3

The Psycholinguistics of Anaphora

Janet L. Nicol and David A. Swinney

1 Introduction

In this chapter, we will review some of the research on the processing of anaphoric elements during language comprehension. We will focus primarily, though not exclusively, on the processing of reflexives and pronouns (which will be jointly referred to as *proforms*) within sentences, rather, than within discourse contexts.

The psycholinguistic study of anaphora is different from the formal analysis of anaphora, which we take to be, primarily, the characterization of the constraints on coindexation and coreference within a syntactic domain. The psycholinguist's approach is to explore how the sentence comprehension mechanism computes coreference during sentence processing. It is important to bear in mind that formal treatments capture anaphoric dependencies as they are represented within complete sentences, but for the processing system, a sentence may still be underway when a proform appears, so dependency relations must often be computed in the context of incomplete sentences. Another major difference is that while the linguist emphasizes the conditions under which coindexation may or may not occur, the psycholinguist is concerned with the consequences of coindexation, specifically: once it is determined that two elements are coindexed, how is this manifested psychologically during sentence processing?

We will start with the following assumptions about how sentence comprehension proceeds: (i) Incoming words are recognized, grammatical categories are identified, words are grouped into constituents, and constituents are linked. (ii) The lexical and phrase structure processes (or *parsing* processes) do not lag far behind the input. (iii) Lexical and phrase structure processing precedes anaphoric processing. For example, if the sentence input is *Mary told Susan about herself over dinner,* by the time *herself* appears in the input, the sentence has been structured so that *Mary* is in subject position (or *Spec of IP), Susan*

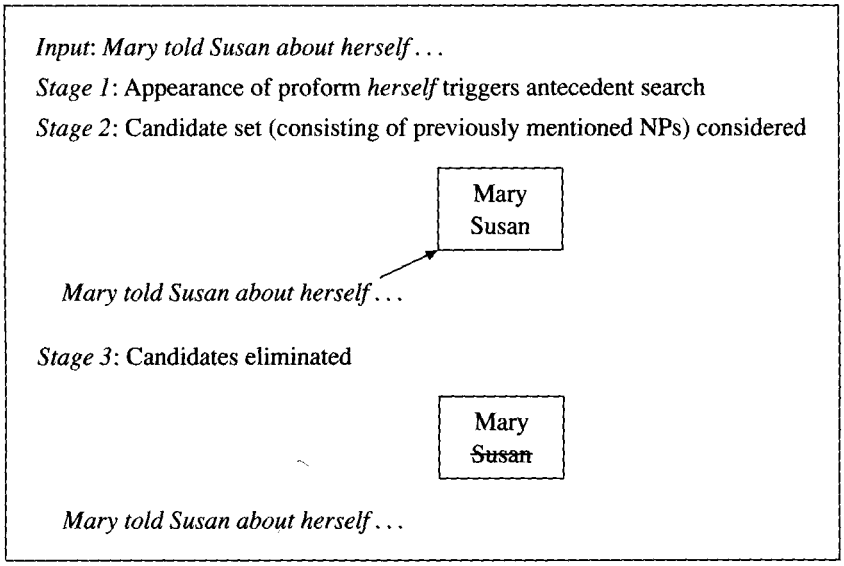

Figure 3.1 *Coreference processes*

is part of the VP, and *herself* is part of a following PP. (iv) The appearance of a proform causes a candidate set of antecedents *(Mary, Susan)* to be considered, from which all but one antecedent (probably *Mary)* is eliminated. We assume that these processes are unconscious except for the end product - e.g. the recognition that *herself* means *Mary*. The coreference processes are represented schematically in figure 3.1.

Associated with this proposed set of processes are a number of empirical questions:

i How quickly after the appearance of the proform is a search initiated?
ii What types of information constrain the initial candidate set?
iii What types of information constrain the elimination of NPs from the candidate set?

There are two observations we should make about this characterization of sentence comprehension. One is that we have omitted a number of details about parsing about which there is a great deal of debate. Some of these are: (a) the immediacy of attachment of words into an ongoing structure (e.g. word by word or phrase by phrase); (b) the level of detail about the attachment of a word or phrase into the already constructed syntactic structure (e.g. in the sentence example *Mary told Susan about herself,* it might initially be postulated that the PP simply *follows* the object NP, and is dominated by a sentence node (such as IP), as proposed by e.g. Gorrell 1995); and (c) the types of information which may affect attachment decisions in the case of an ambiguous structure (e.g. whether semantic factors may be taken into account during the initial construction of a structure).

iv Does coreference processing involve activation or inhibition of previously mentioned NPs?

v What is the nature of the NP representation that comes into play in coreference processing?

We will address these questions in sections 3-5. First, however, we will review some of the methods used to study sentence processing. It is important to understand how the experimental data are obtained in order to interpret some of the contradictory findings that have been reported in the literature. This review includes the techniques that are most commonly used in the study of on-line coreference processing; however, we note that there are a range of other techniques that have been employed, including Event-related Potentials (ERP; see, e.g. Osterhout and Mobley 1995), and various types of metalinguistic tasks, including antecedent identification (e.g. Caramazza, Grober, and Garvey 1977, Oakhill and Yuill 1986), sentence completions (designed to provide information about prior antecedent identification) (e.g. Garvey, Caramazza and Yates 1976), judgments of pronoun naturalness (e.g. Gernsbacher 1991), and other judgments (e.g. Oakhill et al. 1992, Garnham and Oakhill 1988, McKee, Nicol, and McDaniel 1993).

2 Psycholinguistics Methods

2.1 Reading time studies

This method tracks the time it takes to read sentences. There are two major variants: the self-paced reading technique and the eye-tracking methodology. In the former type, a participant in an experiment is presented with sentences on a computer screen. The sentences are presented word by word or phrase by phrase, and the participant controls the rate of presentation with the press of a button. Usually, only one fragment is visible at any one time; when a new sentence region appears, the previous one disappears. There is no looking back. Reading times (i.e. the interval between button presses) for relevant sentence sections are recorded. Clearly, the reading time for a given region of a sentence is in itself uninterpretable; one needs to be able to compare it to something. Typically, minimally different pairs of sentences are constructed, and reading times for similar regions within each sentence are compared. For example, one could compare reading times for the pronoun *she* in the sentences in (1):

(1) a. Sally told Bill that' *she* would be late.
b. Sally told Jill that *she* would be late.

If the reading times for the pronoun were significantly different in the two sentence types, one might draw inferences about the effect of a gender contrast in the set of potential antecedents for the pronoun, and the consequent ease of pronoun resolution. If the reading times differed, but not until several words after the pronoun, one might infer that pronoun resolution in such constructions is not immediate.

The same design principles hold for the second variant on this method, in which readers' eye-movements are recorded. This technique is more natural than the self-paced reading task because participants are able to look back at earlier portions of a sentence and also to get a peripheral glimpse of upcoming material. This technique provides much more information about what readers are doing: time spent in a particular sentence region may be recorded, and also regressive eye-movements from one region to another. Longer reading times and a greater number of backward eye-movements could be indicative of greater processing difficulty. Hence, if readers looked longer at the pronoun or words just following the pronoun, and/or showed a greater number of regressive eye-movements to previous regions in *the Sally-Jill* sentence than in the *Sally-Bill* sentence, then one could infer that pronoun resolution was more difficult in the sentences containing more potential antecedents. In addition, differences in reading times could be informative about the timing of pronoun resolution. For example, as with the self-paced methodology, differences in reading time could start at the point of the pronoun or begin a few words later.

A potential problem for both methods, however, with interpreting the timing of processing disruption is that there is some controversy about the extent to which the eye and mind are synchronized. A tight synchrony means that relative ease or difficulty in processing should be manifest when the eyes are fixated on a critical word, such as a pronoun (see e.g. Just and Carpenter 1980). An asynchrony suggests that eye-movements precede processing; the eye advances rightward to sample additional visual input while the mind continues to process the information (see e.g. Ehrlich and Rayner 1983). In the latter case, a lag in relative difficulty (or ease) of processing would be expected. In the case of a sentence containing a pronoun, if reading time differences do not emerge until several words after the pronoun, this could mean that the appearance of a pronoun does not immediately trigger the coreference process. Alternatively, it could mean that sentence processing in general sometimes lags behind the input: the eye could spend little time fixating on a pronoun because it is a short, high-frequency word.[2] We will return to this issue below.

[2]In fact, Ehrlich (1983) found that pronouns were directly fixated only 13 percent of the time. Hence eye-tracking studies which test pronoun resolution typically use fixations that occur just before and after the pronoun.

2.2 Probe verification studies

This method requires participants to indicate whether a specific word had occurred in a just presented sentence fragment, sentence, or discourse. Sentences are presented either visually (word by word or phrase by phrase) or auditorily, and at some point during the presentation of the sentence, a word appears for verification, such as in (2), where the probe point is indicated with an asterisk and the probe word appears in upper-case:

SALLY

(2) Jane told Sally that she * was in danger.

