Cognitive Mechanisms & Syntactic Theory: Arguments against Adjuncts in the Lexicon

Julie E. Boland
University of Michigan

As a psychologist who studies sentence comprehension and holds a joint appointment in the departments of psychology and linguistics, I have frequent opportunities to observe the interaction, or lack thereof, between the two disciplines. Although cognitive psychology and formal linguistic theory share some common history in Chomsky's (1959) pivotal review of Skinner's (1957) book on language behavior, these two disciplines have not continued to influence one another to the degree one might expect. For example, theoretical developments in syntax have rarely if ever been motivated by an experimental finding about sentence comprehension. In fact, there is good reason for this. I will argue that most psycholinguistic data is irrelevant to formal linguistic theory. Nonetheless, there may be a subset of psycholinguistic data that formal linguists ought to consider. I will attempt to delineate this largely hypothetical subset, using as an example some actual psycholinguistic data on the argument/adjunct distinction. First, let us consider the domains of cognitive psychology and formal linguistic theory.

Cognitive psychology encompasses the study of mental representations and the mental operations for manipulating (creating, accessing, etc.) these mental representations. The central goal is the description of a processing system. In the context of sentence comprehension, for example, a cognitive psychologist might develop a theory of syntactic parsing within sentence comprehension that specifies what types of mental representations are involved, what aspects of linguistic and non-linguistic knowledge are used to create those mental representations, and so forth. In this chapter, I will consider experimentally-oriented psycholinguists like myself to be cognitive psychologists, regardless of their academic credentials.

Formal linguistic theory encompasses the study of what speakers know about their language. The central goal is the development of the simplest and most elegant description of this knowledge (i.e. grammar). By placing a premium on elegance, linguists hope to characterize the properties of the human mind that make language possible and not merely catalogue facts about languages. In addition to elegance, grammars must achieve descriptive adequacy by generating all and only those utterances accepted by native speakers. Descriptive adequacy is evaluated by analyzing human behavior, i.e., linguistic intuitions. However, descriptive adequacy receives somewhat less attention, perhaps because it is difficult to thoroughly evaluate the descriptive adequacy of a grammar for a particular language and completely impossible to fully evaluate its descriptive adequacy for all existing and possible languages—although this is understood to be the ultimate goal. Descriptive adequacy is assessed via the collection and analysis of linguistic intuitions. While this process might be considered a type of psychological study, it is quite unlike the typical experiments carried about by a psycholinguist or some other type of cognitive psychologist.

Formal linguistic theory offers cognitive psychologists a framework for partitioning the subcomponents of language processing (Phonetics, phonology, morphology, syntax, etc.), a vocabulary for describing the mental representations involved in language processing (noun

phrase, empty category, thematic role, etc.), and theories about how the representations are structured (e.g., X-bar theory, Categorial Grammar, Head-driven Phrase Structure Grammar, etc.). In turn, cognitive psychology provides linguistics with methodologies for investigating the mental processes of acquisition, comprehension, and production. Such investigations are clearly of some interest to linguists, in general, but such investigations are relatively unlikely to influence formal linguistic theory for reasons that will be discussed below.

At the level of syntax, the distinction between cognitive psychology and formal linguistics is echoed in the constructs of the human parsing system on one hand and the grammar on the other. In human comprehenders, parsers operate incrementally, analyzing syntactic structure in real time as each word is heard or read. The operation of the parser is subject to performance constraints such as limitations on working memory and the time necessary to complete mental operations. In contrast, grammatical operations do not occur in real time (though they may constitute an ordered sequence of representations in a derivational theory) and working memory is irrelevant. Many confusions arise because psycholinguists and linguists often use similar terminology. For example, a syntactician may assume that a derivation has an input and a sequence of representations leading to an output, but neither the input nor the intermediate representations need align directly with the inputs and intermediate representations within a psycholinguistic theory of sentence comprehension or sentence production. Perhaps one of the most telling contrasts between conceptions of the parser and the grammar is that psycholinguists construct different theories to account for comprehension and production, while such a distinction makes no sense within formal linguistic theory. Rather, the entire syntactic derivation is an atemporal representation of linguistic competence. It represents our implicit knowledge about the language without describing the cognitive operations necessary to understand or produce said language.

The above description assumes weak type transparency between the grammar and the parser (see Berwick & Weinberg, 1984, and Chomsky, 1968, for definitions of weak and strict type transparency, and reasons to avoid strict transparency). Granted, parsing data **is** directly relevant to theories of grammar if linguists assume strict type transparency in mapping from processing to linguistic theory. Fairly strict type transparency was maintained in a 1960s theory of parsing, the Derivational Theory of Complexity (e.g., Miller & Chomsky, 1963), which assumed a one-to-one mapping between the transformations in Transformational Grammar and the mental operations involved in parsing. Many psycholinguists consider the Derivational Theory of Complexity to be a historical illustration of the perils of assuming strict transparency between one's psycholinguistic theory and some syntactic theory. Even ignoring questions of psychological plausibility, the danger of constructing a processing theory around a linguistic theory that will soon be out of date is enough to scare many psycholinguists away from the strict transparency approach. Not all psycholinguists have been convinced of the danger, however. As recently as 1996, Colin Phillips proposed that the parser and the grammar are essentially the same system in his dissertation.

Much more frequently, linguists and psycholinguists assume that there is only weak type transparency between the grammar and the parser. Under weak transparency, some unknown set of linking assumptions maps between linguistic theory and cognitive processes. The output of the parser and the output of the grammar must be roughly compatible, but the two systems may arrive at their respective outputs in very different ways. I say "roughly compatible" because the class of parseable sentences is not equivalent to the class of grammatical sentences. The odd cases are normally explained by "performance factors." For example, working memory constraints might prevent comprehension of a doubly center-embedded, but grammatical, sentence, while the ability to recover from a disfluency could enable comprehension of a superficially ungrammatical sentence.

