JOURNAL OF MEMORY AND LANGUAGE 32, 692–715 (1993)

# The Interaction of Lexical and Syntactic Ambiguity

MARYELLEN C. MACDONALD

*University of Southern California*

Two experiments investigated comprehension of noun/verb lexical category ambiguities such as *trains*, in order to determine whether resolution of these ambiguities was similiar to other types of ambiguity resolution. Frazier and Rayner (1987, *Journal of Memory and Language*, **26**, 505–526) argued that these ambiguities were resolved with a delay strategy that is not used for other ambiguities. Experiment 1's self-paced reading data replicated Frazier & Rayner's results but also showed that evidence taken to support delay had other explanations. Experiment 2 investigated the influence of semantic biases on ambiguity resolution and found that three probabilistic factors influenced lexical category ambiguity resolution: (1) the relative frequency of head vs. modifying noun usage of a biasing noun, (2) the frequency of cooccurrence of a biasing noun and category ambiguous word in English, and (3) the combinatorial semantic information in the sentence. The extent to which alternative models account for the use of probabilistic information in ambiguity resolution is discussed. 

Human language contains ambiguities at many levels of linguistic representation, but two kinds of ambiguity, lexical and syntactic, have received the most attention in recent psycholinguistic research (e.g., Altmann, 1990; Small, Cottrell, & Tanenhaus, 1988). The dominant theory of lexical ambiguity resolution is one in which alternate senses of ambiguous words may be available for a short time after the ambiguous word is encountered, even when the ambiguous word occurs in a biasing context (Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979; Tanenhaus, Leiman, & Seidenberg, 1979). By contrast, "serial" or "single representation" models form the dominant accounts of syntactic ambiguity resolution. The best known of these is the Garden Path model developed by Frazier and her colleagues (Frazier, 1987; Ferreira & Clifton, 1986; Rayner, Carlson, & Frazier, 1983). This model holds that the syntactic component of the language processor has no access to nonsyntactic information; it is aware of only the lexical categories of the input words (Determiner, Noun, etc.) and it assigns a grammatical syntactic structure to this input using only knowledge of the grammar and some independent strategies (e.g., Minimal Attachment, a heuristic that guides the parser to choose the simplest syntactic structure). Semantic and discourse information can influence syntactic ambiguity resolution only at a later stage, according to this model.

Given such disparate accounts of lexical and syntactic ambiguity resolution, it is interesting that many syntactic ambiguities are triggered by lexical ambiguities. One particularly striking case concerns *lexical category ambiguities*, ambiguous words that have multiple meanings across two or more different lexical categories. For example the words *train* and *watch* have both noun and verb interpretations, so that when words of this type are encountered in a sentence, both the meaning and the syntax of the sentence become temporarily ambigu-

This research was supported by NIH BRSG 1711 at the Massachusetts Institute of Technology, and by NSF Grant DBS-9120415. I thank Greta Buck, Alex Rosen, Josh Butler and Beatrice Dupuy for help with experiments and with coding the corpus data. Thanks also to Mark Seidenberg, Neal Pearlmutter, and three anonymous reviewers for helpful comments on an earlier draft. Address correspondence and reprint requests to: Maryellen C. MacDonald, Hedco Neuroscience Building, University of Southern California, Los Angeles, CA, 90089-2520. E-mail: mcm@gizmo.usc.edu.



0749-596X/93 $5.00

ous. If such ambiguities are treated like lexical ambiguities, then it would be expected that multiple syntactic/semantic interpretations are briefly activated, but if these ambiguities are instead handled in the manner that the Garden Path model proposes for syntactic ambiguities, then only one syntactic interpretation would be considered. Lexical category ambiguities therefore present a challenge not only for the comprehender, who must deal with ambiguity at both the semantic and syntactic levels simultaneously, but also for psycholinguistic theories, which have traditionally offered very different explanations for semantic and syntactic ambiguity resolution. This paper investigates the nature of lexical category ambiguity resolution, especially the extent to which it resembles ambiguity resolution at other levels of linguistic representation.