When sentences (and probes) are presented visually, the probe words usually appear in a different font, and/or in a designated position, and/or in a different color. In this example, the probe word *Sally* would be clearly distinguishable from the words in the sentence. Participants would be asked to determine whether *Sally* had appeared in the current sentence and to press the appropriate response key. Response times would be recorded. Again, pairs of sentences need to be compared in order for the response times to be meaningful. Hence, one might compare response times to *Sally* in each sentence of the pair in (3):

(3) a. John told Sally that she * was in danger.
b. Jane told Sally that she * was in danger.

The rationale behind the technique is that verification times will be affected by the prominence of the probed item in memory. Following the pronoun *she, Sally* could be more prominent in the *John-Sally* sentence than in the *Jane-Sally* sentence (if pronoun resolution is immediate) because readers might be more likely to interpret *Sally* as the antecedent of *she* in sentence (3a) than in (3b), since (3b) contains another potential antecedent *(Jane).*

2.3 Sentence-based semantic priming studies

This method allows precise details to be established concerning which words are activated at any particular time during the ongoing comprehension of a sentence. It has come to be particularly useful in examining ongoing auditory comprehension of fluent speech in a way that does not disrupt or change the comprehension process itself. Several variants of this approach exist, but cross-modal semantic priming (Swinney et al. 1979) has proven to be a particularly sensitive measure of moment-by-moment sentence processing, and we will demonstrate the general approach via a description of this task.

In this approach, experimental participants hear sentential or discourse material presented auditorily at normal rates of speech. Listeners are told that their primary job is to understand the sentence(s) or discourse they hear. (They are tested for their comprehension at various points throughout the experiment.) The listener also has an additional task to perform: at some point during the auditory input, a string of letters will appear on a computer screen and he or she has to make a rapid word/non-word decision about that letter string. On some of the trials in which the visual letter string forms a word, that word will be a semantic associate of a key word in the sentence that is being heard.

Thus, for example, the listener might hear the sentence:

(4) The teacher grabbed the frog from the young boy before it could get hurt.

Immediately after hearing the word *frog* in the sentence, the letter string TOAD, a semantically related target, might appear on the screen and participants be asked to make a lexical (word/non-word) decision. (Alternatively, an unrelated "control" word might appear; control words are matched to the related target words along variables that known to affect access time - length and frequency of occurrence - and are typically also pretested in isolation to ensure that a related target word and its matched control elicit equivalent response times.) Based on principles of automatic semantic priming (e.g. Neely 1991), activation of the word *frog* is reflected in decreased reaction time to make a lexical decision about the letter string *toad.* Thus, reaction time to classify *toad* as a word takes less time than to classify the unrelated control word, when participants respond just after hearing *frog.*

This "semantic priming" effect can thus be used to track activation of the words in a sentence. Priming effects caused by a word in a sentence typically last, all else being equal, on the order of 250-700 milliseconds (depending on such things as how rapidly new information is input, or whether the sentence contains other words related to the target item) .[3] Note that 250-700 msec. is roughly equivalent to one to three syllables.

[3]There are several aspects of this technique that we wish to stress. To begin with, presentation of auditory discourse/sentential material *always* continues to the sentence's end; it is never stopped or modified, even when the visual probe/target is presented. Additionally, the visual target never occurs at the end of the sentence. This prevents the visual target from being integrated into the ongoing sentential material (provided, of course, that the sentential material is presented at normal volume levels and speed - see further discussion by Swinney et al. forthcoming). Second, judgments about the visual targets never involve subjects making metalinguistic judgments about anything in the auditory material they hear (e.g. "Was the word `waitress' in the sentence?"). This prevents metalinguistic examination of the auditory sentence, processing which involves conscious (non-automatic) processes in addition to normal comprehension processes. Finally,

It is critical to our interests that this technique can be used to examine when and where the antecedents related to a proform are activated. Thus, for example, one can use this technique to see if and when a decision to *toad* is facilitated (relative to the control target) with respect to the occurrence of the pronoun *it* in the example above (in which the pronoun appears well after the activation of the antecedent would be expected to have decayed).

3 Coreference Processes: Timecourse and Mechanisms

3.1 When is a search for an antecedent initiated?

In principle, it is possible that the occurrence of a proform does not immediately trigger a search for an antecedent, since parsing - the computation of phrase structure - does not typically depend on having computed coreference, so parsing would not be held up by slow coreference processes.' More specifically, a pronoun or reflexive has the category "NP," which is sufficient for phrase structure building. Subsequent inferences about the identity of the proform may be postponed, or they may simply take time.

The literature on the processing of proforms has no single answer to the timecourse question. Several reading studies have shown no effects at the point of the pronoun of manipulations of prior context (e.g. Ehrlich and Rayner 1983, Smyth 1986, Stevenson and Vitkovitch 1986, Vonk 1985). For example, Ehrlich and Rayner (1983) used an eye-tracking methodology to examine coreference processing in short discourses in which the distance between the antecedent and pronoun was varied, as in (5), but in all cases, there was a unique antecedent for the pronoun (underlines have been added here to highlight the distance manipulation):

processing of the sentence is uninterrupted and "normal," at least up to the point where the visual target is presented - which is the final processing point examined in such studies. In this regard, the task differs considerably from most other on-line tasks, which often ask the subject to evaluate each word in a sentence successively, or to hold a search target in mind *while the* sentence is being processed, or try to remember if a search-probe target *was* in the sentence, which involves recall of the sentence in memory while examining its contents for the probe. Overall, the cross-modal semantic priming task is one of the most sensitive, least intrusive behavioral techniques available for the examination of the ongoing comprehension process.

" There are cases in which coreference may interact with parsing, as in examples like *I know the woman with the boy who bought herselhimself a soda.* In the *herself* case, the relative clause must attach high, since it modifies *the woman.* In the *himself* case, the relative clause modifies *the boy*. It is likely, however, that the attachment decision is made even before the reflexive appears; hence, the appearance of the reflexive might either confirm the attachment or trigger a structural reanalysis. In short, coreference could play an after-the-fact role in parsing, rather than a determining role.

(5) ***Intermediate*** A group of people who shared an interest in photography had recently started writing a newsletter of their activities. In fact, in one room Mark was mailing a copy of the paper to Susan. He was very involved in photography and spent every weekend taking pictures.

Far A group of people who shared an interest in photography had recently started writing a newsletter of their activities. Mark wrote most of the copy but the other members did a lot of work as well. In fact, in one room Cathy was mailing a copy of the paper to Susan. He was very involved in photography and spent every weekend taking pictures.

(Ehrlich and Rayner 1983:80 table 1)

Ehrlich and Rayner's hypothesis was that pronoun resolution would take longer when there was more (vs. less) intervening material between the pronoun and antecedent (the "far" vs. "intermediate" cases). They expected that if it were the case that readers fully process each word of an input text before advancing to the next word, the difficulty associated with the "far" condition would manifest itself while the readers were fixating on (or near) the pronoun. If readers do not fully process each word before moving on, then the distance manipulation should not have an immediate effect, but should show up downstream from the pronoun. Their results showed that the effect of resolution difficulty did not appear at the point of the pronoun, but did emerge on just-subsequent words. According to the authors, the "data suggest that although this [pronoun resolution] process may be started when the pronoun is within fixation, there is no evidence that assignment is necessarily completed within the same fixation, especially when the antecedent occurs some distance back in the text" (Ehrlich and Rayner 1983:84).

This description of the timing of pronoun resolution appears to apply to self-paced reading studies as well. Although this task may actually encourage readers to integrate sentence material in word-by-word fashion (because once the response key is pressed, the next word appears and the current word *disappears),* few studies actually show a significant effect of a relevant manipulation of prior context *at the point of the pronoun.* For example, Badecker and Straub (1992) compared word-by-word reading times for sentences such as the following:

(6) a. Grace persuaded Mary that he needed to rest more often than the schedule allowed.

b. Grace persuaded Howard that he needed to rest more often than the schedule allowed.

These sentences differ with respect to whether the matrix clause provides an antecedent for the pronoun (or whether readers have to create one). Clearly, sentences such as (6a) should be harder to process, and indeed, they are. However, in the Badecker and Straub study, a significant processing slowdown emerged after the pronoun, not on the pronoun itself.

A number of probe verification studies also show a delay in pronoun resolution. For example, McElree and Bever (1989) compared the processing of sentences such as (7a) and (7b). Participants read the sentences phrase by phrase (region boundaries are indicated in the examples with "/"). The probe word appeared either closely after the pronoun (probe point #1) or at the end of the sentence (#2):

(7) a. The stern judge/ who met with the defense/ thought he should (#1) / argue about the appeal.(#2)
b. The stern judge / who met with the defense / flatly rejected the (#1) / arguments for the appeal.(#2)

The probe word that McElree and Bever used was not the antecedent noun (e.g. *judge)* but rather the adjective modifying this noun (e.g. *stern).* The rationale, however, is the same as for any probe task: the pronoun should make the antecedent phrase more salient in memory (compared to the no-pronoun condition), hence components of this phrase should be easier (and faster) to verify. McElree and Bever found that verification times for the adjective were faster following a pronoun sentence like (7a) (vs. the no-pronoun counterpart (7b)), but only at the second probe point, not at the first.

