level statistical information: twelve-month-old children can use their newly discovered word boundaries to discover regularities in the syntactic distributions of a novel word-string grammar (Saffran and Wilson 2003). Gerken, Wilson, and Lewis have demonstrated that fifteen-month olds are able to combine multiple cues in order to learn morphophonological paradigmatic constraints (Gerken, Wilson, and Lewis 2005).

It is a lucky thing that children can learn initial aspects of language from statistical features of the input, since **there are no stable formal cues cross-linguistically** to identify word forms, grammatical categories, or relations (Barðdal forthcoming; Croft 2001; Pinker 1984). Are more abstract aspects of language less amenable to learning through exposure? Perhaps. However, the rest of this chapter focuses on children's ability to learn and learn quickly, from the available input, one such aspect of language: linking rules or argument structure constructions.

4.2 Learning argument structure generalizations

Linguists have observed that within a given language, there exist certain formal patterns that correlate strongly with the meanings of the utterances in which they appear. Such correlations between form and meaning have been variously described as linking rules projected from the main verb's specifications (Levin and Rappaport Hovav 1995; Pinker 1989), as lexical templates overlain on specific verbs (Koenig and Davis 2001; Rappaport Hovav and Levin 1998), or as phrasal form and meaning correspondences (*constructions*) that exist independently of particular verbs (Goldberg 1995; Jackendoff 2002).

Of course one way to account for the association of meanings with particular forms is to claim that the association is simply there, biologically determined or innate from the outset (Baker 1988; Chomsky 1982). This claim

generally rests on the idea that the input is not rich enough for the relevant generalizations to be learned; this is the well-known "poverty of the stimulus" argument (Chomsky 1980, 1988; Pinker 1994). On this nativist view, learning a grammar can be likened to customizing a software package: everything is there, and the learner simply selects the parameters that are appropriate, given the input (Jackendoff 2002). Many have criticized this approach for its biological implausibility (Bates and Goodman 1998; Deacon 1997; Elman et al. 1996; Sampson 1997). Moreover, there have been virtually no successful proposals for what any specific aspect of the parameters might look like (for discussion of this failure see, e.g. Culicover 1999; Jackendoff 2002; Newmeyer 1998; Webelhuth and Ackerman 1998).

This chapter joins the growing body of literature that detracts from the poverty of the stimulus argument by presenting evidence that the nature and properties of at least certain patterns in language are learnable on the basis of general categorization strategies (see also e.g. Bybee and Slobin 1982; Bybee and Moder 1983; Jackendoff 2002; Lakoff 1987; MacWhinney forthcoming; Pullum and Scholz 2002; Scholz and Pullum 2002; Taylor 1995; Tomasello 2003).

It is argued that the language input children receive provides more than adequate means by which learners can induce the association of meaning with certain argument structure patterns insofar as well-established categorization principles apply straightforwardly to this domain. Throughout this chapter, I adopt constructional terminology, but the ideas presented are not actually exclusive to a constructionist account. Those who favor one of the other terminologies mentioned above need only construe this account as a proposal for how children can learn linking rules or learn the semantics associated with various lexical templates on the basis of the input. What *is* crucial is the uncontroversial notion that there do in fact exist correlations between formal linguistic patterns and meaning. Chapter 5 focuses on the issue of how we avoid overgeneralizations, and Chapter 6 discusses the motivation for learning the generalizations that we learn. The final part of the book (Chapters 7–9) analyzes how to account for further regularities both within and across languages without resorting to claims that the generalizations must be hard-wired or biologically determined.

Table 4.1 provides a partial list of form and meaning correspondences (lexical templates, combination of linking rules, constructions) along with the labels that are used as mnemonics throughout the chapter to refer to them.

Previous work on the acquisition of constructions has focused almost entirely on the question of whether the constructions (or "linking rules") that exist in a given language have been acquired at a certain age. Findings

TABLE 4.1. Examples of correlations between form and meaning

Form/Example	Meaning	Construction Label
1. Subj V $Obl_{path/loc}$	X moves $Y_{path/loc}$	**Intransitive Motion**
e.g. The fly buzzed into the room.		
2. Subj V Obj $Obl_{path/loc}$	X causes Y to move $Z_{path/loc}$	**Caused Motion**
e.g. Pat sneezed the foam off the cappuccino.		
3. Subj V Obj Obj2	X causes Y to receive Z	**Ditransitive**
e.g. She faxed him a letter.		
4. Subj V Obj RP	X causes Y to become Z_{state}	**Resultative**
e.g. She kissed him unconscious.		

based on the preferential-looking paradigm have been used to argue that children already have certain linking rules at relatively young ages, the implication being that the linking rules are innate and not learned from the input (Naigles 1990; Fisher 1996; Hirsh-Pasek, Golinkoff, and Naigles 1996). That is, linking rules have been claimed to be "near-universal in their essential aspects and therefore may not be learned at all" (Pinker 1989: 248). These candidates for universal linking rules include a mapping of agent to subject; patient to object; and goal to oblique (prepositional phrase) (Pinker 1989: 74). Naigles, Gleitman, and Gleitman (1993: 136–7) likewise suggest that "there is sufficient cross-linguistic similarity in these linking rules to get the learning procedure started.... [T]here is an overwhelming tendency, cross-linguistically, for agents to appear as subjects and themes as direct objects, with other arguments appearing in oblique cases." Baker (1996: 1) likewise notes: "One central task for any theory of grammar is to solve the so-called 'linking problem': the problem of discovering regularities in how participants of an event are expressed in surface grammatical forms." The implication has been that universal aspects of language are innate, proposed specifically to solve the apparent "poverty of the stimulus" problem (Chomsky 1965, 1981, 1988).

On the other side of the debate, the emphasis has been on the conservative nature of children's early learning, with demonstrations focusing on children's failure to generalize beyond the input until learners have been exposed to a vast amount of data at age 3.5 or older (Akhtar and Tomasello 1997; Baker 1979; Bates and MacWhinney 1987; Bowerman 1982; Braine 1976; Brooks and Tomasello 1999; Gropen et al. 1989; Ingram and Thompson 1996; Lieven, Pine, and Baldwin 1997; MacWhinney 1982; Olguin and Tomasello 1993; Schlesinger 1982; Tomasello 1992). The clear implication of this work is that constructions must be learned, since they are acquired so late and in such a piecemeal fashion.

Training studies are required in order to reconcile the issues involved in this debate, since such studies allow the input and the target construction to be manipulated. Yet precious few training studies exist. Moreover, even fewer facilitory or inhibitory factors have been identified. In one training study, Childers and Tomasello (2001) found a single facilitating factor in the acquisition of the English transitive construction: namely the use of pronouns instead of full NP arguments (Akhtar 1999 also found some facilitative effect for pronouns). Abbott-Smith, Lieven, and Tomasello (2004) attempted to look for other facilitory factors in construction learning, including tight semantic homogeneity and shared syntactic distribution in the input, but found null results.

Kaschak and Glenberg (2004) have investigated adults' online processing of the construction exemplified by *This shirt needs washed*, a construction that was novel to their experimental subjects, although it is used by native speakers of western Pennsylvania. They found that speakers were able to read instances of this construction with greater fluency after hearing or reading other instances of the construction. Facilitation was found as well in testing on the same pattern with *wants* after training on *needs*, demonstrating that the facilitation transferred to a related verb. The increased fluency, as measured by shorter reading times, was interpreted to indicate that speakers learned to comprehend the construction; however, the target construction contains familiar words with appropriate inflectional endings and is closely related to familiar expressions such as *This shirt needs to be washed*, with *to be* omitted. A suggestion that subjects were able to comprehend the construction from the outset comes from the fact that they demonstrated increased reading times for semantically inconsistent follow-up sentences even in the initial testing trials. Still the study demonstrates that familiarity with a construction can be acquired quite quickly.

Certain findings indicate that a *failure* to be able to predict aspects of distribution leads learners to generalize the more regular aspects. In order to investigate the processes involved in creolization, Hudson and Newport (1999) taught adult speakers a toy novel syntax, through exposure to sentences that were paired with video clips to provide interpretations. Several different determiners were used in free variation. Adult subjects were then tested on whether they were able to regularize their use of the determiners. The experimenters initially found that subjects failed to regularize and instead produced the various determiners in roughly the proportions they had heard them used. In a follow-up study, Hudson Kam and Newport (MS) have found that adults do regularize (overgeneralizing the use of one determiner) when the input is sufficiently unpredictable. In particular, adults overgeneralized (regularized) a determiner that had appeared 60 per cent of the time, when various other determiners appeared with unpredictable, low frequencies (each of various determiners appearing 2.5 per cent of the time, in free variation). In an artificial grammar-learning task, Gomez (2002) likewise observes that decreasing predictability between adjacent dependencies increased awareness of dependencies between first and third elements (cf. also Monaghan et al. 2005; Valian and Coulson 1988).

To summarize, a few training studies exist, and some have required children to learn a novel word order and/or a novel morpheme, but none has required children to map a novel word order onto a novel meaning: exactly the task that the child faces when naturalistically learning language. The target meaning involved has been simple transitivity in the case of previous novel word-order studies (e.g. Childers and Tomasello 2001), or identifiabilty (in the case of the determiner study by Hudson and Newport 1999), or no meaning at all (in the case of artificial grammar learning by Gomez 2002). Surprisingly little data has been found that has identified particular facilitory or inhibitory factors in learning to map a novel form to a novel meaning, beyond varying overall exposure and levels of predictability.

4.3 Skewed input

In order to examine more closely the input children receive, Goldberg, Casenhiser, and Sethuraman (2004) investigated a corpus of children's early speech. The main language corpus used was the Bates, Bretherton, and Snyder corpus (1988), on the Child Language Data Exchange System (CHILDES) database (MacWhinney 1995). This corpus contains transcripts from the Bates/Bretherton Colorado longitudinal sample of twenty-seven middle-class children–mother dyads, thirteen boys and fourteen girls at age twenty and twenty-eight months. There are transcripts for fifteen minutes, equally divided into three types of mother–infant interaction: free play, reading of the book *Miffy in the Snow*, and snack time.

In analyzing the mothers' speech we found a strong tendency for there to be one verb occurring with very high frequency in comparison to other verbs used in each of the constructions analyzed. That is, the use of a particular construction is typically dominated by the use of that construction with one particular verb. For example, *go* accounts for a full 39 per cent of the uses of the intransitive motion construction in the speech of mothers addressing twenty-eight-month-olds in the Bates et al. (1988) corpus. This high percentage is remarkable since this construction is used with a total of thirty-nine

different verbs in the mothers' speech in the corpus. The figures for the three constructions investigated are given below in Table 4.2.[2]

TABLE 4.2. Fifteen mothers' most frequent verb and number of verb types for three constructions in Bates et al. (1988) corpus[3]

Construction	Mothers	Total Number of Verb Types
1. Subj V Obl	39% *go* (136/353)	39 verbs
2. Subj V Obj Obl	38% *put* (99/256)	43 verbs
3. Subj V Obj Obj2	20% *give* (11/54)	13 verbs

The question arises as to why these particular verbs are used more frequently in these constructions by mothers. One factor of course is that these verbs are among the most frequent verbs in the language at large (Carroll, Davies, and Richman 1971). But this does not in itself predict that these verbs should account for such a high proportion of tokens of any single construction, since most frequent verbs appear in multiple constructions. Zipf long ago noted that highly frequent words account for most linguistic tokens (1935). Although he did not claim that there should be a single most highly frequent word for each clause pattern, nor did Zipf's work prepare us for the fact that a single verb accounts for such a large percentage of the tokens, his observation suggests that we may find a similar pattern in constructions other than argument structure constructions.

In fact, Stefanowitsch and Gries (2003) have refined the notion of relative frequencies in order to take into account the overall frequencies of the verbs in the language. Given that one can expect high-frequency verbs to appear with high-frequency constructions, the more refined question asks, how frequent are these particular verbs in these particular constructions over and above chance? Stefanowitsch and Gries advocate applying measures of association (e.g. $\chi 2$ or Fisher exact test) to the following matrix of data:

[2] Valeria Quiochi (in prep.) has found a similar pattern for the caused-motion and intransitive motion patterns in Italian. (The ditransitive pattern does not exist in Italian.)

[3] The percentage of uses of *give* in the ditransitive is somewhat less striking than the percentages of *go* and *put* in the intransitive and caused motion constructions, respectively. However, that is likely an effect of the small sample size. Bresnan and Nikitina report that *give* accounts for 226/517, or 44% of the instances of the ditransitive in the parsed Switchboard corpus (Bresnan and Nikitina MS). *Tell*, in our small sample of 54 appeared an equal number of times as *give*. We believe this is an artifact of the story-reading context, since all but one instance of *tell* occurred in the story context, and in 8 out of 10 instances, the theme argument is *story*.

	Argument structure construction (ASC) #1	Other ASCs	Total
*Verb*1	a	b	a+b: frequency of *Verb*1
All other verb tokens	c	d	c+d

If you select all constructions involving *Verb*1 (a+b), how likely is it that there will be (*a*) instances of ASC #1 and (*b*) instance of other ASCs? That is, how biased toward ASC #1 is *Verb*1, given the overall frequencies of *Verb*1 *and ASC#1*? Stefanowitsch and Gries have found that one can never expect a single verb to account for more than 10 per cent of the tokens, even if the verb is very frequent, because there are so many different argument structure constructions: that is, skewing of the magnitude that is found in the data is not expected due to the simple frequencies of the verbs and constructions involved.

The same trends noted in mothers' speech to children are mirrored in children's early speech (Goldberg, Casenhiser, and Sethuraman 2004). This is not surprising since we know from previous studies that children's use of verbs is highly sensitive to their mothers' use of verbs (Choi 1999; De Villiers 1985; Naigles and Hoff-Ginsberg 1998). Moreover the reason certain verbs occur so frequently is constant for both children and mothers as described below.

Accounting for the skewed input

The fact that *go*, *put*, and *give* are so frequent in the input raises the question as to why that should be so. There seem to be two reasons. First, if we compare, for example, *go* with *amble*, or *put* with *shelve*, it is clear that *go* and *put* are more frequent because they apply to a wider range of arguments and therefore are relevant in a wider range of contexts (Heine 1993; Zipf 1935).

In addition, each of the main uses of these verbs designates a basic pattern of experience; for example, someone causing someone to receive something (*give*), something causing something to move (*put*), or something moving (*go*). These meanings are readily accessible to the child, since they involve concrete actions. Clearly the verb meanings need to be accessible as well as highly frequent in the input in order to be frequently produced in early child language (Slobin 1985).

General Purpose Verbs and Constructional Meaning

As represented in Table 4.3, the meanings of the most frequent verbs used in particular argument structure constructions bear a striking resemblance to

TABLE 4.3. Main verbs and the constructional meanings they correspond to

go	X moves Y	Intransitive Motion
put	X causes Y to move Z	Caused Motion
give	X causes Y to receive Z	Ditransitive
make	X causes Y to become Z	Resultative

the meanings independently posited for those argument structure constructions (Goldberg 1995).

"General purpose" verbs, including *put*, *go*, *do*, and *make*, are among the first and most frequent verbs in many languages (Clark 1978, 1996). Clark cites data from Bowerman (1973) for Finnish, Grégoire (1937) for French, Sanchés (1978) for Japanese, and Park (1977) for Korean; Ninio (1999), discussed below, provides similar data from Hebrew (cf. also Uziel-Karl 1999).

The generality of the meanings of these verbs and their highly frequent and early appearance in children's speech suggests that they may aid children in generalizing patterns from the input. Speculations about the close relationship between certain verbs and argument structure patterns have been made previously by certain researchers both in linguistics and in language acquisition. Kay (1996) observes that "it is possible to think of the argument structure patterns as in some sense 'derived from' the semantics of their most neutral verbs" (cf. also Goldberg 1998). Clark (1996) likewise speculates that certain early-learned verbs may serve as "templates" for further acquisition on the basis of their semantic characteristics. She demonstrates that children are aware of much relevant semantic knowledge pertaining to their early verbs as evidenced by their discriminating use of various inflectional morphemes. Ninio (1999) has also suggested that syntactic patterns emerge from generalizing the use of particular verbs. With the possible exception of Ninio, these researchers do not attempt to flesh out this idea. As Kay (1996) notes, "when we come to propose this seriously we will have to specify just what sort of 'projection' we are talking about . . . and what the mechanism is according to which the pattern of the verb is projected to the more general pattern."

Ninio observes that children often begin using a single verb with a direct object long before a direct object appears with other verbs; moreover, she notes the overwhelming tendency for these "pathbreaking" verbs to be drawn from the set of general purpose verbs. In particular, children tend to use verbs meaning "want," "make/do," "put," "bring," "take out," or "give" before other verbs are used. In a longitudinal study, Ninio observes that SVO and VO patterns were initially produced with only one or at most a few verbs for a prolonged period. More and more verbs came to be used in an exponentially

increasing fashion; that is, there seemed to be more facilitation after ten verbs than after five, and so on. She suggests that this increase stems from the fact that children gradually abstract a more general and purely syntactic pattern on the basis of the early verbs, and that the growing generalization allows them to use new verbs in this syntactic pattern more and more easily.

On both Ninio's account and the present proposal, patterns are learned on the basis of generalizing over particular instances. As vocabulary increases, so does the strength of the generalization, making it progressively more and more easy to assimilate new verbs into the patterns. The two accounts complement each other in that Ninio proposes that general purpose verbs lead the way in the early acquisition of syntax, and the present proposal emphasizes the role of general purpose verbs in the acquisition of the semantics associated with basic syntactic patterns. It seems that early uses of general purpose verbs provide the foundation for both initial syntactic and semantic generalizations, and thus provide a route to the acquisition of form and meaning correspondences: i.e. constructions.

The accounts differ in their explanations as to *why* general purpose verbs should be learned so early. While Ninio notes that general purpose verbs are highly frequent and pragmatically relevant, she argues that the tendency for general purpose verbs to be used early in the VO and SVO patterns stems largely from a high degree of semantic transitivity in these general purpose verbs. However, as Ninio herself notes, many of the early general purpose verbs are not highly semantically transitive according to traditional criteria; i.e. they do not involve an agent acting on, or changing the state of, a patient argument (Hopper and Thompson 1980). For example, the general purpose verbs *want*, *see*, *get*, and *have* appear among the very first verbs in Ninio's corpus, and yet they are not highly transitive according to traditional criteria (see Goldberg, Casenhiser, and Sethuraman 2004 for discussion).

The present hypothesis is that the high frequency of particular verbs in particular constructions facilitates children's unconsciously establishing a correlation between the meaning of a particular verb in a constructional pattern and the pattern itself, giving rise to an association between meaning and form.

4.4 Experimental evidence for an effect of skewed input in language learning

Casenhiser and Goldberg (2005) designed an experiment to test learners' ability to learn to pair a novel constructional meaning with a novel form: exactly the task that the child faces when naturalistically learning language.

We created a novel pattern involving known nouns arranged in a non-English word order along with a nonsense verb. At the beginning of each scene, participants heard a simple present-tense sentence describing the scene, and at the end of the scene they heard a past-tense version of the same sentence.

GENERAL FORM:
<NOUN PHRASE$_1$ NOUN PHRASE$_2$ NONSENSE VERB + *o* >

EXAMPLE:
The spot the king moopos; The spot the king moopoed

VIDEO CLIP PAIRED WITH EXAMPLE:
a spot appears on the king's nose

The meaning of the phrasal pattern was that of APPEARANCE (a meaning novel for English phrasal patterns): the entity named by the first noun phrase comes to exist in the place named by the second noun phrase. For example, the intended meaning for the sentence *the sailor the pond neebod* was "the sailor sailed onto the pond from out of sight."

Fifty-one children aged 5–7 (mean age 6;4) were randomly and equally assigned to three conditions: the *control* condition, the *balanced frequency* condition, and the *skewed frequency* condition. The number of different novel verbs and the overall number of examples were controlled for: five different novel verbs and sixteen total examples were presented in both training conditions. In the balanced frequency condition, subjects heard the five novel verbs, each with a relatively low token frequency: three novel verbs occurring four times and two novel verbs occurring twice (4-4-4-2-2). In the high token frequency condition, subjects again heard the five novel verbs, but this time one had especially high token frequency of 8, while the other four verbs appeared twice each (8-2-2-2-2). The control condition watched the same film but with the sound turned off; thus any difference among groups can only be attributed to a difference in the linguistic input that subjects were exposed to, as ALL THREE CONDITIONS WATCHED EXACTLY THE SAME VIDEO. Training lasted less than three minutes.

The test was a forced choice comprehension task: subjects saw two film clips presented side-by-side on the screen and heard a sentence describing one of the clips. Sentences included six test trials with the novel phrasal pattern and *new* novel verbs; interspersed were six filler trials with new novel verbs in the familiar transitive pattern. Each test film-clip pair involved the same entity engaged in a similar action, but in only one did the entity appear on the scene within the clip (e.g. in one case, a sailor sails in on a boat from off the screen; in the paired foil clip the sailor sails around in a boat on screen). Half of the foil films involved transitive actions (a clown tossing a ball in the air was matched with a clown falling from out of sight onto the ball), and half involved intransitive actions (a flower growing taller was matched with a flower growing from out of the ground).

Subjects were asked to point to the film clip on the computer screen that corresponded to the description that they heard. Responses were coded for accuracy. The task is reminiscent of the preferential-looking paradigm, the main difference being that our subjects provided an unambiguous behavioral response, pointing to the matching scene instead of simply looking longer at one scene than another. Results are given in Fig. 4.1.

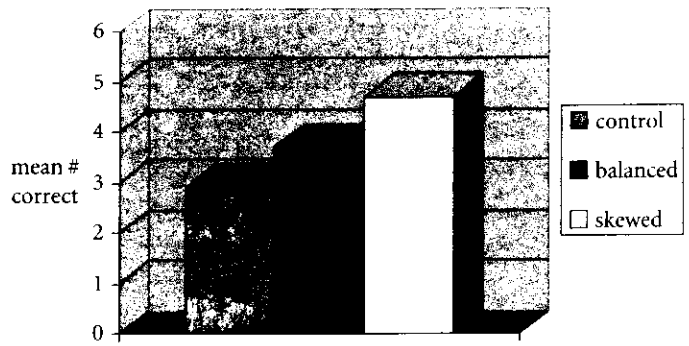

FIGURE 4.1 Experimental results of children's learning of a novel construction with novel verbs (based on Casenhiser and Goldberg 2005)

As expected, subjects in the control (no-sound) condition did no better than chance at choosing the correct scene. The balanced condition showed a statistically significant improvement over the control condition, indicating that they had learned something about the construction's semantics from the training film involving five relatively low token frequency verbs. As predicted by our hypothesis, the skewed frequency condition showed a statistically significant improvement in accuracy over the balanced condition.[4]

The results are striking: after only three minutes of training, children and adults are able to glean the novel abstract meaning that is associated with a

[4] An ANOVA confirmed a significant main effect for group, $F_{(2, 48)} = 11.57$, $p < .001$. Planned comparisons analyzed with Fisher's PLSD show that both the high-frequency and the balanced groups performed significantly better than the control group ($p < .001$ and $p < .05$ respectively). Moreover, the high-frequency group performed significantly better than the balanced group ($p < .01$).

novel formal pattern involving novel verbs and extend what they have learned to *new* utterances that use *new* novel verbs. Particularly facilitative is input in which one verbal token accounts for the balance of utterances, when the number of verb types is held constant.

The learning involved was implicit insofar as (*a*) no direct instructions were given to the subjects regarding what they were expected to learn during the training, and (*b*) subjects were unable to articulate explicitly the meaning of the construction when asked to paraphrase an instance of the new construction afterwards.

Goldberg, Casenhiser, and Sethuraman (2004) found similar results with adult subjects, as shown in fig. 4.2.

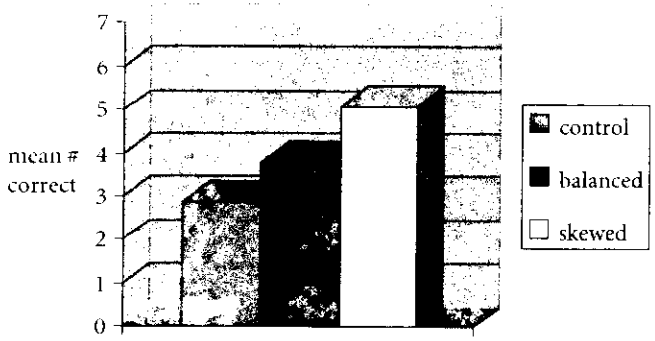

FIGURE 4.2 Experimental results of adults' learning of a novel construction with novel verbs[5]

Thus with less than three minutes of training, both children and adults demonstrated an ability to learn constructional meaning: they were able to extend the semantics of the construction to new novel verbs and new scenes in this forced choice comprehension task. Moreover, the results demonstrate that high token frequency of a single general exemplar does indeed facilitate the acquisition of constructional meaning.

Kidd, Lieven, and Tomasello (2005) have recently found relevant results in an unpublished study. Four-year-old children were asked to repeat sentences which involved complement-taking verbs. Some of the sentences had errors in them (e.g. *I say her give the present to her mom*). They found that 25 per cent of

[5] The experiment on adults, performed before the better-controlled experiment described above on children, did not rule out the possibility that watching the scenes helped the children instead of the language. This is because the videos as well as the language differed slightly in the two adult groups that watched the training films. The control condition in the experiment on adults did not hear the language associated with the films *or* watch either film.

the time, children would substitute a different main verb when repeating the sentence. The particular verb *think* accounted for fully 70 per cent of these substitutions (whether or not the sentence was corrected). In a second study, Kidd et al. found that even when none of the input sentences contained *think*, children nonetheless still displayed just as strong a tendency to substitute *think* as the main verb when repeating the sentences (81 per cent of substitutions by four-year-olds; 70 per cent of substitutions by six-year-olds). Since *think* is the most common verb occurring in the [V S] frame (accounting for almost 40 per cent of the tokens in their sample), this provides a further indication that young children's generalizations about constructions appear to be focused around particular verbs that appear frequently in those constructions.

Fast Mapping as evidence of UG?

It is possible that the quick learning of the mapping could be taken as an indication that the particular mapping is a part of universal grammar and is innately available. A mapping between Subj Obj V and <theme location V>, could be added to the set of mapping principles sometimes claimed to be universal. However, it is not clear that languages encode "appearance" in this way. In particular, the mapping violates at least one of the proposed universal linking rules suggested by both Pinker (1989) and Naigles, Gleitman, and Gleitman (1993): namely the generalization that a locative argument should be expressed by an oblique complement. The location argument used in the experiments would require prepositional marking to be considered an oblique in English; yet locations were expressed as simple noun phrases in the experimental stimuli. Moreover, the child's word-order parameter would presumably already be set by age six, and yet children had no trouble learning the SOV order. Therefore the idea that the mapping must be part of UG is without independent support.

