Empty Categories, Implicit Arguments, and Processing

Empty Categories, Implicit Arguments, and Processing

Maryellen C. MacDonald and Robin Clark

Carnegie Mellon University

1. Introduction

The development of a theory of empty categories permits us to entertain the hypothesis that empty categories have implications for sentence processing. For example, we can consider the hypothesis of Chomsky (1982) and Bouchard (1984) that empty categories do not differ qualitatively from phonologically realized elements. In particular, Bouchard (1984) suggests that an empty anaphor will not differ with regard to its basic properties from an overt anaphoric expression. We might extend Bouchard's hypothesis by positing a distinction between the processing of projected positions, which are structurally realized in a syntactic representation, and unprojected arguments, which may have semantic properties but are not related to a structural position in the syntactic representation. To clarify these points, consider the sentences in (1-2):

1a. John believes himself to be the best student.
b. John $_{i}$ is believed [e] $_{i}$ to be the best student.

2a. The tenement was destroyed for personal gain.
b. Love of money is a corruption of the soul.

In (1a), there are two structurally realized and coreferent positions, John, and the anaphor bound by John, himself. In (1b), there are also two structurally realized and coreferent positions, John and the empty category bound by John, indicated by [e]. Our hypothesis predicts that the language processor should treat himself in (1a) and the empty category in (1b) in a unified manner. In the examples in (2), there are a number of implicit arguments, including the agent of destroyed and the person who gains by the destroying in (2a). Although neither argument is mapped onto structural positions, the standard judgment for this sentence is that these two arguments pick out the same individual. A similar analysis can be developed for (2b). Our hypothesis is that there will be a distinction in the manner in which the language processor treats these structurally unrealized positions and the way that overt and empty anaphors are treated.

In order to formulate the processing theory more fully, let us first consider an analysis of the passive construction. Within the framework of Transformational Grammar, this construction has generated an extensive literature (Jaeggli, 1986; Levin and Rappaport, 1986; Chomsky, 1981; Wasow, 1977, among others). Within this literature, it is relatively uncontroversial that there is a distinction between verbal passives, as in (3), and adjectival passives, as in (4):

3a. The girl was kissed.
b. The hypotheses were believed.
c. The cart was pushed.

4a. The girl was surprised.
b. The hypotheses were complicated.
c. The cart was stained.

Verbal and adjectival passives can be distinguished first by the subtle semantic differences between them: The sentences in (3) describe events, while those in (4) have a more stative reading. The two types of passives can also be differentiated by the syntactic environments in which they can appear. Bare adjectival passives can appear in exactly those environments where adjectives can appear, but the appearance of verbal passives is more restricted (Levin and Rappaport, 1986; Wasow, 1977). For example, both adjectives and adjectival passives can appear as prenominal modifiers, while verbal passives cannot, as can be seen in the phrases in (5), which contain in order, a prenominal adjective, adjectival passive, and verbal passive:

5a. the happy/surprised/*kissed girl
b. the clear/complicated/*believed hypotheses
c. the useable/stained/*pushed cart

Similarly, adjectival passives can appear as the complements of verbs such as "seem" and "appear" that select for adjectival rather than verbal complements. Verbal passives cannot appear in this environment:

6a. the girl appeared happy/surprised/*kissed
b. the hypotheses seemed clear/complicated/*believed
c. the cart looked useable/stained/*pushed

While some adjectives and adjectival passives can be prefixed by a negative un, verbal passives cannot take this prefix, as evidenced by (7). Finally, as the examples in (8) indicate, verbal passives cannot be modified by intensifiers such as "very" and "rather".