This delay in facilitation has also been found when sentences are presented auditorily (and probes are presented visually). For example, MacDonald and MacWhinney (1990) compared response times to a probe in a pronoun vs. no-pronoun clause (as in 8), at three different points: just after the pronoun, 250 msec. post-pronoun, and 500 msec. post-pronoun. In addition, the researchers probed both referents mentioned in the matrix clause:

(8) *1st Clause:* Just before dawn, Lisa was fishing with Ron/Penny in the boat
a. *2nd-clause pronoun:*
1 2 3
and *she* caught a big trout right away.
b. *2nd-clause no pronoun:*
1 2 3
and lots of big trout were biting.

MacDonald and MacWhinney's results showed that the pronoun condition had a facilitatory effect for the antecedent *(Lisa)* compared to the non-antecedent

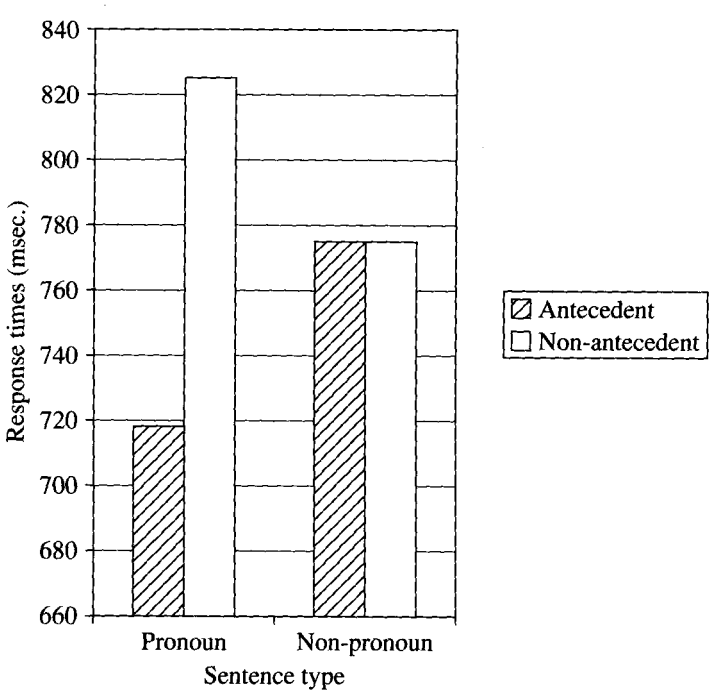

Figure 3.2 *Response times to probes, third probe point (MacDonald and MacWhinney 1990: exp. 2)*

(Ron) only at the last probe point - 500 msec. after the pronoun had occurred. Hence the probe task appears to show a delay in pronoun resolution - processing may not be completed until 500 msec. These results (reformatted for presentation here) are displayed in figure 3.2. In contrast, results obtained with the semantic priming task suggest that resolution processes are set into play *immediately* after a proform is presented.

Nicol (1988) employed the cross-modal priming technique to explore the role of featural constraints in permitting (or restricting) reactivation of antecedents. In one study, she contrasted sentences in which two nouns in the matrix clause were probed, and one of them (balanced for subject/object of the matrix verb) mismatched the pronoun in gender, as in (9), where the probe point is indicated with an asterisk, NPs that were probed for appear in italics, and the antecedent is underlined:

(9) The *ballerina* told the <u>*skier*</u> that the doctor for the team would blame him * for the recent injury.

Target items for lexical decision were either semantically related to each of the relevant nouns (e.g. *dance* for *ballerina* and *snow* for *skier)* or unrelated to these nouns (but matched to the related targets in terms of length, frequency, and other relevant variables). Using a relative difference in response times to related vs. unrelated targets as an index of lexical activation, Nicol found an immediate activation effect *only* for the gender-congruent noun (here, *skier*).

The discrepancy in the timing of pronoun resolution needs to be considered more closely. Examination of the differences between methods and materials suggests a few possible sources for the timing difference. Above, a question about the eye-mind synchrony during reading was raised. If rightward eye movement is triggered by, say, successful word recognition, then a failure to find an immediate effect of coreference processing is simply difficult to interpret. Therefore, reading measures may not always be the most appropriate way to assess timing issues (though they are certainly very useful for examining other kinds of questions).

If we consider only the experiments which involved auditory presentation - the cross-modal probe verification and semantic priming studies - several differences in task demands and stimulus materials are evident:

i Probe verification - but not simple lexical decision - may require a "replaying" of the sentence (from "left to right") in order to perform the task. In fact, the finding of a "first mention advantage" (e.g. Corbett and Chang 1983, Gernsbacher and Hargreaves 1988, Gernsbacher, Hargreaves, and Beeman 1989, McDonald and MacWhinney 1995) is quite robust in this literature; this could reflect left-to-right playback. When a probe item appears, participants in these experiments could initiate sentence replay, which might interfere with the processing of the material just preceding the probe. In other words, pronoun resolution processes could be disrupted or delayed by the presentation of a probe item (and the requirement to verify the occurrence of that probe word).

ii The cross-modal priming experiments have typically used object pronouns, but the probe verification tasks have typically used subject pronouns. One reason why this subject/object difference might be pertinent here is that subject pronouns occur early in a new clause and listeners might need a larger sample of speech in order for lexical recognition and syntactic structuring to occur. There may also be prosodic differences associated with these two pronoun types - for example, subject pronouns are less likely than object pronouns to be lengthened (since object pronouns are more likely to be phrase-final), and more likely to be followed by a pause. Both of these prosodic effects (lengthening and pausing) could facilitate identification of the pronoun.

iii The probe tasks have typically used sentence stimuli in which the pronoun appeared in an adjunct clause, but the cross-modal semantic priming studies used stimuli in which the pronoun appeared in an argument clause. It is possible that in the latter case, the NPs within the matrix clause are more readily activated, since they, along with the embedded clause, satisfy the requirements of the matrix verb's argument structure.

3.2 What constrains the initial candidate set and the elimination of candidates?

The questions concerning constraints on (i) which NPs are included in the initial candidate set and (ii) the elimination of NPs from the candidate set will be addressed together because they are interdependent. Obviously, if gender congruence requirements, for instance, do not constrain the initial set of candidates, these requirements must surely come into play in the elimination process. Hence, any bit of relevant information that does not play a role early in the process must play a role later in the process (except for the relatively rare cases in which there are multiple potential antecedents, and readers and listeners are ultimately uncertain about the identity of the antecedent).

The sources of constraining information that we will consider include the following:

i *Syntactic position information.* This includes information about whether an NP is subject or object, whether it is the head of a phrase or part of a modifier.

ii *Coreference constraints.* Such constraints include (roughly) the requirement that the antecedent of a *reflexive* (as object of the verb) is the subject of the clause in which the reflexive appears, the antecedent of an *object pronoun* must be some NP other than the subject NP, and the antecedent of a *subject pronoun* must be found outside the clause containing the pronoun. These constraints are captured formally and in greater detail by the Binding Principles (Chomsky 1981). These are illustrated in (10), (11), and (12) respectively (note that the asterisked indices below represent ungrammatical cases throughout, but that unasterisked indices represent an obligatory pairing in (10) and a permissible pairing in (11) and (12)):

(10) Sally; told Mary that Sue_k couldn't vote for herself.;,.$_{j,k}$
(11) Sally; told Mary that Sue_k couldn't vote for her;,;rwk
(12) Sally; told Mary that she;$_{j}$,.$_k$ couldn't vote for Sue_k

iii *Gender, number, animacy (or humanness) congruence.* Obviously, the gender, number and animacy (or humanness) features of an antecedent cannot clash with those of a proform (with which it is coindexed):

(13) Sam; told Mary that Sue_k couldn't vote for her.;,$_{j}$,.$_k$
(14) The Smiths; told $Mary_3$ that Sue_k couldn't vote for her.;/,,.k
(15) After looking at the car,,, Sam; told $Mary_3$ that Sue_k liked ith,.;,.j,.k

iv *Prominence in memory*. The syntactic position of an NP could affect its prominence in memory: a subject NP could be quite salient because it is a subject (and so it is natural to refer to it in upcoming discourse), but an NP within a modifier may be readily forgotten.

v *Semantic and pragmatic constraints.* Syntactic and feature constraints do not always eliminate all but one antecedent. Often, other factors come into play; these include knowledge about how the world works (e.g. in *Sam told Bill that he struck/saw/liked Mary, he* most likely refers to *Sam* because Sam would be unlikely to inform Bill about Bill's own activities or states of mind).