Two different accounts have been offered for the relationship between lexical, syntactic, and lexical category ambiguity resolution. The first of these is a "constraint-based" approach to ambiguity resolution in which both lexical and syntactic ambiguities are resolved via the operation of multiple probabilistic lexical, syntactic, and discourse-level constraints (e.g., MacDonald, 1993; MacDonald, Pearlmutter, & Seidenberg, 1993; McClelland, St. John, & Taraban, 1989; Pearlmutter & MacDonald, 1992; Spivey-Knowlton, Trueswell, & Tanenhaus, 1993; Taraban & McClelland, 1988; Trueswell, Tanenhaus, & Garnsey, 1993a; Trueswell, Tanenhaus, & Kello, 1993b). These accounts differ in detail, but the research in this tradition views ambiguity resolution at the lexical and syntactic levels as sharing some basic features: Partially activated alternative lexical forms and meanings, modulated by discourse context and lexical frequency, support alternative syntactic structures to greater or lesser degrees. In the process of activation spreading across several levels of representation, one syntactic structure and one lexical interpretation begin to dominate the alternatives; the time course of this process would depend on the strength and consistency of the constraints that modulate the activation. On this view, lexical category ambiguity resolution should proceed in this same fashion: Multiple factors, including contextual information and the inherent bias of the ambiguous word (e.g., whether *train* is more often a noun or verb), should rapidly affect both the syntactic and the semantic interpretation of the ambiguity. Some support for this view is offered by Tyler and Marslen-Wilson (1977), who found that prior discourse context had a rapid effect on the syntactic interpretation of lexical category ambiguities, as measured by cross-modal naming of a disambiguating word.

The alternative view of ambiguity resolution is that it is not desirable to attempt to unify lexical and syntactic ambiguity resolution, because lexical representations are *stored* and *accessed*, whereas syntactic representations are *constructed* (e.g., Frazier & Rayner, 1987). Indeed, Frazier (1989) has argued that different mechanisms are needed to handle each kind of ambiguity: (1) for purely semantic lexical ambiguities (with no lexical category ambiguity), the alternative interpretations are available in parallel and one interpretation is rapidly chosen on the basis of frequency and/or prior context, (2) for syntactic ambiguities, the serial parser of the Garden Path model constructs only one interpretation, and (3) lexical category ambiguities like *train* are handled by a "delay" mechanism. The delay model (Frazier & Rayner, 1987) holds that the alternative meanings and categories of the ambiguous word are initially accessed in parallel, but no alternative is rapidly chosen. Instead, the language processor is hypothesized to delay selection of an alternative until a definitive disambiguation is reached later in the sentence.

The delay account of lexical category ambiguity resolution conflicts with the constraint-based models in several respects.

First, in arguing that lexical category ambiguities are resolved with delay and other ambiguities are not, Frazier (1989; Frazier & Rayner, 1987) challenges the view that ambiguity resolution at each level of linguistic representation shares the same basic mechanisms. Second, the claim that no action is taken during lexical category ambiguity resolution until a definitive disambiguation is reached challenges the constraint-based claims that many kinds of information can rapidly constrain ambiguity resolution, beginning when the ambiguity is introduced. The delay model is therefore an important part of the debate concerning the relationship between ambiguity resolution at different levels of linguistic representation, and it merits further investigation.

## The Delay Model

Support for the delay model of lexical category ambiguity resolution can be found in several experiments conducted by Frazier and Rayner (1987). They examined two-word phrases containing lexical category ambiguities such as (*the*) *desert trains*. These phrases also contain a second ambiguity that is inherent in all noun phrases (at least in written input), in that at the point of the first noun (e.g., *desert*), is it temporarily ambiguous whether this noun is a modifying noun or the head noun of the noun phrase. Thus in the *modifier noun + head noun* (NN) interpretation of the phrase, *desert* is a noun modifying the head noun *trains*, whereas in the *head noun + verb* (NV) interpretation, *desert* is the head noun of the noun phrase, and *trains* is a verb.

Frazier and Rayner measured eye fixation times for their ambiguities in ambiguous and disambiguation regions, compared to unambiguous control sentences. In the unambiguous conditions, prior disambiguating information was provided by replacing the determiner *the* that preceded the ambiguous words with either *this* or *these*, yielding the unambiguous condition *these desert trains* for the NN interpretation and *this desert trains* for the NV interpretation.

Frazier and Rayner's studies yielded two interesting results. First, reading times in the ambiguous region (e.g., *desert trains*) were shorter in the ambiguous (*the* determiner) condition than in unambiguous (*this/these*) control sentences, but reading times in the remainder of the sentence, containing the disambiguation, were longer in the ambiguous condition compared to controls. Their results can be seen in Fig. 1, which shows the mean reading times per character in these regions, averaged across their Experiments 1 and 2. The effects are quite subtle but were generally reliable in subjects analyses (Frazier and Rayner conducted Experiment 2 with new materials in lieu of items analyses).

