The Role of Structure in Coreference Assignment During Sentence Comprehension

Janet Nicol[1] and David Swinney[2]

This paper examines the role of syntactic constraints on the reactivation and assignment of antecedents to explicit and implicit anaphoric elements during sentence comprehension. Evidence from on-line studies examining the time course of coreference processing supports the view that reactivation of potential antecedents is restricted by grammatical constraints when they are available. When structural information cannot serve to constrain antecedent selection, then pragmatic information may play a role, but only at a later point in processing.

A fundamental issue in understanding the general nature of human sentence comprehension is that of detailing characteristics of the process that assigns coreferential relationships among sentential elements. This paper presents studies examining the temporal course of the assignment of antecedents to both explicit anaphoric elements (such as pronouns and reflexives) and implicit anaphoric elements (the TRACE-related "gaps" identified with NP-movement and the obligatory PRO "gaps" related to certain verb forms) during sentence comprehension. Our focus throughout presentation of this work is on the role of syntactic constraints on this coreferential process.

The basic processing issues concerning coreference assignment revolve around the questions of *when* and *how* (and *if*) such relationships are established between referentially dependent elements and their

[1]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology.
[2]Linguistics and Psychology, Graduate School and University Center, City University of New York, 33 West 42nd Street; New York, New York 10036-8099.

antecedents. Thus, as a first-pass approach to the problem it is necessary to determine whether the processing device determines coreference relationships immediately (during initial perceptual processing) and, if so, whether it does so by actively linking the antecedent to the reference-dependent element, or by some much more subtle (perhaps postperceptual) process.

There is a growing literature concerning coreference processing, a literature filled with both theoretical and methodological issues. However, one consistent finding of recent studies is that reference-dependent sentential elements appear to cause reactivation of the antecedent noun phrases to which they refer (e.g., Corbett & Chang, 1983; MacDonald & MacWhinney, 1988; Bever & McElree, 1988; Tanenhaus, Stowe, & Carlson, 1985; Tanenhaus, Carlson, & Seidenberg, 1985). In general, these studies have tested for evidence of this ''reactivation'' at the end of the sentence containing the reference-dependent element, and hence at a point in time somewhat after that element was processed. (The major exceptions to this trend are studies by Tanenhaus, Stowe, & Carlson, 1985, who used a word-by-word reading task, and Garnsey, Tanenhaus, & Chapman, 1989, who used an evoked potential anomaly technique). Even the end-of-sentence studies, however, have found evidence of increased levels of activation (reactivation) of the antecedent to a reference-dependent element.

Thus, while these studies answer the initial question concerning the psychological and perceptual reality of coreference assignment, they raise an even more detailed and interesting set of issues concerning this process. In particular, these findings are compatible with three quite different accounts of the reactivation process: (1) A referentially dependent NP may cause reactivation of *all* previously mentioned NPs, which would then be considered for their suitability as antecedent. (2) A referentially dependent NP may give rise to reactivation of all and only those NPs that bear the appropriate structural relation to the dependent item. [A variant of this hypothesis, involving the notion of interpretative strategies or heuristics, might suggest that where two or more referents are structurally permissible as antecedents, there will be reactivation of only one—the one that is nearest to the referentially dependent element (see, for example, Frazier, Clifton, & Randall, 1983) or the one that carries a parallel grammatical function (Grober, Beardsley, & Caramazza, 1978)]. (3) A dependent NP causes only *the* best-fitting preceding NP to be reactivated.

These accounts differ primarily in their characterization of the time course of information availability and the nature of information interac-

tion (constraint) on the activation and assignment process. At some point during coreference processing, information about semantics, pragmatics, and discourse context, on the one hand, and sentence structure, on the other, will all come into play. The three alternatives give above make two basically different predictions about when and how these types of information have their effect. Under one view, such information will eliminate inappropriate NPs from a candidate set of activated NPs, and under the other, this information will act to restrict the candidate set by constraining which prior NPs are reactivated.