Even under weak transparency, there is some appeal for linguistic theories that map straightforwardly to processing data. See for example, Crain and Fodor (1985), who maintain that the relationship between the parsing system and Generalized Phrase Structure Grammar (GPSG) is relatively transparent, and consider this an advantage for the GPSG approach. More recently, Jackendoff (2002) stated that a more satisfactory union of linguistics and pycholinguistics was one of his goals in developing a new linguistic framework. From the perspective of a psycholinguist,

there is a big advantage to linguistic formalisms, such as those in Categorial Grammar (e.g., Steedman, 1996), which can be incorporated into a processing model in a fairly straightforward manner. Yet, even though I am more likely to use formalisms from Categorial Grammar than those from Minimalism (Chomsky, 1995) to describe the syntactic representations that are accessed from the lexicon and used during sentence comprehension, I generally remain agnostic as to which theory provides a more optimal account of linguistic knowledge. My agnosticism stems from the belief that, while simple linking assumptions between linguistic and psycholinguistic theories would be ideal, the adequacy of a linguistic theory depends upon linguistically internal criteria and does not hinge upon the linking assumptions.

In contrast to my view, some researchers find formal linguists' lack of interest in psycholinguistic data quite troubling. For example, Edelman and Christiansen (2003, p. 60) argued for "the need to demonstrate the psychological (behavioral), and, eventually, the neurobiological, reality of theoretical constructs" such as the operations merge and move with the Minimalism Program of syntactic theory. However, unless one assumes strict transparency, experimental psycholinguistic data is not needed to test these theoretical constructs. Phillips and Lasnik (2003) did take the strict transparency view and replied to this criticism by providing a long list of experimental and neuroscience papers demonstrating findings such as (a) agreement violations produce a particular kind of electrical brain response, and (b) readers reactivate boxer at the underlined gap location in *The boxer that the journalist questioned __ got angry*. While such results are consistent with particular grammatical formalisms, these data are beside the point. They are not relevant to the theoretical foundations of Minimalism questioned by Edelman and Christiansen, and such data has had no observable impact on the development of syntactic theory. One could take the position that linguistic theory **should** account for the available psycholinguistic data, including "performance factors" such as working memory constraints and garden path recovery strategies. Jackendoff (2002) is an example of a move in this direction. However, such an obligation would dramatically change the goal of most formal linguists from the description of linguistic knowledge to the description of how linguistic knowledge is implemented within a processing system that operates in real time.

This is not to say that psychological data is completely irrelevant to linguistic theory or that linguistic intuitions have a privileged access to the mental representations postulated by syntactic theories. Psychological data is directly relevant if a linguistic theory predicts that constituents of Type X will be processed differently than constituents of Type Y. All grammatical theories make this type of prediction with regards to grammaticality: The word strings that can be generated by the grammar should be exactly those strings judged to be acceptable. These predictions are usually tested somewhat informally via linguistic intuitions. However, they can be tested experimentally by predicting patterns of syntactic anomaly effects in (for example) an event-related potential (ERP) paradigm. There is little justification for such efforts from the point of view of syntactic theory, because linguistic intuitions can be collected much more quickly, easily, and inexpensively. Although there are numerous concerns about the reliability of linguistic intuitions, similar concerns apply to experimental research. Collecting either linguistic intuitions or experimental data requires careful thought and considerable expertise.

The important question is: Are there cases in which linguistic theory predicts that different types of constituents will be processed differently and linguistic intuitions and formal linguistic methods alone have not provided clear data? Both conjuncts might be true if there is a distinction in formal linguistic theory that entails a distinction in how linguistic knowledge is stored in long-term memory. One apparent example is the well-known debate about the verbal past tense. The traditionalist approach assumes that irregular verbs must encode the past tense within the lexicon while regular verbs need not do so because the past tense can be formed by applying a general rule (e.g., Pinker & Prince, 1988). The opposing connectionist view is that both regular and irregular past tenses are formed via the same mechanism (based on the properties of all the individual lexical items), without any explicit rules (Rumelhart & McClelland, 1986). However, even if the connectionist account is right, it doesn't threaten the existence of explicit rules **within** linguistic theory, because the rules still provide an elegant description of our linguistic knowledge. As long as they assume weak transparency, linguists can consider the

connectionist account to be one possible implementation of the rules vs. lexical-specification contrast within the formal theory.

This example illustrates one half of an interesting asymmetry. Psychological evidence that some item-based mechanism can mimic rule-governed behavior is not enough to eliminate rules within linguistic theory. In contrast, psychological evidence that the necessary knowledge is not specified in the lexicon does strongly suggest the use of a general rule. For an example of psychological data that meets this additional constraint, let us consider the linguistic distinction between arguments and adjuncts.

ARGUMENTS VS. ADJUNCTS

Most syntactic theories distinguish between arguments and adjuncts in terms of lexical specification. In the sentence, *Chris gave Kim some candy on Tuesday in the park*, the verb *gave* is the lexical head of the verb phrase (VP). As such, it specifies three arguments and assigns a thematic role to each: *Chris* (agent), *Kim* (recipient) and *candy* (theme). In contrast, *Tuesday* and *the park* are adjuncts, getting their thematic roles from the prepositions that head their phrases. Many syntactic theories have a structural distinction as well: arguments are sisters to the head, while adjuncts are sisters to a phrasal node (e.g., Chomsky, 1995). Processing evidence can't address the structure of the phrase tree as long as we assume weak type transparency, but if arguments and adjuncts are processed differently because arguments are lexically specified and adjuncts are not, processing evidence may be able to distinguish the difficult cases. Critically, we need psychological evidence that adjuncts are **not** lexically specified.