Frazier and Rayner also investigated the effect of semantic information on lexical category ambiguity resolution. Unlike Tyler and Marslen-Wilson (1977), they did not manipulate prior context for each ambiguous phrase but instead assessed the preferred interpretation of each stimulus item with sentence completion norms (e.g., norming subjects interpreted *desert trains* with an NN interpretation 60% of the time). These measures yielded the second interesting finding, that the inherent biases to the NN or NV interpretation across items did not influence the patterns of reading

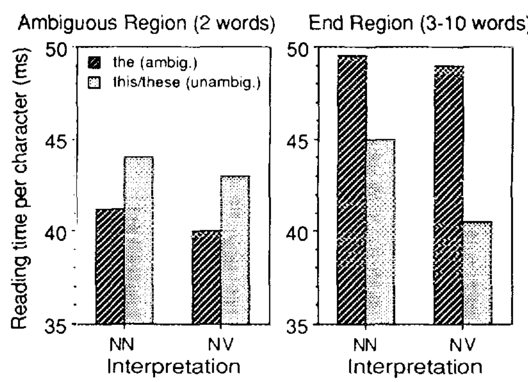


FIG. 1. Reading times in ambiguous and unambiguous regions in Frazier and Rayner's (1987) Experiments 1 and 2.

times in the ambiguous region or in the disambiguation.

The combination of the reading patterns and the semantic bias results led Frazier and Rayner to propose the delay model: During the delay period at the onset of the ambiguity, very little analysis of the ambiguity takes place, resulting in lower processing loads and shorter reading times in ambiguous conditions compared to unambiguous controls, for which there is no delay. When the disambiguation is reached, the delay process is terminated and there is now some catching up to do in analysis of the ambiguous condition, with the result that reading times increase in this condition over the unambiguous one. Frazier and Rayner suggested that very strong prior contexts, such as a preceding paragraph about locomotives for the *desert trains* item, might change this pattern of results, but in the absence of such context, the inherent bias of a phrase like *desert trains* toward one or the other interpretation is not important, because the ambiguity resolution mechanism is one in which processing is delayed until the disambiguation is reached.

The delay model does account for Frazier and Rayner's results, but an alternative interpretation is possible. It is a truism that when ambiguous sentence reading times are compared to unambiguous control conditions, the comparison is valid only to the extent that the control condition does not differ from the ambiguous condition except in its ambiguity. Frazier and Rayner's controls may have introduced another factor in their materials. The concern with using the determiners *this* and *these* to create unambiguous conditions is that these determiners also have a deictic function, directing the comprehender to a previously mentioned entity. That is, the phrase *this desert* is generally felicitous only when a desert has been previously established in the discourse. When deictic determiners are used in isolated sentences with no prior context, as in the Frazier and Rayner experiments, they sound quite awkward. Long reading times immediately following *this* or *these* may therefore simply reflect readers' confusion with the use of infelicitous determiners, and the effects may be wholly unrelated to ambiguity. If so, the data in Fig. 1 could not be used to support a delay model.

These concerns suggest that it is necessary to evaluate the alternative explanation of reading time differences in the ambiguous region in lexical category ambiguous sentences. Experiment 1 below addresses these issues: Reading times for sentences containing lexical category ambiguities were compared to times for unambiguous counterparts such as *deserted trains* and *desert trained*, and the type of determiner was manipulated factorially with the new ambiguity manipulation. If some awkwardness from using *this* and *these* in isolated sentences was producing the negative ambiguous–unambiguous difference in the ambiguous region in Frazier and Rayner's experiments, then reading times should vary in this region only as a function of determiner, not ambiguity. If Frazier and Rayner's effects were really due to ambiguity, however, then there should be an interaction in Experiment 1, such that reading times in ambiguous phrases like (*the*) *desert trains* are shorter than in any unambiguous condition, both when the phrase is unambiguous by virtue of the preceding determiner and also with the new unambiguous conditions that do not involve a determiner change.

## EXPERIMENT 1

### *Method*

*Subjects.* Forty-eight MIT undergraduates were paid for their participation. An additional five subjects were tested but rejected for comprehension question error rates exceeding 20%.