There is some evidence that pragmatic information does *not* constrain initial reactivation of a candidate set of antecedents. Corbett and Chang (1983) presented subjects with two-clause sentences containing two potential antecedents in the first clause and a subject pronoun in the second. The correct antecedent for the pronoun was either the first- or second-mentioned name in the first clause; which of the two was most appropriate depended entirely on pragmatic factors. For example, in the sentence "Ellen aimed a pistol at Harriet, but she did not pull the trigger," "she" most likely refers to "Ellen"; clearly, what guides this judgment about antecedence is our knowledge about the world (e.g., that pistols have triggers and that one must be holding the pistol to pull the trigger). An end-of-sentence name-verification task indicated that reaction times were facilitated to *both* potential antecedents of a pronoun, not just to the appropriate one. This finding, clearly incompatible with the third account described above, supports either of the first two: Either all previously mentioned NPs are reactivated or all and only those NPs that are in an appropriate structural position to act as antecedent are activated.

The analysis of experimental evidence presented below speaks to the question of whether or not syntactic constraints on coreference may act to restrict the reactivation of prior NPs. In all, we feel that the evidence provides strong support for the notion that reactivation is restricted by grammatical constraints: The initial set of candidate antecedents contains all and only those referents that bear the appropriate syntactic relation to the referentially dependent NP.

EMPIRICAL EVIDENCE FOR SYNTACTIC CONSTRAINTS ON COREFERENCE ASSIGNMENT

The research reviewed here is largely restricted to studies that use on-line techniques to measure the temporal course of reactivation of antecedents to referentially dependent elements. Compared to end-

of-sentence probes, on-line measures tend to be less prone to conscious reflection on the part of the subjects, since the subjects are engaged in listening to the sentence while they respond to the task. Further, while postsentence tasks (such as probe recognition) may reflect either reactivation that has occurred automatically in response to the presence of a referentially dependent NP or reactivation brought about by the probe recognition task itself (which promotes a memory search), on-line tasks appear more likely to tap automatic, unconscious processes.

We will first review the reactivation findings for wh-traces, NP-traces, and overt anaphors. These are the cases for which the syntax provides informational constraints that could allow choice of a unique antecedent. Then we will examine the reactivation patterns in sentences in which there are several potential antecedents (i.e., fewer constraints): sentences containing pronouns, and PRO.

Wh-trace

A number of studies have examined the reactivation of an antecedent in sentences containing wh-trace.

Using a cross-modal paradigm, Swinney, Ford, Frauenfelder, and Bresnan (1988) examined reactivation patterns in response to wh-traces contained in relative clause constructions. Sentences were presented auditorily, and subjects were asked to make a lexical decision to visually presented word/nonword targets. The appearance of the targets on a CRT screen coincided with the portion of the sentence in which the trace would be represented—test point 2 in the following sentence—or with control-test positions (t indicates "trace" and * indicates the probe point; subscripts indicate coreference).

(1) The policeman saw the boy$_i$ that$_i$ the crowd at the party
    $\overset{*}{1}$ accused t_i $\overset{*}{2}$ of the $\overset{*}{3}$ crime.

In the critical cases, the visual target was either an associate of one of the previously mentioned NPs or a control word matched in length and frequency and a priori reaction time to the associate. (Re)activation of an antecedent was indicated by a comparison of reaction time to make a lexical decision to a particular associate and to its control word. A significant priming difference between them was taken to indicate activation of that prior referent. (Note that although the antecedent of the trace in this instance is the relative pronoun *that*, rather than the head of the relative, *boy*, it is reasonable to assume that, since they corefer, the semantics of the head are inherited by the relative pronoun. Therefore,

Table I. Priming Scores (Lexical Decision RTs to Control Words Minus RTs to Semantically Related Words for Each Referent) at Each Probe Point

Referent	Probe point		
	1	2	3
Boy	12	27[a]	27[a]
Crowd	44[a]	19	8

[a] $p < .05$.

probing with an associate of *boy* ought to provide a measure of priming for the antecedent.) To ensure that any finding of activation actually signifies reactivation of a referent, there were two additional probe locations: just prior to the verb, *accuse,* and 300 msec following the initial probe point.

Table I displays the results from this experiment. (Values given here and in subsequent tables are magnitude of priming measures given in milliseconds—i.e., the difference between the mean response time for an unrelated matched control target and that for related targets. Significant effects at least at the .05 level are shown by a superscript.)