7a. the girl is unhappy/unsurprised/*unkissed
b. the hypotheses were unclear/uncomplicated/*unbelieved
c. the cart was unuseable/unstained/*unpushed

8a. the girl is rather happy/surprised/*kissed
b. the hypotheses were very clear/complicated/*believed
c. the cart was quite useable/stained/*pushed

The standard analysis of passives within the Government Binding framework (Chomsky, 1981; Jaeggli, 1986) treats the adjectival passives as pure adjectives. A morphological rule changes the category of the stem from Verb to Adjective, absorbs the external theta-role, and externalizes the internal theta-role. An adjectival passive like complicated[1] is thus an intransitive which assigns an external thematic role. Verbal passives, on the other hand, undergo no category change. The external theta-role is absorbed; the resulting verb assigns no external theta-role, and by Burzio's Generalization (Burzio, 1986), cannot assign accusative Case. An NP which is dependent on the verb for case must move to a position which receives Case but not a theta-role. The only candidate landing site is subject position. For example,

---

[1] We note that the absorbed external theta role of the verb stem may still be assigned in a by-phrase (cf. (i)) and may be treated as an implicit argument.

(i)a. People inhabit this island.
b. This island is inhabited (by people).

(9a) is the d-structure for "The girl was kissed," and (9b) is the S-structure.

9a. [e] was kissed [the girl]
 b. [the girl] was kissed [e]
 i i

In (9a) the girl occurs in object position, where it is assigned the THEME theta-role. Subject position is vacant, since it receives no theta-role and therefor cannot be occupied by an argument. Because the object cannot receive Case, the girl must move to a position that receives Case but not a theta-role in order to pass the Case Filter. The result of movement is shown in (9b); the girl receives nominative Case from Infl and an NP-trace appears in object position, transmitting the THEME theta-role to the girl. The structure in (9b) satisfies the Projection Principle, the Theta Criterion, and the requirements of Case Theory.

Given this analysis of passives and our hypotheses concerning the processing of empty categories, we can make several predictions. First, the postverbal empty category in verbal passives should behave like an overt anaphoric element. Second, the processing of adjectival passives should differ from that of verbal passives to the degree that the adjectival passive lacks a post-head empty category. Third, the processing of an implicit agent of the suppressed external thematic role should diverge from the processing of an NP-trace.

2. Psycholinguistic Measures

Since the relationship between the NP-trace and its antecedent is a coreference relationship, we would expect that a psycholinguistic measure which is sensitive to coreference would also be sensitive to the antecedent-trace relationship. Although the psycholinguistic literature on syntactic anaphors is rather sparse, a number of studies on coreference relations involving pronouns can be found (MacDonald, 1986; Leiman, 1982; Chang, 1980; Corbett and Chang, 1983). The most common experimental paradigm in these studies is to present sentences such as the ones in (10), where the pronoun is coreferential with one of two names in the sentence.

10a. At the carnival, Tommy guided Ruth through the maze, and the creepy fake monsters terrified him at almost every turn.

b. At the carnival, Tommy guided Ruth through the maze, and the creepy fake monsters terrified her at almost every turn.

At some point after the subject has heard or read the pronoun in the sentence, the sentence is interrupted and a probe word is shown on a computer screen. The subject must make some sort of response to the probe, such as naming it aloud or pressing a key indicating whether the probe had occurred in the sentence. Subjects are generally faster at making a response to the probe word when the probe word is the referent of the pronoun. Figure 1 presents data from MacDonald (1986), in which subjects read sentences such as in (10). Subjects saw a probe word 1/2 second after they had read the pronoun and pressed a key to indicate whether the probe had been in the sentence. When the probe word was Tommy, subjects were faster making this judgment if they had been reading (10a), which contained the pronoun "him"; they were much slower making the jugment if they had been reading (10b), which contained "her". Conversely, subjects were faster to judge that the probe Ruth had been in the sentence if they had been reading (10b), compared to (10a).[2] The interpretation of these and similar results (Chang, 1980; Leiman, 1982) has been that the presence of a pronoun can strengthen the mental representation of its referent, and this strengthening effect is reflected in faster response times to probes in these experiments.

---

[2] This figure actually represents responses to 90 sentences with structures similar to those in (10), all with different probe names, rather than just responses to the two names Tommy and Ruth.