*Materials.* To best evaluate the Frazier and Rayner (1987) effects, the ambiguous stimuli for Experiment 1 were developed from 24 of the 32 experimental items from

TABLE 1
EXAMPLE STIMULI, EXPERIMENT 1

| | | Point of disambiguation |
|---|---|---|
| NN Interpretation | | |
| 1.* Ambiguous, *the* | I know that the desert trains could resupply the camp. | first end region word (*could*) |
| 2.* Ambiguous, *this/these* | I know that these desert trains could resupply the camp. | usually amb. region Word 2 (*trains*) |
| 3. Unambiguous, *the* | I know that the deserted trains could resupply the camp. | usually amb. region Word 2 (*trains*) |
| 4. Unambiguous, *this/these* | I know that these deserted trains could resupply the camp. | usually amb. region Word 2 (*trains*) |
| NV Interpretation | | |
| 5.* Ambiguous, *the* | I know that the desert trains soldiers to be tough. | last four words |
| 6.* Ambiguous, *this/these* | I know that this desert trains soldiers to be tough. | usually amb. region Word 2 (*trains*) |
| 7. Unambiguous, *the* | I know that the desert trained soldiers to be tough. | usually amb. region Word 2 (*trained*) |
| 8. Unambiguous, *this/these* | I know that this desert trained soldiers to be tough. | usually amb. region Word 2 (*trained*) |

* These items are replications of the Frazier and Rayner conditions.

their Experiments 1 and 2.[1] The words in the sentence up to the ambiguous phrase were kept exactly as in Frazier and Rayner's versions, but new endings were written for most items so that each sentence contained exactly four words after the ambiguous region. Eight versions of each sentence were prepared, manipulating three factors. The first two of these, Interpretation (the ambiguous phrase is disambiguated with the NN or NV interpretation) and Determiner Ambiguity (*the* vs *this/these*), form replications of the original Frazier and Rayner studies. The third factor, which will be termed Inherent Ambiguity, was manipulated by modifying the two words in the ambiguous region to remove the ambiguity in the unambiguous condition. As can be seen in the examples in Table 1, the nature of the change to this region depended on the noun vs verb Interpretation variable. When the ambiguous word (e.g., *trains*) was given a noun interpretation, the preceding word was changed to be either an adjective (*deserted trains* in Table 1), or to a possessive (e.g., *cashier's checks* replaced the ambiguous *cashier checks*). When the ambiguous word was disambiguated as a verb, it was changed from present tense to past tense, and the preceding word was left unchanged (*desert trained, cashier checked,* etc.).

Table 1 also contains information about the point of disambiguation. Because a noun precedes the noun/verb ambiguous words in these stimuli, the lexical category ambiguity is always in close proximity to a second ambiguity concerning which noun is the head of the noun phrase. For this reason, the "unambiguous" conditions in both the present stimuli and in the original Frazier and Rayner items are never completely free of ambiguity in the first region. Consider first the NN items. In all three (determiner and inherent) unambiguous versions, it is unambiguous that the first word (*desert/deserted*) is a modifier of some upcoming head noun, either because of the adjectival nature of this word and/or because of

[1] The eight Frazier and Rayner items that were not used were numbers 3, 9, 14, and 16 from their Experiment 1, and numbers 2, 10, 12, and 16 from their Experiment 2. These numbers reflect the numbering in Frazier and Rayner's (1987) appendices; their extensive appendices obviously made it possible to evaluate their claims in the present paper.

the plurality of the preceding determiner.[2] Some ambiguity is introduced at the next word (*trains*), however, in that this noun could be taken either as the head noun of the noun phrase (the correct interpretation) or as a modifier of an upcoming head noun (as in *the(se) desert(ed) trains inspectors*). Plural nouns like *trains* are rarely modifiers in English, so there is good probabilistic evidence that *trains* is likely to be the head noun. The last column in Table 1 therefore indicates that by the point of *trains*, the unambiguous NN conditions probably, but not certainly, become fully unambiguous (meaning both that *desert/deserted* is identified as a modifier and *trains* as the head noun). In the ambiguous NN condition, by contrast, it is completely ambiguous whether *desert* is a modifier or head noun, and whether *trains* is a noun or verb, until the first word of the end region forces the head noun interpretation of *trains*.