Note that priming for the correct antecedent of the trace *(boy)* becomes significant only after the embedded verb, where the representation contains a trace. There is significant activation of *crowd* only at a point several hundred milliseconds after the actual occurrence of *crowd* in the sentence; after this point, it diminishes. This suggests that the significant priming of *crowd* at the first probe point was simply due to residual activation for that word. A comparison of the patterns of priming for the two referents indicates that a trace does give rise to a reactivation of its antecedent *(boy)*. Note that since there was no probe to examine activation of *policeman* in this experiment, one cannot conclude that a trace reactivates just its antecedent. However, a follow-up study did demonstrate reactivation of *only* the appropriate antecedent for trace. Swinney et al. (1988) used a naming paradigm to examine priming for all three potential referents in these sentences. Here, again, priming for *crowd* was nonsignificant, as was priming for *policeman*; only *boy* was significantly activated.

This result, that a *wh*-trace triggers reactivation of its antecedent, has been demonstrated by a number of other investigators. Tanenhaus, Carlson, and Seidenberg (1985) report a study that used a cross-modal "rhyme priming" methodology. This study looked at reactivation of the

antecedent of *wh*-trace, again in relative clauses. Two versions of each experimental sentence were created, such as "The man was surprised at which beer/wine the judges awarded the first prize to *t*." Each version was followed by a target that rhymed with one of the antecedents of trace; in this case, for example, the word *fear* was displayed. Response latencies to a lexical decision on *fear* was compared for the two sentence versions to obtain a measure of priming. Their hypothesis is this: If the end-of-sentence trace causes the antecedent to be reactivated—specifically, if the trace causes the phonological properties of the antecedent to be reactivated—then reaction times ought to be facilitated when the target rhymes with the antecedent, as compared with instances in which the target does not rhyme with the antecedent. Thus, responses to *fear* were expected to be faster when the antecedent is *beer* than when the antecedent is *wine*. Their results support their hypothesis.

Reactivation of an antecedent triggered by a *wh*-trace obtains even when the antecedent is semantically inappropriate. Using the cross-modal priming paradigm, Nicol and Osterhout (1988) examined reactivation patterns in sentences such as "That's the actress that the dentist from the new medical center in town had planned/invited to go to the party." In both types of sentences, there was priming for *actress* immediately following the embedded verb, suggesting that a *wh*-trace is postulated following verbs such as *invite* (after which a trace must be represented, since this verb is transitive), as well as after verbs such as *plan* (which optionally take an NP complement). Hence, despite the fact that one cannot "plan the actress," this referent is reactivated. These results are presented in greater detail below in the discussion of PRO.

Further support for this general finding of reactivation for antecedents to *wh*-traces comes from a study by Garnsey et al. (1989), who measured evoked brain potentials of their subjects during the controlled serial visual presentation of sentences containing relative clauses. They exploited the fact that the implausibility of a word in context results in a particular EP pattern (N400) by constructing sentence pairs that contained as the antecedent of the *wh*-trace an item that was either contextually plausible or implausible, such as in (2):

> (2) a. The mother found out which book$_i$ the child read t_i in school.
>   b. The mother found out which food$_i$ the child read t_i in school.

They created, in addition, a pair of control sentences that did not contain a trace:

> (3) a. The mother found out whether the child read the book in school.
>   b. The mother found out whether the child read the food in school.

They predict that the appearance of the contextually implausible object, *food,* in (3b) should produce N400, whereas the plausible object, *book,* should not. As for the relative clause sentences (2a) and (2b), they expect that if the antecedent of the *wh*-trace is reactivated immediately, then N400 ought to occur after the verb *read* in (2b) only. Their data show the N400 as predicted. Thus, this study provides additional evidence that *wh*-traces reactivate their antecedents. It also has interesting implications for the time course of information use within the processing system. Given the above finding by Nicol and Osterhout, it would be expected that the implausibility of the antecedent would not prevent it from being reactivated. If it had, there would be no N400 effect in (2b). Yet the presence of N400 at this point indicates that once reactivated, the anomaly is registered, and further, that it is registered relatively quickly.

In sum, these studies indicate that a *wh*-trace triggers (1) immediate reactivation of its antecedent, (2) reactivation of only the correct antecedent (in terms of *structural* constraints), and (3) reactivation of the antecedent independently of plausibility considerations. This suggests that information concerning argument position and the requirement that NPs receive theta roles is used to restrict reactivation of potential antecedents.