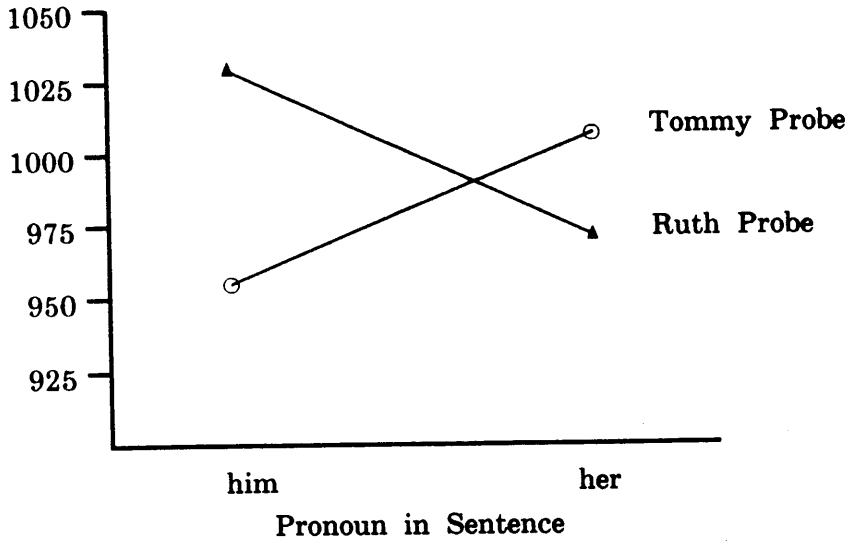

Figure 1. Time to judge if probe word occurred in the sentence

Psycholinguistic measures of the effects of pronominal reference might provide a clue for studying the processing of binding. Although the computation of pronominal reference is presumably carried out at a discourse level of representation, it is possible that the processing of binding, which is a coreference relationship in the syntax, might produce similar effects to those found with pronominal reference. If we can tap processing of binding with the probe task described above, then we can test the putative asymmetries between structurally realized empty categories and implicit arguments. In particular, if the probe task is tapping coreference relations, then we should find strengthening effects for the antecedent of a bound trace analogous to the effects found for the antecedent of a coreferential pronoun. For example, in the sentences in (11), we should find faster responses to the probe word "John" in the verbal passive (11b) compared to the adjectival passive in (11a), because the bound trace in (11b) should strengthen its antecedent.

11a. John was delighted.
b. John was kissed [e]
i i

Continuing the logic outlined above, let us examine how the probe

task would be affected by the presence of implicit arguments. Recall that implicit arguments are not mapped onto structural positions and therefore cannot enter into "strong" syntactic binding (e.g., the binding relation between two syntactic positions). Given the differences between implicit arguments and structurally realized positions, we predict that the strengthening effect triggered by an implicit argument will be qualitatively different from the effects found with overt pronouns and the effects that we hypothesize for empty categories. Consider the sentences in (12) when read as a coherent text:

12a. The furious parents marched right into the classroom.
b. [The new principal] with the modern methods was
i
assaulted [e]
i

The verbal passive assaulted in (12b) has an implicit AGENT thematic role with the NP the furious parents in (12a) as a plausible controller. Because this theta-role is unprojected, we predict that its strengthening effect on parents is relatively weak. In contrast, we predict much stronger strengthening effects for principal, which is the head of the NP that binds the trace in (12b). These predictions are tested in the experiment below.

3. Experimental Method

3.2 Materials

Fourteen word triples were constructed such that one member of the triple was an adjective, one an adjectival passive, and one a verbal passive, such as "furious/surprised/shot" (see Levin & Rappaport, 1986, and Wasow, 1977, for the tests used here to distinguish adjectival and verbal passives). The members of the triples were matched across these three experimental conditions for length and frequency of occurence in English (Kucera and Francis, 1967). Three two-sentence "passages" were written for each adjective/adjectival passive/verbal passive triple, so that any member of the triple could sensibly appear as the last word of each of the three passages for that triple. The topics of the passages were unrelated. The three passages for one triple are shown in sentences 13-15; the last word of each passage was either the adjective incoherent, the adjectival passive embarrassed, or the verbal passive assaulted.