4.5 What exactly did they learn? Examining the role of morphology and word order

We considered the possibility that in the experiments described above, subjects attended only to the morphological cue, which was constant: an *-o* suffix on each of the novel verbs. In a follow-up experiment, we removed this morphological ending, so that there was no stable cue to the construction except for word order (Casenhiser and Goldberg 2005).

We also wanted to make certain that children were able to recognize the novel word order in our novel appearance construction. To investigate this,

we systematically interspersed in the test films three novel transitive scenes with three novel scenes of appearance. A scene of appearance was available for each test trial; if the children learned that the scenes of appearance were paired with the SOV order, they should choose scenes of appearance when queried by instances involving the SOV order, while choosing transitive scenes when queried by the familiar SVO order. This is just what we found. Results confirmed that children were able successfully to learn the novel construction even without the morphological cue, in that they performed significantly different than the controls, as shown in Fig. 4.3.

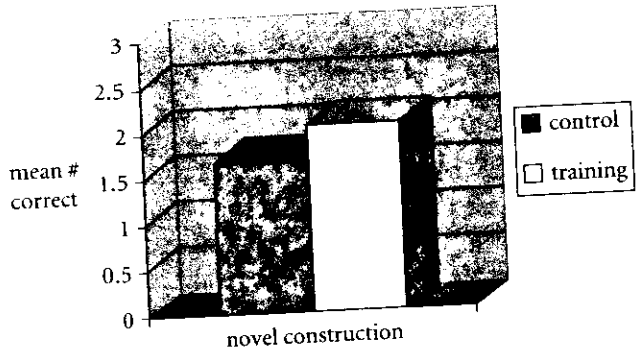

FIGURE 4.3 Average number of trials in which children correctly matched a new instance of the novel construction with a new scene of appearance (out of 3).

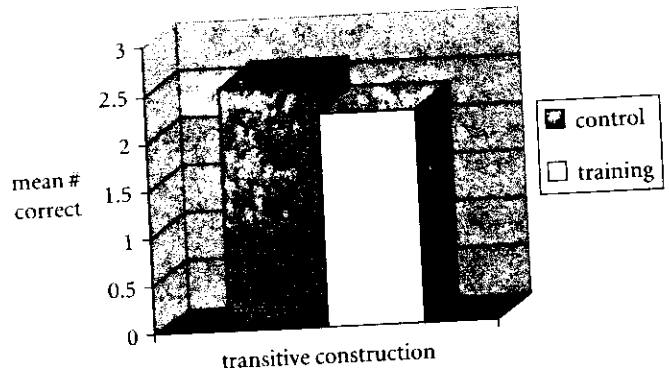

FIGURE 4.4 Average number of transitive stimuli correctly paired with transitive scenes (out of 3)

At the same time, children did not assign a scene of appearance to the transitive stimuli, but instead chose semantically transitive scenes. As expected, children were able to match a transitive stimulus to a transitive scene with or without training, as they were already familiar with the transitive construction: the performance on transitives does not differ significantly across the two groups (Fig. 4.4). See Casenhiser and Goldberg (2005) for specifics.

In the following section we outline why the fact that construction learning is facilitated by high token frequency is in fact expected, given general findings in the non-linguistic categorization literature.

4.6 The role of skewed frequency in non-linguistic categorization

Research in category learning has demonstrated that there is a strong correlation between the frequency with which a token occurs and the likelihood that it will be considered a prototype by the learner (Nosofsky 1988; Rosch and Mervis 1975). Homa, Dunbar, and Nohre (1991) found that token frequency was an important variable at early and intermediate stages of category learning, with increased token frequency facilitating category learning. In learning generalizations about dot patterns, Posner, Goldsmith, and Welton (1967) demonstrated that the rate at which subjects classified patterns correctly was a direct function of the amount of distortion from their respective prototypes: the less variability or distortion, the faster the category was learned.

Elio and Anderson (1984) set up two conditions relevant to the current discussion. In the "centered" condition, subjects were initially trained on more frequently represented, more prototypical instances, with the study sample growing gradually to include more members of the category.[6] In the "representative" condition, subjects were trained on a fully representative sampling from the start. In both conditions, subjects were eventually trained on the full range of instances. Elio and Anderson demonstrated that the order in which subjects received the more prototypical instances played a role in their learning of the category. In particular, they demonstrated that categories were learned more accurately in the "centered" condition; the "representative" condition yielded poorer typicality ratings and accuracy during the test phase on new instances. Elio and Anderson observe, "The superiority of the centered condition over the representative condition suggests that an initial, low-variance sample of the most frequently occurring members may allow the

[6] The study involved descriptions of people belonging to one of two clubs, with members' descriptions varying on five four-valued dimensions.

learner to get a 'fix' on what will account for most of the category members" (p. 25). They go on to note that "a low-variance sample, in which there is a maximum amount of similarity among items, is particularly conducive to forming strong category generalizations" (p. 28).[7] Similar results were found by Avrahami et al. (1997) who demonstrated that subjects learned categories better when presented with several ideal positive cases followed by borderline cases than if they were presented with sequences that emphasized category boundaries from the start.[8]

Gentner likewise notes that processes of analogy required for generalization (her "structural alignment") are facilitated when instances being compared are similar to one another. Gentner, Loewenstein, and Hung (2002) performed an experiment that illustrates this idea: they showed children a particular picture of a Martian to be used as a standard for comparison, and two alternative Martian creatures. The standard Martian and one of the alternatives shared one body part, while the distinct Martian didn't. Children were asked, "This one has a *blick*. Which one of these has a *blick*?" The results demonstrated that if the two alternatives were highly similar to the standard, children were better able to pick out the relevant shared body part; when they were only weakly similar, finding the body part was more difficult. In addition, Gentner et al. demonstrated that children who were tested in the high-similarity condition first, were subsequently more successful on the low-similarity items than children who had the same amount of experience with only low-similarity items.

Moreover, categories that are identifiable with a salient type of stable feature are easier to learn than categories in which the feature is instantiated in different ways, even when the variability is relevant to the feature dimension (Markman and Maddox 2003). The analogy to language is that constructions that are instantiated (to a great extent) by a single verb should be initially easier to learn than constructions that are instantiated by many different verbs.

[7] Interestingly enough, Elio and Anderson (1984) also found that when subjects were explicitly asked to form hypotheses about what criteria governed category membership, the advantage of learning the centered instances first disappeared. They therefore conclude that the advantage is only evident when the learning is *implicit*. Implicit learning involves knowledge that is not accessible to consciousness, is fairly complex and abstract, an incidental consequence of some task demand, and preserved in cases of amnesia (Seger 1994). Relevantly, language learning is an excellent example of implicit learning, since it is largely learned below the level of consciousness, is very complex and ultimately quite abstract, is a consequence of trying to communicate, and is preserved in cases of amnesia.

[8] Stimuli in this experiment also consisted of non-linguistic stimuli: variable-sized semicircles with variably oriented radial lines.

We performed an experiment with a design parallel to that described for construction learning to test whether the advantage of a single high frequency exemplar would hold in a non-linguistic task as well (Goldberg and Casenhiser forthcoming). We created a random dot pattern (with ten dots) to be used as a prototype as well as four systematic variations from the prototype pattern. Subjects in the skewed frequency group saw twice as many instances of the prototype dot pattern as any of the other dot patterns. Subjects in the balanced group were not given this preferential training with the prototype; instead, they saw a more balanced distribution of the prototype pattern in comparison to the other dot patterns. Subjects were again tested with a forced choice to determine if they were able to distinguish a *new* variation of the prototype from a dot pattern generated randomly. New variations used at test differed from the prototype to the same degree as the variations used in training.

Twenty-four college undergraduates were tested, distributed randomly and equally into the two groups. The results demonstrate that subjects in the skewed frequency group were significantly more accurate at test than those in the balanced frequency group, thus confirming the idea that learning of categories generally is facilitated when a prototype is encountered with high frequency as opposed to experience with the same variety of instances, when the prototype does not account for the balance of items (see Fig. 4.5).

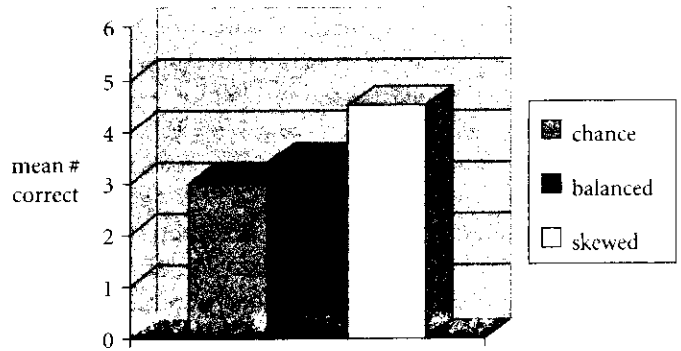

FIGURE 4.5 Mean number of times subjects were able to match correctly the new variation of the random dot pattern (based on Goldberg and Casenhiser, forthcoming).[9]

[9] An ANOVA confirmed a significant main effect for group, such that the high frequency training condition performed significantly better than the balanced frequency training condition $F(1,27)=6.78$, $p < .02$. Chi-square test with 6 degrees of freedom was performed to compare subjects' performance to chance. Subjects in the balanced group did not perform significantly above chance $p = .13$. Subjects in the high frequency group did perform significantly above chance $p < .01$.

These results demonstrate that the learning advantage of skewed frequencies is not specific to language.

To summarize, we know that frequency and order of acquisition play key roles in category formation in that training on prototypical instances frequently and/or early facilitates category learning (Bruner, Goodnow, and Austin 1956; Kruschke 1996; Maddox 1995; Nosofsky 1988). This generalization together with the experimental evidence discussed in the previous section suggests that the very frequent and early use of one verb in a pattern facilitates the learning of the semantics of that pattern. The corpus findings demonstrate that exactly this sort of tailor-made input is available to language learners. We suggest, for example, after using many sentences with *put* in the construction in (1), children come to associate the meaning of *put* with the construction even when the verb is not present as in (2):

(1) She put a finger on that.

(2) He done boots on. (STE, 28 months; Bates, Bretherton, and Snyder 1988)

The result is that the meaning of roughly "X causes Y to move Z_{loc}" comes to be associated with the Subj V Obj $Obl_{path/loc}$ formal pattern.

The implications of this work are potentially far-reaching, since it is very common for constructions to exhibit statistical "skewing" toward a subset of possible data types. That is, tokens of constructions are typically centered around one or a few specific words, or around a semantic prototype, even when they potentially occur with a much broader range of words or meanings (Brenier and Michaelis 2005; Cameron-Faulkner, Lieven, and Tomasello 2003; Deane 2003; Diessel 2001; Goldberg 1996, 1999; Hunston and Francis 1999; Scheibman 2002; Stefanowitsch and Gries 2003; Thompson and Hopper 2001). In the case of argument structure constructions, we have demonstrated a facilitory effect for a high-frequency verb.[10]

In the case of other constructions, relevant skewing of the input could be around a noun, adjective, or complementizer. For example, Brenier and Michaelis (forthcoming) note that the noun *thing* appears in more than half of the tokens of the double *is* construction (e.g. *The thing is, is that...*). Other nouns appear such as *problem, difficulty, question, point, issue,* and *rumor,* but with much less frequency. Thus the general facilitory effect demonstrated in the experiments described above may have a general utility in language learning.

[10] See Del'Orletta et al. (forthcoming), Alishahi and Stevenson (forthcoming), and Borovsky and Elman (MS) for computational models that capture this effect.

The advantage of skewed input: an anchoring effect

One way to think of these effects is that they may involve a type of **cognitive anchoring,** where a high-frequency type of example acts as an anchor, i.e. a salient standard of comparison. Numerical anchoring effects have been demonstrated to be quite robust in cognitive psychology. Tversky and Kahneman (1974) demonstrated that people's numerical estimates are influenced by a salient standard of comparison even if that comparison was completely irrelevant to the task at hand. For example, estimations of the percentage of African nations in the UN was influenced by an arbitrary number determined by a manipulated spin of the "wheel of fortune." That is, estimates were significantly higher among the group exposed to a high anchor (guessing 45 per cent if the wheel stopped at 65) than if the group were exposed to a low anchor (guessing 25 per cent if the wheel stopped at 10); i.e. guesses were assimilated towards the anchor. Anchoring effects persist over time (Mussweiler 2001), have an effect even if the anchor values are obviously irrelevant (Strack and Mussweiler 1995), and even when subjects are warned about the distorting effect of anchors (Wilson, Houston, Etling, and Brekke 1996). At the same time, anchoring effects are markedly stronger when the anchor is perceived to be *relevant* to the subsequent task (Wilson, Houston, Etling, and Brekke 1996). The existence of a type of instance that occurs with high token frequency may well provide a highly relevant cognitive anchor, serving to organize memory and reasoning about other related types.

Single high frequency exemplar facilitory, but not necessary

It is important to observe that it is *not necessary* for there to be a single verb with frequency far greater than other verbs for successful learning to take place. The correlation between form and meaning can be learned by noting their association across several distinct verbs, each with relatively low frequency. This is in fact evident from our data insofar as subjects in the balanced condition outperformed those in the control condition: they clearly did learn something from witnessing several verbs in the construction, each with relatively low frequency. This is important because in naturalistic data, there is not always a single verb that has far greater frequency than other verbs, at least if constructions are defined as generally as possible: the transitive construction may be just such a case (Sethuraman and Goodman 2004).

I need to emphasize that I am *not* claiming that general purpose verbs are necessarily the very first verbs uttered. Neither the cross-sectional corpus data nor the experimental data address this question. Moreover, longitudinal studies have yielded differing answers (see Ninio 1999 for a suggestion that

such verbs are likely to be the very first or second verbs uttered; but see Campbell and Tomasello 2001 for evidence that they are not always the very first verbs).[11]

Reconciling results with the observation that high token frequency is inversely related to productivity

Bybee (1995) has argued that morphological tokens with especially high frequency do *not* lead to generalizations because they are routinized to such an extent that their internal structure is unanalyzed and therefore unavailable for analogical extension. Bybee's argument is based on irregular morphological items such as *went* and *am* which clearly do not lend themselves to generalizations. It is quite possible that such morphological forms are used without internal analysis. However, it is clear that the constructions under discussion represent a different case. They must be analyzed because they contain argument positions that must be filled. The caused-motion pattern with *put*, for example, has slots that are filled by different arguments from one use to the next. The same is true with our constructed stimuli. Psycholinguistic experiment has revealed that even frequent VP idioms such as *kick the bucket* are analyzed in online sentence comprehension (Peterson et al. 2001). In fact, VP idioms have to be analyzed since their internal structure is minimally variable: the verb may be variably marked for tense or agreement (*kick the bucket*; *kicked the bucket*; *kicks the bucket*).

The experiments discussed above had a type frequency of five, not one, and so were unlikely to be treated as fixed idioms. However, we have since run an experiment that includes a training condition in which the only novel verb is *moop*. In this case, subjects do in fact show evidence of less robust generalization (Goldberg and Casenhiser in prep.). Thus a combination of type and token frequencies are relevant.

4.7 Reconciling evidence for fast mapping of phrasal forms and meanings with conservative learning

The finding that the mapping between a new form and meaning pair can be generalized so quickly, with so little input, appears to run counter to the large body of evidence that indicates that children are very conservative learners.

[11] See Tomasello and Stahl (2004) for arguments that extreme care must be taken to avoid confusing high frequency with early acquisition when intermittent sampling techniques are used. Ninio's (1999) data combined parental reports with intermittent sampling, so the data should be reliable (unless high frequency affects the veridicality of maternal reports). See Goldberg, Casenhiser, and Sethuraman (2004) for discussion of Ninio's Pathbreaking Verb analysis.

That is, as discussed in Chapter 3, many studies have demonstrated that the initial production of argument structure patterns is very conservative in that children stick closely to the forms they have heard used with particular verbs (Akhtar 1999; Baker 1979; Bates and MacWhinney 1987; Bowerman 1982; Braine 1976; Brooks and Tomasello 1999; Gropen et al. 1989; Ingram and Thompson, 1996; Lieven, Pine, and Baldwin 1997; MacWhinney 1982; Olguin and Tomasello 1993; Schlesinger 1982; Tomasello 1992, 2000, 2003).

It is possible that the difference between conservative learning found in other tasks and the fast mapping discussed here is age-related. Clearly work on younger children is needed to demonstrate that the categorization strategies at play in the experiments discussed above are operative at younger ages, when children actually begin to acquire language.

However, the reason for the conservative learning has been claimed to be that children—and *adults*—operate with a usage-based model of language (Tomasello 2003). The difference between children and adults has been thought to involve a difference in the amount of experience with the ambient language. Therefore we would not necessarily expect to find quicker generalizations in older children than in younger children if they are both exposed to the same amount of input on a novel construction.

Therefore, it seems more likely that the difference is task-dependent. Studies that have documented conservative learning have used a variety of methods including (*a*) spontaneous production, (*b*) elicited production, and (*c*) act-out tasks. These tasks require *recall* of at least aspects of the pairing of form and meaning. Clearly this is true in the case of production, since in order to produce an utterance, the child must be able to recall its form correctly and use it appropriately. In act-out tasks, children are encouraged to act out scenes that they hear verbal descriptions of; this also requires that the child recall the relevant meaning associated with the given form. The task outlined here, on the other hand, only requires that children *recognize* the relevant meaning from among two given alternatives. This is more akin to the preferential-looking paradigm that has been argued in fact to demonstrate early generalizations (Gleitman 1994; Naigles and Bavin 2001), although somewhat controversially (Tomasello 2000, 2003). It may be that learners are able to make tentative generalizations on the basis of very little input; only if those generalizations are reinforced, however, do they become a stable source of productive language creation (see Section 3.7 for evidence of early tentative generalizations).

It seems that children are both conservative and quick generalizers. Indeed, if in fact language is usage-based, as argued in Chapter 3, then we should expect to find evidence of both item-based learning and the quick grasping

for generalizations. Children must "make sense of" their language so that they are able productively to use it in new situations. Therefore they are likely constantly to seek out generalizations. The advantages to learning constructions are discussed in Chapter 6. But before we turn our attention there, we first address the issue of how generalizations are constrained. This is the subject of the following chapter.

4.8 Summary

To summarize, it appears that the input is structured in such a way as to make the generalization of argument structure constructions straightforward. One particular verb accounts for the lion's share of tokens of each argument frame considered in an extensive corpus study on the Bates et al. (1988) corpus, in both mothers' and twenty-eight-month-old children's speech. The dominance of a single verb in the construction facilitates the association of the meaning of the verb in the construction with the construction itself, allowing learners to get a "fix" on the construction's meaning. Research on construction learning and on the categorization of random dot patterns presented here supports this idea. In this way, grammatical constructions may arise developmentally as generalizations over lexical items in particular patterns. As Elman has suggested, "Knowledge of constructions is a straightforward extension, by generalization, of knowledge of groups of words that behave similarly" (Elman forthcoming: 13).

The present proposal for how the semantics associated with constructions is learnable from the input directly undermines the "paucity of the stimulus" argument as it is aimed at the particular issue of linking generalizations. Before we decide that language-specific properties must be innate, it is worth investigating how they might be learned, given general cognitive processes such as categorization, together with a closer look at the input children receive.

Still, it might be argued that while linking generalizations may be learnable, they are not learned, insofar as they exhibit some striking regularities across languages. We address this argument in Chapters 7 and 9, where it is argued that the cross-linguistic generalizations that exist are readily attributed to general cognitive, pragmatic, and processing constraints and do not require recourse to any genetic, domain-specific linguistic knowledge.

5

How generalizations are constrained

A boulder is a ROCK?... TWO names?! <incredulous>... Is "boulder" a LAST name?

(Zach 3;3)

The previous chapter argued that categorization of attested instances leads learners to generalize grammatical constructions beyond their original contexts. In this chapter, we address the question of how the learner knows whether a pattern can be extended for use with new verbs for the sake of production, that is, the question of productivity. We need to explain how generalizations are constrained: How do children avoid or recover from overgeneralizing their constructions? Children are not strictly conservative, producing only what they have heard, and yet they are not reliably corrected when they produce overgeneralizations or ungrammatical utterances of other kinds. As many have noted, the most obvious possible solutions are not viable, including a reliance on overt corrections or corrective repetitions (Bowerman 1996; Brooks and Tomasello 1999; Pinker 1989). How then can children retreat or avoid overgeneralizations?

There has been much discussion in the literature about productivity and I do not attempt to review it all here (Baker 1979; Bowerman 1988; Brooks and Tomasello 1999; Goldberg 1995; Pinker 1989). At least four factors have been proposed as relevant to predicting a pattern's productivity: (*a*) the number of times an item occurs—its token frequency or degree of *entrenchment;* (*b*) *statistical pre-emption*: the repeated witnessing of the word in a competing pattern (Brooks and Tomasello 1999; Goldberg 1993, 1995; Marcotte 2005; Pinker 1981); (*c*) the absolute number of distinct items that occur in a given pattern or a pattern's *type frequency* (Bybee 1985; Goldberg 1995; MacWhinney 1978; Plunkett and Marchman 1991, 1993); and (*d*) the variability of the items

that occur in a given pattern: a pattern's *degree of openness* (Bowerman and Choi 2001; Bybee 1995; Janda 1990; Pinker 1989).

In order for *any* of these factors to work in constraining generalizations, some memory of how particular words are used in particular constructions is essential. That is, the only way to account for partial productivity is to recognize item-specific knowledge (cf. discussion in Chapter 3).

5.1 Statistical Pre-emption

Several theorists have argued that the process of *entrenchment* or hearing a pattern with sufficient frequency plays a key role in constraining overgeneralizations (Braine and Brooks 1995; Brooks and Tomasello 1999). For example, Braine and Brooks propose a "unique argument-structure preference" such that once an argument structure pattern has been learned for a particular verb, that argument structure pattern tends to block the creative use of the verb in any other argument structure pattern, unless a second pattern is also witnessed in the input. Brooks et al. (1999) demonstrated that children in an experimental setting were more likely to overgeneralize verbs that were used infrequently (e.g. to use *vanish* transitively), and less likely to overgeneralize frequently occurring verbs (e.g. to use *disappear* transitively). The difference was attributed to the difference in frequency. This sort of explanation, however, does not address the fact that verbs that frequently appear in one argument structure pattern can in fact be used creatively in new argument structure patterns, without any trace of ill-formedness as in:

(1) She sneezed the foam off the cappuccino.
(2) She danced her way to fame and fortune.
(3) The truck screeched down the street.

Upon closer inspection, effects that might be ascribed to entrenchment are better attributed to a statistical process of pre-emption, critically involving the role of semantic or pragmatic contrast.

That is, one way that overgeneralizations can be minimized is based on the idea that more specific knowledge always pre-empts general knowledge in production, *as long as either would satisfy the functional demands of the context equally well.* That is, more specific items are preferentially produced over items that are licensed but are represented more abstractly, as long as the items share the same semantic and pragmatic constraints. The idea that more specific information should override more general information when the two are functionally equivalent is one with much precedent (cf. e.g. the Elsewhere Condition of Kiparsky 1973, who attributes the generalization to Panini).

In the case of morphological pre-emption (or blocking), this idea is already familiar. While the agentive nominalizing suffix, *–er*, for example, is partially productive (one who is prone to squabbles can be called a *squabbler*), it does not apply to words for which there already exists an agentive nominal counterpart. For example, while someone can *ref* a game, he is not a *reffer*, because *referee* pre-empts the creation of the new term *reffer*. Similarly, *went* pre-empts *goed*, *children* pre-empts *childs*, etc. The pre-emption process is straightforward in these cases because the actual form serves the identical semantic/pragmatic purpose as the pre-empted form. This idea implies that there is no complete synonymy—while two words may refer to the same thing (*auto* and *car*), they will generally differ in terms of register (as *auto* and *car* do), construal (as *thrifty* and *stingy* do), or in terms of dialect (as *soda* and *pop* do) (Clark 1987).

DiSciullo and Williams (1987) discuss a case of pre-emption in which the existence of a morphological form pre-empts the use of a phrasal pattern. In particular, the ready availability of a lexical comparative pre-empts the formation of a comparative phrase: e.g. *better* renders the adjective phrase *more good* ill-formed. This is expected on a constructionist view, since the morphological form [adj-er] and the phrasal pattern [more adj] are both stored constructions, and they have nearly identical meaning and pragmatics. At the same time, if the instance of the morphological comparative is not stored as an entrenched lexical item, there should be some leeway in choosing the phrasal form, even when the phonology of the adjective would allow it to appear with the morphological comparative. That is, the process of pre-emption requires that an alternative form be more readily available than the pre-empted form. As expected, in fact, a cake can be *moister* or *more moist* than another cake, *fresher* or *more fresh*, *duller* or *more dull*. *Moister, fresher, duller,* being relatively infrequent, appear to be created on the fly from the [adj-er] lexical construction and not stored as independent words that would block the creative formation of *more moist*, *more fresh*, or *more dull*.[1]

The role of pre-emption between two phrasal forms requires explanation, since expressions formed from distinct phrasal constructions are virtually never semantically and pragmatically identical. Any two phrasal constructions will differ either semantically or pragmatically (or both). As we discussed in Chapter 2, the ditransitive construction in (4) is distinct semantically and pragmatically from the prepositional paraphrase. Thus knowledge that the

[1] The role of frequency is sometimes overlooked in discussions of these constructions because the [adj-er] form requires single syllable adjectives and such shorter forms tend to be highly frequent. Therefore, the comparative forms are also often sufficiently frequent to pre-empt the phrasal forms.

prepositional paraphrase is licensed as in (5) (based on positive evidence) should not in any simple way pre-empt the use of the ditransitive (Bowerman 1996) in(4), which is ill formed:

(4) *She explained me the story.
(5) She explained the story to me.

In fact, a large number of verbs do freely appear in both constructions. Goldberg (1995), following up on a suggestion made by Pinker (1984: 400), argued that a statistical form of pre-emption could play an important role in learning to avoid expressions such as (4), once a speaker's expectations are taken into account in the following way. In a situation in which construction A might have been expected to be uttered, the learner can infer that construction A is not after all appropriate if, consistently, construction B is heard instead.

The type of counterfactual reasoning required may seem overly complex. However, it has been demonstrated in another domain that young infants are entirely capable of a very similar sort of logic. Gergely, Bekkering, and Kiraly (2002) divided fourteen-month-old babies into two conditions. In one condition the babies watched an experimenter turn on a light in front of him by using his head instead of his hands. In the second condition, the experimenter also turned on the light with his head, but was simultaneously holding a blanket close to his chest with both hands, feigning chills. They found that when the demonstrator's hands were free, 69 per cent of the babies mimicked his head action to turn on the light (a good example of humans' predilection to imitate). But only 21 per cent of the infants did so when the adult's hands were occupied. In the latter case, the infants' reasoning appeared to be: "if the experimenter were able to use her hands she would have (but they were busy holding the blanket)." In the condition which saw the experimenter use the head action despite having her hands free, the child appeared to reason, "if the experimenter meant to use her hands she would have; since she didn't, perhaps she used her head for some reason." Similarly, the reasoning required for pre-emptive processes to work is "if the person meant to use the other formulation she would have; therefore, since she didn't, perhaps she used the alternate formulation for a reason."