NP-trace

There is some evidence that NP-trace may trigger automatic reactivation of antecedents in a fashion similar to that for *wh*-trace. However, the issue is not nearly as clear-cut as in the case of *wh*-trace, and the pattern of reactivation, if it exists, is considerably different from that found for *wh*-trace. There are two major pieces of evidence about NP-trace—the work by Bever and McElree (1988) and work by Osterhout (1988, personal communication). Bever and McElree have shown that passive sentences yield results similar to those obtained for sentences containing other types of referentially dependent NPs—namely, reactivation of the appropriate antecedent. However, their data involve an end-of-sentence word recognition probe. Osterhout used a somewhat more on-line cross-modal technique in probing for activation of antecedent to the NP trace in a passive sentence (e.g., "The dentist$_i$ from the new medical center in town was invited t_i by the actress to go to the party") and found no significant priming in the trace position. He did, however, find a nonsignificant trend for priming at a test probe point 1,000ms downstream from the trace position. Thus, while there may be some reactivation of possible antecedents, it appears to take place in a manner

Table II. Priming Scores (Lexical Decision RTs to Control Words Minus RTs to Semantically Related Words for Each Referent) for Parallel Anaphor and Pronoun Studies

Referent	Anaphor	Pronoun
Boxer	−1	43[a]
Skier	11	58[a]
Doctor	104[a]	−21

[a] $p < .05$.

quite different from that for *wh*-trace. In all, however, the answer to what is going on in these NP-trace conditions is not clear, and the issue is in need of much further work.

Overt Anaphors versus Pronouns

We now turn to evidence concerning the *overt* counterpart to the NP-trace. There is evidence that an overt anaphor (the reflexive) also triggers reactivation of its antecedent. Nicol (1988) examined the reactivation patterns of referents in sentences such as the following:

(4) The boxer told the skier that the doctor₁ for the team would blame himself₁ * for the recent injury.

A cross-modal priming experiment revealed that immediately after the anaphor there was significant priming of the antecedent, *the doctor*, but nonsignificant priming of *the boxer* and *the skier*. This finding is compatible with two possible explanations: (1) The reactivation of prior referents in response to the anaphor is restricted *only* to the binder of the anaphor; (2) the priming effect found here could simply reflect the residual activation of *the doctor*. However, this second possibility is nullified by the following. Sentence (4) above was contrasted with (5), in which the embedded clause contains a pronoun:

(5) The boxer₁ told the skier₁ that the doctor for the team would blame him₁ * for the recent injury

Again, all three referents were tested. Magnitude of priming for each referent within both the pronoun sentences and the anaphor sentences is given in Table II.

Results for the pronoun sentences indicate *no* priming for *doctor* but significant priming for both *boxer* and *skier*. The pattern obtained for

Table III. Priming Scores (Lexical Decision RTs to Control Words Minus RTs to Semantically Related Words for Each Referent) at Each Probe Point

Referent	Probe point		
	1	2	3
Boxer	23	20	51[a]
Doctor	32	56[a]	42[a]
Swimmer	9	0	20

[a] $p < .05$.

both sentence types suggests that the activation of a candidate set of antecedents occurs in compliance with syntactic (binding) constraints. It is important to note in this instance that Binding theory does not dictate the precise pattern of reactivation in this experiment; hence, a finding of no reactivation of a prior referent is compatible with the constraint that a pronoun must be free in a local domain.

Evidence for this position is also provided by Swinney, Ford, Bresnan, Frauenfelder, and Nicol (cited in Swinney et al., 1988), who examined reactivation patterns in sentences such as (6).

(6) The boxer$_i$ visited the doctor$_k$ that$_k$ the swimmer at the competition had $*_1$ advised $*_2$ him$_i$ $*_3$ to see t_k about the injury.

In this sentence, *boxer* is the only structurally appropriate antecedent of the pronoun *him*. The second NP, *the doctor,* is coindexed with the relative pronoun, which must bind a trace rather than an overt element. Finally, the third NP, *the swimmer,* cannot corefer to the pronoun since it is the subject of the clause in which the pronoun appears.

Patterns of priming for each potential antecedent varied according to the point in the sentence at which it was tested, as shown in Table III.