There are three logical possibilities for when these constraints come into play:

i None constrains the initial candidate set: all NPs mentioned prior to the proform are initially considered.
ii Only a subset of constraints affect the initial candidate set: only certain NPs are initially considered.
iii All constraints affect the initial candidate set: only the correct antecedent is ever considered.

How can an experimenter address the question of which NPs are initially considered? For investigators who use the semantic priming technique, it is clear: any NP that is part of a candidate set of antecedents should show priming. Similarly, any such NP should be easier to verify in a probe verification task. For reading time tasks, variations in set size should manifest themselves in reading time latencies - the larger the set, the greater the processing difficulty and the longer the reading times. In all cases, however, the consideration of a set of potential antecedents (and elimination from this set) is assumed to be entirely unconscious; coreference processing is simply one of a multitude of unconscious processes involved in understanding language.

It may seem most logical for the first option to be true. After all, how can an NP be rejected as a potential antecedent unless it has been considered in the first place? So it would be reasonable to consider all previously mentioned NPs and then eliminate all but one.

In what may seem (in retrospect) like a bold claim, Nicol and Swinney (1989) argued for the second viewpoint, specifically, that syntactic and featural constraints on coreference limit the initial candidate set to just those NPs which could be the antecedent. Evidence for this point of view comes from a set of studies by Nicol (1988) in which these constraints were manipulated. In a first study, she used semantic priming to indicate which word meanings were available after a proform in sentence pairs like those in (16) (italics indicate nouns probed for via the semantic priming task; underlines indicate possible antecedents; and the probe point is indicated with an asterisk):

(16) a. The boxer told the skier that the doctor for the team would blame him * for the recent injury.
b. The boxer told the skier that the doctor for the team would blame himself * for the recent injury.

Nicol presented visual target words immediately after the proform in each of these sentences. The visual target word was either a semantic associate of each of the potential antecedent(s) of the proform *(punch* for *boxer, slope* for *skier,* and *nurse* for *doctor)* or control words which were unrelated to anything in the sentence.' (The different distribution of antecedents that are permitted by syntactic constraints for the pronoun (in (16a)) and the reflexive (in (16b)) in these otherwise identical sentences is shown by the words underlined in each example.) Results showed that for sentences containing a pronoun, as in (16a), significant priming was obtained for target words related to *boxer* and to *skier* but no significant priming occurred for a target word related to *doctor. How-*ever, for the same sentences containing a reflexive, as in (16b), no significant priming was obtained for target words related to *boxer* or to *skier,* whereas significant priming was obtained for the target word related to *doctor.'* Thus, priming was obtained only for nouns which constitute syntactically licit antecedents to the particular proform. Note that the reactivation of these potential antecedents occurs *immediately* once the proform is encountered.

' Note that all effects are evaluated in comparison to lexical decision reaction time to a "control letter string" presented at each of these test points; a "control letter string" is a word that is associatively/semantically unrelated to the key word in the sentence, but which is matched to the "experimental" (related) letter string on the basis of a priori reaction time (lexical decisions taken on the words presented in isolation). Also note that no subject hears a sentence more than once or sees more than one visual target with that sentence.

' These significant effects were obtained on a large number of structurally similar materials in all of these studies, and not just on the example given here; however, for ease of explication, the description of results is in terms of this example.

Table 3.1 *Priming scores and level of statistical significance for probes related to the three potential antecedents (see example (16)) in sentences containing either pronouns or reflexives*

	Referent		
Anaphor type	NP1 (e.g. *boxer)*	NP2 (e.g. *skier)*	NP3 (e.g. *doctor)*
Pronoun	43*	58*	-21ns
Reflexive	-IDS	11	104*

* statistical significance at p<.o5
m not statistically significant

See table 3.1 for priming scores (i.e. differences in response times for the related vs. unrelated targets) for this study.

In all, this study suggested that syntactic constraints on coreference limit the initial candidate set. This finding has been supported by further empirical work. For example, Swinney, Ford, Bresnan, Frauenfelder and Nicol (cited in Swinney et al. 1988) employed the cross-modal priming task to study the activation of antecedents to sentences containing three potential antecedents to pronouns in sentences such as (17) (again, italicized nouns were probed; the underlined noun was the antecedent; probe points are indicated with an asterisk):

(17) The <u>*boxer*</u> visited the doctor who the *swimmer* at the competition had advised him * to see about the injury.

Here, *boxer is* the only structurally permissible antecedent of the pronoun *him*. The matrix object NP in this example *(doctor) is* linked to the relative pronoun, and ultimately connected to a gap after *see,* and cannot be the antecedent of *him*. Immediately following the pronoun, significant priming was found for visual target words related to *boxer* (as it was in the Nicol 1988 study). Similarly, no significant priming was found for words related to the noun *swimmer.* This study also examined activation (or reactivation) of various potential antecedents at a baseline probe position in the sentence - ¡mmediately before the verb *advised.* Results at this probe position showed that there was no significant priming for the visual word targets related to *boxer* (and there was no significant facilitation in responses to targets related to *swimmer).* Hence, the priming found for the structurally appropriate

antecedent *boxer* just after the pronoun suggests that the syntactically appropriate antecedent was reactivated after the pronoun was presented.

In much the same vein, Fodor, McKinnon, and Swinney (cited in Fodor 1993) examined reactivation effects with materials based on those in Nicol's (1988) studies (but with crucial changes), such as:

(18) a. The *boxer* knew that the *doctor* for the team was sure to blame him * for the injury.
b. The *boxer* knew that the *doctor* for the team was sure to blame himself * for the injury.

This study was aimed at examining the effects of a mediating NP-trace before the infinitive (created by the "raising" predicate, such as *be sure* in these materials) on coreference processing. The goal was to discover whether the "invisible" NP-trace would have a real and immediate on-line effect on assignment of an appropriate antecedent to the pronoun or reflexive. However, for our purposes, it is sufficient to examine the results obtained for pronoun and reflexive coreference, independent of other considerations. Using the cross-modal semantic priming technique, Fodor et al. found significant priming for words related to only the structurally appropriate antecedent for each proform *(boxer* for *him* and *doctor* for *himself),* and no other significant priming effects, at a test point immediately after the pronoun/reflexive in the sentence. These results comfortably replicate the Nicol (1988) findings, demonstrating immediate and syntactically constrained reactivation of antecedents for proforms.

Nicol (1988) also employed the cross-modal priming technique to examine the role of featural constraints in permitting (or limiting) reactivation of antecedents. In one study, she contrasted sentences varying the number specification of the antecedents of number-marked pronouns, as in:

(19) a. The *boxers* told the *skier* that the doctor for the team would blame them * for the injury.
b. The *boxers* told the *skier* that the doctor for the team would blame him * for the injury.

Nicol found significant priming only for words related to the potential antecedents which matched in number with the pronoun (only the underlined noun was significantly primed in each of the examples in (19)). (Note that this is so even though the conjoined NP containing both *boxers* and *skier* is an acceptable, but apparently not immediately activated, antecedent of *them.)*

Nicol (1988) also examined the effect of the lexical feature "gender" on the processing of proforms and their antecedents. She contrasted sentences like those in (20):

(20) a. The *ballerina* told the *skier* that the doctor for the team would blame him * for the recent injury.
b. The *ballerina* told the *skier* that the doctor for the team would blame her * for the recent injury.

Again probing (immediately after the pronoun) with visual target words related to each of the possible antecedents, Nicol found significant priming for only those referents that did not clearly clash in gender with the pronoun. This produced an asymmetric pattern of results, in which the masculine pronoun triggered priming of a single referent (e.g. *skier),* but the feminine pronoun triggered priming of both referents (e.g. both *ballerina* and *skier).* The explanation for the asymmetry was that the non-feminine nouns used in the sentence stimuli were not all unambiguously masculine (there are many more nouns in English which unequivocally denote a female referent than there are nouns which unequivocally denote a male referent; hence, *ballerina* is unequivocally female but *skier is* not unequivocally male). Overall, feature congruence (or non-incongruence) was an effective factor in constraining antecedent activation for pronouns.

This has not always been found to be the case for animate (or human) vs. inanimate (or non-human) referents. For example, Marslen-Wilson and Tyler (1980) conducted a cross-modal priming study aimed at determining the effect of lexical features on selection of a correct antecedent during ongoing comprehension, using sentences such as (21):

(21) The *sailor* tried to save the *cat* but it/he fell overboard.

Marslen-Wilson and Tyler found facilitation for lexical decisions to the visual targets *boat* and *dog* when they appeared immediately after *sailor* and *cat,* respectively. However, they found no differential facilitation when these words followed the pronouns *he* and *it.* They took this result to indicate that both structurally possible antecedents were activated immediately after both pronouns, and that featural properties of the pronoun only give weighting to selection of one of the antecedents, but do not actually select for a unique antecedent. We note, however, that the gender-marked pronouns (such as *he)* are commonly used to refer to both human and non-human (but animate) antececents (e.g. both *sailor* and *cat* in these sentences could be appropriate), and if both potential antecedents were animate in most of their materials, then

this would hugely diminish any ability to test for selectivity of priming based on animacy. Hence it is not clear that the failure to find selective activation in this study in any way changes the conclusions we have thus far drawn from the sentential priming studies - that pronominal features can constrain the reactivation of antecedents. Furthermore, there is counterevidence.