The fact that statistically based pre-emption involving related, but non-synonymous constructions, plays a role in avoiding overgeneralizations has in fact been demonstrated empirically. Brooks and Tomasello (1999) found that children aged six or seven were less than half as likely to productively produce a novel verb in a transitive frame when the verb had been modeled in both an intransitive and periphrastic causative construction, than when it was only

modeled in the simple intransitive. For example, if the child had heard both *The ball is tamming*, and *He's making the ball tam*, then they were less likely to respond to "what's the boy doing"? with *He's tamming the ball*, than they were if only the simple intransitive had been witnessed. It seems that hearing the novel verb used in the periphrastic causative provided a readily available alternative to the causative construction, statistically pre-empting the use of the latter. That is, hearing a periphrastic causative in a context in which the transitive causative would have been at least equally appropriate led children to avoid generating a transitive causative in a similar contextual situation.

The Brooks et al. (1999) finding that high-frequency verbs are less likely to be overgeneralized than low-frequency verbs is consistent with the idea that it is pre-emption that prevents overgeneralization, not the frequency of the verb per se. That is, the pre-emptive context in which *disappear* might have been expected to occur transitively but instead is witnessed intransitively (in a periphrastic causative construction) occurs more frequently than the same pre-emptive context for *vanish*.

It might seem that the conditions for applying statistical pre-emption are not common enough to constrain generalizations. However, defending the same type of pre-emptive process, Marcotte (2005) observes that instances of such indirect negative evidence are readily available to the child:

> children's ability to detect shared meanings between their own utterances and adult ones is an unspoken but crucial precondition to obtaining even positive evidence. Both positive and negative evidence are the outcome of a process of comparison between the child's parse of an adult utterance in its context, and a child-generated representation expressing the same meaning in that context. Matches yield positive evidence, mismatches yield negative evidence.

In a computer modeling study designed to analyze how spatial terms can be learned on the basis of positive evidence, Regier (1996) found that learning was dramatically improved if each positive use of a spatial term was taken as a statistical indication by the system that all of the other possible spatial terms were inappropriate. The inference is necessarily only statistical: the model would have been ill-advised to assume that every use of *above* indicated that *over* had not been a possible alternative, since the two are often both equally applicable. Still, by treating each use of *above* as a tentative indication of "not *over*" and "not *across*" and "not *under*," etc., the model was able to learn to map words successfully onto an impressive variety of spatial configurations. Alishahi and Stevenson (forthcoming) have successfully applied statistical pre-emption in a computational model of argument structure learning; the system recovers from overgeneralizations such as *She falled the cup*, by con-

sistently hearing *fall* used intransitively even when an actor argument was part of the message to be conveyed (as in *She made the cup fall*).

As is discussed more in Chapter 7, there is a clash in information-structure properties that results in a preference for (7) over (6). In learning to avoid examples like (6), the child may be aided by statistical preemption in the input:

(6) ??Who did she give a book?
(7) Who did she give a book to? (preferred, despite prescriptive injunction against stranded prepositions)

That is, when a learner might expect to hear a form like that in (6), she is statistically overwhelmingly more likely to hear a form such as (7). (In online data, actual occurrences of questioned prepositional goals outnumber questioned ditransitive recipients by roughly forty to one; see Chapter 7 for discussion). This statistical pre-emption may lead the child to disprefer questions such as that in (6) in favor of ones such as (7).

The pre-emptive process, unlike the notion of simple high token frequency, predicts that an expression like (8) would *not* be pre-empted by the overwhelmingly more frequent use of *sneeze* as a simple intransitive (as in (9)) because the expressions do not mean at all the same things.

(8) She sneezed the foam off the cappuccino.
(9) She sneezed.

At the same time, frequency does play some role in the process of statistical pre-emption exactly because the pre-emption is statistical. Only upon repeated exposures to one construction in lieu of another related construction can the learner infer that the second construction is not conventional. As noted above, this requires that a given pattern occur with sufficient frequency.

5.2 Type frequency/Degree of Openness of a pattern

The process of statistical pre-emption is a powerful way in which indirect negative evidence can be gathered by learners. At the same time, it cannot account fully for children's lack of overgeneralizations. Constructions may be either too low frequency or too semantically or pragmatically specialized for another construction effectively to pre-empt their appearance (cf. discussion in Goldberg 1995: ch. 5). Moreover, upon first encountering a novel verb, speakers presumably know something about the contexts in which it can appear and the contexts in which it cannot, without there being a possibility of a pre-emptive context (since it is a new word).

Several authors have proposed that *type frequency* correlates with productivity (Bybee 1985, 1995; Clausner and Croft 1997; Goldberg 1995). Constructions that have appeared with many different types are more likely to appear with new types than constructions that have only appeared with few types. For example, argument structure constructions that have been witnessed with many different verbs are more likely to be extended to appear with additional verbs. To some extent, this observation has to be correct: A PATTERN IS CONSIDERED EXTENDABLE BY LEARNERS ONLY IF THEY HAVE WITNESSED THE PATTERN BEING EXTENDED.

At the same time, it is clear that learners do not generate new instances on the basis of type frequency alone. In fact, a training study by Childers and Tomasello (2001) sought to find a role for increased type frequency in the productive use of the transitive construction, but failed to find an effect. The degree of semantic relatedness of the new instances to instances that have been witnessed is likely to play at least as important a role as the simple type frequency. Constructions that have been heard used with a wide *variety* of types are more likely to be extended broadly than constructions that have been heard used with a semantically circumscribed set of types. That is, learners are fairly cautious in producing utterances based on generalizing beyond the input; they can only be expected confidently to use a new verb in a familiar pattern when that new verb is relevantly close in meaning to verbs they have already heard used in the pattern.

In writing about how non-linguistic inductions are made, Osherson and colleagues (Osherson et al. 1990) propose a relevant notion of COVERAGE. They demonstrate that speakers are more confident about generalizing a property to a new instance to the degree that the new instance fits within the category determined by entities known to exhibit the property. For example, speakers are more confident of the soundness of the conclusion in (A), that rabbits have some property X,[2] than they are of the same conclusion in (B):

(A) assumption 1: Lions have property X.
assumption 2: Giraffes have property X.

Conclusion: Rabbits have property X.

(B) assumption 1: Lions have property X.
assumption 2: Tigers have property X

Conclusion: Rabbits have property X.

[2] The property X in the above syllogisms is typically filled in by a "blank" predicate, i.e. a predicate of which subjects have no knowledge.

Intuitively this is because the assumptions in (B), that lions and tigers have some property X, tells us something only about large felines, and says nothing at all about rabbits, whereas the assumptions in (A) lead us to suspect that the generalization may hold of all mammals, in which case there is more reason to believe the property holds also of rabbits as an instance of the category mammal. Osherson et al. characterize coverage as the degee to which the premise categories are similar to members of the lowest-level category that includes both the premise and the conclusion categories. In the case of language, the analogy is clear: The greater the degree to which previously attested instances fill a semantic space that includes the potential target instance, the more confident speakers will be in using the target instance.

5.3 Applying pre-emption and openness to particular examples

Let us consider again the novel utterances in (1)–(3) repeated below:

(10) She sneezed the foam off the cappuccino.
(11) She danced her way to fame and fortune.
(12) The truck screeched down the street.

The reason *sneeze* can readily appear in the caused-motion construction as in (10) is because *sneeze* can be construed to have a meaning relevantly like other verbs that readily appear in that construction, as a verb that effects a causal force. Other verbs that appear in the construction indicate that the causal force may involve air (*blow*), and need not be volitional (*knock*). Since *sneeze* has not been pre-empted in this use—given that this meaning has only rarely if ever been expressed—(10) is fully acceptable. Example (11) of the *way* construction is fully acceptable since a large variety of verbs have been attested in that particular construction. *Dance* is relevantly like any of a number of these attested verbs including in particular various verbs of performance (e.g. *sing*). The intended meaning in (11) is also not pre-empted by another construction, since the *way* construction is both relatively infrequent and has a very specialized meaning, roughly "to (metaphorically) travel despite difficulty or obstacles" (Goldberg 1995). That is, since it can only very rarely be expected with any degree of confidence, its non-occurrence in a given context cannot be taken as evidence that it is infelicitious. Finally, (12) is an acceptable use of *screech* because other verbs of sound emission are attested in the intransitive motion construction with a similar meaning (e.g. *rumble*, used to mean "to move causing a rumbling sound") and since again, it is not likely that the meaning "to move causing a screeching sound" could have been systematically pre-empted by another construction.

Thus a combination of both conservative extension based on semantic proximity to a cluster of attested instances, together with statistical pre-emption can go along way toward an avoidance of overgeneralizations in the domain of argument structure.

5.4 Combining multiple cues

Recent work in automated learning algorithms has made possibly relevant discoveries. It has been shown that any collection of rules of thumb ("weak hypotheses") that are imperfect but correct more often that chance can be combined to yield a highly accurate predictive system through a "boosting" algorithm (Freund and Schapire 1999; Meir and Rätsch 2003; Schapire 1990). Boosting has been applied to both traditional symbolic and to connectionist network systems. The algorithm works roughly as follows. It takes a representative subset of data to use as training and formulates the first weak hypothesis (h1), where $h1(x_i)$ is the function that classifies each x_i in the training set according to h1. The training items x_i, for $i = 1$ to n, are then weighted more heavily if their classification was incorrect according to h1, less heavily if their classification was correct. The relevance of this sort of procedure to language learning clearly requires some sort of indirect negative evidence, such as the statistical preemption just discussed. This procedure provides a distribution for all the x's.

The value of a weak hypothesis is measured by its error rate, taking into account the distribution. That is, it is equal to the probability that x is misclassified by h, where x is chosen randomly on the basis of the distribution of weights of x's.

A second weak hypothesis (h2) is formulated on the basis of the revised data set (the data set is distinct from the original data set insofar as the distribution assigned to x's is distinct, with the system treating those examples with the strongest weights as the most important to get correct). Again the system iterates, slightly increasing the weight of all incorrectly classified examples and decreasing the weight of all correctly classified examples. These weak hypotheses are then combined via an algorithm "ADABoost" to yield H(x), where H(x) is the weighted sum of all weak hypotheses, where the weight of each hypothesis is determined by its degree of value as described above. Boosting has been shown to yield an algorithm with an arbitrarily small error rate on the basis of such a collection of rules of thumb that are simply above chance.

This sort of algorithm appears to be well suited for determining the conditions in which one construction is chosen over an alternative, when

several factors are involved. For example, it is well known that the choice of the ditransitive construction over its prepositional paraphrase is conditioned by the semantic class of the verb, the phonology of the verb (Germanic sounding versus Latinate sounding), semantic characteristics of the recipient argument, and the length, complexity, and focus structure of the two non-subject arguments. Each of these factors can be identified as a heuristic—the boosting algorithm would provide a relevant tool to identify interactions and weights among the factors in order to predict the preferred form for new instances.

5.5 Conclusion

A pattern can be extended to a target form only if learners have witnessed the pattern being extended to related target forms, and if the target form has not systematically been pre-empted by a different paraphrase. Statistical pre-emption provides indirect negative evidence to learners—allowing them to learn to avoid overgeneralizations.

In choosing between allowable alternatives that are conditioned by multiple factors, speakers must combine multiple cues. Successes in machine-learning algorithms that involve iterative refinements of combinations of cues provide promising avenues to explore.

In some ways the task of learning would seem to be made easier if speakers never generalized beyond what they had positive evidence for. For example, it would seem to simplify the task if languages used each particular verb in its own particular frame, without generalizing across verbs or using verbs in novel ways. However, in the following chapter we investigate the benefits of generalizing beyond the input to the level of argument structure constructions.

6

Why generalizations are learned

> People do not rely on simple observation or feature learning in order to learn new concepts. They pay attention to the features that their prior knowledge says are the important ones.
>
> (Murphy 2002: 63)

As many psychologists have emphasized, human categorization is generally driven by some functional pressure, typically the need to predict or infer certain properties on the basis of perceived characteristics (Anderson 1991; Holland et al. 1989; Kersten and Billman 1997; Leake and Ram 1995; Murphy 2002; Ross and Makin 1999; Wisniewski 1995). That is, cognitive systems do not generalize randomly or completely. Holland et al. (1989), in their monograph on induction, emphasize that generalizations are constrained in that "the inferences drawn by a cognitive system will tend to be . . . relevant to the system's goals" (p. 5). In the case of language, the language learner's goal is to understand and to be understood: to comprehend and produce language. There is ample functional pressure to predict meaning on the basis of given lexical items and grammatical characteristics (comprehension); conversely, there is pressure to predict the choice of lexical items and grammatical characteristics given the message to be conveyed (production). Since the sentences the child is learning to understand and produce form an open-ended set, it is not sufficient simply to memorize the sentences that have been heard. The child must necessarily generalize those patterns at least to some extent in order to understand and produce new utterances.

In the first part of this chapter, it is argued that the predictive value of constructions encourages speakers to learn them. A second motivation for representing generalized constructions is suggested in Section 6.10, namely that constructions are primed in production.

6.1 Background: the predictive value of verbs in argument structure patterns

There is a long history in the field of linguistics of considering the main verb to be the key word in a clause (Chomsky 1965; Grimshaw 1990; Lakoff 1970; Levin and Rappaport Hovav 1995; Pinker 1989). This has also been true in the field of language acquisition (e.g. Tomasello 1992, 2000). A critical factor in the primacy of verbs in argument structure patterns stems from their relevant predictive value. If we compare verbs with other words (e.g. nouns), verbs are much better predictors of overall sentence meaning, where by "overall sentence meaning" we basically intend "who did what to whom," a level of generalization that is uncontroversially required for adequate sentence comprehension.

Experimental evidence for the idea that verbs play a key role in semantic interpretation is provided by Healy and Miller (1970). Healy and Miller compared the relative contribution of verbs and subject arguments to overall sentence meaning. These two candidates, verb and subject, were presumably chosen because they appear to be the best candidates for representing overall sentence meaning. The subject argument is often referred to as the "topic" argument in a sentence or what the sentence is "about" (Kuno 1972; Lambrecht 1994; Reinhart 1982). At the same time, the verb provides a great deal of information about who did what to whom. Healy and Miller constructed twenty-five sentences by crossing five subject arguments (*the salesman, the writer, the critic, the student, the publisher*), five verbs (*sold, wrote, criticized, studied, published*), and one patient (*the book*). Participants were asked to sort the sentences into five piles according to similarity in meaning. Results showed that participants reliably sorted sentences together that had the same verb much more often than sentences that had the same subject argument. That is, for example, all five sentences with the verb *criticized* were categorized together much more often than five sentences with the subject *the critic*. Given these results, Healy and Miller concluded that the verb is the main determinant of sentence meaning.

Another source of evidence for the idea that the verb is a good predictor of sentence meaning comes from work on analogy. It has been richly documented that relational aspects of meaning are fruitful sources of analogy and similarity judgments (Gentner 1982). Markman and Gentner (1993), for example, found that in making non-linguistic judgments, similarity is judged to be greater when two representations share the same relations between the entities in each representation. That is, the entities are aligned based on the structure that relates them, rather than on the basis of independent characteristics of the entities. The relevance to language is straightforward. In comparing two sentences, the main relational predicates, the verbs, are more likely to be used than the independent characteristics of the arguments (cf. also Tomasello 2000). Why should this be so? The purpose of analogies is generally one of drawing inferences and making predictions: what can be predicted on the basis of one situation about another situation (Gentner and Medina 1998). Thus it is the value of verbs as good cues to sentence meaning that results in the child's early learning of verb-centered argument structure patterns ("verb islands").

This chapter focuses on the question of why learners generalize beyond the verb to the more abstract level of argument structure constructions.

6.2 The value of constructions as predictors of sentence meaning

In the first part this chapter (Sections 6.2–6.9) we demonstrate that generalizing beyond a particular verb to a more abstract pattern is useful in predicting overall sentence meaning, more useful in fact than knowledge of individual verbs. Bates and MacWhinney (1987) have stressed the importance of weighting different cues, dependent on how reliable and available each cue is. Goldberg, Casenhiser, and Sethuraman (2005) hypothesize that the predictive value encourages speakers to generalize beyond knowledge of specific verbs to ultimately learn the semantic side of linking generalizations, or constructional meaning.

Precedent for this idea comes from work in the non-linguistic categorization literature. Kruschke (1996) and Dennis and Kruschke (1998) discuss how learners shift attention away from less reliable (i.e. less distinctive) cues toward more reliable cues, when learning overlapping instances that belong to distinct categories. For example, if two diseases share one symptom but have their own distinctive symptom, subjects will attend more to the distinctive symptoms than the shared one.[1]

[1] Kruschke discusses how this tendency is so strong that it can account for the neglect of base-rate information (the *inverse base-rate effect* discovered by Medin and Edelson (1988)). In Medin and Edelson's study, subjects were taught to classify diseases on the basis of symptoms. Subjects were trained to discriminate disease C (for common) and R (for rare), with C being presented three times as often as R. In the training, every instance of C had two symptoms, I and PC, and every instance of R had two symptoms, I and PR. Since I appeared with both diseases, it was an imperfect predictor; PC was a perfect predictor of the common disease, PR was a perfect predictor of the rare disease. When tested with the ambiguous symptom I, subjects acted in accord with the disease's overall frequency and chose the common disease, C. But when presented with conflicting symptoms PC and PR, people

Consider the construction types below in Table 6.1:

TABLE 6.1. Construction types, defined formally

Label	Form
VOL:	(Subj) V Obj $Obl_{path/loc}$
VOO:	(Subj) V Obj Obj2

It is clear that constructions are sometimes better predictors of overall meaning than many verbs. For example, when *get* appears in the VOL pattern, it conveys caused motion, but when it appears in the VOO pattern, it conveys transfer:

(1) a. Pat got the ball over the fence.
get + VOL pattern → "caused motion"
b. Pat got Bob a cake.
get + VOO pattern → "transfer"

As quantified below, *get* in isolation has low cue validity as a predictor of sentence meaning. Since most verbs appear in more than one construction with corresponding differences in interpretation, speakers would do well to learn to attend to the constructions. As an indication of the fact that the construction is at least as good a predictor of overall sentence meaning as the verb, we consider the actual predictive value of verbs versus formal patterns in a corpus of speech to young children.

Clearly if we compare the contribution of verb and construction to subtle aspects of meaning involving manner or means, the verb would be more predictive than the construction. This is necessarily true since constructions rarely encode specific meanings: compare "X causes Y to receive Z," the meaning of the ditransitive construction with the meaning of the verbs *hand* or *mail*. At the same time, both verbs and constructions have the potential to convey the overall event-level interpretation, roughly "who did what to whom." Since the event-level interpretation (who did what to whom) is clearly a *necessary* component of interpretation, we chose to compare the relative contribution of constructions and verbs at this level. Clearly, in order to arrive at a *full* interpretation of a sentence, the specifics contributed by only the verb (and its arguments) are required as well.

tended to choose the rare disease, contrary to what would be expected if base rates were taken into account in a simple way. Krushke (1996) suggests that subjects learn that the shared symptom I tends to be misleading for the less frequent disease, inferring that PR is a correspondingly better cue.

6.3 Corpus evidence of the construction as a reliable predictor of overall sentence meaning

Goldberg, Casenhiser, and Sethuraman (2005) examined the Bates corpus (Bates et al. 1988) on the Child Language Data Exchange System (CHILDES) database (MacWhinney 1995), as described in Chapter 4.

Results for the VOL pattern

We first examined whether the formal pattern VOL predicted the semantic caused-motion meaning. "Cue validity" is the conditional probability that an object is in a particular category, given that it has a particular feature or cue (Murphy 1982). Two coders classified mothers' utterances as either entailing caused motion or not; those that we judged not to entail caused motion were separated and further analyzed as discussed below. Agreement between the two independent coders was 99 per cent for classifying utterances as instances of the VOL pattern. Agreement for classifying VOL utterances as entailing literal caused motion, metaphorical caused motion, caused location, or not was 97 per cent. Disagreements were resolved through discussion.

P(A|B) is the probability of A, given B. As detailed below, the cue validity of VOL as a predictor of "caused-motion" meaning, or P("caused motion" | VOL), is somewhere between .63 and .83, depending on how inclusive we take the notion of caused motion to be, and how inclusively we define the VOL formal pattern. We found that 63 per cent (159/256) of the mothers' instances of the construction clearly entail literal caused motion. The following examples are representative:

(2) a. get some more in it
b. bring 'em back over here
c. stuff that all in your mouth
d. put 'em in the box

Another twenty instances involve the verbs *keep, have, get,* or *leave* as in the following type of examples:

(3) a. keeping these people in the garage
b. leave it right there

The utterances in (3) entail that the subject argument acts to keep or allow the theme argument to stay in a particular location. The subject argument is agentive and the locative phrase is predicated of the direct object argument just as in instances that entail prototypical caused motion. Many researchers

have related these instances to cases of caused motion independently (e.g. Goldberg 1995; Matsumoto 1992; Talmy 1976). If we include these cases in the final tally, 70 per cent of VOL utterances imply caused motion or caused location. Another 2 per cent (5) of instances involve the verb *want* as in:

(4) Oh, you want them in a cup?

These instances convey possible future caused motion. If we include these in the tally, the percentage of instances that are related to caused motion increases to 72 per cent. Another 5 per cent (13) of instances involve the verbs *read* or *say*, which could be argued to encode metaphorical caused motion (Ackerman and Webelhuth 1998; Goldberg 1995; Pinker 1989; Reddy 1979). Including these cases would raise the total number of VOL utterances whose meanings are related to caused motion to 77 per cent.

Of the remaining VOL utterances, 25, or 10 per cent involved locative adjuncts. If we exclude these, the total number of utterances included as VOL utterances would be reduced from 256 to 231. The total percentage of VOL utterances that involved caused motion would be 85 per cent.

The remaining 34 tokens include examples such as the following, which do not convey caused motion:

(5) a. What is your foot doing on the table? (The WXDY construction: 16 instances)
b. What did Ivy do to her arm? (1 instance)
c. find the bird in the snow (utterances with find: 2 instances)
d. get Papa at the airport (move-from interpretation: 1 instance)
e. stand it up (verb particle interpretations: 14 instances)

To summarize, the cue validity of the VOL pattern as a predictor of caused-motion meaning is provided in Table 6.2 below and in the pie chart of Figure 6.1. on the following page.

We then investigated the extent to which individual verbs predicted caused-motion meaning. Table 6.3 shows our calculations of the cue validities of individual verbs that appear at least five times in mothers' speech in the Bates et al. (1988) corpus.[2] For reliability, a second coder independently classified a

TABLE 6.2. Cue validity of VOL construction as a predictor of caused motion

	Strict encoding of caused-motion meaning	Inclusive encoding
VOL	.63	.85

[2] Verbs appearing in the VOL pattern less than 5 times each accounted for 56 instances. These verbs were: *eat, pick, set, throw, want, clean, close, cook, dump, find, give, hold, keep, leave, let, make, move, park, pour, push, send, stack, stick, stuff, try, wake, walk, wear, wipe.*

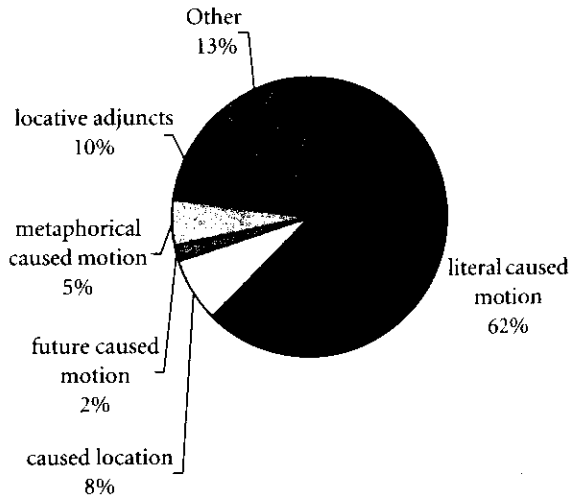

FIGURE 6.1 Proportion of utterances that were related to caused motion (all but those indicated by "other"or "locative adjuncts")

sample of 34 per cent (404/1195) of the total number of utterances, including all of the tokens in the corpus, for the verbs *put, get, take, read, see, stand,* and *turn.* The two coders agreed on whether the verb determined the overall sentence meaning reliably 89 per cent (360/404) of the time. The other

TABLE 6.3. Cue validity of verbs as predictors of whatever meaning they predominantly have in the VOL pattern

Verb	A. Number of instances of stable verb meaning	B. Total number of times verb appears in corpus	C. Cue validity (A/B)	D. Number of times in VOL pattern	E. Weighted cue validity (C × D/200) High (low)
put	113	114	1	99	.50
do	147	601	.24	26	.03
have	15	115	.13	16	.01
get	18	108	.17	14	.01
take	9	44	.20	14	.01
read	14	24	.58	7	.02
see	30	86	.35	7	.01
stand	8	8	1	7	.04
say	6	65	.09	6	.00
bring	10	10	1	6	.03
turn	10	20	.5	6	.02
Total		1195		200	**.68**

utterances were also classified independently by two coders, with disagreements resolved through discussion.

Column A of Table 6.3 lists the number of utterances in which the verb's meaning as it is determined to appear most frequently in the VOL pattern, appears overall in the corpus (including appearances in the VOL pattern). Column B lists the total number of times the verb appears in the corpus. Simple cue validity was calculated by dividing column A by column B; this is recorded in Column C.

In order to determine the cue validity of verbs as they appear in the VOL pattern, it is necessary to weight the cue validities of each verb according to how often the verb occurs in the VOL pattern. We therefore multiply the cue validity obtained in Column C by the number of times the verb occurs in the VOL pattern (recorded in Column D) and divide by the number of VOL tokens in the corpus (200). Summing over the weighted cue validities in Column D provides us with the overall cue validity of .68 for verbs in the VOL pattern in our corpus. If we compare the .63—.85 cue validity for the VOL pattern as a predictor of caused-motion meaning, we can see the construction is roughly as valid a cue.

To see that the weighted average of cue validities is more revealing than the simple average, it is illustrative to consider the following hypothetical situation. Imagine that there were one verb that accounted for 90 per cent of the tokens of a particular construction and had a cue validity of 1. In this particular case, the predictive value of verbs in the construction would clearly be quite high. Yet if there happened to be ten other verbs appearing in the construction, each accounting for 1 per cent of the tokens, and each having low cue validity (in the limiting case close to 0), the average cue validity would only be approximately 0.09. This is not the number we are after, since it does not reveal the fact that the most likely verb to appear is highly predictive. The weighted average in this circumstance would be roughly 0.9, accurately reflecting the predictive nature of verbs in this hypothetical construction. Thus it is the weighted average of cue validities that more accurately reflects the predictive value of verbs.[3]

It is clear from Table 6.3 that there is a wide variability of cue validities across verbs. While a few verbs have perfect or near perfect cue validities (*put, bring, stand*) in our corpus, other verbs' cue validities were low (*do, get, have, let*, and *take*). For the latter verbs, relying on the construction in conjunction with the verb is essential to determining sentence meaning. This fact in itself is

[3] The simple average of cue validities would be .5 instead of .68, which would only strengthen our claim that the cue validity of constructions is at least as high as that of verbs.

sufficient to conclude that attention to the semantic contribution of constructions is required for determining overall sentence meaning.