13. The furious parents marched right into the classroom.
    The new principal with the modern methods was _____
    Probes: PARENTS, PRINCIPAL

14. The obnoxious reporters made the news conference a very rowdy affair.
    The famous politician from the tiny country was _____
    Probes: REPORTERS, POLITICIAN

15. The customers demanded an explanation for the milk shortage.
    The store owner with the wireless microphone was _____
    Probes: CUSTOMERS, OWNER

Note that the verbal passive participle and possibly the head adjective of the adjectival passive contain a suppressed external thematic role (Chomsky, 1981; Jaeggli, 1986; Levin and Rappaport, 1986). It is logically possible, contrary to our hypothesis, that the implicit argument might produce strengthening effects for, e.g. the agent in the first sentence of our passages, even though the implicit argument is not structurally projected. For example, in the passage in (15) ending in the verbal passive "assaulted," "customers" is the most plausible antecedent of the implicit argument. If our probe task can distinguish between implicit arguments and structurally projected positions, then we should find asymmetries between responses to the antecedent of the NP-trace and the antecedent of the implicit argument.

The first sentence of each passage varied between 7-11 words in length, while the second sentence was always 9 words in length. The third word of the second sentence was always a noun, and the ninth (last) word was always one member of the word triple for that passage. Two probes were selected for each passage, the subject of the first sentence (Word 2 or 3), and the subject of the second sentence (always Word 3 of that sentence). For ease of exposition the probe from the second sentence will be termed the binding probe, as any effects of binding should be seen with this probe. The first sentence probe is called the agent probe, because in all passages the implicit role probed was a suppressed agent when the passage ended in a verbal passive.

Sixty-six practice and filler passages were prepared, each containing a pair of sentences with a variety of syntactic constructions. All probes for filler passages either did not occur in the passage (the majority of probes), or were contained in the passage in some position other than the subject position of either sentence. Every probe that was not contained in a passage was an associate of a word that had occurred in the passage. This control on "false" probes forced subjects to pay close attention to the passages and refrain from making the probe recognition judgment merely on the familiarity of the concept being probed. A

yes/no comprehension question was prepared for every passage, with "yes" as the correct answer for half of the questions for both experimental and filler passages. The questions for the experimental passages avoided reference the last word in the passage so that the same question could be used for a passage regardless of whether the passage ended in an adjective, adjectival passive, or verbal passive.

3.2 Procedure

Stimuli were presented on a computer screen such that subjects were able to read one sentence at a time. At the start of each trial, the display was filled with three lines of dashes representing all nonspace characters of the stimulus sentences, probe word, and comprehension question. When subjects pressed a key with the left hand, the first sentence replaced the dashes on the top line of the display. A subsequent keypress replaced the first sentence with dashes and revealed the entire second sentence on the second line. A third keypress removed the second sentence and displayed the probe word in capital letters four spaces to the right of the end of the second sentence. Prior to this keypress, the probe word was always indicated by 11 dashes, regardless of word length, so that subjects would not have early information about the length of the word to be probed.

Subjects judged whether the probe had occurred in the preceding two sentences and pressed a "YES" or "NO" key with the right hand. This response removed the probe and displayed a comprehension question on the bottom line. When subjects had pressed the "YES" or "NO" key to the question, the screen was cleared, and a new set of dashes indicated a new trial. Subjects were encouraged to read at a normal rate, and speed and accuracy were stressed for responses to both the probes and comprehension questions. Response times to the probe and comprehension question were recorded, as were reading times for each sentence. The purpose of collecting reading time and question-answering data was to ensure that all three constructions were equally easy to read and comprehend.

Subjects saw an equal number of sentences in each of the 6 combinations of probe and sentence type conditions. They completed the experiment without a break in one 30 min session.