Shillcock (1982), in one of the earliest cross-modal priming studies to carefully examine activation of a single antecedent at various points throughout the sentence, performed an experiment which was similar in design to that of Marslen-Wilson and Tyler (1980). In this study, participants heard sentences such as (22):

(22) The *teacher* [1] did not board the train, [2] for the simple reason that it/he [3] was not going to the South Coast of England.

The associated visual target word (e.g. *school)* or an unrelated control word (e.g. *street)* was presented at each of the numbered probe positions. Shilicock found results that stand in opposition to those reported by Marslen-Wilson and Tyler (1980): significant priming was found for *school* (compared to the unrelated *street)* at the test point immediately after the pronoun *he,* but not after the pronoun *it.* Further, Shilicock found this to represent a change in activation level from the previous test position in the case of *it,* but not in the case *of he.* In all, this study demonstrated selective and immediate priming of the correct antecedent for a pronoun upon the occurrence of that pronoun. In this case, animacy constrained the set of NPs which were active at the point of the pronoun. Hence, it appears that pronominal feature information does indeed play a role in constraining activation of potential antecedents.

We now return to the puzzle raised above: how can an NP be excluded from a candidate set without first being activated? We propose that as a sentence is processed, grammatical information about each NP (including hierarchical and featural information) is noted and stored separately from its lexical semantic information. Presumably, the activation of word meaning declines over the course of a sentence, although it may be retrieved (or reactivated) at relevant points throughout the sentence. When a proform appears, the grammatical status of each NP is considered and the lexical semantics associated with each NP is retrieved only for those NPs which have some chance of being the antecedent.[1] If there is more than one grammatically appropriate NP within the candidate set of antecedents, the semantics of each

[1] The coreference process described here is meant to characterize coreference only in constructions in which a proform appears in an argument clause.

Input: The ballerina told the queen that the princess had introduced her ...

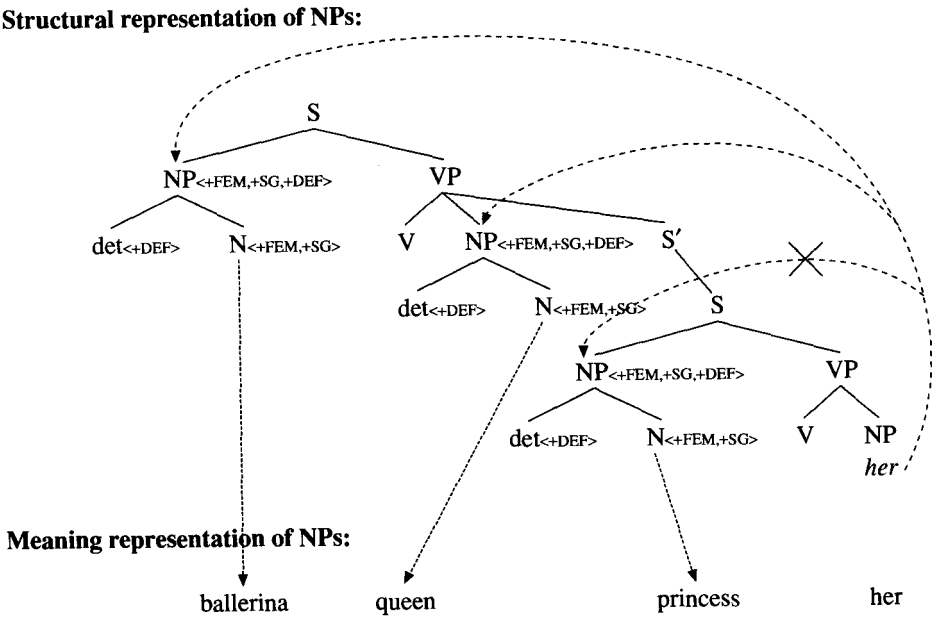

Figure *3.3* *Schematic representation of coreference processing*
A proform initiates a process whereby the syntactic information (both structural and featural) associated with previously mentioned NPs is used to eliminate inappropriate NPs (note the "X" in the arrow to the embedded subject NP). Once permissible NPs are identified, their semantics may be retrieved.

NP will be necessary for the elimination process because semantic/pragmatic information plays a key role in this process. This proposed set of processes is outlined in figure 3.3.

An alternative explanation is that the puzzle is not a puzzle: perhaps semantic priming actually reflects a later stage of processing, in which all but one NP has been eliminated. This explanation was initially resisted by Nicol (1988) because of the findings for the pronoun/reflexive study, which showed that *both* referents in the matrix clause were primed following the pronoun, suggesting that grammatically viable antecedents had not yet been eliminated. A counterargument might be that this pattern of priming could simply indicate that the sentence materials, which contained matrix NPs which were meant to be equiprobable antecedents of the subsequent pronoun, were successful in this respect: each matrix NP might have been construed as the antecedent about half the time, and so - according to the counterargument - both matrix

NPs only *appeared* to be active simultaneously. However, this possibility seemed less likely for the gender study, in which both matrix NPs were primed following the feminine pronoun *her.* It was assumed that given a sentence such as *The ballerina told the skier that the doctor for the team would blame 1erfor the recent injury,* participants in the experiment would be more likely to interpret *her* as coreferent with the gender-congruent *ballerina* than with the gender-neutral *skier.' If* the cross-modal priming task is sensitive only to the final resolution *of* a proform, then it must be explained how it is that the gender-ambiguous *skier is* so often preferred as the antecedent *of her,* when the gender-marked *ballerina* is also a candidate. This alternative possibility - that cross-modal priming taps into the endpoint *of* pronoun resolution - is compatible with a number *of* studies reported by Badecker and Straub (1992, 1994, 1995, 1997), who examined sentence pairs such as those in (23), in which the two initial NPs either were the same gender or differed in gender. In both cases, the embedded subject (e.g. *John)* was the antecedent *of* the reflexive:

(23) a. Bill; thought that John$_s$ owed himself$_j$ another opportunity to solve the problem.
b. Beth; thought that John$_s$ owed himself$_j$ another opportunity to solve the problem.

The researchers used a self-paced reading task in which sentences were presented word by word, and participants pressed a response key to advance to the next word (which overwrote the preceding word, so that participants saw only one word at a time). The results showed that for the region of text following the reflexive, participants read sentences like (23a) more slowly than (23b). Since the only difference between the two sentences is the gender difference, this must be the source *of* the reading time difference. But if, as argued by Nicol (1988), proforms reactivate only syntactically licit NPs, then the reflexive in these constructions should reactivate only the embedded subject. The gender *of* the matrix subject should not matter because this NP is not reactivated anyway. So why is the same-gender case (as in 23b) producing a reading slowdown? Badecker and Straub (1997) have suggested that the gender of structurally irrelevant NPs matters because the coreference process does not initially take syntactic information into account. They propose that any NP that matches the pronoun in gender is initially considered a potential antecedent, and that syntactic information then acts (as a filter) to eliminate NPs. The more NPs in the candidate set, the greater the processing load.

[s] It is important to bear in mind in considering the materials that a range of gender neuter nouns were used, including nouns that, although not obligatorily masculine (e.g. *football player, are* certainly stereotypically so.

This is an interesting hypothesis, which is clearly problematic for our account. However, it is compromised by the results of other studies. One study was conducted by Badecker and Straub themselves (Badecker and Straub 1997). They examined reading times in sentence pairs such as those in (24):

(24) a. John thought that Bill's cousin owed herself another opportunity to solve the problem.
b. John thought that Beth's cousin owed herself another opportunity to solve the problem.

In these sentences, there is one structurally constrained antecedent of the reflexive - *cousin.* With respect to gender-congruence, there are no other gender-matching NPs in (24a). In (24b), however, there is another gender-matching NP *(Beth).* On Badecker and Straub's account, reading times should be slower following the reflexive in. (24b) than in (24a) because the structurally illicit NP *Beth* would need to be eliminated from the candidate set. However, the researchers found no reading time differences at all. This suggests that either both possessors *(Bill* and *Beth)* are initially considered (which is in direct conflict with Badecker and Straub's other results), or neither are considered. This latter possibility seems more feasible; after all, the possessor is structurally more deeply embedded than the other NPs in the sentence. Hence, perhaps such nouns are not automatically considered as antecedents. (Though obviously, they can be antecedents, as in *John thought that Beth;'s cousin owed her; another opportunity to solve the problem.)*

In fact, there is evidence that structural position does affect the accessibility of an NP as the antecedent of a pronoun. Matthews and Chodorow (1988) used reading times (specifically, reading times to the last word of the sentence) to compare the processing of sentence pairs such as the following:

(25) a. After the bartender, prepared the patron's drink$_{k,}$ he; got a big tip.
b. After the bartender; prepared the patron's drink$_{k}$, he$_{j}$ left a big tip.