There are many difficult cases in which linguistic intuitions fail to provide clear data concerning which phrases are adjuncts and which phrases are arguments. One such example is the underlined prepositional phrase (PP) in *Kim changed the tire <u>with a monkey wrench</u>*. Although numerous tests have been devised for soliciting the critical intuitions, instrument PPs remain difficult to categorize (e.g., Larson, 1988; Marantz, 1984; Schutze & Gibson, 1999). Following linguistic tradition, ungrammatical sentences will be preceded by an asterisk in the following examples. Like typical arguments, instruments can't normally be iterated (**John cut the meat with a knife with the sharp end.*), but they can be extracted from weak islands (*With which key do you deny that the butler could have opened the door?*). However, like typical adjuncts, they allow pro-form replacement (*John will eat the cake with a fork and Mary will do so with a spoon.*). The sheer number of linguistic tests for argument status is a sign that none of them work very well. One might conclude that there is no sharp distinction between arguments and adjuncts—such a possibility has certainly been considered, as discussed below. Whether or not a sharp distinction exists, the argument/adjunct contrast may be a situation in which psycholinguistic data is more informative than intuitive data. The degree to which the argument/adjunct distinction is unique in this respect will be discussed in the final section of this chapter.

Although the argument/adjunct distinction figures prominently in many linguistic and psycholinguistic theories, there have been attempts to eliminate the distinction entirely or to reshape it. For example, in Steedman's (1996, p. 77) Categorial Grammar, he assumes that "all PPs, even those that would normally be thought of as modifiers rather than subcategorized, are in fact arguments." Similarly, in some versions of Head-driven Phrase Structure Grammar (Pollard and Sag, 1994), both arguments and adjuncts are lexically specified by the lexical head, and Pesetsky (1995) argued that arguments and adjuncts can occur in the same syntactic environments. Alternatively, some linguists have argued that the distinction is not binary (e.g., Grimshaw, 1990). Within cognitive psychology, MacDonald, Pearlmutter, & Seidenberg (1994) envisioned an argument/adjunct continuum, with arguments and adjuncts differing only in the frequency of co-occurrence with the lexical head. It's worth noting that the argument/adjunct status of instruments and other phrases may differ from one verb to another.

In sum, there are at least two controversies within formal linguistic theory that psycholinguistic data may speak to. The first is whether there is in fact any distinction between the

lexical specification of arguments and adjuncts. Secondly, if such a distinction is to be maintained, psycholinguistic data may help resolve the debate over problematic cases such as instrument PP's.

In order to address the argument/adjunct distinction from a psycholinguistic perspective, I will adopt many of the assumptions of constraint-based lexicalist theories of sentence processing (e.g., MacDonald et al., 1994). I will assume that much of syntactic knowledge is stored lexically and accessed via word recognition. I will assume that syntactic structures are built incrementally during sentence comprehension, and that new constituents are attached to the developing structure via competition between lexical alternatives. Constraints from any level of representation can influence competition, but the relative frequency of lexical forms is an especially powerful constraint. Just as more frequent meanings of semantically ambiguous words are accessed more easily than less frequent meanings, so more frequent syntactic forms are accessed more easily. This means that lexically specified structures exhibit lexically specific frequency effects. Consider the following example. Both *delegate* and *suggest* can head either a dative or a simple transitive structure, but the dative form is relatively more frequent for delegate. This is illustrated in Figure 1, with the more frequent structure in boldface. Lexicalized versions of both structures are accessed by recognition of either verb, but weighted by frequency. An alternative that is more strongly available is assumed to be easier and/or faster to integrate with the developing structure. Thus, *to the students* would be attached more easily following *delegate* than *suggest*, because the PP is specified by the dominant structure in the former case. The subcategorization preference effects reported by Stowe, Tanenhaus, and Carlson (1991), Trueswell (1996), and others provide evidence for this type of lexical frequency effect.
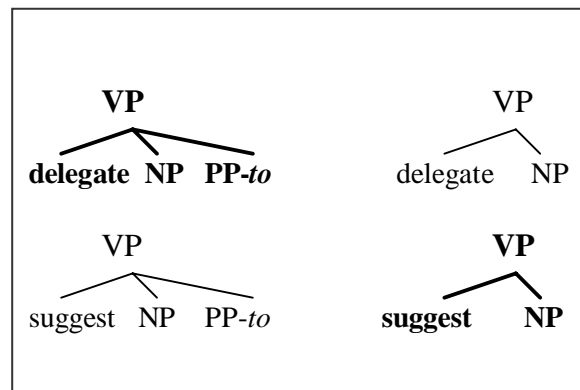

FIG 1. Alternative syntactic forms of *delegate* and *suggest*.

If argument slots are represented in the lexical entries of their heads, but adjunct slots are not, only arguments could be attached using the tree-adjoining mechanism summarized above and illustrated in Figure 2. Given the structures in Figure 1, attaching an adjunct such as *during the meeting* would have to be accomplished using some other mechanism such as an attachment rule that is not associated with any particular lexical head. Under this type of a two-mechanism account, lexical frequency effects would be predicted for arguments, but not for adjuncts. This Argument Structure Hypothesis is relevant to linguistic theory because it makes the prediction that argument phrases will be processed differently than adjunct phrases.
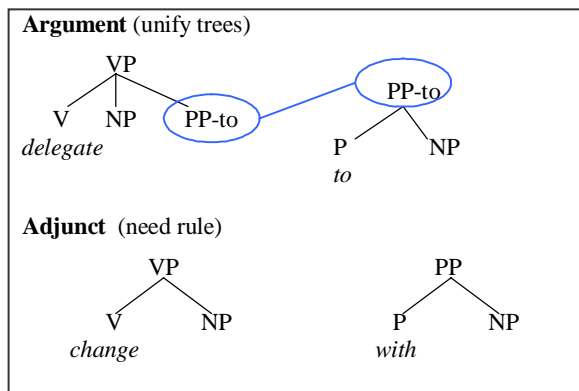
FIG 2. The upper portion of the figure illustrates a lexical unification mechanism for argument attachment. The lower portion of the figure illustrates that same mechanism cannot be used for adjunct attachment if adjunct slots are not lexically specified by the head (i.e., *change*).

Testing the Argument Structure Hypothesis with reading paradigms

In the previous section, I suggested that argument status can be diagnosed by the presence/absence of a certain type of lexical frequency effect. One challenge for testing this Argument Structure Hypothesis is distinguishing lexical frequency effects from plausibility effects and other factors that might influence our dependent measure, such as reading time on the phrase of interest.   In the current section, I will illustrate this problem using a finding from Spivey-Knowlton and Sedivy (1995), and suggest a solution offered by Boland, Lewis, & Blodgett (2004).