6.4 Results of an analysis of the VOO pattern

Comparable results exist for the VOO pattern. We first examined whether the formal pattern predicted the meaning of "transfer." There were a total of 54 VOO utterances in our database. After initial discussion of criteria to be used, agreement was 100 per cent for classifying mothers' VOO utterances as entailing transfer: literally, metaphorically, or not (n=54). If we include instances that involve metaphorical transfer (specifically those involving the verbs *read* and *tell*), 51 of those 54 (94 per cent) convey transfer. If we exclude instances of *read* and *tell*, 34/54 (61 per cent) code transfer. These figures are given in Table 6.4, with the breakdown given in Figure 6.2.

Goldberg, Casenhiser, and Sethuraman (2005) also investigated the extent to which the verbs predicted overall sentence meaning. Agreement among the two coders was 97 per cent for classifying mothers' uses of verbs as predictive of overall sentence meaning (who did what to whom) (n=307).

We determined for each utterance whether the predominant meaning (who did what to whom) conveyed by the verb's use in the VOO pattern held in each of the remaining utterances involving the same verb. For example, we decided that *tell* involved a speaker, a listener, and some kind of content

TABLE 6.4. Cue validity of the VOO pattern as a predictor of the meaning of transfer

	Strict encoding of transfer	Inclusive encoding of transfer
VOO	.61	.94

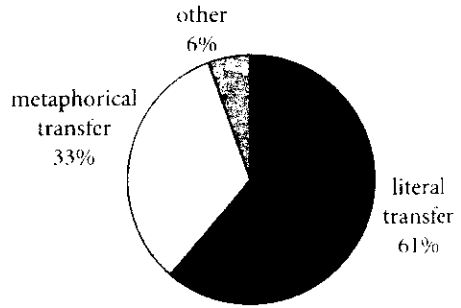

FIGURE 6.2 Proportion of utterances involved literal or metaphorical transfer

conveyed in all uses of the VOO pattern and also in every other use in our corpus; therefore *tell* has a perfect cue validity of 1. On the other hand *read* involved a reader, a listener, and some kind of content in its three appearances in the VOO pattern and in eleven other utterances (phrased using "someone read something to someone"); in the remaining ten utterances in our corpus, *read* did not involve a listener. Therefore the cue validity of *read* was determined to be: (3+11)/24=.58.

Table 6.5 shows our calculations of the cue validities of each of the thirteen verbs that appeared in the VOO pattern in our corpus. As in Table 6.4, Column A lists the number of utterances in which the verb's meaning as it is determined to appear most frequently in the VOO pattern, appears overall in the corpus (including appearances in the VOO pattern). Column B lists the total number of times the verb appears in the corpus. Simple cue validity is recorded in Column C and was calculated by dividing column A by column B. To determine cue validity across verbs, it is necessary to weight this number according to how often the verb occurs in the VOO pattern. We therefore multiply the cue validity obtained in Column C by the number of times the verb occurs in the VOO pattern (recorded in Column D) and divide by the number of VOO tokens in the corpus (54). Summing over the weighted cue validities in Column D provides us with the overall cue validity of .61 for verbs in the VOL pattern in our corpus.

TABLE 6.5. Cue validities of verbs in the VOO pattern as predictors of whatever meaning they predominantly have in the VOO pattern

Verb	A. Number of times verb has stable meaning	B. Number of instances in corpus:	C. Cue validity (A/B)	D. Number of times in VOO pattern	Weighted cue validity (C × D/54)
Give	14	14	1	11	.20
Tell	37	37	1	11	.20
Get	7	108	.06	7	.01
Build	6	24	.25	5	.02
Make	3	56	.05	4	.00
Read	14	24	.58	3	.03
Show	6	6	1	3	.06
Bring	2	10	.2	2	.01
Feed	2	2	1	2	.04
Pour	2	4	.5	2	.02
Buy	1	9	.11	1	.00
Fix	1	7	.14	1	.00
Call	4	6	.67	2	.02
TOTAL		307		54	**.61**

The weighted average cue validity for verbs is .61. We again see that the overall cue validity of constructions is at least as high as the cue validity of verbs. If we are generous in deciding what utterances involve transfer, the cue validity of the construction is markedly higher than the cue validity for verbs.

Once again there is a wide variability of cue validities across verbs. While a few verbs had perfect cue validities (*feed, give, show, tell*) in our corpus, other verbs' cue validities were quite low (*fix, get, make*). Again, regardless of the overall cue validity of verbs, this fact in itself indicates that attention to the construction's contribution is key to determining who did what to whom.

What about verbs in other constructions? It may be that verbs are more predictive for some constructions than others. For example, in the simple intransitive construction, the verb supplies almost all of the lexical content. There is a large difference in sentence meaning between *The vase broke* and *She shouted.* Still, even in these cases, the verbs involved are far from perfect predictors of overall sentence meaning. *Break* can appear both transitively and intransitively—to know whether an agent is known or relevant, one needs to know which construction was used. *Shouted,* too, can be used as a verb of communication (e.g. *She shouted the directions*) or simply a verb of sound emission (e.g. *She shouted for joy*).

The fact that the cue validities calculated (.68 for verbs in the VOL and .61 for verbs in the VOO) are close is intriguing. It is possible that there is a generalization about the overall cue validity of verbs: they may be predictive of sentence meaning roughly two-thirds of the time in English. Calculations on other constructions and other corpora are needed to confirm this figure.

6.5 Experimental evidence for constructions as predictors of sentence meaning

Bencini and Goldberg (2000) conducted an experiment inspired by the Healy and Miller (1970) sorting experiment described in Section 6.1, which had been titled, "The Verb as the Main Determinant of Sentence Meaning." In this earlier experiment, stimuli were created by crossing subject arguments with verbs, since it was assumed that the two best candidates for determining what the sentence was about were the verb and the subject argument. We aimed to compare the semantic contribution of the construction with that of the morphological form of the verb. The stimuli were sixteen sentences created by crossing four verbs with four different constructions.

Undergraduate students were asked to sort these sixteen sentences, provided in random order, into four piles based on "overall sentence meaning." They were instructed that there was no right or wrong answer, that the

experiment was only intended to determine how people sorted sentences according to sentence meaning. Subjects could sort equally well by verb: e.g. all instances of *throw* (1a-d) being put into the same pile, regardless of construction; or subjects could sort by construction: all instances of e.g. the VOO (ditransitive construction) (1a, 2a, 3a, and 4a) being put into the same pile.

It would of course be possible to design stimuli with a great deal of overlapping propositional content such that we could a priori predict either a verb or constructional sort. For example, the sentences *Pat shot the duck* and *Pat shot the duck dead* would very likely be grouped together on the basis of overall meaning despite the fact that the argument structure patterns are distinct. Conversely, *Pat shot the elephant* and *Patricia stabbed a pachyderm* would likely be grouped together despite the fact that no exact words were shared. The stimuli were designed to minimize such contentful overlap contributed by anything other than the lexical verb. No other lexical items in the stimuli were identical or near synonyms.

The use of the sorting paradigm is a particularly stringent test to demonstrate the role of constructions. Medin, Wattenmaker, and Hampson (1987) have shown that there is a strong, domain-independent bias towards sorting on the basis of a single dimension, even with categories that are designed to resist such one-dimensional sorts in favor of a sort based on a family resemblance

TABLE 6.6. Stimuli for sorting experiment

1a. Pat threw the hammer.	(VO) Transitive
b. Chris threw Linda the pencil.	(VOO) Ditransitive
c. Pat threw the key onto the roof.	(VOL) Caused Motion
d. Lyn threw the box apart.	(VOR) Resultative
2a. Michelle got the book.	(VO) Transitive
b. Beth got Liz an invitation.	(VOO) Ditransitive
c. Laura got the ball into the net.	(VOL) Caused Motion
d. Dana got the mattress inflated.	(VOR) Resultative
3a. Barbara sliced the bread.	(VO) Transitive
b. Jennifer sliced Terry an apple.	(VOO) Ditransitive
c. Meg sliced the ham onto the plate.	(VOL) Caused Motion
d. Nancy sliced the tire open.	(VOR) Resultative
4a. Audrey took the watch.	(VO) Transitive
b. Paula took Sue a message.	(VOO) Ditransitive
c. Kim took the rose into the house.	(VOL) Caused Motion
d. Rachel took the wall down.	(VOR) Resultative

structure (Rosch and Mervis 1975). One-dimensional sorting has been found even with large numbers of dimensions (Smith 1981), ternary values on each dimension (Anh and Medin 1992), holistic stimuli, and stimuli for which an obvious multidimensional descriptor was available (Regehr and Brooks 1995). The stimuli presented subjects with an opportunity to sort according to a single dimension: the verb. Constructional sorts required subjects to note an abstract relational similarity involving the recognition that several grammatical functions co-occur. Thus we would expect verb sorts to have an inherent advantage over constructional sorts.

Six subjects produced entirely construction sorts, seven subjects produced entirely verb sorts, and four subjects provided mixed sorts. In order to include the mixed sorts in the analysis, results were analyzed according to how many changes would be required from the subject's sort to either a sort entirely by verb (VS) or a sort entirely by construction (CS). The average number of changes required for the sort to be entirely by the verb was 5.5; the average number of changes required for the sort to be entirely by construction was 5.7. The difference between these scores does not approach significance. That is, subjects were just as likely to sort by construction as they were to sort according to the single dimension of the morphological form of the verb. If verbs provided equally good cues to overall sentence meaning, there would be no motivation to overcome the well-documented preference for one-dimensional sorts: subjects would have no motivation to sort by construction instead of by verb. Bencini and Goldberg hypothesize that constructional sorts were able to overcome the one-dimensional sorting bias to this extent because constructions may be better predictors of overall sentence meaning than the morphological form of the verb.

This experiment was performed with adults, but the implications for language learning are clear. Insofar as constructions are at least as good predictors of overall sentence meaning as any other word in the sentence, learners would do well to learn to identify construction types, since their goal is to understand sentences.

Kaschak and Glenberg (2000) demonstrate that subjects rely on constructional meaning when they encounter nouns used as verbs in novel ways (e.g. *to crutch*). In particular they show that different constructions differentially influence the interpretations of the novel verbs. For example, *She crutched him the ball* (ditransitive) is interpreted to mean that she used the crutch to transfer the ball to him, perhaps using it as one would a hockey stick. On the other hand, *She crutched him* (transitive) might be interpreted to mean that she hit him over the head with the crutch. Kaschak and Glenberg suggest

that the constructional pattern specifies a general scene and that the "affordances" of particular objects are used to specify the scene in detail. It cannot be the semantics of the verb alone that is used in comprehension because the word form is not stored as a verb but as a noun. Ahrens (1994) conducted an experiment with a novel verb form. She asked 100 native English speakers to decide what *moop* meant in the sentence *She mooped him something.* Sixty percent of subjects responded by saying that *moop* meant "give," despite the fact that several verbs exist that have higher overall frequency than *give* and could be used in that frame, including *take* and *tell.* Similarly, Kako (2005) finds that subjects' semantic interpretations of constructions and their semantic interpretations of verbs that fit those constructions are highly correlated, concluding as well that syntactic frames are "semantically potent linguistic entities."

A question arises as to why constructions should be at least as good predictors of overall sentence meaning as verbs. The answer I believe stems from the fact that in context, knowing the number and type of arguments tells us a great deal about the scene being conveyed. To the extent that verbs encode rich semantic frames that can be related to a number of different basic scenes (Goldberg 1995), the complement configuration or construction will be as good a predictor of sentence meaning as the semantically richer, but more flexible verb.

6.7 Increased reliance on constructions in second-language acquisition

Liang (2002) replicated the sorting task of Bencini and Goldberg (2000) with Chinese learners of English of varying proficiencies. Learners categorized as advanced passed the Chinese national test for non-English majors, a test that is generally recognized in China as indicating advanced English ability. Learners categorized as at the intermediate level had passed the national entrance examination to college, which indicates an intermediate level of proficiency with English. Beginning English learners had only two years of English instruction. Liang found that subjects produced relatively more construction-based sorts as their English improved. Her data is provided in Fig. 6.3 on the following page.

For early learners (n=46), the average deviation from an entirely verb-based sort was 5.8; the deviation from an entirely construction-based sort was 6.2. For intermediate learners (n= 31), the average deviation from a verb sort was 6.2; from a construction sort, 5.3. For advanced learners (n=33), the average deviation from a verb sort was 8.2, from a constructional sort 4.9.

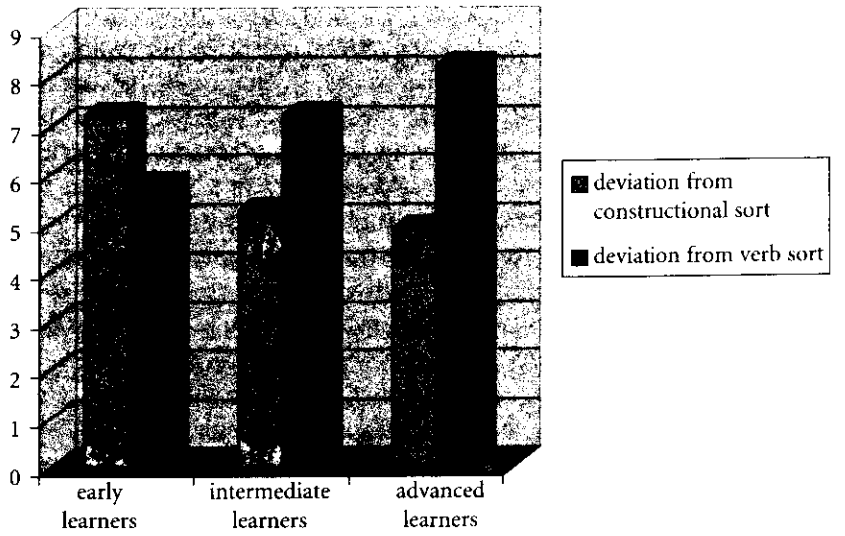

FIGURE 6.3 Results from Liang (2002)

These results indicate that the ability to use language proficiently is correlated with the recognition of constructional generalizations.[4]

6.8 Category Validity

We have discussed cue validity, the probability that an item belongs to a category, given that it has a particular feature: P(cat | feature), and we have found that when the category is taken to be overall sentence meaning, constructions have roughly equivalent cue validity compared with verbs. There is also a second relevant factor. CATEGORY VALIDITY is the probability that an item has a feature, given that the item belongs in the category: P(feature | cat). Thus category validity measures how common or available a feature is among members of a category. The relevant category is, again, sentence meaning.

The category validity of particular verbs as a feature of the semantic category caused motion was determined by hand-coding each utterance in a

[4] Gries and Wulff (2004) also replicated the sorting study, this time with advanced German learners of English, finding similar results to those found for advanced learners by Liang: the average deviation from a verb sort was 8.9, and from a constructional sort was 5.4, a significantly stronger reliance on constructions. These studies raise an interesting issue as to why it might be that learners of a second language appear to rely more heavily on constructions than do native English speakers, at least in this sorting paradigm.

randomly selected subset of the Bates et al. (1998) corpus for whether each of the mothers' utterances expressed caused motion or not. In our sample of all of the utterances of four mothers, there were 47 utterances that conveyed caused motion involving twelve different verbs. As discussed in Chapter 4, there is often one verb that accounts for the lion's share of tokens of particular constructions (Goldberg 1999; Goldberg, Casenhiser, and Sethuraman 2004; Sethuraman 2002); in particular, *put* accounts for many the tokens of the caused-motion construction, but since the transitive and resultative constructions can also convey caused motion (with verbs such as *send, bring, carry*), the category validity for even *put* is not particularly high; in our corpus, 29/47 expressions conveying caused motion involved the verb *put*, resulting in a category validity of .62. The probability that a sentence with caused-motion meaning contained the verb *bring* was only .02, since only 2 per cent of the utterances expressing caused motion used *bring* (1/47). Similarly low probability was found for another seven verbs in the corpus (*drive, make, open, ride, stack, leave, wear*). The category validity for *dump* was .04, *stick* was .06, and *turn* was .11.

The average category validity of all verbs that may convey caused motion is equal to 1/n, where n = the number of verbs that express caused motion, or in our sample, 1/12 = .08. Clearly as the sample size increases, the average category validity for verbs is lowered. The actual average category for verbs approaches 0, since more than a hundred different verbs can be used to convey caused motion (n > 100; average cue validity = 1/100' < .01). Another relevant number is the MAXIMUM CATEGORY VALIDITY, since the maximum category validity provides an estimate of the category validity associated with the "best guess" of a relevant verb. In our sample, *put* had the highest category validity of .62; the other eleven verbs conveying caused motion had markedly lower category validities.

The category validity of a construction as a feature of the semantic category caused motion is the probability that the particular construction will be involved, given the interpretation of caused motion. There were only three constructions used to convey caused motion (the VOL, the resultative, and the transitive construction). If we make the very conservative assumption that these three constructions are independent, we find an average category validity of .33 for constructions. The average category validity for constructions may also go down as the sample size increases; but since there are less than a handful of constructions that can be used to convey caused motion, the average category would not dip below .20. The VOL pattern has the maximum category validity for "caused motion" meaning of .83 (39/47 utterances that expressed caused motion involved the VOL pattern). *Put* is only grammatical

in the VOL pattern (**She put on the table; *She put the book*): P (*put* | "caused motion") = P (*put* & VOL | "caused motion"). Since hundreds of verbs in addition to *put* can appear in the VOL pattern with caused-motion interpretation, P (*put* & VOL | "caused motion") is necessarily less than P (VOL | "caused motion"). Therefore, P (*put* | "caused motion") < P (VOL | "caused motion"). That is, the VOL pattern must have a higher category validity than *put*; it is more *available* as a cue to caused-motion meaning. The comparison between verbs and constructions is given in Tables 6.6 and 6.7. Table 6.6 provides the AVERAGE CATEGORY VALIDITY for the VOL pattern and verbs that may be used to convey caused motion. Table 6.7 provides the MAXIMUM category validity of the VOL pattern and verbs for caused motion meaning.

On both measures, the average category validity and the maximum category validity, the construction has a higher score than the verb. All things being equal, if two cues have roughly equal validity, the higher category validity of one cue will naturally result in a greater reliance on that cue in categorization tasks (Bates and MacWhinney 1987; Estes 1986; Hintzman 1986; Nosofsky 1988). Thus constructions are better cues to sentence meaning than verbs insofar as they are as reliable (with equivalent cue validity) and more available (having higher category validity).

TABLE 6.6. Average category validity for the VOL pattern and verbs as features of the category of caused motion meaning in a sample of 47 utterances (to illustrate) and more generally, based on rational assumptions

	In sample of 47 utterances	Asymptotic category validity
n	(n = 12)	(n > 100)
$\sum_{i=1}^{n} P(\textbf{verb}_i \mid \text{"caused motion"})/n$ where verb_i is a verb that may encode caused motion	=.08	<.01
n	(n=3)	(since n < 5)
$\sum_{i=1}^{n} P(\textbf{construction}_i \text{ "caused motion"})/n$ where construction_i is a construction that may encode caused motion	=.33	>.20

TABLE 6.7. Maximum category validity for the construction and verbs as features of the category of caused motion meaning in a sample of 47 utterances (to illustrate) and more generally, based on rational assumptions

	In sample of 47 utterances	In general:
Verb: P (*put* | "caused motion")	=.62	P (*put* | "caused motion") = P (*put* & VOL | "caused motion")
Construction: P (VOL | "caused motion")	=.83	P (*put* & VOL | "caused motion") < P (VOL | "caused motion")
		Therefore, P (*put* | "caused motion") < P (VOL | "caused motion")

6.9 Languages in which verbs are more predictive

The verbs in many languages are more restrictive than they are in English, only appearing in constructions that match their meanings. Verbs in Latinate languages, Turkish, and Hindi, for example, do not appear in anything like the range of constructions that they do in English even though they have quite parallel meanings (see, e.g., Narasimhan 1998). Therefore it would seem that the verbs in e.g. Turkish have much higher cue validity than they do in English. And yet it seems unlikely that they fail to form argument structure constructions in such languages. Dan Slobin (personal communication, Feb. 14, 2004) has found in unpublished experimental work that speakers of Turkish readily interpret novel verbs presented in familiar constructions, indicating that they are at least able to construct information about an abstract argument structure construction for the purpose of comprehension. The fact that the category validity for constructions is generally higher than that for verbs—insofar as there are more verbs that can be used to convey a particular event frame than there are constructions—may be responsible in part for yielding constructional generalizations.

In addition, there is a second factor that may well play a role in encouraging speakers to form argument structure constructions, even when the cue validity of the verbs in the language is consistently high. This factor involves the phenomenon of CONSTRUCTIONAL PRIMING.

6.10 Structural Priming and its relation to constructions

A second type of motivation for learning constructions, outlined in this section, is that constructions are primed in production. That is, saying or hearing instances of one grammatical pattern primes speakers to produce

other instances of the same. Kathryn Bock and colleagues (Bock and Loebell 1990; Bock, Loebell, and Morey 1992; Bock 1986; Loebell and Bock 2003) have shown in a number of experimental studies that passives prime passives, ditransitives prime ditransitives, and datives prime datives (cf. also Bock and Griffin 2000; Loebell and Bock 2003; Branigan et al. 1995; Chang et al. 2000; Friederici, Schriefers, and Lindenberger, 1998; Hare and Goldberg 1999; Nicol 1996; Potter and Lombardi 1998; Saffran and Martin 1997; Savage et al. 2003; Scheepers 2003; Smith and Wheeldon 2001; Tomasello 2003; Yamashita, Chang, and Hirose 2003).

This sort of priming provides a useful tool to investigate the mental representation of linguistic expressions (Branigan et al. 1995; Bencini 2002). The naturalness of the priming paradigm is supported by the fact that a tendency towards structural repetition occurs in natural unmonitored speech or text (Kempen 1977; Tannen 1987; Weiner and Labov 1983; Levelt and Kelter 1982). This suggests that structural priming is not simply a laboratory-induced phenomenon.

Priming has been argued to represent implicit learning in that its effect is unconscious and long-lasting (Bock and Griffin 2000; Chang et al. 2000). Thus the existence of structural priming may be an important factor underlying the fact that there are generalizations in languages. The same or similar patterns are easier to learn and produce. At the same time, priming of course is not particular to language—repetition of the same motor programs also leads to priming effects.

Bock's original claim was that syntactic tree structures, not constructions with associated meanings, were involved in priming (Bock and Loebell 1990; Bock, Loebell, and Morey 1992; Bock 1986; Loebell and Bock 2003). In recent work, the question of whether constructional priming exists has been investigated. That is, can abstract pairings of form with meaning be primed? Chang, Bock, and Goldberg (2003) conducted a simple experiment in which syntactic structure was controlled for, while two different constructions were used as primes. Sample prime target sentences are given below:

Sample Primes:

(6) a. She loaded the wagon with hay. ("load with")
b. She loaded hay onto the wagon. ("load onto")

Sample Targets:

(7) a. He embroidered the shirt with flowers.
b. He embroidered flowers onto the shirt.

Subjects were asked to recall a sentence as it was presented after a short distractor task. Such rapid serial visual presentation (RSVP) tasks have been

shown to yield priming effects (Potter and Lombardi 1998). If semantics matters in priming, then we should see "load with" structures priming other "load with" structures more than "load onto" structures. In fact this is exactly what was found.

Also, as predicted by constructional priming, subjects produced more load-with types of sentences after load-with type primes than after load-onto primes as shown in fig. 6.4; and more load-onto sentences after load-onto primes than after load-with primes as shown in fig. 6.5.

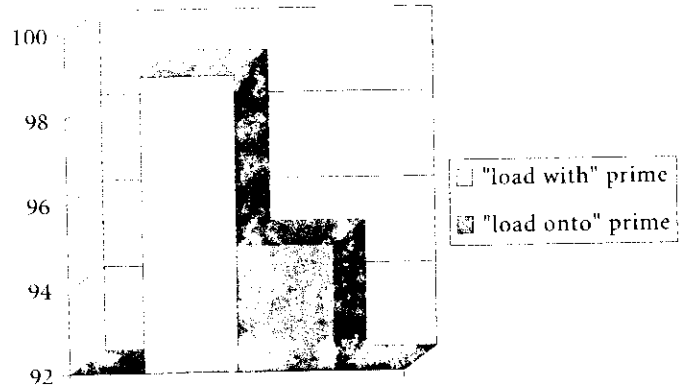

FIGURE 6.4 Percentage of "load with" responses to "load with" target (n=80)

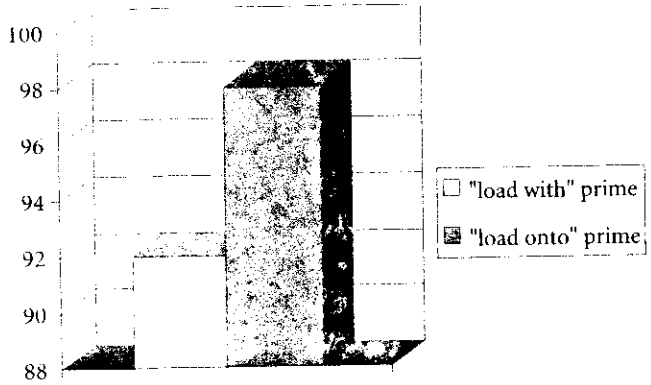

FIGURE 6.5 Percentage of "load onto" responses to "load onto" target (n=80)

The constructional priming found cannot be wholly attributed to an overlap in prepositions, since the "load onto" sentences used a variety of prepositions including *over, onto, into, around, on*. Moreover, Griffin and Weinstein-Tull (2003) have shown that object-raising sentences prime object-raising sentences more than object-control sentences, despite a lack of

any shared morphology. This also suggests constructional priming; the results are unexpected if only formal cues were taken into account, since the two constructions arguably have the same form.

Given these results, it is worth returning to the original motivation for earlier claims that syntactic constituent structure, not constructions (form–meaning pairings) are primed. Bock and Loebell (1990) made perhaps the strongest case for this claim with a series of experiments. In one experiment, they showed that both datives and locatives primed dative descriptions of (unrelated) pictures equally well. Example primes are given below:

Primes:

(8) a. The wealthy widow gave her Mercedes to the church. (dative)
b. The wealthy widow drove her Mercedes to the church. (locative)

The constructional interpretation of this result stems from a claim discussed in Chapter 2, namely that so called "dative" and "locative" expressions are actually both instances of the same caused-motion construction (cf. also Goldberg 1995).