Note that activation levels for the appropriate antecedent, *boxer,* reach significance only at the third probe point—i.e., just after the pronoun to which it refers. There is significant priming for *doctor* just after the embedded verb, with continued priming after the pronoun. There is no significant priming for *swimmer* at any point. For the referents *boxer* and *swimmer,* this pattern of priming conforms to the hypothesis that syntactic constraints restrict which referents are reactivated. However, the priming found for *doctor* at probe points 2 and 3 require an explanation. Given the results of studies that examined reactivation of referents at the point of a *wh*-trace, and the fact that there is priming for *doctor* prior to the occurrence of the pronoun, it seems reasonable to

assume that a trace has been erroneously postulated following the embedded verb, and that the only appropriate antecedent for this trace is consequently reactivated. Suppose that the experimental sentences had been constructed as follows: "The boxer visited the doctor_i that the swimmer at the competition had advised t_i about his backstroke." In such cases, there would be a trace following the verb, and the trace would be linked to *doctor*. If traces are postulated immediately upon processing the embedded verb, no matter what follows, then in sentences such as (6) one would expect a trace to be represented after the verb, *advise,* and thus would expect the antecedent of the trace to be reactivated. The overall pattern of priming for these referents raises an interesting question. While the account presented here concerning the activation of *doctor* at points 2 and 3 provides a sufficient explanation for the data, it implicitly assumes that different syntactic constraints are simultaneously taken into account. Consider the following. Suppose that a pronoun triggers reactivation of all prior NPs except the local c-commanding NP. In the sentence above, both *boxer* and *doctor* would be reactivated at point 3. Suppose also that a trace is erroneously postulated following the verb: *Doctor* would be reactivated in response to the postulation of the *wh*-trace. The overall pattern of priming would look exactly like what was obtained. Hence, it is unclear, at this point, whether a pronoun gives rise to the reactivation of all prior NPs except the local c-commanding NP or to the reactivation of only those NPs that are not coindexed to some other referentially dependent element. In other words, this result leaves open the question whether different types of constraints are consolidated prior to reactivating referents.

PRO

The results for the pronoun cases suggest that there is multiple activation of referents in sentences in which there is more than one referent that may be coindexed with the pronoun. In other words, where semantic or pragmatic information is necessary to determine which of two or more referents is the actual antecedent of the pronoun, all candidates are reactivated for consideration. Recall that in the discussion of PRO, above, the claim was put forth that coreference of PRO with an antecedent is often dependent on semantics. Hence, one would expect activation of multiple referents in response to PRO. The following studies address this question.

Using the cross-modal technique, Nicol and Osterhout (1988)

Table IV. Priming Scores (Lexical Decision RTs to Control Words Minus RTs to Semantically Related Words for Each Referent) at Each Probe Point for Transitive and Intransitive Verbs

Referent	Transitive verbs (e.g., *invite*) Probe point			Intransitive verbs (e.g., *plan*) Probe point		
	1	2	3	1	2	3
Actress	30	75[a]	57[a]	12	71[a]	48[a]
Dentist	6	4	12	28	36	21

[a] $p < .05$.

examined priming of both previously mentioned referents in sentences such as (7).

(7) a. There is the actress$_i$ that$_i$ the dentist from the new medical center in town $*_1$ had invited t_i PRO$_i$ to $*_2$ go to $*_3$ the party.
   b. There is the actress$_i$ that$_i$ the dentist$_k$ from the new medical center in town $*_1$ had planned PRO$_k$ to $*_2$ go to $*_3$ the party with t_i.

In (7a), there is a transitive verb in the embedded clause for which the direct object is a trace of the relative pronoun. PRO in this sentence is coindexed with the nearest NP, the trace. Thus, PRO is identified with *the actress*. In (7b), however, the embedded clause contains an intransitive verb, and PRO is therefore identified with the preceding subject, *the dentist*. These sentences were designed to test the on-line application of processing strategies in determining the antecedent of PRO, in response to Frazier et al. (1983). They hypothesized that control information would be unavailable to the parser during the initial analysis of such sentences. Hence, a strategy, which they call the Most Recent Filler Strategy, is invoked to assign PRO an antecedent: PRO is linked to the most recent potential antecedent. They would expect the results of this experiment to show priming of *dentist*, and only *dentist*, in both sentence types at the appropriate point in the sentence (i.e., probe point 2). The patterns of priming displayed in Table IV were obtained.