Some potential counter-evidence to the Argument Structure Hypothesis was reported by Spivey-Knowlton and Sedivy (1995). Using stimuli like those in (1), they found that VP-attached PP adjuncts were read more quickly than noun phrase (NP)-attached adjuncts following an action verb, while the reverse pattern was found for psych/perception verbs. Because action verbs are more likely to be modified by a PP headed by *with* (see Table 1), this data pattern might represent a lexical frequency effect, with the co-occurrence frequency between the adjunct and its lexical head influencing the ease of attachment. If so, it demonstrates that PP adjuncts are lexically specified, contrary to the Argument Structure Hypothesis. However, an alternative account is based upon a difference in local plausibility. Upon reading *The mechanic changed a tire with…* it might seem more appropriate to say what or who the tire was changed with, rather than to further define the tire as being one with some property. In contrast, for the psych/perception verb example, noticing with someone or something is less plausible than the customer being defined by some property that can be expressed in a *with*-PP. The plausibility account is consistent with the Argument Structure Hypothesis, because both NP-attached and VP-attached adjunct options could be generated by rule rather than lexically specified.

(1) Spivey-Knowlton and Sedivy (1995) stimuli
Action Verbs
  The mechanic changed a tire…
    …with a faulty valve   (NP-attached, slow)
    …with a monkey wrench   (VP-attached, fast) [Instrument]
Psych/Perception Verbs
  The salesman noticed a customer…
    …with ripped jeans   (NP-attached, fast)
    …with a quick glance   (VP-attached, slow)

TABLE 1
Spivey-Knowlton and Sedivy (1995) Normative Data for action verbs and
psych/perception verbs, concerning the VP-attachment bias for PP's headed by *with*

Verb Class	*The N V'd the N with…* **VP-attached Sentence Completions**	**Brown Corpus: Number VP-attached** *with***-PP's**
Action	90%	40
Psych/Percept	24%	4

In a recent paper, some colleagues and I testd the lexical frequency hypothesis to determine whether the adjuncts were in fact lexically specified (Boland et al., 2004). Unfortunately, no dependent measure provides a pure index of lexical frequency effects, uncontaminated by other variables. However, it may be possible to distinguish lexical frequency effects from other influences on reading time. In a theoretical approach that I have advocated (Boland, 1997; Boland & Blodgett, 2001), lexical frequency has a privileged status in influencing syntactic analysis: lexical frequency, but not plausibility, influences the initial generation of syntactic structure(s), while both lexical frequency and plausibility influence syntactic ambiguity resolution, as shown in Figure 3. This approach maintains a distinction between the generation of syntactic structure and selection processes that operate when multiple grammatical structures are possible. The distinction between the generation of syntactic structure and syntactic ambiguity resolution is explicit in some parsing theories (Altmann & Steedman, 1988; Boland, 1997; Lewis, 1993) and acknowledged as functionally necessary in others (Spivey-Knowlton & Tanenhaus, 1998).
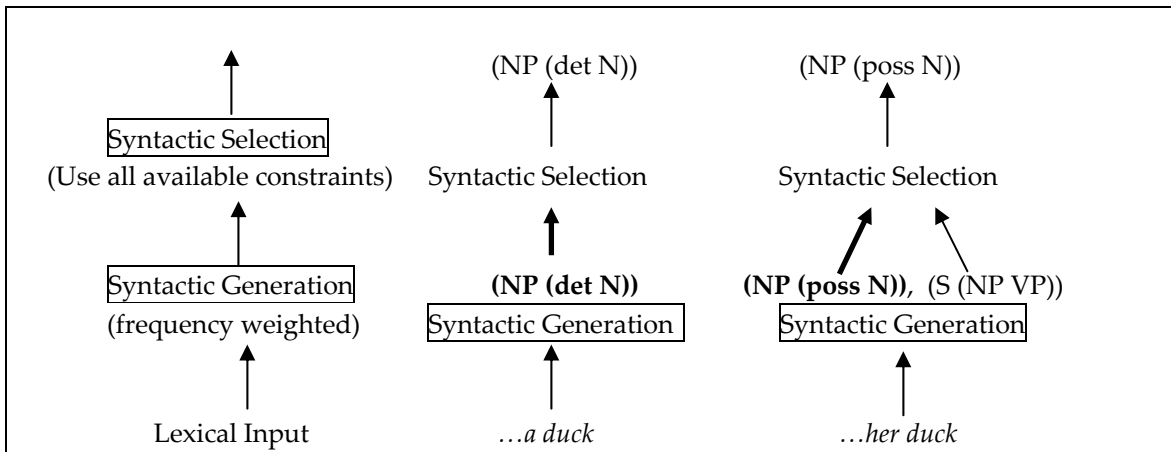

FIG 3. The architecture of the parsing model is given on the left, the representations generated by the model are provided for an unambiguous example (center) and an ambiguous example (right).

Frequency effects in syntactically unambiguous sentences are the strongest possible evidence for the lexicalization of syntactic knowledge, because the effects must arise during lexical access and generation of syntactic structure. To illustrate, compare two noun-verb homographs, as in (2): *play* occurs most frequently as a verb, while *duck* occurs most frequently as a noun. Boland (1997) and Corley (1998) each found that encountering a lexically ambiguous word in its less frequent syntactic form increased reading time compared to encountering its more frequent form. Thus in (2), reading times for *duck* are faster than for *play*, because the syntactic context is consistent with the dominant form of *duck*, but the subordinate form of *play*.

Importantly, reading times in syntactically ambiguous sentences (i.e., *She saw her play)* are influenced by high level constraints like discourse congruency, but reading times in unambiguous sentences like (2) are not (Boland, 1997). Boland and Blodgett (2001) found additional evidence that lexical frequency constraints and discourse constraints impact sentence comprehension in different ways. In an eye tracking experiment that used unambiguous target sentences like those in (2), we found lexical frequency effects only in first pass measures of processing difficulty, while discourse congruency effects were limited to second pass measures. In the absence of any alternative structures, discourse congruency had no impact on syntactic analysis. Rather, the second pass effects were presumed to reflect an anomaly within the discourse level representation. Together, these findings suggest that lexical frequency affects lexical access and syntactic generation, but discourse congruency does not. Instead, discourse congruency plays a role in ambiguity resolution (syntactic selection) and relatively late discourse coherence processes.