(9) Caused-motion construction:

CAUSE-MOVE	(cause	theme	path)
|	|	|	|
V	Subj	Obj	Obl

(10) Examples:
a. She drove the box to Missouri.
b. She drove the box to Mary.
c. She threw the box to Mary.
d. She gave the box to Mary.

Therefore, the findings are that caused-motion expressions prime caused-motion expressions, a result that is expected by a constructionist account of priming. In fact, Bock and Loebell also acknowledge that locative and dative expressions are semantically similar. They therefore performed a second experiment in which they investigated whether intransitive locative expressions such as found in Table 6.8 (example b) primed passives: a

TABLE 6.8. Stimuli types used in Bock and Loebell (1990)

Primes Type	Example
a. Passives:	*The construction worker was hit by the bulldozer.*
b. Intransitive Locatives:	*The construction worker was digging by the bulldozer.*
c. Actives (control):	*The construction worker drove the bulldozer.*

construction with the same syntax as locatives but with clearly distinct meaning.

Bock and Loebell (1990) found that in fact intransitive locatives did prime passives. This is the strongest evidence for purely syntactic, non-constructional priming. Yet all of the intransitive locative primes used as stimuli included the preposition *by* and the auxiliary *be*. A question naturally arises, was it the shared morphemes, and not the shared syntactic structure, that produced the priming (Hare and Goldberg 1999)? In order to address this question, Bencini, Bock, and Goldberg (in preparation) attempted to replicate the Bock and Loebell findings while adding a fourth condition in which intransitive locatives without shared morphology were used as primes (prime type d in Table 6.9).

TABLE 6.9. Stimuli types used in Bencini, Bock, and Goldberg (in prep.)

Primes Type	Example
a. Passive	*The construction worker was hit by the bulldozer.*
b. Locatives w/shared morphology	*The construction worker was digging by the bulldozer.*
c. Actives (control)	*The construction worker drove the bulldozer.*
d. **Locatives w/o shared morphology**	***The construction worker might dig near the bulldozer.***

Bock and Loebell was replicated, demonstrating that locatives with the shared morphology prime passives and as also expected, passives prime passives. At the same time, a significant difference between the passive condition and the locative condition without shared morphology was also found, and the locatives without shared morphology condition did *not* prime passives significantly more than the control group. This finding is intriguing because it may indicate that shared syntactic structure is not sufficient to induce priming.[5]

Hare and Goldberg (1999) designed a different test of the idea that pure syntactic tree structure and not some sort of form–meaning pairing was involved in priming. Recall that it has been well established that ditransitives prime ditransitives, and also that instances of the caused-motion construction prime other instances of the caused-motion construction. We attempted to

[5] At the same time, the data are a bit ambiguous, because there is a stepward trend such that passives and by-locatives prime passives (significantly), and in addition, numerically more passives were produced after non-by locatives than after controls; however the latter difference was non-significant despite the running of 130 subjects.

determine whether a third sort of prime, "provide with" primes, would differentially prime either caused-motion expressions ("datives") or ditransitive descriptions of scenes of transfer. Examples of the "provide with" sort of primes are given in Table 6.10.

TABLE 6.10. Key priming condition in Hare and Goldberg (1999)

Prime type	Examples
"provide with" condition	*The government provided the troops with arms.* *His editor credited Bob with the hot story.* *The father entrusted his daughter with the keys.*

"Provide with" sentences arguably have the same syntactic form as caused-motion expressions: NP [V NP PP], and yet the order of rough semantic roles involved parallels the ditransitive: Agent Recipient Theme. Results demonstrated that "provide with" expressions prime ditransitive descriptions of (unrelated) pictures as much as ditransitives do. There was no evidence at all of priming of caused-motion expressions, despite the shared syntactic form (Hare and Goldberg 1999). Thus when order of semantic roles is contrasted with constituent structure, the order of semantic roles shows priming, with no apparent interaction with constituent structure.[6]

What do the structural priming facts mean? First of all, constructions can be primed, which means that the level of generalization involved in argument structure constructions is a useful one to acquire. It is further possible that priming of structure may not be independent of meaning. Thus the priming mechanism may encourage speakers to categorize on the basis of form *and* meaning.

6.11 Conclusion

We have offered two factors that likely encourage speakers to form the argument structure generalizations they do. Children initially generalize at

[6] One interpretation of the Hare and Goldberg findings is that it was order of animate participants that effected priming, not the order of semantic roles. This possibility is exists because animacy has been shown to induce priming, even when the overall construction is held constant (Bock, Loebell, and Morey 1992). In fact, it remains to be shown that ditransitive vs. dative priming is not induced by differing order of animate participants as well. But, Yamashita, Chang, and Hirose (2003) have shown that dative sentences with the order AGENT-*wa* RECIPIENT-*ni* PATIENT-*o* (*wa-ni-o*) prime other *wa-ni-o* ordered productions, even though the animacy of recipients and patients were controlled for. These results suggest that structural priming can be sensitive to the order of syntactic functions or thematic roles.

the level of specific verbs plus argument slots (Tomasello's "verb islands") because the verb in an argument frame is the best single word predictor of overall sentence meaning. We argue further that children generalize beyond specific verbs to form more abstract argument structure constructions because the argument frame or construction has roughly equivalent cue validity as a predictor of overall sentence meaning to the morphological form of the verb, and has much greater category validity. That is, the construction is at least as reliable and much more available. Moreover, given the fact that many verbs have quite low cue validity in isolation, attention to the contribution of the construction is essential.

Motivating constructional generalizations in a different way is the simple fact that hearing or producing a particular construction makes it easier to produce the same construction. Instead of learning a myriad of unrelated constructions, speakers do well to learn a smaller inventory of patterns in order to facilitate online production.

Part III
Explaining Generalizations

7

Island constraints and scope

This chapter explores the relevance of information structure and processing demands in an explanation of constraints on movement ("island constraints") and constraints on relative scope assignment. We build on previous accounts in order to give a motivated account of a wide range of facts, emphasizing an analysis of the ditransitive that predicts many of the special properties that were observed in Chapter 2.

At the outset, certain terminology is introduced that will be critical to the discussion that follows. Gundel (1985: 35) makes a useful distinction between two types of given/new distinctions: referential givenness/newness and relational givenness/newness. Each is discussed below in turn.

Referential givenness/newness is a relationship between a linguistic expression and a non-linguistic entity. It is correlated with degrees of cognitive activation or identifiability (Ariel 1990; Chafe 1987, 1994; Lambrecht 1994; Prince 1981). There are strong correlations between NP marking and given or new status as follows:

- **Given arguments** are typically expressed by weakly accented pronouns or not at all. They can be established by linguistic or extra-linguistic context. They are considered to be cognitively "active" in the mind of the addressee.
- **Accessible arguments** are normally expressed by definite full NPs. They are considered to be semi-active in the mind of the addressee.
- **New arguments** are normally expressed by indefinite NPs. They are considered to be inactive in the mind of the addressee at the time of utterance.

Relational givenness/newness refers to the assumed informational contribution made by a particular utterance to the knowledge state of the addressee: the INFORMATION STRUCTURE of a sentence (Allen 1999; Halliday 1967; Lambrecht 1994). The choice of particular constructions often determines the information structure of a sentence, including its topic and potential focus domain. Differences in the packaging of information are perhaps the most

important reason why languages have alternative ways to say the "same" thing.

A sentence **topic** is a "matter of [already established] current interest which a statement is about and with respect to which a proposition is to be interpreted as relevant" (Lambrecht 1994). The topic serves to contextualize other elements in the clause (Chafe 1994; Kuno 1972; Langacker 1987a; Strawson 1964).

In languages that have uncontroversial subjects, the **subject argument is the default topic** in the clause (Chafe 1987; Lambrecht 1994; Langacker 1987a; MacWhinney 1977).

The **potential FOCUS DOMAIN** of a sentence is that part of a sentence that is interpretable as being asserted.

TEST for being within the focal domain: propositions expressed within the potential focus domain can be understood to be negated by sentential negation.

The focus domain is thus "one kind of emphasis, that whereby the speaker marks out a part (which may be the whole) of a message block as that which he wishes to be interpreted as informative" (Halliday 1967). Similarly Lambrecht (1994) defines the focus relation as relating "the pragmatically nonrecoverable to the recoverable component of a proposition [thereby creating] a new state of information in the mind of the addressee."

The primary topic and the focus domain together do not exhaustively identify the information units of a sentence. We will refer to elements of a sentence that are neither the primary topic nor part of the focus-domain as BACKGROUNDED elements (corresponding roughly to the TAIL of Vallduví 1993).

BACKGROUNDED elements: constituents that do correspond neither to the primary topic nor to part of the potential focus domain.

Elements that are a part of presupposed clauses are backgrounded. For example, consider the restrictive relative clause in (1):

(1) I read the book that Maya loaned me.

We can see that the proposition conveyed by the relative clause is not part of the focus domain because it is not negated by sentential negation. As a presupposition, it is implied by both the positive and the negative form of the sentence:

(2) I read the book that Maya loaned me. → Maya loaned me the book.

(3) I didn't read the book that Maya loaned me. → Maya loaned me the book.

Of course, aspects of presuppositions can be negated with "metalinguistic" negation, signaled by heavy lexical stress on the negated constituent (*I didn't read the book that Maya gave me because she didn't GIVE me any book!*). But then metalinguistic negation can negate anything at all, including intonation, lexical choice, or accent. Modulo the possibility for metalinguistic negation, the presupposed parts of a sentence are taken for granted, and are not understood to be part of what is asserted by the sentence. Contrastive focus, like metalinguistic negation, is also an orthogonal dimension: it is marked by stress on a constituent in any position—even on a topic or on a backgrounded element. Contrastive focus is often indicated by a fall–rise pitch accent. It will not be central to concerns discussed in this chapter.

The three relevant categories of relational information structure are exemplified in Table 7.1.

TABLE 7.1. General categories of relational information structure

	Example (relevant constituent underlined)
primary topic	She hit a pole.
within the potential FOCUS DOMAIN	George met her.
backgrounded elements	The man who she told him about called.

7.1 "Island" Constraints

John Robert Ross, in his stunning dissertation, noticed that it was not possible to create unbounded dependencies involving just any aspect of a sentence (Ross 1967). In particular, certain syntactic constructions are "islands" to unbounded dependency relations or "extraction." These include complex noun phrases, complex subjects, complements of manner-of-speaking verbs, and adjunct clauses as illustrated in Table 7.2.

The judgments in the case of the complex NPs and subject islands are more robust, and less dependent on context, than in either of the latter two instances, each of which is marginally acceptable (despite the fact that all of these types are classified as "strong" islands in the generative literature).[1] Exploring these subtle differences in judgments requires us to look in a more detailed way at the discourse functions of each of the constructions involved, and is beyond the scope of the present chapter.

[1] Baltin (1982) has suggested that manner-of-speaking complements are adjuncts, not arguments, in that they are generally omissable. Under that interpretation, this case would be an instance of the generalization about adjuncts.

TABLE 7.2. Classic examples of "island" constraints

*Who did she see the report that was about? (cf. She saw the report that was about x)	**Complex NPs** (both noun complements and relative clauses)
*Who did that she knew bother him? (cf. That she knew x bothered him)	**Subjects**
??What did she whisper that he left? (cf. She whispered that he left x)	**Complements of manner-of-speaking verbs**
??What did she leave the movie 'cause they were eating?[2] (cf. She left the movie because they were eating x)	**Presupposed adjuncts**

In a constructional approach, "movement" phenomena are understood to involve the combination of some construction with an unbounded dependency construction (e.g. a question, relative clause, topicalization). The constructions that are combined each have particular information-structure properties, and those properties must be consistent in order to avoid a pragmatic clash. Most if not all of the traditional constraints on "movement"—i.e. the impossibility of combining a construction involving a long-distance dependency with another construction—derive from clashes of information-structure properties of the constructions involved. In Section 7.10, we will address the role of processing in judgments of ill-formedness.

Direct replies are sensitive to islands

Largely ignored are other types of discourse-level phenomena that are sensitive to islands. For example, Morgan (1975) long ago observed that DIRECT REPLIES TO QUESTIONS are sensitive to islands. Let us assume that the answer to (4) is that Laura was dating someone new. None of the replies in (5)–(8) is an appropriate answer to the question posed in (4), since the proposition that would answer the question (namely, that Laura was dating someone new) is expressed within an island. The answer cannot felicitously be expressed in a relative clause (5), in a sentential subject (6), nor in the sentential complement of a manner-of-speaking verb (7), nor by a presupposed adverb (8):

(4) Why was Laura so happy?

Relative clauses are islands to appropriate answers

(5) #The woman who thought she was dating someone new lives next door.
(cf. The woman who lives next door thought she was dating someone new.)

[2] 'Cause is used here instead of *because* to encourage the presupposed interpreation of the adjnct.

Sentential subjects are islands to appropriate answers

(6) #That she's dating someone new is likely.
(cf. It's likely that she's dating someone new.)

Complements of manner-of-speaking verbs are islands to appropriate answers

(7) #John shouted that she was dating someone new.
(cf. John said she was dating someone new.)

Presupposed adverbials are islands to appropriate answers:

(8) #John was hysterical 'cause she was dating someone new.
(cf. John left Manhattan in order that she could date someone new.)

Through Gricean implicatures of relevance, contexts can be found in which the sentences marked as infelicitous above, or closely related ones, are much improved, interpreted as *indirect* responses to the question in (4). But as direct responses to the question posed, each of the responses above is markedly odd.[3] Since this island phenomenon exists across sentences, indeed, across interlocutors, it strongly raises the possibility that constraints on islands are fundamentally related to discourse; the phenomenon is not easily described in purely syntactic terms.

Exclamative ah! *is sensitive to islands* Another source of island effects that cries out for a discourse-level explanation is James' (1972) observation that certain discourse particles such as EXCLAMATIVE *ah*! cannot be used to remark on propositions within islands. For example, the following formulations in (9) cannot be used to convey an exclamation about the fact that Laura was dating someone new, since that fact is conveyed within a complex NP, an island. The examples in (10)–(12) demonstrate that other islands are similarly outside the scope of the exclamation.

Relative clauses are islands to scope of *ah!*

(9) Ah! the woman who thought Laura was dating someone new lives next door! (exclamative cannot refer to the proposition that she was dating someone new)
cf. Ah! the woman who lives next door thought she was dating someone new!

[3] It is possible directly to answer the question in (4) with a sentence such as *I heard a rumor that she was dating someone new*, where the answer appears to be within a RC. But such expressions are equally well interpreted such that the entire NP, *a rumor that she was dating someone new*, is the answer, and of course the entire NP can be extracted (*What did you hear?*).

Sentential Subjects are islands to scope of *ah!*

(10) Ah! That she is dating someone new is likely! (exclamative cannot refer to the proposition that she was dating someone new)
cf. Ah! it is likely that she was dating someone new!

Complements of manner-of-speaking verbs are islands to scope of *ah!*

(11) Ah! John shouted that she was dating someone new! (exclamative cannot refer to the proposition that she was dating someone new)
cf. Ah! John said she was dating someone new!

Complements of presupposed adjuncts are islands to the scope of *ah!*

(12) Ah! John was hysterical 'cause she was dating someone new. (exclamative cannot refer to the proposition that she was dating someone new)
cf. Ah! John left Manhattan in order that she could date someone new.

The facts about indirect questions and *ah* exclamatives cannot naturally be accounted for by purely syntactic accounts (Cole et al. 1977). Both phenomena cry out for an explanation in terms of discourse properties of the constructions involved.

7.2 Backgrounded constructions are islands

Several researchers have suggested that the extraction site must be a *potential focus domain* (Erteschik-Shir 1979, 1998a; Takami 1989; Van Valin 1998; Van Valin and LaPolla 1997). That is, the constituent in which the gap exists (i.e. the constituent containing the canonical position for the fronted element) must be within the part of the utterance that is asserted. It cannot be presupposed. In accord with this observation, notice that none of the constructions in Fig. 7.1 is part of the focus domain—they are all pragmatically

Complex NPs
1. She didn't saw the report that was about him. → The report was about him.

Sentential subjects
2. That she knew it didn't bother him → She knew it.

Complements of manner-of-speaking verbs
3. She didn't whisper that he left. → He left.

Presupposed adverbials
4. She didn't leave the movie after they ate it. → They ate it.

FIGURE 7.1 Islands that involve presupposed information

PRESUPPOSED. That is, their propositional content is implied by both the positive and negative form of the sentence.

Elements within these presupposed clauses are not part of the potential focus domain. In accord with this idea, Erteschik-Shir (1979), Takami (1989), and Van Valin (1998; Van Valin and LaPolla 1997) each propose some type of a **negation test**, as independent verification that constructions such as those identified in Figure 7.1 are not part of the potential focus domain; cf. the test for being part of the potential focus domain above. When an assertion is negated, only elements within the potential focus domain are contradicted.

The idea that the extraction site must be within the potential focus domain can be used to explain why these particular constructions are islands to unbounded dependency relations. But a constraint that extraction can only occur from potential focus domains does not explain how it is that the SUBJECT argument (the whole subject constituent), which is also not part of the focus domain, is readily available for unbounded dependencies. In fact on Keenan and Comrie's (1977) accessibility hierarchy, subjects are the *most* likely candidates for potential extraction:

(13) subject > direct object > oblique object > object of comparison

If a language permits a relative clause to be formed on a noun phrase associated with a grammatical function low in the hierarchy, it will permit relativization on NPs representing all grammatical functions above it. Work in Centering Theory, a computational linguistics tool, has found that the same hierarchy predicts the likelihood of subsequent mention (Grosz, Joshi, and Weinstein 1983, 1995).

It is possible to account for both facts, that unbounded dependencies are normally available to clausal subjects or to elements within the focus domain with the following generalization:

(14) Backgrounded constructions are islands (BCI)

Given the definition of backgrounded elements provided at the outset of this chapter, this claim entails that only the primary topic in a clause, or elements within the potential focus domain are candidates for unbounded dependencies. Notice that elements *within* clausal subjects are backgrounded in that they are not themselves the primary topic, nor are they part of the focus domain.

The restriction on backgrounded constructions is clearly motivated by the function of the constructions involved. Elements involved in unbounded dependencies are positioned in discourse-prominent slots. *It is pragmatically anomalous to treat an element as at once backgrounded and discourse-prominent.* A critical role for processing comes into play as well, as discussed in Section 7.10.

The definition of backgroundedness implicitly acknowledges that the notions of topic and focus are not opposites: both allow for constituents to be interpreted as having a certain degree of discourse prominence (see, e.g., Arnold (MS) for experimental and corpus evidence demonstrating the close relationship between topic and focus). One sentence's focus is often the next sentence's topic. That is, once new material is introduced into the discourse, it is available to persist as a continuing topic during subsequent discourse: i.e. it may have high *topic persistence*. Centering Theory captures the relationship between topic and focus very naturally (Grosz, Joshi, and Weinstein 1983, 1995). In the theory, discourse referents in the speaker's focus of attention are called *centers*. All arguments in each utterance are forward-looking centers, which become potential antecedents for referential terms in a subsequent utterance. A special member of the forward-looking centers is also a backward-looking center, corresponding roughly to "topic," in that it indicates what the utterance is "about" and serves to link the utterance to the preceding utterance.

The fact that topic and focus are not opposites goes some way toward explaining the pervasive confusion about the terms. Both topic and focus are sometimes described as the "locus of attention." There is likely a sense in which that is exactly accurate: both the primary topic and focus are centers of cognitive attention (Deane 1991). These stand in contrast to elements that are neither the primary topic nor within the focal domain of a sentence: *backgrounded* elements.

The formulation that "extracted" elements cannot be backgrounded predicts that certain aspects of sentences heretofore unclassified as traditional "islands" should in fact resist unbounded dependency relations. This is true, for example, of parentheticals, which are not part of the focus domain as evidenced by the fact that they are not understood to be negated by sentential negation, and neither are they the primary topic:

(15) I just read—stop me if I have already told you about this—a great new book.

(16) I didn't just read—stop me if I have already told you about this—that great new book.

(17) *Who$_i$ did I just read—stop me if I have already told [i] about this—a great new book.

The generalization that extracted elements cannot be backgrounded accounts for a wide range of facts. Since backgrounded propositions are not part of what is asserted, they do not provide felicitous answers to questions; this

explains the fact that answers are sensitive to island phenomena. Since the exclamative marker *ah!* has scope only over what is part of the focal domain, it follows that it, too, cannot refer to backgrounded propositions, since they are by definition not part of the focus domain. Many other facts about the resistance of certain constructions to unbounded dependencies are accounted for as well. The facts surrounding the ditransitive construction are discussed below.

7.3 The ditransitive recipient argument resists unbounded dependencies

A statistical generalization about the information structure of ditransitives can be used to account for how the ditransitive construction interacts with unbounded dependency constructions. In particular, recall from Chapter 2 that the recipient argument of the ditransitive resists unbounded dependency relations (Erteschik-Shir 1979; Oehrle 1975):

(18) ??Who did Chris give the book?

(19) ??The boy who Mary had already given the key let himself in.

These judgments are somewhat subtle. Therefore, a Google search was performed to attempt to quantify the dispreference. Results showed that when a recipient was questioned (e.g. *who did she give...*), prepositional paraphrases (e.g. *who did she give the money **to**?*) outnumbered ditransitives by forty to one. In particular, only three questioned recipients of a ditransitive were returned out of the first 120 examples, accounting for only 2.5 per cent of the examples considered.[4] This skewing of the data towards questioning the recipient of the prepositional paraphrase and not the ditransitive recipient, exists despite the prescriptive injunction against stranding prepositions. Moreover, *give* is lexically biased to appear in the ditransitive over the prepositional paraphrase when there is no long-distance dependency

[4] The three attested instances of questioned recipients of ditransitives found are provided below. Two of these involved particularly long theme phrases, strongly motivating the use of the ditransitive (Wasow 2002).

(i) When Julia left the Valley, who did she give control of her interest in Falcon Crest?

(ii) In Paul's report to James and to the elders, who did he give credit for the work among the Gentiles?

(iii) Jack: Yes, but who did she give the eye? ☺

Example (iii) involves an idiomatic phrase *to give someone the eye*, "to look seductively at someone." The expression with *to* (*she gave an eye to him*) only has a literal interpretation.

relation.[5] Therefore it is fair to say that there is a systematic and strong dispreference for questioning the recipient argument of the ditransitive construction. In order to account for the judgments in (18) and (19), we take a brief detour, outlining the discourse properties of the recipient argument of the ditransitive.

The ditransitive recipient as secondary topic

While the subject argument is generally agreed to be the default primary topic in a clause, the recipient argument of the ditransitive construction has been described as a secondary clausal topic (Dryer 1986; Givón 1979, 1984; Langacker 1987a; Van Hoek 1995; see also Nikolaeva 1991 for a general discussion of secondary topic status). Since it is not the primary topic, and as we shall see, it is generally not within the potential focus domain, it qualifies as being backgrounded.

In both corpus and experimental studies, it has been demonstrated that argument status as new or given plays a role in conditioning whether the ditransitive construction is chosen over the dative paraphrase (Arnold et al. 2000; Bresnan and Nikitina ms; Dryer 1986; Givón 1979, 1984; Wasow 2002). In contrast to the ditransitive recipient argument, the theme argument of the ditransitive strongly tends to be new information, rarely being already given in the discourse context. The proposed generalizations with example sentences are provided below:

(20)	Ditransitive:	Subj	Obj1	Obj2
		agent	recipient	theme
		topic	secondary topic	new/accessible

(21) She gave him a book.

(22) ??She gave a man them.

The paraphrase with *to* is not constrained in this way, as indicated below:

(23)	"Dative":	Subj	Obj	PP
		agent	theme	goal
		topic		

(24) She gave a book to him.

(25) She gave it to a man.

[5] Wasow (2002) cites the lexical preference for ditransitives over prepositional paraphrases when *give* is used to be roughly four to one.

What is the evidence for the claim that the recipient argument of the ditransitive is a secondary topic? Corpus studies have demonstrated that the recipient argument is typically pronominal and if it is not expressed pronominally, it tends to be expressed with a definite NP description. That is, the recipient argument of the ditransitive construction rarely introduces a new argument into the discourse. These strong trends have been observed by Thompson (1990) with a database of 162 tokens, and by Collins (1995) who analyzes a corpus of spoken and written language containing 108 instances of the ditransitive construction and 57 prepositional paraphrases. In both studies, the ditransitive recipient is overwhelmingly already given in the discourse.[6] Because these corpus studies were relatively small, and because statistics are known to differ somewhat across different corpora, I replicated these earlier studies on two distinct databases as described in the Appendix to this chapter. The first database consisted of spoken child-directed language from the Suppes corpus in the CHILDES database (MacWhinney 1995).[7] The second database involved primarily written texts gathered from the web using the Linguists' Search Engine on Alta Vista. In fact, the ditransitive recipient argument is highly likely to be pronominal as compared with either the goal argument of the prepositional paraphrase or the theme argument of the ditransitive. In neither corpus were there significant instances of new recipient arguments. See also Ruppenhofer (2004) for additional corpus analysis of the ditransitive.

Referential givenness is not identical to topicality, but the two are correlated in the following way. Continuing topics are given in that they have to have been mentioned in order to be continuing as topics, and even newly established topics tend to be accessible or anchored in the discourse as opposed to brand new, insofar as they typically appear with a definite determiner or are explicitly related to a given entity by means of a possessive determiner or relative clause (Francis, Gregory, and Michaelis 1999; Lambrecht 1994).

The ditransitive recipient has other hallmarks of topicality as well: it is animate and its existence is presupposed (Polinsky 1998). Animate referents are more likely to be topical than inanimate referents (e.g. Bock, Loebell, and Morey 1992), and topical referents are generally presupposed to exist (Strawson 1964). The fact that both topics and ditransitive recipient arguments tend

[6] In data from an experiment reported in Wasow (2002), with the relative weight of recipient and theme controlled for, we see a strong preference for the recipient argument to be given as opposed to accessible. The discourse context the experimenters created was not amenable to instances of new recipients.

[7] I am grateful to Mike Tomasello for making these data available to me.

to be given, animate, and presupposed to exist is explained if we assume that the recipient argument is a secondary topic.

Moreover, designating the recipient argument of ditransitives as a secondary topic predicts that it should not be part of the focus domain. This prediction holds; notice that it is not understood to be negated by simple negation (Erteschik-Shir 1979):

(26) She gave her a ball.
#No, him.

Only if there is lexical stress (contrastive focus) on the recipient argument can a given ditransitive recipient be construed to be part of the focus domain:

(27) She gave HER a ball. (recipient is contrastive focus)
No, him.

This is in contrast to the prepositional goal argument which *can* be interpreted as part of the focus domain without contrastive stress.[8]

(28) She gave a ball to her.
No, (to) him.

Since the recipient argument is not the primary topic and is not within the focal domain, it qualifies as being a backgrounded element. An independent indication of the backgrounded nature of the recipient argument comes from the fact that a full 26 per cent (50/190) of examples involving *give* on a search using the Linguists' Search Engine (Alta Vista) were non-idiomatic examples such as the following in which the recipient or goal argument was unexpressed:

(29) But he *gave the glad tidings* that despite all the negative propaganda, Islam was spreading rapidly in America.
(30) Furthermore we *give exact algorithms* for interval graphs . . . and graphs of bounded asteroidal number.