3.3 Subjects

Subjects were 30 undergraduates enrolled in psychology classes at Carnegie Mellon University who participated as part of a course requirement. All subjects were native speakers of

English. An additional 5 subjects were tested but were not used in the final analyses because of error rates on comprehension questions over 20%.

3.4 Results

Response times to the probe, to the question, and reading times to the second sentence were analyzed. Prior to analysis, all incorrect probe responses (error rate 4%) and incorrect comprehension question responses (6.8% error rate) were removed. Extremely long responses in each task (reading, probe judgment, question comprehension) were also removed, using a two-step procedure. First, all responses over three times the grand mean for each task were removed. Next, means and standard deviations were calculated for each subject for each task, and all scores more than 3.5 standard deviations over the subject's mean were removed. This procedure removed less than 3% of the responses for each task.

Reading times in the second sentence did not differ across the adjective (2.951 seconds), adjectival passive (3.005 seconds), and verbal passive (2.941 seconds) conditions, $F < 1$.[3] Response times to the comprehension questions (1.717, 1.748, and 1.666 seconds to the adjective, adjectival passive, and verbal passive conditions, respectively) also did not differ across sentence type, $F(2, 58) = 2.25$, $p > .11$.

Response times to the probes are shown in Figure 2. Statistical analyses indicated that responses to the agent probe following the adjective sentence, the adjectival passive sentence, and the verbal passive sentence did not differ from one another, $F < 1$. That is, the line indicating agent probe responses in Figure 2 cannot be distinguished from a flat line. Responses to the binding probe revealed a different story: Responses in the verbal passive condition were faster than in the adjective condition, $F(1, 29) = 9.63$, $p < .005$, and there was a trend for the binding probe responses to be faster in the verbal passive condition compared to the adjectival passive condition, $F(1, 29) = 3.46$, $p <$

---

[3] An F of 1 or less indicates results entirely consistent with chance variations. If F is greater than 1, then p, the probability that the results would be obtained by chance, is reported. The accepted value for statistically significant results is $p < .05$; that is, the obtained results would be expected to occur by chance less than 5% of the time.

.08. Responses to the second sentence probe did not differ in the adjective and adjectival passive conditions, $F(1, 29) = 1.35$, $p > .25$.