(2)   a. *She saw a play.*
        b. *She saw a duck.*

Under this approach, we can minimize the influence of factors that affect selection processes by using maximally unambiguous contexts, as in (3). This should increase the role of lexical frequency relative to plausibility in syntactic processing. For PP adjuncts like those used in Spivey-Knowlton and Sedivy (1995), there should no longer be a difference between action verbs and psych/perception verbs. In contrast, lexical frequency effects should be found for VP-attached PP arguments, and we used some new dative sentences to demonstrate a true lexical frequency effect.

(3) Boland et al. (2004) stimuli structurally biased toward VP-attachment
    <u>VP Adjuncts</u>
        High Lexical Frequency of VP Attachment
            The tire that the mechanic changed with a monkey wrench...
        Low Lexical Frequency of VP Attachment
            The customer that the salesman noticed with a quick glance...
    <u>VP Arguments</u>
        High Lexical Frequency of VP Attachment
            The chores that the parents delegated to their kids …
        Low Lexical Frequency of VP Attachment
            The chores that the parents suggested to their kids …

Consider the context *The mechanic changed the tires…*. If the next word is *with*, English syntax allows for two possible adjunct attachments of the PP headed by that preposition: modification of the VP or modification of the direct object NP. If these attachment alternatives are both rule-generated and thus equally available, one must use pragmatic knowledge or some other mechanism to select the most likely attachment site. In contrast, pragmatic knowledge and plausibility would play a reduced role if we decreased the availability of NP attachment by fronting the direct object: *The tire that the mechanic changed with…*. In this case, structural factors such as recency and complexity make VP attachment more accessible and would likely swamp the selection process.

Even in such relatively unambiguous structures, if the verb takes an argument PP, we should see effects of lexical frequency. Consider the context *The chores that the parents delegated/suggested…*. A PP like *to the children* should be read more quickly following *delegate* compared to *suggest* because the dative syntactic structure (shown on the left half of Figure 1) is more strongly available for *delegate* than for *suggest*. The lexical frequency effect arises because

access to the competing argument structures is weighted by relative frequency.

These predictions were tested by Boland et al. (2004) and confirmed using both self-paced, phrase-by-phrase reading and eye-fixation measures. The Spivey-Knowlton and Sedivy (1995) contrast between action verbs and psych/perception verbs was replicated in locally ambiguous structures like those in (1), but was greatly reduced in the maximally unambiguous versions shown in (3). That is, self-paced reading times for the VP-attached PPs were faster in the action verb condition than the psych/perception verb condition in the locally ambiguous structures like those in (1), but not in the versions in (3) that were strongly biased toward VP attachment. Importantly, lexical frequency effects were obtained in the unambiguous structures for dative argument PP's: Self-paced reading times for the PP following a high-frequency dative like *delegate* were faster than after a low-frequency dative like *suggest*. The eye-tracking data were particularly informative because they offered additional details about the relative timing of the argument and adjunct effects. The lexical frequency effects for the dative arguments were apparent in the early eye movement measures such as the first-fixation and the first-pass reading times over the PP. In contrast, the attachment site by verb class interaction, replicating the Spivey-Knowlton and Sedivy (1995) finding on locally ambiguous adjuncts, was found only in the total time on the PP region. There were no first-pass effects during the PP for the adjunct stimuli in either the locally ambiguous condition (1) or the maximally unambiguous condition (3).

In sum, Boland et al. (2004) found lexical frequency effects in argument attachments, but not adjunct attachments. This suggests that arguments are attached using detailed lexical information that is weighted by frequency, while adjuncts are attached using more global syntactic knowledge. In our eye-tracking replication, the adjunct effects analogous to those reported by Spivey-Knowlton and Sedivy (1995) occurred later than the lexical frequency effects and were most likely caused by the influence of pragmatics on syntactic selection. The contrast between the argument and adjunct stimuli observed by Boland et al. suggests that the PP's assumed to be adjuncts, including the instrument PP's, are not lexically specified by the verbal heads. This finding ought to be considered, along with traditional linguistic tests, when evaluating the argument status of instrument PP's.

Implicit introduction of arguments in listening paradigms

Frequency effects are one consequence of the lexical specification of arguments. Such effects are an empirical marker of argument status that can be investigated in psycholinguistic experiments. Another consequence of the Argument Structure Hypothesis is that recognition of a lexical head provides access to the thematic roles associated with frequently occurring arguments. This prediction is supported by reading experiments that have demonstrated that verbs implicitly introduce their arguments into the discourse, without the arguments being explicitly mentioned (e.g., Carlson & Tanenhaus, 1988; Mauner, Tanenhaus, & Carlson, 1995). Converging evidence can be found within a listening paradigm.

When we are listening carefully, we tend to look at things as they are mentioned if the mentioned items are in the visual environment. This phenomenon extends to items that have not (yet) been explicitly mentioned. For example, Sedivy, Tanenhaus, Chambers, and Carlson (1999) found that when listeners were asked to *Pick up the tall glass…*, they often planned an eye movement to the glass during the adjective, prior to hearing *glass*. Instead of waiting for acoustic evidence that the target object is being mentioned, listeners used the current visual context (i.e., the presence of a tall/short contrast set) to select the pragmatically appropriate referent. In another study, Sussman, Campana, Tanenhaus, and Carlson (2002) found that listeners made an eye movement to an appropriate instrument (a pencil) upon hearing *Poke the dolphin* but not *Touch the dolphin*. Even though no instrument was mentioned, listeners used their knowledge about the two verbs to decide whether to manipulate the dolphin with their finger or a pencil in a real-world environment. Listeners in the Sussman et al. study were also sensitive to contextual factors that altered verb meaning. For example, they looked at a potato peeler when asked to *Peel the potato*, but not when asked to *Peel the banana*.