Backgrounded elements are often candidates for omission (see Chapter 9). Thus it is safe to conclude that the recipient argument is generally backgrounded (as a secondary topic) in the clause.

[8] Parallel judgments hold for the following version of the negation test:

(i) She didn't give him a ball. It was a bat.
(ii) She didn't give him a ball. #It was her/Mary.

Again, with contrastive stress, it is possible to interpret the scope of negation to include the recipient:

(iii) She didn't give HIM a ball. It was Mary (she gave it to).

Statistical generalization, not hard and fast constraint

Despite the strong statistical trend for the recipient argument of the ditransitive construction to be non-new and not part of the focus domain, it is important to acknowledge that introspective grammaticality judgments are less robust. Constructed examples such as the one below are acceptable:

(31) The US committee hoped to give an American the award.

In this example, the recipient argument, *an American*, is new and may receive either a specific or a non-specific interpretation: either the committee wanted to give a specific American the award or they wanted to give some American or other the award. In the same vein, while the focus domain tends not to include the recipient argument, when the recipient argument is new, it can be part of the potential focus domain, as it is within the scope of simple negation:

(32) The US committee hoped to give an American the award.
No, a Canadian.

To account for this fact, it may be necessary to introduce degrees of backgroundedness. If backgrounded arguments correspond to a lack of cognitive attention, it would make sense that backgroundedness might be a gradient term. We return to this idea below.

The analysis of the ditransitive in (20) finds additional support in that it directly predicts both the fact that the recipient argument resists unbounded dependencies and the scope facts concerning the ditransitive, as we shall see in Section 7.12. These are two aspects of the construction that have been emphasized but not adequately explained by formal accounts (recall the discussion in Section 2.4).

As defined above, elements that do not correspond to the primary topic and are not within the focus domain are backgrounded and therefore are not candidates for extraction. Because the status as backgrounded in the ditransitive is actually a statistical generalization, not a categorical constraint, the fact that unbounded dependencies involving the recipient argument are occasionally found is also expected. In addition, this information structure account of the ditransitive facts serves to unlock an often overlooked puzzle: **passive recipients can be questioned and relativized.** The active ditransitive recipient resists unbounded dependencies because it strongly tends to be backgrounded: although topical, it is not the primary topic, and it is not within the focus domain. Nonetheless, the recipient argument can be freely questioned or relativized if it is already the subject of a passive:

(33) Who was given the book? passive
(cf. ??Who did John give the book?) active

(34) The man who was given a book left early. passive
(cf. ??The man who she gave a book left early.) active

That is, if the recipient argument is a subject (via passivization), then it is free to be involved in unbounded dependency relations that are otherwise only marginal. This makes sense since passivized recipients are the primary topics in a clause, and primary topics are readily involved in unbounded dependencies.

Empirical evidence exists to support the idea that the recipient argument is not backgrounded when it is in subject position. It turns out that the information structure generalizations about the recipient argument differ in the active and passive. The passive-ditransitive recipient patterns like other passive subjects, not like recipients in the active voice. In particular, the subject-recipient argument of the passive ditransitive is markedly less likely to be already given, i.e. marked with a pronominal, than the recipient argument of active ditransitives. Of the first ninety-two passive ditransitives found by the Linguists' Search Engine, the subject-recipient argument was marked as indefinite 10 per cent of the time, a definite lexical NP 70 per cent of the time, and a pronominal only 20 per cent of the time. Similar percentages were found for passive sentences generally (12 per cent, 67 per cent, and 21 per cent, respectively of the first 100 passive sentences returned). On the other hand, the recipient argument of an active ditransitive was only marked as indefinite 1.5 per cent of the time, and definite 32.5 per cent of the time. It was marked with a pronominal a full 66 per cent of the time. These comparisons are shown in Fig. 7.2.

Given this empirical finding, the fact that the recipient argument is available for unbounded dependency constructions when it is passivized is expected on an account that predicts that unbounded dependency restrictions stem from information structure properties of the constructions involved.[9]

The statistical generalization that the recipient strongly tends to be given in the discourse and not part of the focus domain is used to explain the scope relations of the ditransitive as well in Section 7.12.

[9] The percentages of lexical NP subjects are known to vary greatly across different types of corpora (Francis, Gregory, and Michaelis 1999). However, since I used the same corpus to search for all of the arguments under discussion, it is fair to compare the *relative* percentages.

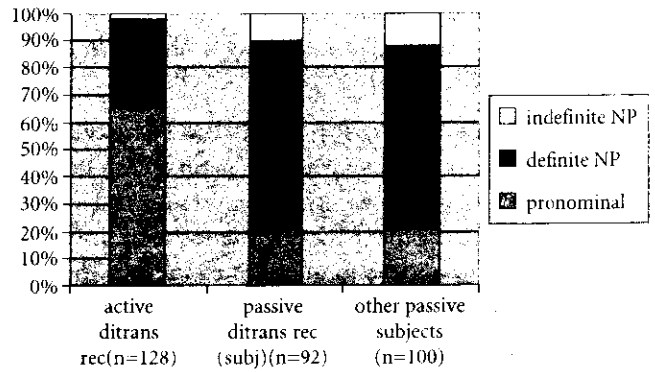

FIGURE 7.2 **Comparison of recipient argument in active ditransitives and in passive ditransitives, and passive subjects generally, in data collected from Alta Vista using the Linguists' Search Engine.** Active ditransitive recipients are different from passive ditransitive recipients ($\chi^2(2, N = 211) = 67.0$, $p < .001$). Passive ditransitive recipients and other passive subjects are not different from each other ($\chi^2(2, N = 192) = .26$, $p = .88$).

7.4 Subordinate Clauses

Are subordinate clauses backgrounded? It might seem that negating the main clause would not negate the subordinate clause, so that subordinate clauses should uniformly be backgrounded. In fact as will be discussed below, many subordinate clauses are backgrounded. However, in other cases, sentential negation can be used to imply a negation of the proposition expressed by the subordinate clause. For example,

(35) She doesn't think they sent the letter.

(36) He didn't say they went to the market.

(37) They don't believe she left early.

The complement clauses in (35)–(37) are within the potential focus domain of the main clause negation. That is, (35) for example, will be true if she thinks they didn't send the letter (or if she didn't think anything at all). On the other hand, negating other main verbs much less readily conveys the implication that the complement clause is within the scope of negation. We have already seen that complements of manner-of-speaking verbs are generally islands. As predicted, the complement of manner-of-speaking verbs is generally not within the potential focus domain (without special context):

(38) I didn't grumble that they sent the letter. (manner-of-speaking verb) (asserts that I didn't grumble; there is typically no implication about whether they sent the letter)

This is because the main focus of the clause is on the manner of speaking, since otherwise a simpler verb such as *say* could have been used. The subordinate clauses in (35)–(37) are not necessarily construed to be backgrounded, whereas the subordinate clause in (38) strongly tends to be. In accord with this idea is the fact that it is acceptable to question elements within subordinate clauses of the types in (35)–(37) but not those in (38), as illustrated below:

(39) What did you think that they sent?

(40) What did you say that they sent?

(41) What did you believe that they sent?

(42) ??What did you grumble that they sent? (manner-of-speaking complement)

At the same time, if it is clear in the context that the manner of speaking is already topical, the complement clause can be construed to convey the focus domain of the sentence. In this case, long-distance dependencies from within the manner-of-speaking complement are improved.

7.5 Reason clauses

Reason clauses may be backgrounded or not, depending on context, position, intonation, and choice of connector. The reason clause may be part of what is asserted, particularly if the main sentence accent falls within the reason clause. This interpretation is clear in (43):

(43) She didn't travel to Memphis because she wanted to see ELVIS. (can be used to imply that while she did travel to Memphis, it was not to see Elvis)

Under this interpretation, unbounded dependencies from within the *because* clause are acceptable.

(44) ?Who did she travel to Memphis because she wanted to see?

As a type of adjunct, reason clauses are considered to be "strong" islands within mainstream generative grammar (Cinque 1990). This predicts that examples such as (44) should be unacceptable, and yet they are not (cf. also Kluender 1998). Reason clauses involving *in order to* are even more problematic for the notion that all adjuncts are "strong" islands. *In order to* reason clauses differ from those expressed by *because* phrases in that the scope of negation tends to include the *in order to* verb phrase even without special intonation (45):

(45) She didn't travel to Memphis in order to see Elvis (can be used to imply that the traveling was done but not in order to meet Elvis)

That is, *in order to* clauses are not presupposed and are not backgrounded, even though they are adjuncts (in that they are omissible and apply to a wide range of expressions). In accord with the BCI, questioning elements from within *in order to* clauses is fully acceptable:

(46) Who did she travel to Memphis in order to see?

Similarly the following attested cases, all from Santorini (2005), involve unbounded dependencies from within adverbial clauses that are not presupposed (the dependent filler is boldfaced, the adjunct phrase underlined):

(47) Why did this admirable man turn to **the very tyranny** he sacrificed so much <u>to overthrow</u>? (J. Blatt and S. Jacobs (2002), "Things fall apart", *The Nation* 274, no. 13: 30)

(48) Enron . . . ingratiated itself with **those very politicians** it gloried <u>in mocking in its ads</u>. (T. Frank (2002), "Shocked, shocked! Enronian myths exposed", *The Nation* 274, no. 13: 17)

(49) . . . Bush declared in **a victory speech** that he was forced to wait 36 days <u>to deliver</u>. (Kathy Kieler, "We can unite," *USA Today*, December 12, 2000, p. 1A)

(50) **a scenario** that government agencies are spending billions of dollars <u>preparing for</u> (*The World*, NPR, February 10, 1999)

Thus it is clear that adjuncts are not islands across the board; only backgrounded (presupposed) adjuncts are islands.

7.6 Non-restrictive relative clauses

The BCI account predicts straightforwardly that definite restrictive relative clauses and noun complements should be unavailable for long-distance dependencies, because both are backgrounded. *Non*-restrictive relative clauses might seem to pose a problem insofar as they often convey new information. Yet, the negation test still clearly demonstrates that they are backgrounded: they act as informational asides, not unlike the parentheticals mentioned above.

(51) I saw John, who I told you about last week.

(52) I didn't see John, who I told you about last week. → I told you about John last week.

(53) *Who did John, who I told _ about last week leave early?

Thus the BCI correctly predicts that restrictive as well as non-restrictive relative clauses should be islands.

7.7 Presentational relative clauses

Not all relative clauses are equally backgrounded. "Presentational" relative clauses can serve to convey the main assertion in a clause. In Danish, and to some extent, in English as well, these relative clauses are, as predicted by the BCI, available for unbounded dependencies:[10]

(54) Hvad for en slags is er der mange born der kan li? Danish (Erteschik-Shir 2002)
What kind of ice cream are there many children who like?

(55) We have a visitor who there's no one who's willing to host. (Chung and McCloskey 1983)

(56) John is the sort of guy that I don't know a lot of people who think well of. (Culicover 1999: 230)

(57) That's the article that we need to find someone who understands. (Kluender 1998)

Notice that in these cases, the content of the relative clause is negated by sentential negation. For example:

(58) There are not many children who like this kind of ice cream. → many children don't like this kind of ice cream.

(59) That's not the article that we need to find someone who understands. → we don't need to find someone who understands that article.

These "exceptional" cases of presentational relative clauses that allow long distance dependencies are also expected on the BCI account.

7.8 Factive complements

Complements of factive verbs presuppose the truth of their complements by definition, as in (60):

[10] Relative clauses in French can be used presentationally, serving to convey the main assertion in the clause (Lambrecht 2004), and yet they are unavailable for unbounded extraction for reasons that remain mysterious.

(60) It bothered Sue that the mayor smoked cigars. → The mayor smoked cigars.

The BCI account predicts, then, that the complements of factive verbs should be islands, and in fact, long-distance dependencies from within the complement of a factive verb are not fully acceptable (Kiparsky 1971; Ross 1967):

(61) ??What did it bother Sue that the mayor smoked? (factive complement)

Factive complements are sometimes discussed as "weak" islands insofar as complements are more readily available for unbounded dependencies than are adjuncts:

(62) *Why did it bother Sue that the mayor smoked cigars_?
(the sentence is acceptable if the question is understood to ask why Sue was bothered, but it is unacceptable under the interpretation that the reason for the mayor smoking is being queried)

The negation test would seem to indicate that factive complements should be quite robust islands, whereas in fact it is clear that unbounded dependencies from within these complements are less than crashingly bad. This fact may be taken as undermining the account, or it might instead suggest that the negation test is not a perfect indicator of backgrounded status. Intuitively the proposition expressed by the complement clause of factive verbs is more central to the discourse than that expressed in relative clauses. I leave this question aside until a better gauge of backgrounded status is found.

There are other unsolved mysteries that exist on an account that relies purely on backgrounded status. These include indefinite full relative clauses and wh-complements, as discussed below.

7.9 Tricky cases

Indefinite Relative clauses

In the case of relative clauses with *in*definite heads, the content of the relative clause appears to be within the focus domain, at least according to the negation test. That is, the content of relative clauses that are headed by indefinite NPs is not presupposed (Hooper and Thompson 1973):

(63) She didn't meet a boy who resembled her father... the boy resembled her mother.

Thus the BCI predicts that indefinitely headed relative clauses should not be islands, insofar as they are within the focus domain and are therefore not

backgrounded. In fact, we see the predicted contrast between definite and indefinite head nouns in the case of "reduced" relative clauses:[11]

(64) Who did she see a report about?

(65) ??Who did she see the report about?

(see Deane 1991 for discussion and analysis of such cases).

However, contrary to the prediction of the BCI, full relative clauses on indefinitely headed arguments *are* islands:

(66) *Who did she see a report that was about?

An explanation of this fact remains elusive; in order to investigate the contrast between (64) and (66) it is clearly necessary to gain a better understanding of the distinct discourse properties of reduced versus full relative clauses. It seems, however, the contrast is in the right direction—intuitively, full relative clauses are more backgrounded than reduced relative clauses. Again, this suggest that the negation test may not be a refined enough measure of background status. Also it is clear that if anything, definitely headed RCs are more island-like than indefinitely headed RCs. Postal (1998), for example, argues that indefinitely headed RCs are "selective" islands, whereas definitely headed RCs are "absolute" on the basis of the following types of differences in judgments:

(67) a. *It was the distance to the chasm that Frank liked the man who was able to determine.
b. ?It was the distance to the chasm that Frank knew someone who was able to determine.

Clearly more work is needed to understand fully the restrictions on long-distance dependencies out of indefinitely headed relative clauses, but at least the preference for such dependencies when compared with those out of definitely headed relative clauses is in the predicted direction.

Wh-complements

Wh-complements are part of the potential focus domain, as is evidenced by the fact that they may be negated by sentential negation without special intonation:

[11] On the basis of their surface structure, these cases are probably better understood to be noun complements instead of relative clauses. Their classification makes no difference to the issues discussed here—as part of complex noun phrases, traditional linguistic theory predicts that they should be islands.

(68) She wasn't wondering whether she would meet William, she was wondering whether she would meet Bob.

This then predicts that wh-complements should *not* be islands, although they are generally taken to disallow unbounded dependencies. In fact, judgments about the following examples are somewhat variable, with a minority of speakers rating them fully acceptable:[12]

(69) ?What did Bush ask whether he could do?

(70) ?Which man is he wondering whether she's met?

However, as was the case with factive islands, unbounded dependencies from within wh-clauses are clearly degraded when the dependent element is an adjunct instead of an argument.

(71) *When did he ask whether he could eat dessert _?
(on the interpretation that the timing of desert is at issue)

It is clear that the BCI account predicts a wide variety of facts in a straightforward way: the fact that subjects, definite relative clauses, full noun complements, the ditransitive recipient argument, presupposed adjuncts, complements of manner-of-speaking verbs, and factive verbs are all islands is predicted. The fact that unbounded dependencies are sometimes allowed from presentational relative clauses, (reduced) indefinite relative clauses, non-presupposed adjuncts, and verbs of saying is also expected. In addition, the fact that direct questions and exclamative *ah!* is sensitive to islands is also expected on an account that relies on the discourse properties of the constructions involved.

The first two columns in Table 7.3 summarize the predictive power of the BCI generalization. Also included, simply for reference and a point of comparison, in the final column are the predictions made by the "subjacency" constraint that is typically appealed to in generative accounts of island phenomena. The subjacency constraint is a purely formal generalization that states that there may not be more than one "bounding node" (NP or S) between a dependent (filler) and its canonical position (gap).

The BCI generalization arguably has more predictive power than the syntactic subjacency generalization insofar as it allows for finer-grained distinctions than are possible on a purely syntactic account. The BCI generalization

[12] Informal judgments were solicited from twelve naive native speakers based on a six-point scale, where 1 was "terrible" and 6 was "perfect." The average score across all subjects for extraction out of wh-clauses was 3.2; the average for the grammatical extractions (e.g. *Who did he tell her he met?*) was 4.8; for extraction out of definite relative clauses was 1.75.

TABLE 7.3. Predictions made by BCI generalization and "subjacency" constraint

	Islands	Predicted to be islands by BCI	Predicted to be islands by subjacency constraint
Arguments within subjects	Yes	Yes	Yes
Ditransitive recipient:			No
Active	Yes	Yes	
Passive	No	No	
Adjuncts	Sometimes (when presupposed)	Sometimes (when presupposed)	Yes
Relative clauses and Noun Complements			Yes (except by stipulation)
Definite	Yes	Yes	
Presentational	No	No	
Indefinite:			
"Reduced"	No	No	
Full	Yes	No	
Wh-complements	Yes (marginally)	No	Yes
Sentential Complements:			No (except by stipulation)
say, tell	No	No	
manner-of-speaking verbs	Yes (unless manner-of-speaking verb is made neutral by context)	Yes (unless manner-of-speaking verb is made neutral by context)	
Factive verbs	Yes (marginally)	Yes	

predicts that the active ditransitive recipient argument should disfavor long-distance dependencies, while the passive recipient does not; the BCI also predicts in a principled way the fact that many adjuncts are not islands. Moreover, the BCI generalization predicts, again, in a principled way, the difference between *say* and *tell* on the one hand and manner-of-speaking and factive verbs on the other. It also correctly predicts that presentational relative clauses and "reduced" relative clauses (or complements) of nouns should not be islands, while definite relative clauses should be. The preponderance of the evidence indicates that island constraints exist because backgrounded elements are not available for unbounded dependencies.

At the same time, Table 7.3 includes two mismatches between information structure status and island status: unbounded dependencies from within full

indefinite relative clauses are not possible and are only marginal from within wh-complements. These judgments hold despite the fact that neither indefinite relative clauses nor wh-complements are backgrounded according to the negation test. A full account will require an explanatory analysis of these cases. In the case of wh-complements, processing demands have been demonstrated to play a crucial role, as outlined in the following section.

7.10 A critical role for processing demands

Kluender and Kutas (1993) provide compelling experimental evidence that processing demands coupled with lexical factors lie at the root of variable acceptance of unbounded dependencies within wh-complements. To demonstrate a role for lexical factors, they report that judgments are systematically variable, depending on the choice of complementizer, even in yes/no questions that are considered by linguistic theory to be fully and equally grammatical. The relative judgments are indicated in (72)–(74):

(72) Isn't he sure [that the TA explained it to them in the lab]? >

(73) Isn't he sure [if the TA explained it to them in the lab]? >

(74) Isn't he sure [what the TA explained ___ to them in the lab]?
(Kluender and Kutas 1993)

That is, questions with embedded complements were judged to be less acceptable when the complement clause contained an *if* as compared with *that*, and less acceptable still if they contained *who/what*. Kluender and Kutas observe that the complementizer *that* requires little semantic analysis, whereas the complementizer *if* demands that a hypothetical context be constructed, and the complementizers *who/what* requires that a referent be activated.

There clearly is an extra processing load involved when arguments appear in displaced positions relative to their canonical expression. Just and Carpenter's (1992) Capacity Constrained Comprehension Theory suggests that the system's goal is to release elements from the memory store as quickly as possible (cf. also Gibson 2000; Gibson et al. 2005; Kitagawa and Fodor 2003; Kluender 1998; Kluender and Kutas 1993; Ellefson and Christiansen 2005). Kluender and Kutas (1993) in fact demonstrate a processing cost to holding a filler in working memory before finding its canonical position in the embedded clause (the "gap"). They find a marked decrease in acceptability ratings between (74) above and example (75), even though (75) is supposed to be fully grammatical. That is, while (74) was judged grammatical 90 per cent of the

time (with mean acceptability rating of 32/40), (75) was judged grammatical only 55 per cent of the time (with mean acceptability rating of 19/40).

(75) Who$_a$ isn't he sure [that the TA explained it to ____$_a$ in the lab]? >

(76) ?Who$_a$ isn't he sure [if the TA explained it to ____$_a$ in the lab]? >

(77) *Who$_a$ isn't he sure [what the TA explained ____$_b$ to ____$_a$ in the lab]? (Kluender and Kutas 1993)

Thus displaced constituents across clause boundaries place a burden on working memory. At the same time, accounts based on simple processing difficulty do not predict the full range of facts in Table 7.3. Processing accounts do not explain *why* it is that clause boundaries should present a problem, nor why clause boundaries involving factive and manner-of-speaking verbs should present more difficulty than clause boundaries of other verbs. Moreover, accounts based on the difficulty of crossing clause boundaries do not explain the facts surrounding the ditransitive, since a strong dispreference for unbounded dependencies exists, but no clause boundary is crossed; i.e. the recipient argument resists being the filler in an unbounded dependency despite the fact that it is an argument of the main clause. Finally, processing accounts also do not account for the fact that direct questions and exclamative *ah!* are sensitive to islands. The present proposal relies on recognizing the functions of the constructions involved in addition to recognizing the existence of processing constraints.

7.11 Cross-linguistic facts

The discourse properties of the constructions involved combine with processing factors to yield the witnessed patterns of acceptability. A critical role for processing demands is evident in languages that allow question words to appear in canonical position (wh-*in-situ* languages). These languages are much freer in their unbounded dependency relations (Huang 1982). For example, the following examples in Japanese (78) and Korean (79) are judged to be acceptable, despite the fact they involve questioning within complex NPs:

Japanese (examples from Norvin Richards (1999), Japanese Islands, http://web.mit.edu/norvin/www/papers/)

(78) Nakamura-san-wa yakuza-ga dare-o korosita tatemono-o
Nakamura-HON-TOP gangster-NOM who-ACC killed building-ACC
kaimasita ka?
bought
(*)Who did Nakamura-san buy a building [where gangsters killed _]?

Korean (examples from Sung Ho Hong (2000), "A Non-movement approach to wh-*in situ*", http://privatewww.essex.ac.uk/~shong)

(79) Mary-ka [[t$_i$ nwuku-lul bipanha-n] chaek$_i$]-lul choaha-ni?
Mary-NOM what-ACC criticize-REL book-ACC like-Q
Mary likes books [that criticize what]?

The present account predicts that there should be some evidence of the dispreference for asking questions within backgrounded constructions, even in *in situ* languages, since questions within islands should be dispreferred due to the function of the constructions involved. In fact, judgments about wh-words within islands are quite variable in some *in situ* languages. Examples often require special context to be judged acceptable (Kitagawa and Fodor 2003). Thai is an *in situ* language for which I was able to collect judgments from five native speakers; as indicated, there was little agreement among speakers about whether questions within islands are acceptable[13](*=unacceptable; #=requires special context; √=acceptable). All five judgments are indicated below by a string of *s, #s, and √s:

Complex NPs:

(80) ***√√ khaw hen raayngaan thii kiawkap khray?
(s)he see report that about who
She saw that report that was about who?

Subjects:

(81) ****√ kaan thii Mary ruucak khray thamhay John ramkhaan?
NOM that Mary know who make John be-bothered
She knew who bothered him?

Complements of manner-of-speaking verbs:

(82) * #√√√ Mary krasip waa John luum aray?
Mary whisper that John forget what
She whispered that he forgot what?

Presupposed adjuncts:

(83) **√√√ Mary ook caak rongnang phro phuun-phuun kin aray?
Mary exit from theatre because friends eat what
She left the movie because friends were eating what?

[13] I thank Theeraporn Ratitamkul for collecting these judgments for me.

Lakhota is another wh-*in-situ* language, and Van Valin and La Polla (1997) report a systematic avoidance of questions from within islands. An example is cited below (Van Valin and LaPolla 1997: 617):

(84) *Wičháša ki [[šųka wa táku *ϕ*-*ϕ*-yaxáke] ki le] wa-o-o-yáka he?
Man the dog a 3sgU-3sgA-bite the this 3sgU=3sgA-see Q
The man saw the dog which bit what?
(acceptable under yes–no question interpretation)

There are other types of evidence for island effects within other *in situ* languages, as well. For example, while it is possible to have a wh-word within a relative clause in Japanese (85), it is not possible to put *ittai* ("the hell") within the island (86); *ittai* must appear just outside the island (Pesetsky 1987). This is despite the fact that in main clause questions, *ittai* generally appears directly before the wh-word (87) (examples from Pesetsky 1987):

(85) Mary-wa ittai [John-ni nani-o ageta hito-ni] tatta-no?
Mary-TOP the.hell John-DAT what-ACC gave man-DAT met-Q
What the hell did Mary meet the man who gave t to John?

(86) *Mary-wa [John-ni ittai nani-o ageta hito-ni] atta-no?
Mary-TOP John-DAT the.hell what-ACC gave man-DAT met-Q
What the hell did Mary meet the man who gave t to John?
(Pesetsky 1987: 112, 126)

(87) Taroo-wa Hanako-ni ittai nani-o agemasita ka? (main clause: non-island)
Taroo-TOP Hanako-to the.hell what-OBJ gave Q
What the hell did Taro give to Hanako?

Quite parallel facts have been cited for Sinhala, another wh-*in-situ* language. The Sinhala question particle *də* cannot appear within an island but must appear just outside as well (Gair 1998; Hagstrom forthcoming).

Thus I hypothesize that constructions that serve to convey backgrounded information are *dispreferred* for containing question words cross-linguistically, even in *in situ* languages that do permit wh-words within backgrounded constructions. Suggestive evidence comes from the variable Thai judgments, and the data reported on Lakhota by Van Valin and La Polla. Systematic corpus evidence and sentence processing experiments would be required to confirm fully the claim more generally, however. It is predicted that examples of questions within islands are relatively rare as compared with other sorts of questions of similar complexity. It is also expected that there should be an

increase in processing load required for comprehending question words within islands. Still, it seems that the pragmatic clash is not necessarily sufficient to result in clear judgments of ungrammaticality unless there is displacement involved. **Displacement from canonical position creates additional processing load and this combines with the pragmatic clash to result in unacceptability.**

On the basis of different data, Deane likewise comes to the conclusion that information structure combines with processing load to result in ill-formedness. He puts the idea this way: "Extraction consists in the establishment of a long-range grammatical relation. An obvious prerequisite to establishing a relation between two concepts is that one be paying attention to both concepts at the same time... But in long-distance extraction, the two concepts to be linked are separated far enough from one another that some means must be provided to focus attention on both. And what means would be more natural than if the two concepts were ones which commanded attention anyway? [i.e. topic and focus]" (Deane 1991).