The researchers found longer reading times when the more deeply embedded NP turned out to be the more plausible antecedent: reading times were longer for (25b) than for (25a). In other words, if the more plausible antecedent is part of a modifier, reading time is slowed. Matthews and Chodorow found a similar pattern when the more plausible antecedent was part of a PP modifier, as in (26):

(26) a. Because the owner; of the restaurant$_{i}$ prepared the food$_{k}$, it$_{k}$ was always delicious.

b. Because the owner; of the restaurant; prepared the food$_k$, it; was always crowded.

In this experiment, reading times were longer when the more plausible antecedent turned out to be the NP within a modifying phrase *(restaurant)* than when the antecedent was the head of its maximal projection *(food).* In both cases, the syntactic position within a structure appears to affect how accessible in memory an NP might be. This, in turn, should have an effect on coreference processes.

One could argue that there is an initial phase of coreference processing in which only salient feature-congruent NPs are considered. These are then filtered out by other types of constraints, including syntactically based coreference constraints and semantic/pragmatic information (such as the fact that patrons leave tips and bartenders get tips, that food is delicious but a restaurant is not, etc.). Sometimes the filtering process will produce a null set, in which case other strategies will need to come into play. It would be at this later stage of processing that the less accessible NPs would be considered. In fact, Badecker and Straub (1997) do argue that only salient NPs are part of the initial set of candidates. The problem with this argument is that it cannot quite account for reading time patterns for sentences in which a possessor is the antecedent of a pronoun, as in (27):

(27) a. John thought that Bill's cousin owed her another opportunity to solve the problem.

b. John thought that Beth's cousin owed her another opportunity to solve the problem.

In (27a), there is no antecedent for *her,* but in (27b), the possessor *Beth* could be the antecedent. Badecker and Straub's reading time data suggest that differences between the sentence types emerge early - at the point of the pronoun. Therefore, readers cannot be excluding the (non-salient) possessor from the candidate set; if they did, reading times for the pronoun would be identical in the two sentences. Overall, although the notion that NP-accessibility plays a role in pronoun resolution receives independent empirical support, the Badecker and Straub account cannot benefit from this because their conflicting reading time data for pronouns and reflexives, which show access to the possessor in one case but not the other, may signal that the task is unreliable. Indeed, the results of another experiment suggest that something else may be going on in their reading studies.

A followup study was conducted by Nicol (1997), who compared reading times for sentences such as those in (28) in a "phrase-by-phrase" self-paced

reading task. These were embedded in a larger study, so that participants were presented with not only these types of sentences but with sentences containing reflexives and a variety of other constructions. These also provided some variety in the number of words appearing at once:

(28) a. My aunt / heard that / the congresswoman / would contact / her about / the complaint.
b. My aunt / heard that / the congressman / would contact / her about / the complaint.
c. My aunt / heard that / the congresswoman / would contact / me about / the complaint.
d. My aunt / heard that / the congressman / would contact / me about / the complaint.

Stimulus sentences contained two contrasts: the embedded subject was either congruent with the matrix subject (as in (28a) and (28c)) or it was incongruent ((28b) and (28d)); and in order to examine the effects of gender congruency of NPs with and without coreference, the embedded object was either a third person pronoun (which was congruent with the matrix subject) or a different NP (either a first or second person pronoun or a short name (e.g. *Ruth)* or phrase (e.g. *the boys)).* This followup to Badecker and Straub's work differed from theirs in two ways. First, common nouns were used instead of proper nouns. It was felt that participants might engage more fully in the task if sentences described more concrete scenarios. Second, every single experimental item was followed by a comprehension probe that required the reader to have identified the object pronoun (e.g. *The congressmanlcongresswoman contacted my auntlme* or *The congressmanlcongresswoman contacted my uncle).* The systematic inclusion of comprehension probes allowed a partitioning of the data according to correctness of the response.

Let us first consider all the data for the two critical regions: the penultimate region, which contained either the pronoun or control, and the next, final region (these are displayed in figure 3.4). At the pronoun region, there is a clear difference in reading time for the embedded object, with third person pronouns being read faster than the control pronoun or NP. However, the next region shows in interaction of type of embedded object (third person vs. other) and congruence: reading times are fastest when the third person pronoun is preceded by incongruent NPs. This latter result suggests that NP congruence does indeed make it more difficult to compute coreference, despite the fact that one of the NPs is not a syntactically permissible antecedent.

However, consider now the reading times for only those sentences for which participants provided a correct response (shown in figure 3.5). Again, for the

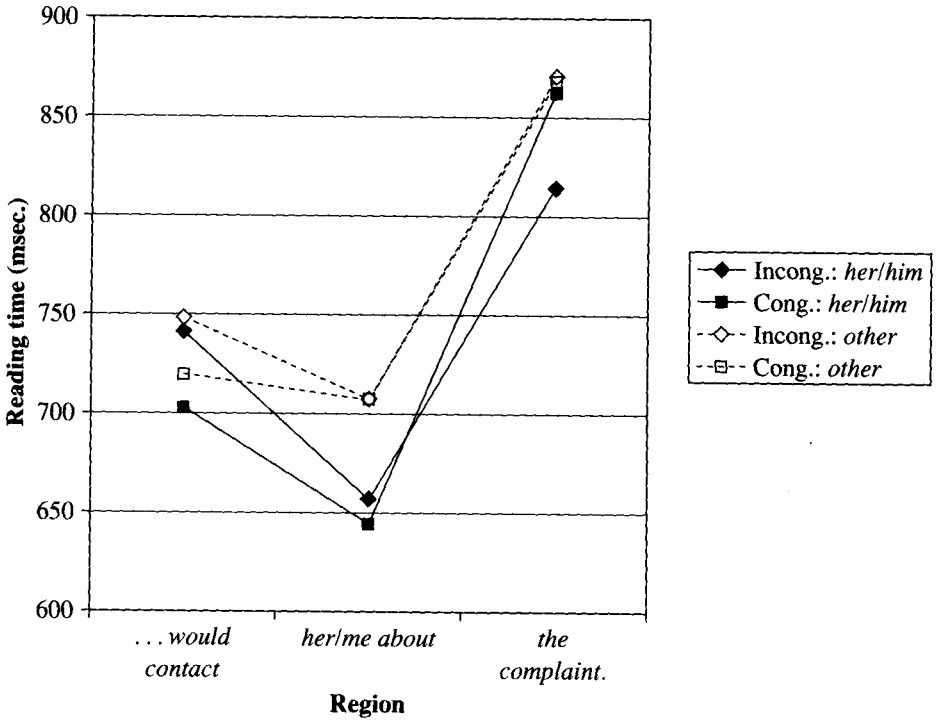

Figure 3.4 *Reading times for last three regions of test sentences (all data) (Nicol* 1997)

The regions correspond to regions in the example sentence. The four sentence types reflect two crossed variables: the inclusion of referents which have congruent pronominal features ("Cong.") vs. incongruent pronominal features ("Incong.") and the presence of a third person pronoun ("her/him") vs. a control NP ("other").

pronoun region, there is a main effect of embedded pronoun type. However, for the following region, there is no interaction: NP congruence has no effect. Hence, it appears that manipulating the gender of the embedded subject (which cannot be the antecedent due to its position in the sentence) does not have an effect on coreference processing when readers are fully comprehending the sentence.

We do not believe that this finding entirely undermines the Badecker and Straub results. Rather, we believe that there may be different modes of reading for which comprehension varies enormously: how well readers maintain sentence structure, and how fully they compute argument relations and coreference, very likely depends on the task demands (including how often comprehension questions are asked and what the comprehension questions probe), the nature of the non-experimental distractor sentences, and more generally, the reader's own level of arousal and interest in the material (for