Directed action tasks such as that used by Sedivy et al. (1999) and Sussman et al. (2002) raise the concern that the eye-movement patterns are caused by strategies the listeners adopt to guess. Arguably, normal conversation involves a great deal of strategic guessing about the speaker's intent, so this is not a problem if the goal is to study the output of the complete comprehension process. However, if there are some partially or fully automatized aspects of syntactic and semantic processing, the directed action paradigm is not ideal for studying the representations that result from those automatized processes alone. For example, one might question whether the recognition of *poke* **obligatorily** introduces an instrument into the discourse model.

Encouragingly, there is converging evidence for the automatic activation of thematic role information from passive listening tasks. In one study, Altmann and Kamide (1999) had people listen to a sentence like *The boy will move/eat the cake* while looking at a semi-realistic scene with a boy, a cake, and some inedible (but moveable) toys. Altmann and Kamide found faster looks to the cake following *eat* compared to *move*. In fact, participants often looked at the cake in the *eat* condition prior to the onset of the noun. Altmann and Kamide concluded that the verb's thematic roles were used to pro-actively restrict the domain of subsequent reference.

Even in a passive listening task, it is difficult to identify the cause of the anticipatory fixations, because both linguistic and general world knowledge could have contributed to the effect.[1] An important question is whether the discourse elements that can be introduced by a verb are limited to members of its thematic grids. In other words, do a verb's arguments hold a privileged status or are all related words and concepts accessed in the same way? If it is solely the verb's argument structure that is driving eye movements, then listeners should not look at a bed upon hearing *The girl slept* because *bed* cannot be an argument of *slept*. Alternatively, listeners might look at a bed because beds are part of a prototypical sleeping event and are thus conceptually related to sleep. Furthermore, discussions about sleep often include mention of a bed, so linguistic co-occurrence frequency is high and the co-occurrence of sleeping and beds in participants' actual experience is likely to be extremely high. One might consider an account of Altmann and Kamide's (1999) effect that is akin to semantic priming--a conceptual, essentially intra-lexical, process. However in more recent work, Kamide, Altmann, and Haywood (2003) found that combinatory semantics rather than simple lexical relationships influenced eye movements. For example, when viewing a carnival scene, listeners looked at a motorcycle upon hearing *The man rode…* and looked at a merry-go-round upon hearing *The girl rode….* Thus, knowledge higher-level than simple lexical associations must have influenced gaze. Was it

---

[1] A different type of concern is that the visual display changes the way that linguistic material is processed. For example, if the affordances of the pictured objects lead one to expect particular verbs, then the objects introduce a verb, rather than the verb introducing its arguments. Such a circumstance may in fact be consistent with a more nuanced version of the Argument Structure Hypothesis, but it is not clear that the visual display alone would activate argument structure knowledge. This can be illustrated by Kamide, Altmann, and Haywood's (2003) experiment using Japanese, in which the verb is clause-final. Translated example stimuli are in (4). Kamide et al. found that a dative particle on the second noun (4a) prompted anticipatory looks to a potential direct object (a hamburger). Clearly, these anticipatory looks were not due to the introduction of an argument by the verb *bring*. Rather, the case-marking particle facilitates thematic role assignment, which in turn constrains the class of verbs that can occur. In (4a), a verb of transference can be inferred from the dative particle. A verb of transference requires that something be transferred, and real world knowledge about waitresses and customers may lead to the expectation that some food will be the transferred substance as opposed to a trash can, the other pictured item. Crucially with respect to the concern that the visual display drives linguistic expectations, real world expectations about waitresses and customers alone did not prompt looks to the hamburger, demonstrated by the lack of looks to the hamburger in (4b). Thus, argument structure constraints introduced a new entity into the discourse (the transferred substance) and real world knowledge prompted looks to the most plausible argument.

(4)   a. Waitress-nominative customer-dative merrily hamburger-accusative bring.
        b. Waitress-nominative customer-accusative merrily tease.

argument structure or real world knowledge, or both?

I investigated this question using a passive listening paradigm (Boland, 2004). Across three experiments, effects of both argument status and real world knowledge were found. The first experiment manipulated both the argument structure of the verb and the typicality/co-occurrence frequency of the target argument/adjunct. Example stimuli are in (5); the typical/atypical target is underlined. The goal was to distinguish between anticipatory looks to target pictures representing potential arguments and anticipatory looks to pictures that were strongly associated with the verb, but did not have the linguistic status of argument. The intransitive-location stimuli provide a clear case of an adjunct target (bed/bus), the dative-recipient stimuli provide a clear case of an argument target (teenager/toddler), and the action-instrument stimuli provide an intermediate case in which the targets are arguably adjuncts (stick/hat). Acceptability ratings insured that sentences with typical targets were judged to be more acceptable than sentences with atypical targets. Furthermore, typical targets were more likely to co-occur with their verbs.[2] Importantly, there was no evidence that typical recipients had a higher co-occurrence frequency than typical locations—if anything, the opposite was true.

> (5) Example stimuli from Boland (2004), Experiment 1.
> Intransitive-Location.
>> *The girl slept for a while on the bed/bus this afternoon.*
>> (pictures: girl, bed/bus, pillow, toy car)
> Action-Instrument.
>> *The donkey would not move, so the farmer beat it vigorously with a stick/hat every day.*
>> (pictures: donkey, farmer, stick/hat, grass)
> Dative-Recipient.
>> *The newspaper was difficult to read, but the mother suggested it anyway to her*
>> *teenager/toddler last week.* (pictures: newspaper, mother, teen/toddler, dictionary)

The primary finding in Experiment 1 was that dative verbs prompted more anticipatory looks to potential recipients than transitive action verbs prompted to potential instruments or intransitive verbs prompted to potential locations. That is, listeners were more likely to fixate the teenager or toddler in the dative example than the bed/bus or stick/hat from the intransitive and action verb examples. The relevant time window for examining these anticipatory looks was from verb onset to the onset of the PP that mentioned the target. The argument status effect began about 500 ms after verb onset, suggesting that it occurred soon after lexical access of the verb. Interestingly, listeners were just as likely to fixate the atypical recipient (toddler) as they were to fixate the typical recipient (teenager). In both the typical and atypical conditions, the potential referent met the lexical constraints on recipients for that particular verb. If verbs specify the syntactic and semantic constraints on their arguments, recognizing a verb would make available knowledge about that verb's arguments, and entities that satisfy the syntactic and semantic constraints could be identified in the current discourse model or the situational context.