Similar temporary ambiguities exist in the "unambiguous" NV interpretations. In all conditions, the head vs modifier status of *desert* is initially unclear, and the ambiguity persists to some degree into the next word. In the unambiguous determiner condition (*this desert trains*), it is fairly unlikely that *trains* will itself be considered a modifier for an upcoming head noun, because, as noted above, plurals are rare as modifiers, and so the NV interpretation is likely, but not certain, at this point. In the inherently unambiguous conditions, the modifier interpretation for *trained* is even more remote, because *desert–trained* should be hyphenated to indicate a modifier status, and also because most phrases were highly implausible or ungrammatical as a pair of modifiers for an upcoming noun (e.g., *poor stated, summer flew*). The ambiguous condition is still fully ambiguous at this point, and its disambiguation comes late in the end region, because the words after the ambiguity do not necessarily force the verb interpretation of the noun/verb ambiguous word. The NV interpretation is not forced immediately because English allows the possibility of a postnominal relative clause without a complementizer like *that* (e.g., *The summer flies I caught were big, The desert trains soldiers attacked were destroyed*). The category ambiguity thus persists for some words in the NV interpretation condition, until the relative clause interpretation becomes extremely unlikely or the sentence ends.

Eight lists were prepared to fully counterbalance all factors. Each list also contained 10 practice items and 66 fillers, some of which were test items for an unrelated experiment. A yes/no comprehension question was prepared for each item.

*Procedure.* Subjects read sentences in a Moving Window display in which all nonspace characters of the sentence initially appeared as dashes on a computer screen (Just, Carpenter, & Woolley, 1982). Subjects pressed a computer key to see each word of the sentence. The first keypress revealed the first word, and with the second keypress, the first word reverted to dashes and the second word was revealed, and so on. Following the keypress terminating the last word of the sentence, the comprehension question appeared. Subjects pressed a key marked Yes or No to answer and did not receive feedback. Subjects were tested individually and completed the experiment without a break in a 30 min session.

### *Results*

Only sentences with correct comprehension question responses were included in reading time analyses. Reading times were adjusted for length using a procedure described by Ferreira and Clifton (1986): For each subject, a linear regression equation was calculated to predict reading times in each word from word length in all experimental sentences. Reading time at each word was then expressed as a difference

[2] The exception, as a reviewer pointed out, is the *the* determiner followed by a possessive, e.g., *the cashier's*, which might be interpreted as a contraction. The possessive (i.e., modifying) interpretation is reinstated in the next word, however.

score from the predicted reading time for that subject. Negative reading times for a word indicate that this word yielded shorter reading times than would be expected for a word of that length, relative to other words in the sample. As reading times generally increase through a sentence, the length-adjustment regression commonly results in negative values early in the sentence. The length-adjusted reading times were trimmed at each word by removing all data points that were over 2 sd above the condition mean for that word, affecting less than 4% of the data. Both raw and length-adjusted reading times for individual words are contained in Appendix A.

Figure 2 presents the reading times in the ambiguous region for the replication conditions and the new conditions; a constant of 100 is added to all length-adjusted reading times to make them all positive values, which aids in reading the figure. Figure 2 reveals two important results. First, Frazier and Rayner's effects of Determiner Ambiguity were replicated: Reading times were shorter in the ambiguous determiner *the* conditions than in the *this/these* conditions. Second, the same effect of Determiner also obtained in the inherently unambiguous conditions. This pattern produced a main effect of Determiner in an Analysis of Variance with Ambiguity, Determiner, and Word as factors [$F_1(1,47) = 5.75$, $MS_e = 4681$, $p < .05$, $F_2(1,23) = 3.09$, $MS_e = 2907$, $p < .10$], but no interactions with any other factors ($F$'s $< 1$).

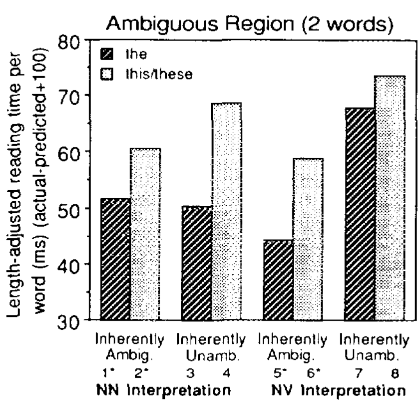


FIG. 2. Length-adjusted reading times in the ambiguous region in Experiment 1. The numbers 1–8 refer to the condition numbers given in Table 1. Asterisks indicate conditions that are replications of Frazier and Rayner (1987).

Not only does Fig. 2 look remarkably like Frazier and Rayner's data for the ambiguous region that was shown in Fig. 1, but the magnitudes of the effects are very similar across experiments. In the ambiguous region in the present study, the Determiner Ambiguity manipulation had an 11 ms/word effect in the inherently ambiguous conditions that formed the replication of Frazier and Rayner, and there was a 12 ms/word effect of Determiner in the new inherently unambiguous conditions. The effect of Determiner in this region in Frazier and Rayner's two eye fixation experiments was 14 ms/word. The change of reading paradigm from eye monitoring to self-paced reading therefore seems to have had little effect in this region.