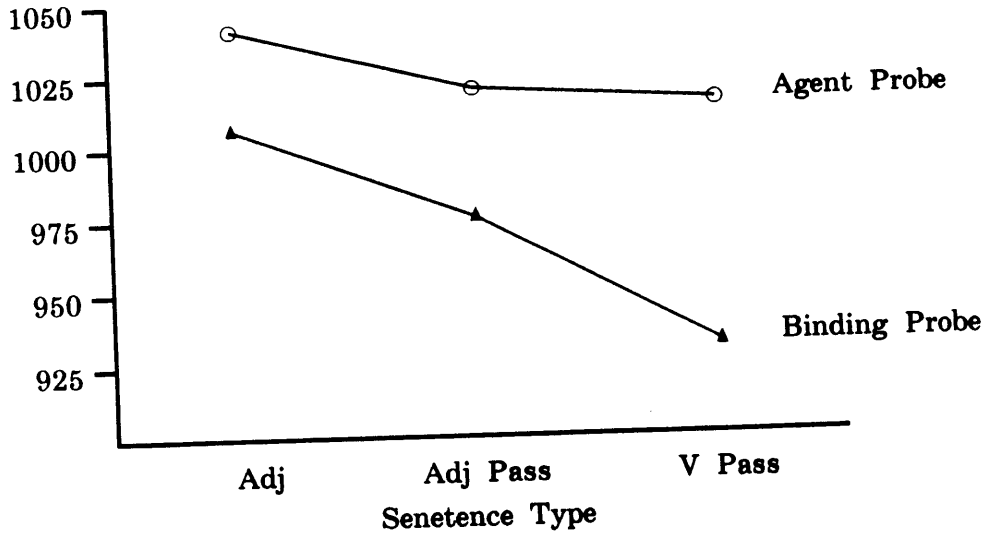

Figure 2. Response Times to Binding and Agent Probes

4. Discussion

When the sentence contaied a verbal passive, responses to the binding probe were different than responses to the same probe in the other two constructions. This result suggests that the language processor treats verbal passives differently than it treats the superficially similar adjective and adjectival passive constructions. The direction of the effect (that is, that responses were faster in the verbal passive condition) is as we would expect given the analysis of verbal passives incorporating noun phrase movement and a bound trace. Because the trace is structurally realized, it strengthens its antecedent, resulting in faster response times in the probe task. Thus we have seen that the probe task is sensitive to coreference relations with both overt anaphoric expression and phonologically null anaphors.

Consonant with our hypothesis concerning the asymmetry between structurally realized positions and implicit arguments, the results in the agent probe condition showed no differences in response times across any of the three sentence types. These

results follow most naturally if structurally realized positions and implicit arguments are handled at different levels of language processing. Only structurally realized positions are handled at the level of processing tapped by the probe task, and thus only the antecedent of the NP trace showed faster response times in the verbal passive condition.

While conclusions from any single experiment must remain tentative, we can rule out several alternative explanations of the results with an examination of some of the experimental controls employed. First, the data cannot be explained by general differences in processing load across the three sentence types, because reading time and question-answering time did not differ across the conditions. The overall reading rate for the second sentence in the passage was about three words per second, a very normal rate. This result discounts any explanations based on strategies in two ways. First, the reading rate indicates that the subjects' reading was not disturbed by the fact that they expected a probe word at the end of the second sentence. Second, the fact that the subjects read the sentences at a normal rate indicates that they they did not spend time developing special strategies to predict what probe might appear.

Finally, one might appeal to some principle of the discourse that could produce these effects for the binding probe. Such an explanation is refuted by the lack of any effect for the agent probe. The implicit agent is presumably highly salient in the discourse representation of the passage, and therefore it is extremely difficult to imagine a discourse representation in which the theme but not the agent was strengthened in the verbal passive condition. Given all these controls, it is highly unlikely that the obtained results could stem from differences in processing difficulty for the three sentences, from any peculiar reading or responding strategies that subjects might have developed, or from some principle governing the construction of discourse representations.

These results are quite promising and suggest several avenues for future research. First, we can investigate further how the language processor initially assigns a structure to adjectival passives--recall that in the research reported here, response times in the adjectival passive condition were intermediate between adjective and verbal passive response times. It is tempting to suggest that subjects sometimes posit a verbal passive structure (with a bound element) for the adjectival passives. Such speculation based on null results is rather dangerous; recall that adjectival passives did not differ significantly from the other two conditions. It is not clear that this strategy of positing the structure containing the trace as a first resort (and later deleting the trace if necessary) is necessarily preferable to a strategy where a trace is not assumed initially (so that in some cases the language processor would

later have to go back over the structure and insert a trace).[4] If the processor does use a first resort strategy, however, it may be possible to change the processor's performance by giving it early information indicating that a passive construction is unambiguously adjectival or verbal. For example, in Hebrew adjectival passives and verbal passives are morphologically distinct (Borer and Wexler, 1987). If perceivers have morphological information available to distinguish the two types of passives, then we should not see any differences between adjectives and adjectival passives. If the slight differences found in the study reported here were reproduced in Hebrew, then some other aspects of processing of passive constructions would need to be explored.

Second, the effect should be investigated with binding in other syntactic constructions, and in fact recent work has found similar effects with NP-trace in raising and control constructions, as well as the verbal passive construction (Bever and McElree, 1988). The effect should not be limited to empty categories produced through NP-movement, and so the null pronominal PRO and the trace of wh-movement should also produce effects. Bever and McElree (1988) found that PRO produced similar, but weaker, effects compared to constructions containing an overt pronoun or an NP-trace. Using a somewhat different experimental method, Clifton and Frazier (to appear) have found tentative evidence that wh-gaps also produce strengthening effects.