In the first experiment, the argument structure of the dative verbs introduced an abstract recipient, but there was only one potential referent pictured—the same one that was explicitly mentioned. A second experiment used the same sentences, but presented both typical and atypical targets (the recipients, instruments or locations) on each trial. This experiment produced clear

---

typicality effects, suggesting that when more than one potential referent is pictured, real world knowledge is used to focus attention on the most appropriate referent. This account is consistent with prior evidence that pragmatic constraints influence ambiguity resolution, but not the generation of linguistic structure (Boland, 1997).

The argument status effect was replicated in a third experiment, in which a single animate NP (and the corresponding picture) served as an argument in the dative condition (6a) and as an adjunct in the action verb condition (6b). No instrument was mentioned in the critical trials, though a prototypical instrument for the action verb was always pictured, and in filler trials, pictured instruments were mentioned. There were more looks to the target picture when it was an argument (recipient) than when it was an adjunct (benefactor, instrument) during the interval 500 - 1000 ms after the onset of the verb. There were very few fixations on the pictured instrument during this time-frame, and there was no difference in the probability of a look to a prototypical adjunct (*fix*-tools) and an improbable adjunct (*mention*-tools). Co-occurrence frequency does not provide an alternative explanation. There were no reliable differences in co-occurrence frequency among the dative-recipient, action-benefactor, and action-instrument pairs.

> (6) *One window was broken, so the handyman…*
> [pictures: window, handyman, couple, tools]
>   a. *mentioned it right away to the owners.* (recipient-Argument)
>   b. *fixed it hurriedly for the owners.* (benefactor-Adjunct)

Together, these findings demonstrate that linguistic constraints play a privileged role in guiding visual attention in this passive listening paradigm. Furthermore, these argument status effects suggest an important distinction between adjuncts and arguments in terms of how verbs introduce entities into the discourse. A verb implies its arguments, but not adjuncts, before they are explicitly mentioned. In addition, these results suggest another experimental test of argument status.

CONCLUSIONS

In sum, the results from reading and listening paradigms converge to support the view that arguments and adjuncts have a different status in parsing. In the reading experiments summarized above, there were lexical frequency effects for PP arguments but not PP adjuncts, suggesting that only the arguments were syntactically analyzed using a lexicalized mechanism. In the listening experiments summarized above, verbs implicitly introduced their arguments, but not adjuncts, and visual attention was drawn to likely referents of those arguments. This is to be expected if only arguments are represented in the lexical entries of their heads. These findings are relevant to two major issues in parsing theory: How is syntactic knowledge stored and accessed? What are the mechanisms for attaching new constituents to the developing syntactic representation?

Are these results also relevant to formal syntactic theory? The psycholinguistic focus on arguments and adjuncts in the discussion above is obviously motivated by the argument/adjunct distinction in formal linguistic theory. In this case and many others, psychologists who study sentence comprehension rely on linguistic theory for insight into the nature of our mental representations and vocabulary for describing them. However, the insights don't flow as freely in the other direction. Formal linguists don't often try to account for phenomena that psychologists discover about the mental representations involved in language processing. This may be because formal linguistics has little to gain from cognitive psychology under weak transparency assumptions. But what about the exceptional cases?

I have suggested that assertions about lexical specification within syntactic (and morphological) theory are in fact claims about how linguistic knowledge is stored, accessed, or acquired. As such, some of these assertions may be tested more definitively with experimental

methods than with linguistic intuitions. If the experimental data are clear, and if linguistic theory makes note of them, the experimental paradigms reviewed above may be able to resolve some of the debates about the distinction between arguments and adjuncts.

In contrast, psycholinguistic research cannot resolve purely structural debates about the geometry of the phrase structure tree or the nature of a derivation within syntactic theory, because these constructs do not generate straightforward predictions about processing. An example is the extensive line of experimental research (e.g., Clahsen & Featherston, 1999) investigating the psychological reality of "traces" left behind by movement in certain theories of syntax. The Trace Reactivation Hypothesis is usually stated as the prediction that an antecedent will be reactivated at its trace site. For example, in the sentence, *In which box did you put the cake__?*, the fronted PP *in which box* is the locative argument of *put*. The PP is said to have moved out of its canonical position, leaving behind a trace, which is represented by the underline. During comprehension of such a sentence, *in which box* would be coindexed with the trace, and as a result, the PP could then be interpreted as the appropriate argument of *put*. The fundamental problem is that recognition and coindexing of the long distance dependency is a complicated **processing** issue that has not been carefully addressed in the trace reactivation literature. Researchers generally assume that coindexing—and therefore reactivation—occurs at the linear position of the trace. In our example, coindexing would take place after the offset of *cake,* so priming of *box* would be predicted at that point in the sentence. Unfortunately, because traces are phonologically null, the listener or reader does not perceive a trace directly. Therefore, recognition and coindexing of the purported trace need not coincide with its linear position in a sentence. If they are psychologically real, traces must be postulated on the basis of cues that may or may not be adjacent to the trace site. For example, recognition of *put* could initiate projection of a VP with slots for a direct object and a locative PP. If so, a trace could immediately be posited and coindexed with *in which box*, leading to priming of *box* at *put*. Depending upon the strategy adopted by the parser, other alternatives are also possible. In short, a syntactic theory of traces makes no predictions about when or if priming should occur unless it is wedded to well-articulated processing theory that specifies how and when traces are postulated, as well as how previously encountered phrases will persist or decay in working memory. Because these processing questions are themselves controversial, it is difficult to see how psycholinguistic research can resolve syntactic debates over traces.