Considering, first, sentences containing transitive verbs in the embedded clause, only *actress* becomes reactivated after the embedded verb. There is no significant priming for *dentist* at any of the probe points, indicating that there is no evidence that the Most Recent Filler Strategy is applied on line. The result for transitive verb sentences, taken in isolation, suggests that only the correct antecedent for PRO is

reactivated. However, this conclusion is not supported by the data for sentences containing intransitive verbs. In the intransitive verb sentences, only *dentist* should be activated, since *dentist* is the antecedent of PRO. Instead, exactly the same pattern of priming was obtained for the intransitive verb sentences as for the transitive verb sentences. The activation of *actress* following the transitive verb is understandable if the effect is interpreted as a response to the *wh*-trace that is represented after this verb. As shown above, *wh*-traces do give rise to reactivation of the head of the relative clause—in this case, *actress*. However, there is no trace following the intransitive verb, so *actress* should not be reactivated in such sentences. Considering only probe points 2 and 3, a possible explanation is that *actress* is in a salient position within the sentence; hence, the activation of *actress* is maintained throughout. However, the lack of significant priming for *actress* at the first probe point does not support this hypothesis. It seems reasonable, instead, to argue that a trace is postulated following intransitive verbs as well as transitive verbs. Notice that verbs such as *plan* do take an NP complement, but they normally take inanimate NPs, not animate ones (as in "plan the trip"), so that in sentence constructions such as those used in this experiment, these verbs appear to be intransitive. The set of intransitive verbs that appear in the experimental sentences actually consist of equal numbers of quasi-intransitive verbs like *plan,* and true intransitive verbs such as *hesitate,* which do not take NP complements of any kind. A follow-up study (Nicol, 1988) that contrasted these two types of intransitive verbs indicates that it is only after the quasi-intransitives that there is significant priming of the head of the relative (e.g., *actress*). Thus, it appears that the activation of *actress* following intransitive verbs in this study is actually due to postulation of a trace in a subset of these verbs. Since PRO is always linked to the *wh*-trace in such constructions, it is unclear whether or not the activation of *actress* also corresponds to the presence of PRO. It is necessary to look at sentences in which PRO and trace are not contiguous in order to examine reactivation in response to PRO. The next study addresses this issue.

　　　Again using the cross-modal technique, Osterhout and Nicol (1988) presented sentences that were composed of a matrix clause containing two potential antecedents and an infinitival clause. The matrix clause was either in the active voice or the passive voice; hence, the linear position of the correct antecedent varied. Within each sentence, referents were probed at five points, as shown in (8). The first point was just after the *to* in the infinitive; with respect to this point, the next four probes

Table V. Priming Scores (Lexical Decision RTs to Control Words Minus RTs to Semantically Related Words for Each Referent) at Each Probe Point

Referent	Test point				
	1	2	3	4	5
	Test of materials of the type exemplified in sentence 8a				
Actress	21	0	28	24	31[a]
Dentist	−10	15	37[a]	77[a]	6
	Test of material of the type exemplified in sentence 8b				
Actress	19	−8	29	32[a]	40[a]
Dentist	19	31	48[a]	2	−3

[a]$p < .05$.

occurred 500 msec, 1,000 msec, 1,200 msec, and 1,500 msec downstream.

(8) a. The actress invited the dentist$_i$ from the new medical center PRO$_i$ to $\dagger$ go to the pa$\overset{*}{2}$rty at the $\overset{*}{3}$ mayor's $\overset{*}{4}$ house $\overset{*}{5}$.

   b. The actress$_i$ was invited by the dentist from the new medical center PRO$_i$ to $\dagger$ go to the pa$\overset{*}{2}$rty at the $\overset{*}{3}$ mayor's $\overset{*}{4}$ house $\overset{*}{5}$.

In (8a), the antecedent of PRO is *dentist*; in (8b), it is *actress*. The activation patterns for these two referents within each sentence type are displayed in Table V.

At the first two probe points, there is no significant priming for either antecedent. At the third point, only the "recent filler" (*dentist*) is significantly primed in either sentence type. At the fourth position, only the actual antecedents are significantly activated: *dentist* in the active sentence, and *actress* in the passive sentence. At the final probe point, only the "distant filler" is primed in either type of sentence. First, these results confirm the above finding that the Most Recent Filler Strategy does not appear to be invoked during the processing of sentences containing empty subjects. In addition, these results support the hypothesis that all structurally appropriate referents are reactivated. With respect to activation of multiple referents, PRO is similar to pronouns, and this is exactly what was predicted. Yet reactivation patterns in response to PRO differ in two ways from the pattern found for pronouns: (1) For PRO, it appears that referents are reactivated in right-to-left fashion, while there was simultaneous activation of potential antecedents of the pronoun; (2) there was no significant activation of an antecedent for PRO until a position 1,000 msec after the initial probe. Unlike the ambiguity

in structure of the passive construction, which could account for the lag in the reactivation of an antecedent, these constructions are not ambiguous. It is unclear why there should be this difference in the time course of reactivation patterns in response to pronouns versus PRO. There are obvious differences between the two NP types, in terms of syntactic distribution, phonological realization, and constraints on coreference; yet it is not at all evident why such differences would give rise to this particular effect.