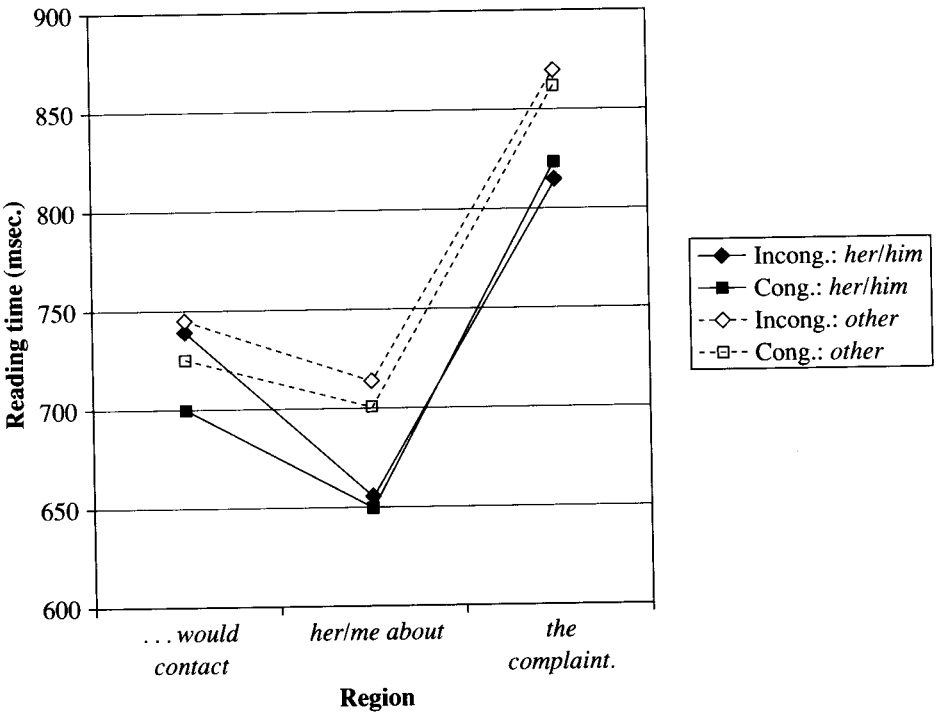

Figure 3.5 *Reading times for last three regions of test sentences (data for correctly understood sentences) (Nicol 1997)*
The regions correspond to regions in the example sentence. The four sentence types reflect two crossed variables: the inclusion of referents which have congruent pronominal features ("Cong.") vs. incongruent pronominal features ("Incong.") and the presence of a third person pronoun ("her/him") vs. a control NP ("other").

support of this notion, see Foertsch 1994, Garnham, Oakhill, and Cruttenden 1992). In short, people can (and do) read more or less carefully, and the various modes of reading merit study.'

We believe too that although auditory comprehension may also be variable, it is at least immune from strategies that may arise from the desire to

v One reservation we have, however, is that reading studies which use a word-by-word presentation could impair performance to the extent that readers have to be inordinately attentive in order to fully understand a sentence containing coreferring elements. In other words, the technique could create a circumstance that arises rarely in the real world. It would be useful to know whether readers in eye-tracking studies can also be found to ignore syntactic constraints in computing coreference.

press forward through the sentence (as is possible when participants pace themselves in reading tasks). In addition, auditory verbal memory may simply be superior to visual verbal memory (Jakimik and Glenberg 1990). Hence, listeners may be better comprehenders, overall, than readers (especially readers who cannot re-read earlier text, as is typically the case in self-paced reading studies). This notion has some experimental support. Nicol (1993) conducted a series of studies that sought to replicate and extend the cross-modal priming work. In these studies, she used an all-visual presentation, in which sentences were presented word by word in the center of the screen. In these studies, participants did not pace themselves; rather, the presentation was automatic (each word appeared for approximately 230 msec.). At a critical point in the sentence, a target item appeared for lexical decision. This target appeared in a different font, in a different color, higher than screen-center, and enclosed in angled brackets. After 1000 msec., the target disappeared and the sentence continued. A portion of the sentences was followed by a comprehension question. Nicol examined sentences such as the following (probed items are italicized, the correct antecedent is underlined, and the probe point is indicated with an asterisk):

(29) a. The boxer said that the *skier* in the hospital had blamed himself for the recent injury.
b. The boxer said that the *skier* in the hospital had blamed him * for the recent injury.

In these sentences, the embedded subject *(skier) is* the correct antecedent only for the reflexive sentence (29a), not for the pronoun sentence (29b), so it was expected that priming would be obtained for the embedded subject *(skier)* in sentences like (29a), and not in sentences like (29b). In fact, that is exactly what was found: there was significant facilitation for the associate of *skier* only when *skier* was the appropriate antecedent.

Although the results of this first study suggest that participants were fully comprehending the sentences, the results of a followup study suggest otherwise. The followup study was conducted in order to determine whether these results might arise due to a "comprehension strategy" by which the reflexive would simply be linked to the nearest preceding NP, which would result in a correct interpretation in a majority of instances. To test this, sentence pairs like those in (30) were examined:

(30) a. The boxer talked to the *skier* in the hospital and blamed him * for the recent injury.
b. The boxer, talked to the *skier* in the hospital and blamed himself for the recent injury.

In this study, priming for the matrix object *(skier)* was expected for sentences containing a pronoun, such as (30a), but not for sentences containing a reflexive, such as (30b). The results of this experiment showed just the *opposite* pattern from what was expected: the matrix object was primed in the reflexive sentence (30b) but not in the pronoun sentence (30a). Clearly, the pattern of priming suggests that participants were not properly computing the syntactic structure of the sentence or the coreference relationships, or both." This mode of presentation might have been particularly difficult for participants; however, even in reading experiments which present sentences in full, readers often show backward regressions when they encounter pronouns (Carpenter and Just 1977, Ehrlich 1983). Since regressions are disruptive to the reading process, it is likely that readers regress only when they are having serious difficulty.

In sum, the experimental work suggests that with auditory presentation, the occurrence of a proform triggers immediate reactivation of a candidate set of antecedents. This initial set is limited to just those referents which are grammatically appropriate as antecedents (with respect to syntactic position and feature attributes). With visual presentation, it is possible that coreference constraints do not initially restrict the candidate set if participants are not questioned closely about their comprehension. In short, there may be different modes of reading, and the early effect of syntactic constraints on coreference may only be seen when participants are fully engaged in the reading task.

Finally, let us consider briefly some of the constraints that are likely to come into play in the elimination of potential antecedents.

(1) There may be a possible preference for *parallel function (i.e.* a preference for coreference with the NP in the matrix clause which has the same thematic role as the pronoun), so that for a sentence such as *Sally, told Mary$_f$ that she$_{if,*k}$ couldn't vote for Sue$_k$,* there might be a preference to construe *Sally* as *she* because both of these NPs are subjects. This effect has been reported by a number of investigators (e.g. Crawley, Stevenson, and Kleinman 1990, Grober, Beardsley, and Caramazza 1978, Sheldon 1974, Smyth 1994, Stevenson, Nelson, and Stenning 1995), though it may be overridden by other factors (such as gender (Matthews and Chodorow 1988) and semantic factors (McDonald and MacWhinney 1995)).

[10] This conforms to our own experiences as participants. Prior to launching new experiments, a great deal of self-testing is routine (in order to optimize presentation rates, exposure times of targets, font readability, etc.); in this case, visual presentation certainly seemed to make coreference more difficult to compute (and note that this was so for the very creator of the sentences being presented).

(2) In addition, there may be a preference for coreference with a "topic": in *Sally; told Mary$_1$ that she;,; couldn't vote for Sue, Sally* could be preferred as the antecedent of *she* because subject NPs are typically topics, and as such are more salient than object NPs (Crawley, Stevenson, and Kleinman 1990, Stevenson, Nelson, and Stenning 1995). Salience may also play a role in structures in which NPs vary with respect to their *depth* within the matrix clause, as in the Matthews and Chodorow (1988) research discussed above (recall, in a sentence fragment like *Because the owner of the restaurant prepared the food, it. ..,* although both *restaurant* and *food* are possible antecedents, *restaurant* may be less preferred because it is embedded within the subject NP, as a part of a phrase which modifies the subject head)."

(3) Another obvious factor in pronoun resolution is real-world sensibility. The semantic context in which referents appear is obviously relevant to selecting an antecedent. For instance, Hirst and Brill (1980) correlated strength of contextual bias of a lead-in sentence (predetermined by pilot testing) with choice of antecedent and with response latencies to determine the antecedent.

One well-studied instance of real-world plausibility effects is the *implicit causality* associated with the matrix verb in constructions in which two clauses are causally linked. For example, in *Sally impressed Sue because she. .. ,* there may be an antecedent preference for *Sally,* but in *Sally scolded Sue because she .... Sue* would be preferred. Implicit causality has been examined by a number of investigators (e.g. Caramazza, Grober, and Garvey 1977, Garvey and Caramazza 1974, Garnham, Oakhill, and Cruttenden 1992, MacDonald and MacWhinney 1995, McKoon, Greene, and Ratcliff 1993), and it is clear that implicit causality plays a robust role in pronoun resolution (in sentences in which "cause" plays a central role).