Even if some psycholinguistic data influence some corners of formal linguistic theory, I do not believe we are on the brink of a revolution in linguistic methodology. Psycholinguistic data--and data from cognitive neuroscience for that matter--will always play a secondary role in formal linguistic theory, adjudicating between linguistic theories that are equally elegant and account for the traditional data (linguistic intuitions from a variety of languages) equally well. This is as it should be, as long as we accept the assumptions of weak transparency. Linguistic theory does not attempt to describe neural or behavioral patterns, but rather the knowledge state that gives rise to those neural and behavioral patterns. Linguistic assertions about lexical specification are unusual in that these assertions concern the linking assumptions between formal theories of linguistic knowledge and processing theories of how linguistic knowledge is stored, accessed, and used. For the most part, the linking assumptions among the knowledge state, the behavior, and the neural activity remain underspecified in both linguistic and psycholinguistic theories.

ACKNOWLEDGMENTS

References

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition, 73*, 247-264.

Altmann, G. T. M. & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30*, 191-238.

Berwick, R. C. & Weinberg, A. S. (1984). *The grammatical basis of linguistic performance: Language use and acquisition.* MIT Press, Cambridge, MA.

Boland, J.E. (1997). The relationship between syntactic and semantic processes in sentence comprehension. *Language and Cognitive Processes, 12*, 423-484.

Boland, J.E. (2004). Visual Arguments.To appear in *Cognition*, pending revisions.

Boland, J.E., & Blodgett, A. (2001). Understanding the constraints on syntactic generation: Lexical bias and discourse congruency effects on eye movements. *Journal of Memory and Language, 45*, 391-411.

Boland, J. E., Lewis, R., & Blodgett, A. (2004). *Distinguishing generation and selection of modifier attachments: Implications for lexicalist parsing and competition models.* Submitted.

Carlson, G. N., & Tanenhaus, M. K. (1988). Thematic roles and language comprehension. W. Wilkins (Editor), *Syntax and Semantics, Volume 21: Thematic Relations*. San Diego, CA: Academic Press.

Chomsky, N. (1959). Review of Skinner's Verbal Behavior. *Language, 35*, 26-58.

Chomsky, N. (1968). *Language and Mind*. New York: Harcourt, Brace, and Jovanovich.

Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.

Clahsen, H. & Featherston, S. Antecedent priming at trace positions: Evidence from German scrambling. *Journal of Psycholinguistic Research, 28*, 415-437.

Corley, S. (1998). *A statistical model of human lexical category disambiguation.* Doctoral dissertation, University of Edinburgh.

Crain, S., & Fodor, J. D. (1985). How can grammars help parsers ? In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds), *Natural language parsing: Psychological, computational, and theoretical perspectives.* Cambridge: Cambridge University Press.

Edelman, S. & Christiansen, M. (2003). How seriously should we take Minimalist syntax? *Trends in Cognitive Sciences, 7*, 61-62.

Grimshaw, J. (1990). *Argument Structure*. Cambridge, MA: MIT Press.

Jackendoff, R. (2002). *Foundations of language: Brain, meaning, grammar, evolution.* Oxford: Oxford University Press.

Kamide, Y., Altmann, G. T. M. & Haywood, S. (2003). Prediction and thematic information in incremental sentence processing: Evidence from anticipatory eye-movements. *Journal of Memory and Language, 49*, 133-156.

Larson, R. (1988). On the double object construction. *Linguistic Inquiry, 19*, 335-392.

Lewis, R. (1993). *An architecturally-based theory of human sentence comprehension.* Doctoral dissertation, Carnegie Mellon University.

MacDonald, M. C, Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*, 676-703.

Marantz, A. P. (1984). *On the nature of grammatical relations.* Cambridge, MA: MIT Press.

Mauner, G., Tanenhaus, M. K., & Carlson, G. N. (1995). Implicit arguments in sentence processing. *Journal of Memory and Language, 34*, 357-382.

Miller, G. & Chomsky, N. (1963). Finitary models of language users. In R. D. Luce, R. R.

Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology, volume 2*. New York: Wiley.

Pesetsky, D. (1995). *Zero Syntax: Experiencers and Cascades*. MIT Press, Cambridge, MA.

Phillips, C. (1996). *Order and Structure.* Doctoral dissertation. Department of Linguistics and Philosophy, Massachusetts. Institute of Technology.

Phillips, C. and Lasnik, H. (2003). Linguistics and empirical evidence. *Trends in Cognitive Sciences*, *7*, 62-63.

Pinker, S. & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28*, 73-193.

Pollard, C. & Sag, I. (1994). *Head driven phrase structure grammar*. Stanford, CA: CSLI Publications.

Rumelhart, D. E. & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, Vol 1: Foundations.* Cambridge, MA: MIT Press.

Schutze, C. & Gibson, E. (1999). Argumenthood and English prepositional phrase attachment. *Journal of Memory and Language*, *40*, 409-431.

Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition, 7*, 109-147.

Skinner, B. F. (1957). *Verbal Behavior.* Appleton-Century-Crofts, New York.

Spivey-Knowlton, M., & Sedivy, J. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition, 55*, 227-267.

Spivey-Knowlton, M., & Tanenhaus, M.K. (1998). Syntactic ambiguity resolution in discourse: Modeling effects of referential context and lexical frequency within an integration-competition network. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24*, 1521-1543.

Steedman, M. (1996). *Surface structure and interpretation.* Cambridge, MA: MIT Press.

Stowe, L., Tanenhaus, M. & Carlson, G. (1991). Filling gaps on-line: Use of lexical and semantic information in sentence processing. *Language and Speech*, *34*, 319-340.

Sussman, R. S., Campana, E., Tanenhaus, M. K., & Carlson, G. M. (2002). *Verb-based access to instrument roles: Evidence from eye movements.* Poster presented at the 8th Annual Architectures and Mechanisms of Language Processing Conference, Tenerife, Canary Islands, Spain.

Trueswell, J.C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of Memory and Language, 35*, 566-585.