These data suggest that the use of *this* and *these* to disambiguate the sentences introduces some factor that does not exist in the *the* conditions. Comparisons across levels of Determiner Ambiguity therefore do not provide a clean basis for assessing the effects of ambiguity. The Inherent Ambiguity manipulation, however, provides another opportunity to assess ambiguity effects, when combined with the *the* determiner. Comparisons across inherent ambiguity yielded a mixed pattern of results. Figure 2 shows no differences between the inherently ambiguous and unambiguous conditions in the NN interpretation (the dark bars in conditions 1 vs 3), but in the NV interpretation (5 vs 7), ambiguous reading times were shorter than unambiguous times. This pattern produced an Inherent Ambiguity × Interpretation interaction (marginal in the subjects analysis), $F_1(1,47) = 3.65$, $MS_e = 4092$, $p < .10$, $F_2(1,23) = 4.69$, $MS_e = 1834$, $p < .05$. The pattern of cell means suggests that a major cause of this difference is in the unambiguous reading times, which were nonsignificantly longer in the unambiguous NV interpretation than in the unambiguous NN interpretation, $F_1(1,47) = 2.31$, $MS_e = 6358$, $p > .10$. The long reading times for the NV inherently unambiguous conditions could have resulted from some infelicity associated with switching from present tense at the beginning of the sentence to past tense (*trained*) in many of these items, or perhaps from some confusion over whether a modifier like *desert-trained* was intended for some items.

Finally, reading times for the second word of the ambiguous region were a mean of 42 ms longer than for the first, $p$'s $< .001$, but Word did not interact with any other factor. Frazier and Rayner also found robust effects of Word here, and they also found some interactions with Word and other factors that were not consistent across their Experiments 1 and 2.

The end region in Fig. 3 shows that the two inherently unambiguous conditions produced an effect of Determiner Ambiguity that is similar to that found in the ambiguous region. The differences here were not significant ($p$'s $> .15$), but there is some suggestion that the effects of the *this/these* determiners carried over to the end region. In the inherently ambiguous conditions, the pattern of reading times was different. The NN interpretation condition had longer reading times in the *the* determiner than in the *this/these* determiner, replicating the pattern that Frazier and Rayner found for this region. The NV interpretation condition, however, produced a nonsignificant difference in the opposite direction, $F < 1$. An examination of the individual words in these conditions (given in Appendix A) reveals that this pattern is primarily due to extremely long reading times for the first word of the end region (the word immediately following the ambiguous region) in the *this/these* condition; when this word is removed, the pattern more closely resembles the ambiguous NN interpretation condition, with reading times in the *the* condition nonsignificantly longer than in the *this/ these* condition. The awkwardness of the *this/these* condition may thus have spilled over to the end region, masking beneficial effects of the disambiguation category that *this/these* can provide. The NN interpretation conditions may have yielded more stable data because the point of disambiguation was more uniform across items in this interpretation than in the NV interpretation.

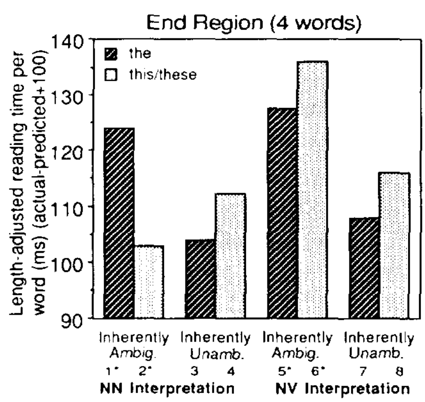


FIG. 3. Length-adjusted reading times in the end region in Experiment 1. The numbers 1–8 refer to the condition numbers given in Table 1. Asterisks indicate conditions that are replications of Frazier and Rayner (1987).

Examining only the *the* conditions in Fig. 3, we can observe clear effects of Inherent Ambiguity in the end region. In both interpretations, reading times were about 20 ms per word longer in the inherently ambiguous conditions than in the unambiguous conditions, $F_1(1,47) = 10.78$, $MS_e = 6989$, $p < .01$, $F_2(1,23) = 8.07$, $MS_e = 7804$, $p < .01$. This result suggests that there really is a disadvantage for sentences containing lexical category ambiguities, compared to unambiguous sentences, even though the head/modifier ambiguities that are inherent in noun phrases may not have been completely removed by the end of the first region in these items.