Finally, we can further examine the asymmetry between implicit arguments and structurally realized positions. In the experiment reported above, the implicit argument was always in the first sentence, while the antecedent of the null anaphor was always in the second sentence. Our hypothesis predicts similar results when the overt argument and the implicit argument are in the same sentence. For example, in (16a) the NP "an attempt at the climb" contains an implicit agent; that is, the agent thematic role is not mapped onto a structural position (Williams, 1987; Roeper, 1985). Note that in this example, "John" is understood as the agent of "an attempt". In its strongest form, our hypothesis predicts no strengthening effect for "John" by virtue of its

---

[4] The Marcus Parser (Marcus, 1980) uses this strategy, where morphological information indicates the presence of a passive, so that traces are not posited after adjectives. For a discussion on the general desirability of positing gaps as a first (or last) resort, see Clifton and Frazier (in press) and references cited there.

controlling the implicit agent of "attempt." In contrast (16b) contains a structurally realized PRO, which is controlled by "John". Our hypothesis predicts a measurable strengthening effect for "John". Bever and McElree (1988) have found this effect for controlled PRO; however the processing of sentences such as those in (16a-b) have yet to be compared.

16a. John made an attempt at the climb.
b. John wanted PRO to attempt the climb
i i

To conclude, we have found an asymmetry between the processing of structurally realized and structurally unrealized (implicit) arguments. The psycholinguistic evidence provides interesting new empirical support for trace theory. Furthermore, the data provide some evidence in support of the linguistically interesting hypothesis that empty categories differ from phonologically realized positions only in lacking phonetic content.

References

Bever, T.G., and McElree, B. 1988. "Empty Categories Access their Antecedents During Comprehension," Linguistic Inquiry 19, 35-45.

Borer, H. and Wexler, K. 1987. "The Maturation of Syntax," In T. Roeper and E. Williams, (Eds.) Parameter Setting. Dordrecht, The Netherlands: Reidel.

Bouchard, 1984. On the Content of Empty Categories. Dordrecht, The Netherlands: Foris.

Burzio, L. 1986. Italian Syntax. Dordrecht, The Netherlands: Reidel.

Chang, F.R. 1980. "Active Memory Processes in Visual Sentence Comprehension: Clause Effects and Pronominal Reference," Memory and Cognition 8, 58-64.

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

Clifton, C., and Frazier, L. To appear. "Comprehending Sentences with Long-distance Dependencies," In G. Carlson and M. Tanenhaus, (Eds.) Linguistic Structure and Language Processing. Dordrecht, The Netherlands: Reidel.

Dell, G.S., McKoon, G., and Ratcliff, R. 1983. "The Activation of Antecedent Information During the Processing of Anaphoric Reference in Reading. Journal of Verbal Learning and Verbal Behavior 22, 121-132.

Kučera, H., and Francis, W.N. 1967. Computational Analysis of Present-day American English. Providence: Brown University Press.

Leiman, J.M. 1982. A Chronometric Analysis of Referent Assignment

to Pronouns. Unpublished Ph.D. dissertation, Wayne State University, Detroit.

Levin, B., and Rappaport, M. 1986. "The Formation of Adjectival Passives," Linguistic Inquiry 17, 623-661.

MacDonald, M.C. 1986. Priming in Sentence Processing: Facilitation of Responses to a Noun from a Coreferential Pronoun. Unpublished Ph.D. dissertation, UCLA.

Marcus, M. 1980. A Theory of Syntactic Recognition for Natural Language. Cambridge, MA: MIT Press.

Roeper, T. (1985). "Implicit Arguments, Implicit Roles and Subject/Object Asymmetry in Morphological Rules". Unpublished ms. UMass.

Wasow, T. 1977. "Transformations and the lexicon" In P. Culicover, T. Wasow, and A. Akmajian, (Eds.) Formal Syntax. New York: Academic Press.

Williams, E. 1987. "Implicit arguments, the Binding Theory, and Control," Natural Language and Linguistic Theory 5, 151-180.