Hundreds of volumes have been dedicated to island phenomena and I do not pretend to have done the topic full justice here. For insightful discussion of the role of processing and information structure in "superiority" effects, see Arnon et al. (2004); for "that"-trace effects see Homer and Michaelis (2005); and for negative clauses Oshima (2005). For discussion of the complexities involved in parasitic gaps, see Levine and Sag (2003); for a non-syntactic account of certain definiteness effects, see Deane (1991) and Kluender (1998). For arguments that long-distance dependencies receive a natural treatment on monostratal approaches without traces, see Sag and Fodor (1994).

What I hope to have accomplished is to provide some new data and arguments for the idea that information structure properties of the constructions involved combine with processing complexity to account for many island phenomena. For additional arguments see also Deane (1991); Erteschik-Shir and Lappin (1979); and Erteschik-Shir (1998a,b). We now turn to issues related to quantifier scope, where again, we see a critical role for recognizing construction-specific discourse properties.

7.12 Topicality and Quantifier Scope

Quantifier scope plays a role in semantic interpretation when there are two quantifiers of different types. For example, if a sentence contains an existential quantifier (e.g. *a, one, some*) and a universal quantifier (e.g. *all, every, each*), then the existential quantifier may have wide scope over the universal: "There

exists (at least) one x, such that all y..."; or the universal quantifier may have wide scope over the existential: "For all y, there exists some x (or other)..." In the first case, the existential quantifier's referent is determined once and for all at the outset; in the latter case, the existential quantifier's referent may vary according to which one of the set of referents referred to by the universal quantifier is considered. It turns out that there are general tendencies that hold across constructions, once the particular discourse functions associated with constructional arguments are understood.

Ioup (1975) long ago observed that relative scope is determined by two interacting hierarchies. One hierarchy, slightly simplified here,[14] is described in terms of grammatical relations (see also Kuno 1991):

Topic	>	subject	>	IO	>	Obl/DO
Greatest inherent tendency for wide scope	>>>					least inherent tendency for wide scope

FIGURE 7.2 Hierarchy of topicality of grammatical relations

The hierarchy is interpretable as a scale of topicality. Topicalized expressions are likely to be informationally topical in that they tend to be related to a frame or set already introduced in the discourse (Chafe 1976; Lambrecht 1994). As noted above, subjects are the default clausal topic (Chafe 1987; Michaelis and Francis forthcoming; Givón 1984; Gundel 1998; Halliday 1967; Lambrecht 1994). Also as discussed in Section 7.3, ditransitive recipients (Ioup's "IO"s) are secondary topics. At the other end of the spectrum, obliques and direct objects are available to be part of the focus domain; they readily accept new or focal information and are generally non-topical (Du Bois 1987; Du Bois, Kumpf, and Ashby 2004).

Ioup provides support for the hierarchy from many diverse languages. It appears that languages generally obey a hierarchy of topicality in determining scope, irrespective of word order. For example, in Arabic and Japanese, a topicalized NP, appearing at the end of the string, necessarily has wide scope (examples from Ioup 1975: 53):

Arabic topicalization:

(88) Khabaz Hanna li kol li binaat feteer.
Baked John for all the girls cake
A cake, John baked for all the girls. a > all

[14] Ioup places Oblique arguments above DO in the hierarchy, but for our purposes it is sufficient to assume that they are both at the low end of the wide-scope hierarchy.

Japanese topicalization:

(89) Keiki-o Taro-wa nisannin-no onnanoko-ni tsukuri-mashita.
cake Taro a few girls for made
A cake, Taro made for a few girls. a > all

(90) Nissanin-no onnanoko-ni Taro-wa keiki-o tsukuri-mashita.
A few girls for Taro cake made
For a few girls, Taro made a cake. few > a

Similar effects can be seen in English (cf. Kuno 1991):

(91) a. A chocolate cake, John baked for all the girls. a > all (preferred)
b. For all the girls, John baked a chocolate cake. all > a (preferred)

(92) a. All of us have read many of his articles with great enthusiasm. all > many
b. Many of his articles, all of us have read with great enthusiasm many > all

The idea that topicality is intimately related to quantifier-scope interpretation can be seen to be intuitive once one examines what it means to have wide scope. A wide-scope interpretation of a variable is one in which the variable is given or fixed, and a second variable is interpreted with respect to it. That is, a variable x that is given a wide-scope interpretation provides the "anchor." It is within the context provided by the wide-scope operator that variables with more narrow scope are interpreted. This is clearly reminiscent of what topics are: the topic is given or fixed, playing the role of an anchor, while the comment is predicated of it. Chafe notes that topics typically set up "a spatial, temporal or individual framework within which the main predication holds" (Chafe 1976).

As Francis and Michaelis (forthcoming) point out, it is the scale of topicality that is important, not particular grammatical relations per se. They observe that particular semantic frames can constrain their arguments to be more or less topical, with the concomitant differences in scope possibilities. For example, they account for the following facts this way:

Creation frames:

(93) She made a canoe from it.
She made a canoe from every log. every > a

(94) An oak grew out of it.
An oak grew out of every acorn. every > a

Transformation frames:

(95) *She made a log into it.
*She made a log into every canoe. *every > a

(96) *A acorn grew into it.
*An acorn grew into every oak. *every > a

While the creation frame allows the material argument to be more topical than the created entity—presumably because a creation scene is typically specified in order to introduce a new entity that is created, the transformation frame requires that the material argument be more topical than the transformed entity, presumably because a scene of transformation is typically specified in order to introduce the new, altered entity into the discourse.

A second interacting factor in determining quantifier scope concerns the particular lexical choices of quantifiers. As Ioup notes, the hierarchy places quantifiers that refer to more members of a given set higher than those that refer to fewer members. A slightly elaborated version of Ioup's quantifier hierarchy is provided in Fig. 7.3.

Each > *every* > *all* > *most* > *many* > *several* > *a few* > *at least one* > *someone or other*

Greatest inherent tendency for wide scope >>> least inherent tendency for wide scope

FIGURE 7.3 Hierarchy of quantifiers

What is the source of the lexical effects on quantifier scope? The lexical effects are arguably related to topicality as well. An unidentified new referent serves as a poor choice of topic, and correspondingly non-specific quantifiers such as *some x or other* or (non-specific) *a* tend to have narrow scope. On the other hand, universal quantifiers, referring to an identifiable-in-context set of entities are like definites, and are thus relatively more likely to serve as topics. Differences among universal quantifiers stem from lexical semantic differences. For example, *each* is lexically a distributive quantifier, and is therefore a wide-scope operator, since a collective reading is required for narrow scope. See also Carlson (MS) for relevant discussion.

Kuno (1991) suggests two other interacting factors that play a role in quantifier-scope interpretation, both of which are interpretable in terms of topicality in a straightforward way. His factors and illustrative minimal pairs are given in Fig. 7.4.

More discourse-linked Q > Less discourse-linked Q
Many Democrats distrust some of these Republicans. some > many (easier than for b)
Many Democrats distrust some Republicans.

1st or 2nd person pronouns (local) Q > 3rd person (non-local) Q
Many of them know some of us. some > many (easier than for b)
Many of us know some of them.

FIGURE 7.4 Other factors relevant to topicality and quantifier scope (Kuno 1991)

"Discourse linking" implies that an NP is anchored in the discourse, and is therefore accessible as opposed to brand new; i.e. an NP that is discourse-linked is therefore more given, and givenness is correlated with topicality. First- and second-person pronouns are always available for topic status.[15]

The Ditransitive Construction and Scope facts

In order to put the arguments of the ditransitive construction on a linear scale of relative topicality, the following scale emerges from the characterization of ditransitives given earlier in (20).

Ditransitive topicality hierarchy:

(97) Agent (subject) >> recipient >> Agent (oblique) >> Theme
High topicality >>> Low topicality

These relative topicality facts predict all of the relevant scope facts that hold of the ditransitive. Since the subject is the primary topic, it should take wide scope over the recipient and theme arguments. This prediction holds:

(98) Everyone gave a girl a book. Everyone > a girl (preferred)

(99) Everyone gave him a book. Everyone > a book (preferred)

That is, in (98), the preferred reading is that everyone gave some girl or other some book or other.

In passive ditransitives, the recipient argument as subject has wide scope over the agent, since the recipient in the case of passive is the primary topic of the clause. The oblique agent remains more topical than the focal theme argument, and as expected, the former tends to be interpreted as having wider scope over the latter:

[15] Further support for the argument made above that the ditransitive recipient is a secondary topic comes from the fact that ditransitive recipients show a marked preference for 1st- and 2nd-person pronouns when compared with recipients expressed pronominally in the prepositional dative (57% for the former as compared with 20% in the latter in a study of the entire parsed Switchboard corpus) (Bresnan and Nikitina MS).

(100) A girl was given a book by everyone. A girl > everyone (preferred)
Everyone > a book (preferred)

Moreover, as observed in Section 2.1, the recipient argument of the ditransitive strongly tends to have wide scope over the theme argument, as in (101):

(101) The teacher assigned one student every problem. One > every
(Bruening 2001; Ioup 1975; Larson 1990: 604)

The goal argument of the prepositional phrase, on the other hand, need *not* take wide scope over the theme argument, and the following sentence is ambiguous between a reading in which one particular problem was given to all the students and one in which all the students received one problem:

(102) The teacher assigned one problem to every student. (ambiguous)

Since the recipient argument is more topical than the theme argument in the ditransitive, we predict the scope facts evident in (101). Since the two non-subject arguments of the prepositional paraphrase are not specified for relative topicality, we also predict that sentences such as (102) will be ambiguous with respect to scope assignment.

We saw above that brand-new recipient arguments do not necessarily strike native speakers as ungrammatical. Therefore, the quantifier hierarchy predicts that it should be possible to invert the most prevalent scope interpretation of recipient and theme if the right quantifiers are chosen. This is also the case:

(103) They gave at least two people each Nobel prize. Each > a least two
(preferred)

Example (103) is naturally interpreted as meaning that each Nobel prize was awarded to at least two recipients: the theme argument taking wide scope over the recipient argument.

Similarly, as expected, the recipient argument can be interpreted as having wide scope over the subject argument when particular quantifiers are chosen. The phrase "some x or another" encourages a narrow scope reading for x, even in positions that normally have wide scope such as subject position:

(104) Some person or other gave each child a book. Each > some or other
(preferred)

Thus the scope properties of the ditransitive construction can be seen to follow naturally from the combination of lexical facts and a topicality scale. Proposals that rely on the structural description of ditransitives for an explanation of scope effects (e.g. Bruening 2001) do not predict the acceptability of the inverse scope in (103) and (104).

To summarize, then, because scope is strongly correlated with topicality, **the information-structure properties of constructions predict their predominant assignment of scope.** Since lexical content also plays a role in scope assignment, it is to be expected that the right choice of lexical elements allows for occasional inverse scope assignment as well.

7.13 Conclusion

In this chapter, we have extended existing accounts that relate apparently syntactic phenomena to the domain of information structure, in order to account for a range of seemingly mysterious facts. By proposing an elaborated version of the insight that unbounded dependencies cannot involve backgrounded elements (Erteschik-Shir 1979; Takami 1989; Van Valin 1998), we have addressed the majority of standard constraints on movement. The discussion surrounding the ditransitive construction provides evidence that a close analysis of the particular constructions involved yields answers to otherwise recalcitrant problems. In particular, we have been able to explain the recipient argument's statistical tendency to resist being involved in unbounded dependencies unless it is the passive subject.

We have also built upon previous observations (particularly that of Ioup, Kuno, and Francis and Michaelis) to offer a general account of how a wide-scope interpretation correlates with topicality. We have demonstrated that an information-structure analysis of the ditransitive predicts the fact that the recipient argument tends to have wide scope over the theme argument (and the fact that it doesn't when it is passivized or when particular lexical items are chosen).

Both explanations benefited from a careful analysis of lexical semantics and the particular constructions involved. Which arguments are backgrounded or topical depends critically on the choice of constructions. Not all sentential complements serve the same function (compare the complement of *say* with the complement of *shout*); neither do all postverbal NPs (compare the direct objects of transitive verbs which are within the potential focus domain with the recipient arguments of ditransitives which are backgrounded). Determining which arguments are topical likewise depends on which constructions one is considering (compare the verbs of creation with the verbs of transformation). FUNCTIONAL EXPLANATIONS REQUIRE REFERENCE TO THE FUNCTION OF THE CONSTRUCTIONS INVOLVED (including the lexical semantics of the words involved).

While the field of information structure is complex and it requires recognizing statistical regularities, a case has been made that information structure and processing are absolutely central to the investigation of issues that lie at the center of linguistic theorizing.

Appendix: Corpus evidence for the proposed information structure of the ditransitive

A replication of earlier findings on the strong tendency for the recipient of ditransitives to be non-new was found in the speech of three mothers, to their children (Nina, Adam, and Eve) in the CHILDES database:

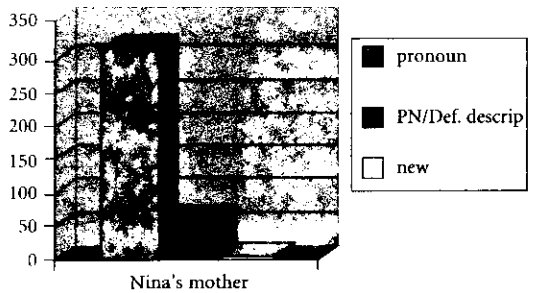

Number of pronouns (given), proper names/definite descriptions (accessible), and indefinite (new) arguments in active recipient position (n = 381) in Nina's mother's speech

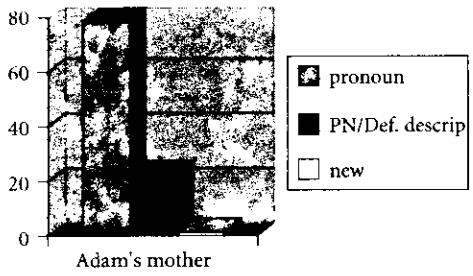

Number of pronouns (given), proper names/definite descriptions (accessible), and indefinite (new) arguments in active recipient position (n = 103) in Adam's mother's speech

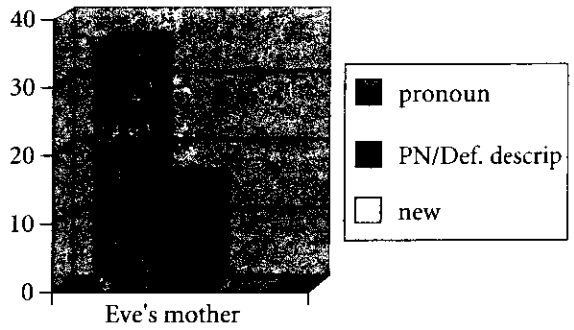

Number of pronouns (given), proper names/definite descriptions (accessible), and indefinite (new) arguments in active recipient position (n = 52) in Eve's mother's speech

The trends are the same across mothers. That is, the distribution of NP types across the three mothers is not significantly different: $\chi^2(4\ N = 536) = 8.9$, $p = .06$.

Data from Alta Vista is again similar, with the ditransitive recipient showing even more of a tendency to be pronominal than the ditransitive actor argument. The theme argument of the ditransitive, on the other hand, overwhelmingly appears as a lexical NP.

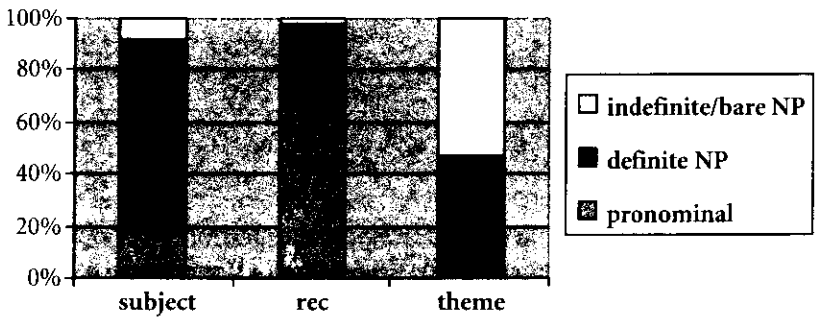

Percentage of various nominal marking in each argument of active ditransitives, found in the first 110 instances involving the verb *give*, gathered on Alta Vista using the Linguists' Search Engine. The distribution of NP types is distinct in subjects and recipients: $\chi^2(2\ N = 211) = 49$, $p < .001$; recipients and themes $\chi^2(2\ N = 181) = 51.5$, $p < .001$; and themes and recipients $\chi^2(2\ N = 190) = 106$, $p < .001$. ($N < 220$ because not all instances of *give* expressed each argument).

The fact that the prepositional goal argument is not constrained to be given is evident from the frequency of proper names and definite descriptions among the corpora of each mother. The data are given below.

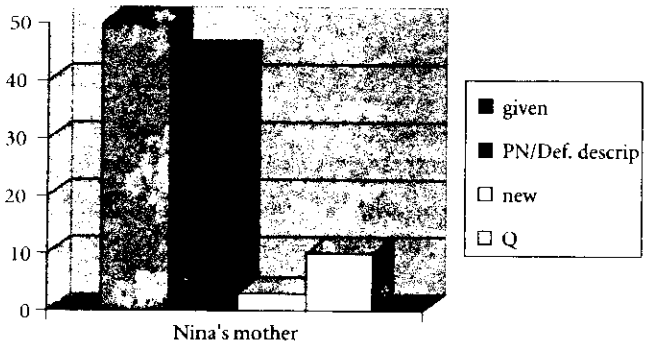

Number of pronouns, definite descriptions, question words, and new arguments of Nina's mother's prepositional goal argument (n = 107)

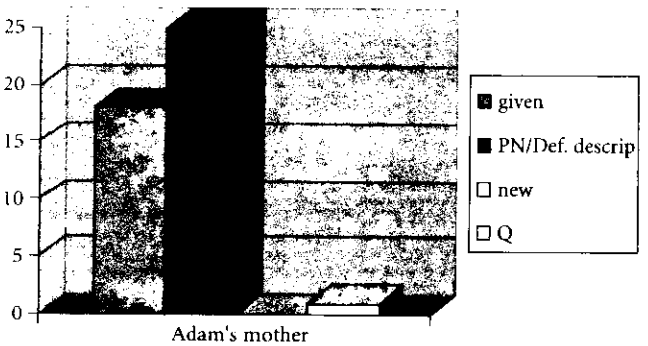

Number of pronouns, definite descriptions, question words, and new arguments of Adam's mother's prepositional goal argument (n = 44)

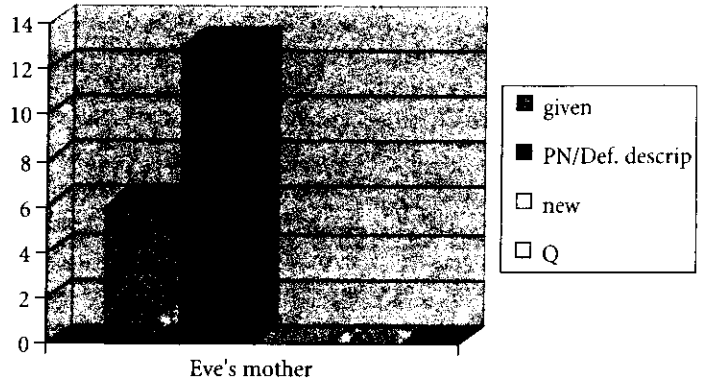

Number of pronouns, definite descriptions, question words, and new arguments of Eve's mother's prepositional goal argument (n = 19)

The trends are again the same across mothers.[16] A comparison of the frequencies of NP types appearing as ditransitive recipient and prepositional goal indicates a significant difference: $\chi^2(3\ N = 706) = 106.6$, $p < .001$.

[16] The distribution of NP types across the three mothers is not significantly different: $\chi^2(6\ N = 170) = 9.9$, $p = .13$.

8

Grammatical categorization: Subject–Auxiliary Inversion[1]

8.1 Introduction

A recognition of the general nature of categories can be used to illuminate the functional underpinnings of aspects of grammar that appear at first to be brute syntactic facts. This chapter investigates one such case, that of Subject–Auxiliary Inversion (SAI). SAI is found in a variety of utterances including yes/no questions, non-subject wh-questions, counterfactual conditionals, sentences with initial negative adverbs, exclamatives, comparatives, negative conjuncts, and positive rejoinders:

(1)	Did she go?	Y/N questions
	Where did she go?	(Non-subject) wh-questions
(2)	Had she gone, they would be here by now.	Counterfactual conditionals
(3)	Seldom had she gone there...	Initial negative adverbs
(4)	May a million fleas infest his armpits!	Wishes/Curses
(5)	Boy did she go!	Exclamatives
(6)	He was faster at it than was she.	Comparatives
(7)	Neither do they vote.	Negative conjunct
(8)	So does he.	Positive rejoinder

These, or a subset of these widely varying contexts have been cited as evidence that there is no functional generalization associated with the syntactic pattern (Green 1985); SAI has been used in this context as evidence for the existence of purely syntactic generalizations or autonomous syntax

[1] This chapter is based on a jointly authored paper with my undergraduate student at UIUC, Alex Del Giudice (forthcoming), in *The Linguistic Review*. I wish to thank Bill Croft, Seizi Iwata, and Kunie Miyaura for very helpful comments on an earlier draft.

(Newmeyer 2000).[2] SAI therefore provides a strong challenge to the idea that formal patterns are generally associated with and motivated by functions, a claim that distinguishes cognitive and functional approaches from otherwise parallel generative approaches (e.g. Jackendoff 2002).

This chapter builds on previous accounts that have offered insightful generalizations about the function of SAI constructions (Lakoff and Brugman 1987; Lambrecht 1994; Michaelis and Lambrecht 1996a; Diessel 1997). In addition, it is argued that a case that is apparently somewhat exceptional, that of exclamatives (e.g. (5)), is actually strongly motivated as well. Apparently formal restrictions on SAI, are in fact demonstrated to be better accounted for by a functional account. Finally, the particular internal syntactic form of SAI is motivated by appeal to its semantic/pragmatic function. The implication of this work is that synchronic functional motivations often lurk behind seemingly syntactic brute facts and can be used to explain many aspects of grammar that appear otherwise to be wholly idiosyncratic.[3]

The relevant functional categories may only be found if researchers are aware of insights gained from cognitive psychology. In particular, the claim that SAI bears no functional generalization is based on a false assumption about what such a generalization would look like. The implicit assumption is that there should be a single feature or set of features common to the category, and yet, as discussed below, this assumption is widely recognized to be false. As Murphy notes, "The groundbreaking work of Eleanor Rosch in the 1970s essentially killed the classical view [of necessary and sufficient definitions], so that it is not now a theory of any actual researcher in this area" (Murphy 2002: 16).

It is argued below that the set of constructions that exhibit SAI naturally form a coherent functional category of the type familiar from general categorization and lexical concepts.

[2] Fillmore (1998) takes a similar stand, providing a detailed discussion of the formal properties of SAI and its various specific subconstructions, but offering no semantic or pragmatic properties to relate the subtypes. He is explicitly agnostic on the question of whether such properties exist (note 11, p. 121) and in fact goes so far as to state that many basic constructions are not pairings of form with function, but are simply formal generalizations. Fillmore's position is somewhat surprising given his central role in the development of Construction Grammar, an approach that is explicitly designed to deal with pairings of form with function.

[3] It should be made clear that the failure to find a unified functional category for SAI or some other pattern would only indicate that certain brute syntactic facts exist that are not synchronically motivated. While I do accept the existence of occasional instances of synchronically unmotivated syntactic facts (normally motivated by diachronic developments), these appear to be the exception rather than the rule.

Newmeyer (2000) claims that SAI is a prime example of a formal generalization with no functional underpinnings. That is, SAI has been claimed to display unified syntactic properties, but not to be associated with any unified function. For example, Newmeyer suggests that no instances of SAI, regardless of type, appear in embedded questions or subordinate clauses:

(9) a. I wondered whether you had been working late.
b. *I wondered whether had you been working late.

This is claimed to be a *formal idiosyncrasy*—since the same restrictions appear to hold of each separate use of the structure despite the fact that they are associated with a range of functions.

The present chapter argues on the contrary that there is a coherent category of functions associated with SAI, and moreover that there are no unified syntactic properties, with the exception, of course, of the surface form of SAI itself, which I will argue is motivated by its function. In this way, the present chapter makes a case study of SAI, the point being that a focus on form to the neglect of function is like investigating a human organ such as the liver, without attending to what the liver does: while this is not impossible, it is certain to fail to be explanatory.

8.2 A semantic/pragmatic generalization

Several semantic/pragmatic generalizations have been proposed in the literature. Many SAI constructions are framed negatively or at least non-positively (Lakoff and Brugman 1987). Many SAI constructions in both German and English are non-declarative or non-assertive speech acts (Lambrecht 1994; Michaelis and Lambrecht 1996a; Diessel 1997). Finally, many SAI constructions are often narrow-focus or sentence-focus constructions—that is, they often do *not* have more typical, predicate focus (i.e. topic-comment) information structure (Michaelis and Lambrecht, 1996a).

All of these accounts implicitly or explicitly note that SAI is a deviation from the prototypical sentence form in one way or another, the prototypical sentence being a declarative, positive assertion with predicate-focus information structure. Another feature of prototypical sentences is that they may stand alone and they are not dependent on another clause. We will see that many SAI constructions differ from prototypical sentences in this respect as well.

At the same time, no SAI construction simultaneously has all the attributes of non-prototypical sentences; each SAI construction has its own subset of the attributes. While this would seem unhappily to lead to too many disjunctive characterizations, such a situation is the norm for lexical items and for categorization generally: a subset of attributes that hold of the prototypical case is instantiated in each conventionalized extension (e.g. Lakoff 1987). That is, as we saw in Chapter 3, learners retain quite a bit of item-specific linguistic information at the same time that they form generalizations. As discussed in Chapter 4, instances with high type frequency provide a cognitive anchor or prototype that helps the learner assimilate new cases to the category.

For example, consider the word *home*: prototypically, it evokes a house where an intact family unit lives and sleeps and where children grow into adulthood, and where one feels comfortable and a sense of belonging. Yet it can be used for many of these aspects in isolation from the others (Fillmore 1992):

(10) a. Prototype: 138 FitzRandoph Rd. is Aliza's and Zach's home. (house, where they grow up, live with their family, feel comfortable, and belong)
b. She owns 14 homes. (houses)
c. She went home to her dorm room. (place where someone lives and sleeps)
d. She traveled home to see her family. (place where she grew up)
e. She's at home in the mountains. (place where one feels a sense of belonging)

Home clearly has a prototypical meaning that can be extended to highlight a particular attribute (or attributes) of the prototypical meaning.