### Discussion

The principal goal of Experiment 1 was to assess one piece of critical evidence taken by Frazier and Rayner (1987) to support a delay model of lexical category ambiguity resolution, namely the finding that reading times in the ambiguous region were shorter when the preceding determiner (*the*) permitted the ambiguity than when the determiner (*this/these*) provided a disambiguation. Experiment 1 replicated this result, but the effects of the determiner manipulation also obtained in the inherently unambiguous sentences. These results suggest that Frazier and Rayner's data reflect something about the change of determiners across conditions rather than any ambiguity effects. The present experiment does not itself identify the source of the *this/these* difficulty—the hypothesis concerning the discourse infelicity of *this/these* is still only an hypothesis—but the important point for theories of ambiguity resolution is that a pattern of reading times thought to implicate a delay period appears to derive from other sources.

A similar conclusion is likely for Frazier and Rayner's (1987) Experiment 3, which also found evidence for delay, but with a different unambiguous condition. In that experiment, Frazier and Rayner examined adjective/verb ambiguities with words like *landing* and *sinking,* in sentences like *Not everyone agrees that sinking ships is/are unpleasant.* Unambiguous controls for most items were created by adding an adjective before or in the middle of the ambiguous region, e.g., *fast sinking ships* and *sinking fast ships.* The rationale for these manipulations is based on two assumptions: (1) that the *-ing* word has exactly two interpretations (adjective or verb), so that the placement of the added word (*fast*) will remove all ambiguity, and (2) that the other word in the region (*ships*) is unambiguously a noun. Neither assumption is wholly justified, however. First, the ambiguous *-ing* words may sometimes have a third interpretation as a noun, and so the addition of the word *fast* does not provide a disambiguation. Thus *fast sinking,* which is intended to unambiguously indicate the sequence *adjective + adjective,* could also be an *adjective + noun* sequence (e.g., *The fast sinking made the diver ill*). In the unambiguous verb condition, the addition of *fast* after *sinking* does guarantee that *sinking* is a verb, but it also adds a temporary ambiguity concerning whether *fast* modifies the word to follow it (the intended interpretation) or whether it modifies *sinking* (e.g., *Sinking fast isn't recommended when diving*). Second, the assumption that the next word is unambiguously a noun as is also not justified—*ships* may also be a verb, introducing another lexical category ambiguity. Because of the ambiguities and the awkward phrasing in many of the unambiguous items, it is again likely that reading patterns attributed to processing delay have other explanations in this experiment.

These results and analyses are mostly negative, insofar as they offer alternative interpretations of "delay" reading patterns; they say little about the processes that comprehenders do use to resolve lexical category ambiguities. The major goal of Experiment 2 is therefore to provide a more direct test of the different predictions derived from the delay and constraint-based models. The best hope of providing stronger tests seems to lie in semantic manipulations: A constraint-based account, with extremely rich and rapid interaction between different levels of linguistic representation, predicts that semantic bias effects concerning the relative plausibility of the alternative interpretations should be observable rapidly, by the end of the ambiguous phrase or by the start of the disambiguation, given strong enough biases. The claim from the delay model, however, is that such inherent biases will not influence the initial ambiguity resolution, because all analyses are delayed until the disambiguation is reached. Experiment 2 investigates these alternative predictions.

## Experiment 2

The source of the delay model's argument for the lack of semantic bias effects is the results from Frazier and Rayner's

(1987) analysis of semantic bias effects in their own materials. They found that reading times did not vary as a function of whether the ambiguity was resolved with the preferred interpretation or the nonpreferred interpretation. Frazier and Rayner interpreted this result as owing to the delay process, but the lack of a semantic bias effect in their data might instead reflect the fact that the semantic biases in their items were quite weak; only half of the items were strongly biased (over 80%) to one interpretation. A stronger bias manipulation might therefore reveal that semantic information does rapidly influence ambiguity resolution.

Experiment 2 investigates this hypothesis with a strong semantic bias manipulation, combined with an ambiguity manipulation. If the bias of the ambiguous phrase to one interpretation does influence ambiguity resolution, then ambiguous reading times should differ from their unambiguous controls only when the ambiguity is resolved with the disfavored interpretation. When the semantic bias strongly promotes the interpretation that turns out to be correct, however, ambiguous reading times should not differ from the unambiguous condition. The delay model, by contrast, predicts a main effect of ambiguity, independent of bias. According to this model, an effect of bias might appear in later regions after the disambiguation, but crucially, the semantic bias manipulations should not be able to affect the initial stage of processing, during which analyses are delayed. The experiment has two parts: First, reading times in different bias and ambiguity conditions are measured, then these reading times are evaluated in light of normative data concerning the semantic biases.