SUMMARY

The results of on-line studies show that in sentences in which the grammar dictates a unique antecedent, there is immediate reactivation of only that antecedent, and in sentences where more than one referent may bind the referentially dependent element, all such referents are reaccessed. Thus, wh-traces and overt anaphors immediately trigger reactivation of only the referent that must bind the antecedent. The findings concerning NP-traces in passive sentences are considerably less clear at this time, particularly given the lack of effects when examined with on-line tasks. It was also found that, with respect to reactivation patterns, pronouns cluster with PRO: They both cause reactivation of multiple referents. This is an expected result for any theory holding that coreference processing is subject to structural constraints, in that structure does *not* dictate a unique referent in these conditions. In this regard, it is important to note that the subject of the clause in which the pronoun appears is not reaccessed (reactivated), something that *is* in keeping with structural constraints of binding theory. As for PRO (the condition for which structure provides the weakest information concerning the actual antecedent of the gap), both potential antecedents are reactivated.

In all, the experiments reviewed here strongly support the view that the reactivation of prior referents is restricted by grammatical constraints. In the case where such information does not sufficiently constrain the list of potential antecedents to a single one, the pragmatic and other sentence/discourse processing procedures undoubtedly come into play, but, given the present evidence, only at a *later* point in processing. In many ways the conditions where multiple antecedents are activated remind one of the conditions the sentence processor faces in dealing with multiple meanings of words; it appears that initial perceptual processing activates all viable candidates (here, those that conform to grammatical

constraints) and the choice among the candidates is relegated to later, perhaps nonmodular (see Fodor, 1983) language processing.

ACKNOWLEDGMENT

REFERENCES

Bever, T.G., & McElree, B. (1988). Empty categories access their antecedents during comprehension. *Linguistic Inquiry, 19,* 34–43.

Corbett, A.T., & Chang, F.R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory and Cognition, 11,* 283–294.

Fodor, J.A. (1983). *The modularity of mind.* Cambridge, MA: M.I.T. Press.

Frazier, L., Clifton, C., & Randall, J. (1983). Filling gaps: Decision principles and structure in sentence comprehension. *Cognition, 13,* 187–222.

Garnsey, S.M., Tanenhaus, M.K., & Chapman, R.M. (1989). Evoked potentials and the study of sentence comprehension. *Journal of Psycholinguistic Research: Special Issue on Sentence Processing, 18,* 51–60.

Grober, E.H., Beardsley, W., & Caramazza, A. (1978). Parallel function strategy in pronoun assignment. *Cognition, 6,* 117–133.

MacDonald, M.C., & MacWhinney, B. (1988). *Effects of pronouns on nonreferents: When inhibition looks like facilitation.* Unpublished manuscript, Carnegie-Mellon University.

Nicol, J. (1988). *Coreference processing during sentence comprehension.* Doctoral dissertation, M.I.T.

Nicol, J., & Osterhout, L. (1988). *Reactivating antecedents of empty categories during parsing.* Unpublished manuscript.

Osterhout, L., & Nicol, J. (1988). *The time-course of antecedent activation following empty subjects.* Unpublished manuscript.

Swinney, D., Ford, M., Frauenfelder, U., & Bresnan, J. (1988). On the temporal course of gap-filling and antecedent assignment during sentence comprehension. In B. Grosz, R. Kaplan, M. Macken, & I. Sag (Eds.), *Language structure and processing,* Stanford, CA: CSLI.

Tanenhaus, M., Carlson, G., & Seidenberg, M. (1985). Do listeners compute linguistic representations? In A. Zwicky, L. Kartunnen, & D. Dowty (Eds.), *Natural language parsing: Psycholinguistic, theoretical, and computational perspectives.* London and New York: Cambridge University Press.

Tanenhaus, M., Stowe, L., & Carlson, G. (1985). The interaction of lexical expectation and pragmatics in parsing filler-gap construction. In *Proceedings of the Seventh Annual Cognitive Science Society Meetings,* pp. 361–365.