A study by MacDonald and MacWhinney (1995) has shown that such information is more quickly available than gender information in the resolution of pronouns, which is not what we would have predicted. The researchers used the cross-modal probe verification paradigm to examine pronoun resolution processes in sentences such as (31) and (32). Two verb types were used, either a "stimulus-experiencer" verb such as *amaze* (as in (31)) or an "experiencer-stimulus" verb such as *admire* (as in (32)). When these are followed by a "because" clause, perceivers expect to hear more about the stimulus, and so construe a pronoun as coreferent with the stimulus *(Gary* in (31) and *Craig* in (32)). In addition, the experiencer and stimulus either matched in gender or

[11] We assume that such less-accessible NPs are part of the initial candidate set, though this has yet to be tested empirically with the cross-modal priming technique.

did not. Both the stimulus and experiencer were probed, and four different points were used (these points appeared at particular temporal distances from the second name, the pronoun, and the sentence end, and are indicated with asterisks at their approximate points):

(31) a. Gary amazed Alan time * after time because he * was * so talented at the juggling competition.*
b. Gary amazed Ellen time * after time because he * was * so talented at the juggling competition.*

(32) a. Ross admired Craig quite * a bit because he * was * so courageous during the debate championship.*
b. Ruth admired Craig quite * a bit because he * was * so courageous during the debate championship.*

MacDonald and MacWhinney found that just after participants heard a pronoun, speed of probe verification was affected by implicit verb causality *(amaze* ("stimulus-experiencer") vs. *admire* ("experiencer-stimulus")) but not by a gender distinction among the participants *(Gary amazed AlanlEllen).* In other words, it appears that verb semantics has a more immediate effect than NP gender. At first, this seem counterintuitive; surely, antecedent identification should be surer and faster if a gender mismatch *eliminates* one NP than if verb semantics to do with causality *merely disfavors* one NP as a potential antecedent. However, it is likely that causality relations have already been computed by the time the pronoun appears, [12] whereas the pronoun, itself bearing gender information, must appear before any elimination process occurs. This would account for the timecourse difference.

4 Activation vs. Inhibition

Above, it was concluded that a proform reactivates its antecedent. But in fact, there is some controversy about whether a proform raises the activation level of an antecedent (activation), lowers the activation level of non-antecedents (inhibition), or both. The study by MacDonald and MacWhinney (1990,

[12] The authors attempt to discount this possibility by conducting a control experiment in which a subordinate clause such as ". . . because he was so talented" is replaced with a PP such as . . at the juggling competition." They found that the verb type (e.g. *amaze vs. admire)* had no effect on probe verification responses. However, it may well be the case that the computation of a causal relationship only occurs under certain discourse conditions, e.g. with the appearance of conjunctions such as *because.*

discussed in section 3.1 above) suggested that pronoun resolution involves activation and inhibition mechanisms. Their critical result is for their unambiguous condition (e.g. the *Lisa-Ron* condition, in which referents differed in gender) at the third probe point (500 msec. post-pronoun) (see figure 3.2). At this point, response times were significantly faster to *Lisa* than to *Ron* in the pronoun condition. In the no-pronoun condition, response times were equivalent. Even more important, a comparison of the response times in the pronoun vs. no-pronoun conditions showed a relative speedup in response times to *Lisa* and a relative slowdown in response times to *Ron.* What is absolutely critical for the resolution of this issue is the nature of the baseline condition. In the MacDonald and MacWhinney study, the sentence versions are increasingly different at the later probe points. For instance, in (33a) and (33b) below, the third probe point is preceded by quite different sentence fragments (notably, an additional NP - *trout* - has likely been processed by this point in (33b) but not (33a)):

(33) a. She caught a big trout right away.
b. Lots of big trout were biting.

It is not at all clear, then, that a no-pronoun condition like (33b) provides a neutral sentence condition.

Note that a major difference between the probe verification task and the cross-modal priming task is that the priming task includes an independent measure of "baseline" activation, through the use of the frequency-matched control word. Response times to related target words are always compared to the response times for unrelated words. The use of *sentence* baseline conditions (e.g. the early probe point in the study discussed above by Swinney et al.) provides an additional measure - by showing that priming is present just after a pronoun but not before, it suggests that the significant priming effect is not due to continued activation of the probed noun. It also allows for moment-by-moment changes in sentence difficulty to be controlled for, since any difficulty should result in increased response latencies for both related and unrelated targets. The probe verification task is missing this independent target baseline, and necessarily so, because there is simply no equivalent within this task to the "unrelated target" condition. This is because the probe of interest (e.g. *Lisa, Ron)* requires a "YES" response from participants, but an equivalent type of target - matched to, say, *Lisa,* and in the same sentential position and also requiring a "YES" response - is not possible. As a consequence, the only possible comparisons are between (i) two NPs within the sentence, and/or (ii) two difference sentences. Hence, the finding that a non-antecedent of a pronoun is responded to faster in the no-pronoun condition

than in the pronoun condition could be due to *other* differences between the two sentence types. In particular, it is possible that the pronoun increases processing load, increasing response times to make judgments about *any* noun that is not directly related to the pronoun. Given the lack of an independent measure of activation in the probe verification task, we will accept the results of the cross-modal studies and assume that, at least in within-sentence coreference contexts, proforms reactivate potential antecedents.

5 Representational Issues

Finally, Love and Swinney (1995) examined the nature of the representations involved in coreference. In doing so, they examined a number of test points throughout the comprehension of a sentence containing a pronoun, with the goal of attempting to obtain more detail about the timecourse of the initial processing and (re)activation of such antecedents. However, their focus was on using cross-modal semantic priming to uniquely distinguish whether, at the point of the pronoun, what is reactivated is a verbatim (form) representation or a semantic or conceptual representation (for other work in this area, see Cloitre and Bever 1988, Waters, Caplan, and Leonard 1992). To do this, they utilized lexical ambiguities as antecedents for pronouns.

The reasoning behind the use of lexical ambiguities is the well-established finding that *all* meanings of a lexically ambiguous word are initially activated when the word is heard in a sentence (see, e.g. Swinney 1979, Tanenhaus, Leiman, and Seidenberg 1979, among others). This is known as *form-driven lexical access* - access of all meanings is triggered by hearing the (acoustic) form of the word. Thus, if all meanings of an ambiguous word are found to be reactivated immediately after the pronoun for which the ambiguous word is the antecedent, this would constitute strong evidence for the view that the form of an NP is maintained in memory and that this form is retrieved during coreference processing. However, if only the contextually relevant meaning of the antecedent ambiguity is reactivated after the pronoun, the representation of the antecedent in memory is likely to be a semantic one (corresponding to a single meaning of the ambiguous word).

In this study, subjects heard sentences of the following form:

(34) Jeff had read about problems with savings and loan institutions, so he went to the *bank* to ask about the safety that it provided with respect to CD investments.

1 2 3

Priming for each meaning of the ambiguous word *bank (i.e. financial institution* and *river edge)* was examined at each of three target presentation points (marked numerically). Significant priming for words related to both meanings of the ambiguity was found at the first test point immediately following initial occurrence of the ambiguous word in the sentence.[13] At test point 2, no significant priming was found for words related to either of the meanings of the ambiguity. (Thus, after being initially activated and processed, the lexical ambiguity (or some subset of its meanings) was then cleared from the immediate level of processing.) Finally, at the critical test point (point 3) immediately following the pronoun, a significant facilitation effect was found for only the target word related to the primary (and contextually relevant) meaning of the ambiguity (here, the *financial institution* interpretation); there was no significant facilitation for the target word related to the other meaning. Thus, only the contextually relevant meaning of the ambiguity was immediately reactivated once the pronoun was heard. This study demonstrates that a pronoun triggers reactivation of the semantic representation of an NP, not the form itself.

6 Summary

When a proform occurs within an auditorily presented sentence, there is immediate reactivation of the candidate antecedents which do not mismatch the pronoun with respect to agreement features and which obey syntactic constraints on coreference. The representation that is reactivated is semantic/conceptual. When there is more than one candidate antecedent, other information may be invoked in order to eliminate all but one antecedent; such information includes real-world knowledge and discourse prominence. Probe tasks do not typically show an *immediate* reactivation effect; however, this may be due to the task demands: participants must check their mental representation of the ongoing sentence for a particular lexical item, and doing so could disrupt coreference processing if the probe appears just after a pronoun. Although the salience of NPs (correlated with depth of embedding) has not, to our knowledge, been tested with priming or probe techniques, it is possible that a deeply embedded NP is less likely to be part of the initial candidate set.

[13] This provided yet another demonstration of exhaustive access of word meanings for lexical ambiguities, even in the presence of strong prior biasing contexts.

When a proform appears within text, coreference processing may not be immediate, because the early visual processing of text may occur when the eye is fixated, but complex processing such as the establishment of coreference may lag behind. It is also possible that the computation of coreference is inherently more difficult in the visual domain (especially when the experimental procedure involves the overwriting of earlier text). Because the priming and probe tasks provide meaningful data only if the timing of processing operations is relatively uniform across materials and participants, the reading lag makes it difficult to examine the activation of word meanings associated with pronoun resolution within written text. Therefore, questions concerning constraints on the initial candidate set are likewise difficult to address. It is certainly plausible, however, that there are inherent differences in pronoun resolution processes in reading and listening.

Acknowledgments

The authors gratefully acknowledge support from NIDCD (Grant DC-01409, to the first author, as part of the National Center for Neurogenic Communication Disorders, University of Arizona), NIH DC02984, and NIH D000494 for some of the research reported herein, and the National Center for Neurogenic Communication Disorders at the University of Arizona for supporting the writing of this chapter.