The word *baby* is another example. Prototypically, *baby* refers to a small human, who is cute, needs to be taken care of, is immature emotionally, and is the youngest in the family. Any single one of these attributes can be used as the reference of *baby* in its various senses, however:

(11) a. Prototype: She had a baby (small, human, cute, immature, needs to be taken care of, youngest in a family)
b. Baby carrots (small)
c. Hey, baby! (cute)
d. Don't baby me. (to take care of as a parent would a baby)
e. He's such a baby. (emotionally immature)
f. Mr Platt is the baby in his family (youngest in the family)

Are we to say that the words *home* and *baby* have no meaning? Or that their various senses are unrelated to one another? To assume so would be to claim

that it is an accident that all the senses are referred to by the same morpheme. However there exists evidence from processing (Klein and Murphy 2002; Pylkkanen and Murphy 2004) and from language acquisition (Casenhiser 2004) that systematic relationships among senses of a polysemous word are recognized by speakers. It is therefore preferable to view the prototypical sense as extended such that one or more particular attributes are highlighted in each extension.

A natural category relating the various conventionalized uses of "baby" is represented in Fig. 8.1.

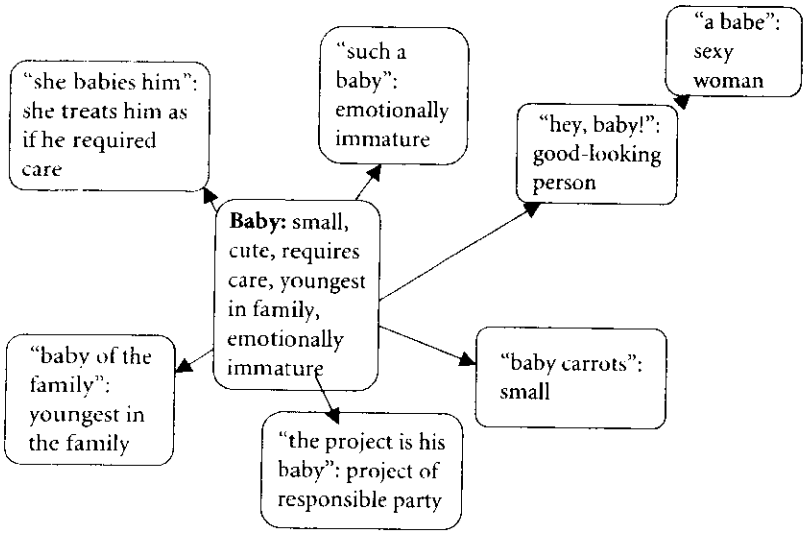

FIGURE 8.1 Polysemous senses of the lexical item "baby"

Each of the links from the prototype can be interpreted as indicating a metonymic relationship whereby an attribute of the prototype is referred to by the same word as the prototype itself.

The function of SAI can be analyzed in similar terms: there are attributes that hold of the prototypical case, and conventional extensions of the prototype systematically differ from the prototype in displaying only a subset of the relevant attributes. On the basis of the distribution of SAI constructions, it is further suggested that the dominant attribute of SAI is NON-POSITIVE; this attribute of SAI constructions serves to motivate the form of the construction.

Questions

Yes/no questions are clearly non-declarative speech acts whose propositional content is non-positive (it does not assert or presuppose the truth of the proposition). For example, (12) does not assert or presuppose that he left. Wh-questions such as (13) do presuppose that he did something, but are non-declaratives.

(12) Did he go?

(13) What did he do?

The fact that (non-echo) *subject* wh-questions do not invert (14) is attributable to the fact that the only way to invert the subject and auxiliary would be to position the auxiliary before the wh-word subject. Yet the wh-question construction requires the wh-word to appear sentence-initially.[4]

(14) *Did who leave?

Counterfactual conditionals

SAI is required in counterfactual conditionals, as in (15a):

(15) a. Had he found a solution, he would take time off and relax.
b. *He had found a solution, he would take time off and relax.

This construction presupposes (not asserts) a *negative* word-to-world fit—that is, the hypothetical antecedent is presupposed to designate a non-actual state of affairs in the world. Thus this construction involves non-positive, non-asserted propositional content. The inverted clause is also not a stand-alone sentence, but instead is dependent on the following clause. Thus counterfactual conditionals have three attributes of non-prototypical sentences: they are non-positive, non-asserted, and dependent.

Initial Negative Adverbs

SAI can be seen to be obligatory with initial adverbs that are construed negatively, as demonstrated in (16) (see also Green 1985; Lakoff and Brugman 1987):

(16) a. Not until yesterday did he take a break.
b. *Not until yesterday he did take a break.

[4] In a default inheritance hierarchy making use of a usage-based model of grammar, these facts are naturally accounted for since the non-subject question construction, as expected, inherits from both the general wh-construction and the SAI construction. The conflict between the two constructions is accounted for by the non-subject wh-question construction, which specifies that its order is [Wh_{subj}] [VP].

At the same time it may not appear with a positively framed adverb such as in (17 a-b):

(17) a. *Everywhere has he found a solution.
b. *Yesterday did he take a break.

The negative implication conveyed by SAI can be seen by comparing (18a) and (18b). Example (18a) implies that even with money offered as incentive she would not quit, while (18b) expresses that she would quit with the slightest incentive (Jackendoff 1972; Lakoff and Brugman 1987; Newmeyer 2000):

(18) a. For no money would she leave. (she wouldn't)
b. For no money she would leave. (she would)

Example (19) can be used to demonstrate that what is important is not the factual word-to-world fit in determining whether SAI is appropriate, but instead whether the propositional content is *construed* positively or not. That is, (19a) is logically equivalent to (19b), and yet (19b), since it is framed positively, cannot be phrased with SAI (19c):

(19) a. Not until yesterday did he take a break. (SAI: framed non-positively)
b. He took a break yesterday and not before. (positively framed equivalent proposition)
c. *Did he take a break yesterday and not before. (SAI with positively framed proposition)

May you live to be 100!

SAI also appears in utterances like (20).

(20) May a million flies infest his armpits!

This special type of statement is an appeal to some unspecified forces about some uncontrollable event. Thus this pattern, like other uses of SAI we have seen, is not a declarative speech act, but rather a curse or wish (an "expressive" speech act in Searle's (1979) terminology). Moreover, no positive state of affairs is asserted or presupposed: whether or not the curse or wish will be fulfilled is left open.

The use and function of statements such as that in example (20) are quite limited. The only auxiliary that can be placed in this construction is "may":

(21) a. May you live a good life!
b. *Should you live a good life!

In addition, this construction differs from other uses of SAI in that it doesn't provide a negative polarity context:

(22) *May you ever lift a finger!

Thus this case must be listed as a special subconstruction (see also Fillmore 1998). No general syntactic account would predict these idiosyncrasies without stipulation. At the same time, the construction obeys the generalization that utterances with SAI are non-declarative and non-positive. Thus its existence is motivated.

Exclamatives: making sense of a recalcitrant case

The exclamative construction seems to pose the greatest threat to a functional account of SAI, since instances of this construction can be positively framed with topic-comment information structure. It has been observed that exclamatives are at least non-assertive in that they presuppose, rather than assert the truth of the propositional content, and in that way have one feature associated with SAI (Diessel 1997; Michaelis 2001).

It is possible to motivate further the use of SAI in the exclamative construction by observing that exclamatives are closely related to a subtype of questions: rhetorical questions. This is not the received wisdom. While Lambrecht (1994), for example, suggests a functional restriction on SAI, namely that utterances with SAI are not declaratives, he nonetheless suggests there is no direct relationship between questions and exclamatives (p. 30). Fillmore (1999) likewise states without qualification that exclamatives "are clearly not questions" (p. 122). This chapter argues that it is possible to relate exclamatives to questions in a quite strong, direct way, therefore more strongly motivating the fact that exclamatives, like questions, are expressed with SAI. Exclamatives can be seen to be closely related to rhetorical questions: that they should have the form of questions is therefore unremarkable.

Cross-linguistically, exclamatives and questions often bear formal similarities (Zanuttini and Portner 2003). Note that many exclamative utterances are at the same time rhetorical questions, as indicated by the phrase *or what* which is generally restricted to appear with rhetorical yes/no questions:

(23) a. Do you want to go? (request for information)
b. Do you want to go or what? (rhetorical question: it is assumed that you want to go)

(24) Boy, are you tired or what?!

Moreover, there is evidence that speakers sychronically perceive a relationship between exclamatives and questions. Alex Del Giudice collected examples of SAI using the Google Search Engine, searching for the following randomly selected four instances of <expletive aux pronoun> as exact phrases (that may or may not continue)(Goldberg and Del Giudice forthcoming):

(25) a. "boy is this"
b. "boy are you"
c. "wow does this"
d. "man does this"

Irrelevant cases such as instances of the string that crossed sentence boundaries (e.g. *First Boy: Is he doing that on purpose?*) were ignored. The first fifty instances of exclamatives returned by Google for each of the four strings were combined into a master list of 200 attested tokens. A full 13 per cent (26/200) of the exclamative sentences included the specific phrase *or what*. A few examples of these are given below. All spelling and punctuation have been preserved:

(26) a. "Boy, is this the summer of angst-ridden Australian tough guys, or what?"
b. Boy is this a bit more involved than I expected, or what?!
c. Boy, is this an awesome picture or what?!?
d. Hoo boy, are you dropped into a pot of boiling water or what.
e. "Wow! Does this guy own stock in Microsoft or what?"

That this many examples should bear the specific lexical phrase *or what* is quite striking and can be taken as evidence that speakers relate exclamatives to rhetorical questions. In fact positively framed exclamatives in general are compatible with the phrase *or what*, just as positively framed rhetorical questions are, indicating that an even greater number of instances may have been construed as rhetorical questions, even though they were not explicitly marked as such.[5]

Thus exclamatives are actually strongly motivated by the fact that they are a minimal extension from rhetorical questions. Rhetorical questions are well suited to being used as exclamatives because rhetorical questions are by

[5] Notice also that exclamatives that do not have subj–auxiliary inversion are not compatible with *or what*?

(i) What a big boy he is! (exclamative, no SAI)
(ii) *What a big boy he is or what?!

definition questions that are posed that require no response because the answer is already clear. Exclamatives likewise are used to express states of affairs that are clearly true (and remarkable).

Comparatives

SAI optionally occurs in comparatives, in a formal or archaic register:

(27) a. He has read more articles than have his classmates.
b. He has read more articles than his classmates have.

Unless the subject is a full lexical NP, inversion sounds decidedly archaic, and it is downright unacceptable with pronominal *it* as subject:

(28) a. (archaic) He has read more articles than have they.
b. *Our library has more articles than has it.

The preference for lexical NP subjects in comparatives follows from the fact that comparatives with SAI require a narrow focus on the subject argument, instead of the more typical topic-comment interpretation.

The use of SAI in comparatives may be motivated independently by a tendency to position heavier or longer constituents clause finally (see, e.g., Hawkins 1994; Wasow 2002). In support of this idea is the fact that SAI is only used in comparatives when there is VP ellipsis in the inverted clause, thus allowing the lexical subject to appear clause-finally:

(29) a. *He has read more articles than have his classmates read.
(cf. He has read more articles than his classmates have read.)
b. He has read more articles than have his classmates.

The inverted clause in comparatives is dependent on the preceding clause and cannot stand alone as an independent sentence. Thus SAI in comparatives is used to mark dependent clauses that have narrow focus on the subject argument: two characteristics of non-prototypical sentences.

Nor is this one unmotivated

Sentences beginning with *neither* or *nor* require SAI:

(30) a. Neither is this construction unexpected.
b. *Neither this construction is unexpected.

Such sentences are negatively framed. They are also pragmatically dependent on another contextually given proposition, insofar as *neither* and *nor* are conjunctions and thus require reference to a second conjunct. Thus these

cases have two attributes of non-prototypical sentences: they are non-positive and they are dependent.

So is this one

A final use of subject–auxiliary inversion to be discussed here is found immediately after initial conjunct "so/as/likewise," as in:

(31) Context: "His girlfriend was worried."
a. So was I.
b. *So I was.

These are distinct from the negative conjuncts just discussed in that they must involve VP ellipsis, as in:

(32) *So was I worried.

These cases are positively framed assertions. But they have two features associated with prototypical SAI: they have narrow focus on the subject argument as opposed to topic-comment information structure, and they are pragmatically dependent on another evoked proposition.

The majority of constructions that license SAI have at least two features of non-prototypical sentences. The case of exclamatives is additionally motivated by its close relationship with rhetorical questions, and wh-questions are also motivated by their close relationship to yes/no questions as well.

A functionally motivated category of SAI is diagrammed in Fig. 8.2. Independent evidence for the naturalness of this category comes from the fact that several of the most central functions have been found to cluster together in cross-linguistic work on an entirely different construction: indefinite pronouns (Haspelmath 1997). Haspelmath finds a tendency for indefinite pronouns to be expressed alike in questions, conditionals, comparatives, and indirect negations. In the case of SAI constructions, the following are roughly analogous: questions, counterfactual conditionals, comparatives, initial negative adverbs, and negative rejoinders: i.e. extensions A,B,C,D,E,G, as indicated in Fig. 8.3.

The SAI category proposed and the category for the senses of *baby* described above are quite similar in having conventionalized extensions radiating out from a central core. One important difference between them, however, is that the prototype suggested for SAI: "non-prototypical sentence," is actually a generalization that is not directly instantiated (as indicated by the dashed circle in Figure 8.2.) That is, while we frequently encounter prototypical sentences, we do not encounter "non-prototypical sentences" as instances

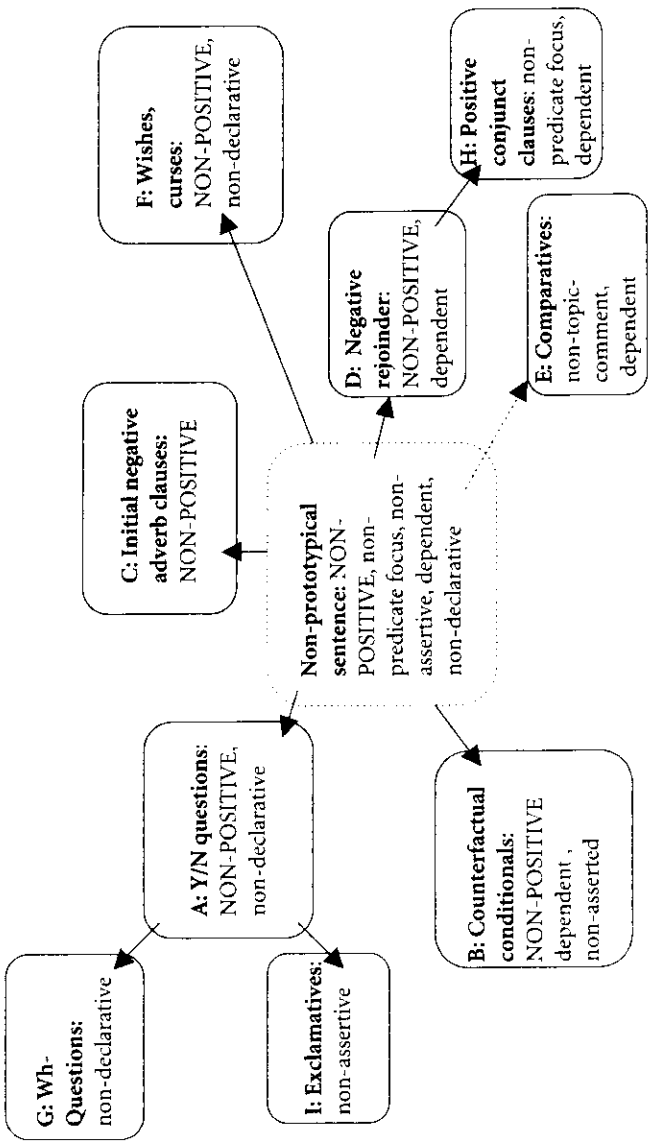

FIGURE 8.2 Functional category of SAI constructions with "non-prototypical sentence" as its prototype

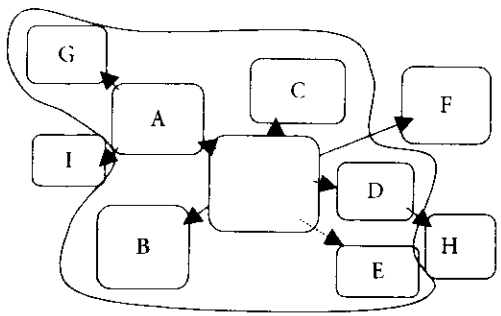

FIGURE 8.3 SAI functions that cluster together in indefinite pronouns

of a non-prototypical sentence category. Therefore it is possible to reconstrue the category of SAI as a halo of constructions that stand in *contrast* to prototypical sentences. The systematic difference in form (subject–auxiliary inversion) signals a systematic difference in function (a distinction from prototypical sentences). This analysis recognizes an additional link-type to those suggested in Goldberg (1995): a markedness link, indicated by a curved link in Fig. 8.4.

8.3 Motivating the form of the construction

The attribute of being non-positive is the dominant feature of SAI constructions. This is evident in Fig. 8.2 and 8.4 insofar as more subconstructions have this attribute than any other single attribute. Moreover, each of the constructions that do not share this attribute can be seen to be immediately motivated by a construction that does, with the exception of comparatives. SAI in comparatives is then by hypothesis the least motivated, and it turns out, is also the least stable. As noted above, it is strictly optional and slightly archaic sounding. It may well drop out of use entirely in the coming decades.

Why should non-positive contexts be indicated by an inversion of subject and auxiliary instead of, say, by placing the subject in sentence-final position or some other arbitrary feature? Newmeyer (2000) suggests that there is no motivation for its particular form, and notes that no one has previously suggested one. Moreover, he further suggests that unless we are able to motivate the particular form that SAI takes on the basis of some proposed function, then this would provide additional evidence for the thesis of autonomous syntax (2000: 46). Let us take up this challenge.

Auxiliaries carry information about polarity as well as tense and aspect (Langacker 1991). By positioning the auxiliary in a non-canonical position,

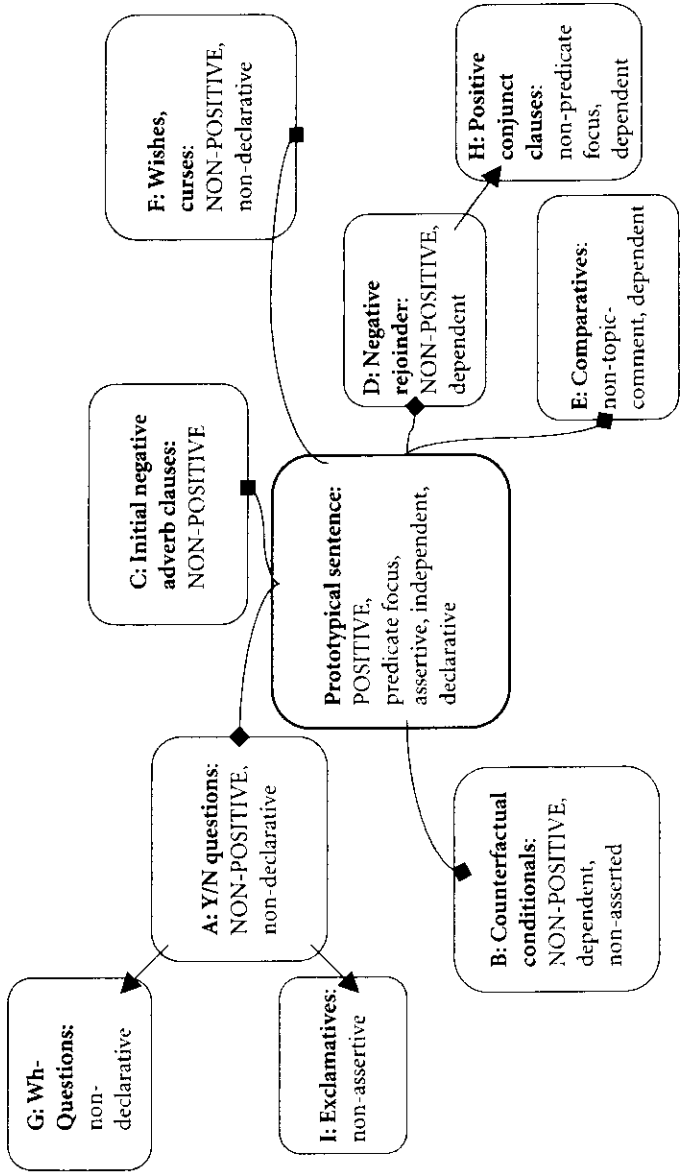

FIGURE 8.4 Functional category of SAI constructions with prototypical sentence as its prototype and markedness links motivating each of the extensions from the prototype

the construction conveys that the polarity involved is not the canonical, positive polarity; i.e. no positive word-to-world fit is asserted. This motivation can be used to account for the fact that only the first auxiliary is inverted (Newmeyer 2000: 48):

(33) a. Have you been working late?
b. *Have been you working late?
c. *Could have been you working late?

Only the first auxiliary serves to indicate polarity of a sentence; the broad-scope negative, for example, can only appear after the first auxiliary:

(34) a. You couldn't have been working late.
b. You could have not been working late.
c. You could have been not working late.

Examples (34b,c) are acceptable only under the narrow-scope reading in which *not* predicates the word or phrase immediately following, and not the entire sentence: that is, only (34a) is paraphrasable by "It's not the case that you could have been working late." Notice that *not* can only cliticize on the first auxiliary (as in 34a):

(35) a. *You have been't working late.
b. *You could haven't been working late.

Putative Main Clause Restriction

Is SAI actually uniformly restricted to main clauses, as has been suggested (Emonds 1970; Newmeyer 2000)? It turns out it is not (Hooper and Thompson 1973; Green 1976; Bolinger 1977b), as the examples in (36) illustrate:

(36) a. They knew that *had she* left on time, she'd be here by now. (counterfactual conditional)
b. She reflected to herself that *never had she* seen such a beautiful sight. (clause-initial negative adverb)
c. Junie B. knew that *boy, was she* in trouble! (exclamative)
d. I knew that they would spend millions on defense, but I knew equally that *not one cent would* they give for tribute. (Bolinger 1977b: 515) (Negative NP preposing)

Certain uses of SAI are restricted to convey particular non-declarative speech acts, including, in particular, questions and wishes/curses. Even these SAI constructions can appear in subordinate clauses, but only when the main speech act is conveyed by the subordinate clause, as in examples (37a,b) (from Green 1976: 385):

(37) a. I guess John didn't come in, did he? (tag question)
b. We ought to assign Postal, because who can understand *Aspects?* (rhetorical question)

There are many subtleties related to which subordinate clauses can appear with SAI, but the restrictions have been widely argued to be pragmatic, not syntactic. (For more specific formulations of where putative main clause phenomena can occur, see Hooper and Thompson 1973; Green 1976; Bolinger 1977a.)[6] As a first approximation, the following generalization holds: (**Only**) **SAIs that are restricted to conveying particular speech acts are restricted to main clauses, or to subordinate clauses that convey speech acts.**

Cross-linguistic rarity

The use of subject–auxiliary inversion to indicate a non-prototypical sentence is rare cross-linguistically. This fact raises a challenge to an autonomous syntax account, at least as traditionally construed by generative grammarians: if SAI were based on a universal system of syntactic knowledge, one should presumably expect the generalizations to be universal or at least common cross-linguistically. On the present account, on the other hand, SAI is a motivated device for indicating deviations from prototypical sentences, particularly for non-positive expressions. It is certainly not the only possible device: overt negatives, discourse particles, and other special constructions could do the job equally well. Once we recognize that SAI and its range of conventional uses are learned on the basis of positive input by each new generation of speakers, there is no expectation that it should necessarily be common across languages.

8.4 Conclusion

Given parallel evidence for radial categories in the concepts designated by words (*home, baby*) and grammatical categories, and the convergent understanding of how human categories are formed generally, the type of category of functions exhibited by SAI can be seen to be quite natural.

The implication of this chapter is that functional motivations often underlie seemingly idiosyncratic facts of grammar. The distribution of SAI in

[6] Hooper and Thompson suggest that SAI can appear in subordinate clauses that are asserted as opposed to presupposed. Green similarly argues that SAI can appear in subordinate clauses in which the speaker appears to agree with the proposition expressed by the subordinate clause. Bolinger observes additional idiosyncratic pragmatic facts about the appearance of SAI constructions in subordinate clauses.

subordinate clauses and its restriction to the first auxiliary, are explained by attention to the functions that SAI conveys. In contrast, a purely syntactic account of the phenomenon can only stipulate the form of SAI—that subject and first auxiliary are inverted—without making any accurate further predictions or generalizations. The formal account would moreover still have to stipulate a list of all of the subconstructions that require or allow SAI (and whether each requires or allows SAI), as well as each of the constructions' other special properties, without motivating the existence of SAI in English in any way. It is fair to say that while the formal approach may be descriptively adequate, it does not, in the case of SAI, have any explanatory force.

In seeking out functional categories we need to be cognizant of the sort of categories we should expect to find. Categories of language, like most human categories, are much more flexible than those defined rigidly by necessary and sufficient conditions.

9

Cross-linguistic generalizations in argument realization

In Chapters 4–6, it was argued that argument structure constructions are learnable, but I acknowledged that we needed an account of the cross-linguistic tendencies that exist. If the mappings between form and meaning were universal and *not* attributable to general cognitive mechanisms, then one could legitimately argue that while they are learnable, they are not learned.

Claims that linking rules are universal are widespread (Baker 1988; Pinker 1989; Gleitman 1994; Levin and Rappaport Hovav 1995). The implication has been that universal aspects of language are innate, proposed specifically to solve the apparent "poverty of the stimulus" problem (Chomsky 1957, 1965). Pinker (1989: 248), for example, suggests that "Linking rules... seem to be near-universal in their essential aspects and therefore may not be learned at all.... Linking rules can be universal and innate in the current theory..." He offers a very explicit proposal for the mappings from semantic roles to surface syntactic positions, as follows (Pinker 1989: 74):

1. Link the agent to SUBJECT
2. Link the patient to OBJECT
3. Link the theme argument (first argument of BE or GO) to SUBJECT unless SUBJECT is already linked; to OBJECT otherwise
4. Link the goal to an OBLIQUE (prepositional phrase) argument
5. Link the theme argument in a CAUSE TO HAVE predicate to the second object in a ditransitive construction

Naigles, Gleitman, and Gleitman (1993: 136–7) suggest that "there is sufficient cross-linguistic similarity in these linking rules to get the learning procedure started.... [T]here is an overwhelming tendency, cross-linguistically, for agents to appear as subjects and themes as direct objects, with other arguments appearing in oblique cases." In the same vein, Gleitman (1994: 203) suggests that "at least some of the mapping rules [between syntax and semantics] have to be in place before the verb meanings are known, or else the whole game is over."