*Method*

*Subjects.* Forty-four MIT undergraduates were paid for their participation. Two additional subjects were rejected for high comprehension question error rates.

*Materials.* The stimuli were constructed around 16 lexical category ambiguous words with both a plural noun and a third person singular verb meaning (*fires, benefits, guards,* etc.). A sentence ending was written for each ambiguous word that resolved the ambiguity in favor of the NV interpretation. This interpretation was chosen because the Experiment 1 results for this interpretation were weaker than for the NN interpretation, and Experiment 2 offers an opportunity to clarify ambiguity effects in this construction. Examples can be seen in Table 2; all sentences are contained in Appendix B.

In contrast to Frazier and Rayner (1987), who compared bias across items, the semantic bias of the ambiguous phrase was manipulated within items in the present experiment. Two biasing nouns were chosen for each of the noun/verb ambiguous words. These biasing words appeared immediately before the noun/verb ambiguous word and were designed to combine with it so that the two words together strongly favored either the NV interpretation or the NN interpretation. For example, the biasing noun *corporation* in the phrase *corporation fires* promotes the NV interpretation

TABLE 2
EXAMPLE STIMULI, EXPERIMENT 2

| | |
|---|---|
| | Supportive bias |
| Ambiguous | The union told reporters that the corporation fires many workers each spring without giving them proper notice. |
| Unambiguous | The union told reporters that the corporations fire many workers each spring without giving them proper notice. |
| | Unsupportive bias |
| Ambiguous | The union told reporters that the warehouse fires many workers each spring without giving them proper notice. |
| Unambiguous | The union told reporters that the warehouses fire many workers each spring without giving them proper notice. |

of *fires* (in which *corporation* is the head noun of the noun phrase and *fires* is the verb), because it seems more plausible for a corporation to fire people than for a corporation to have a fire. Because all sentences are disambiguated with the NV interpretation, the bias that promotes this interpretation is termed the *supportive bias* condition. The *unsupportive bias* condition promotes the incorrect NN interpretation. For example, *warehouse fires* promotes the noun interpretation of *fires* (in which *warehouse* is a modifier of the head noun *fires*) because a fire in a warehouse seems more plausible than a warehouse firing people.

Ambiguity was manipulated by changing only the plurality of the critical pair of words, so that the ambiguous *corporation/warehouse fires* became *corporations/warehouses fire* in the unambiguous condition. By the time the verb (*fire*) is encountered in these conditions, it is virtually unambiguous that this word is the verb and that the preceding word is the head of the noun phrase. Eight words preceded the ambiguous word, and nine words followed. As in the NV interpretation conditions of Experiment 1 and in Frazier and Rayner (1987), there is no one point at which the ambiguous condition is definitively disambiguated.

All of the noun/verb ambiguous words had a higher frequency noun interpretation than verb interpretation (Francis & Kucera, 1982). The bias of this word therefore goes *against* the eventual NV interpretation in the sentence. If the supportive bias condition can remove all ambiguity effects, then this result would provide evidence that the combinatorial information from the biasing word and the next word (e.g., that it is more plausible for a corporation to fire someone than to have a fire) can override simple lexical frequency information (e.g., that the noun interpretation of *fires* is more frequent than the verb interpretation). The delay model, which predicts no analyses until the disambiguation is reached, does not predict any effects of bias. There is also a third hypothesis that is intermediate between the constraint-based prediction of "rich analysis" during initial processing and the "no analysis" prediction of the delay model. This "local analysis" hypothesis is that only the frequency information for the noun/verb ambiguous word influences the early stages of ambiguity resolution, but that combinatorial information is not available early. As the frequency information of the noun/verb ambiguous word always favors the incorrect NN interpretation in these sentences, the local analysis hypothesis predicts that ambiguous sentences should be harder than their unambiguous counterparts, independent of bias. Experiment 2 thus distinguishes hypotheses about combinatorial semantic effects in the constraint-based account, which does predict an effect of bias, from simple frequency effects or no semantic effects (as in the delay model), both of which do not predict a bias effect. The experiment does not distinguish these latter two possibilities, however.

A yes/no comprehension question was written for each sentence. Four stimulus lists were constructed to balance all factors.

*Procedure.* The procedure was identical to that in Experiment 1. Subjects saw 5 practice sentences and one version of each of the 16 experimental sentences, randomly intermixed with 60 fillers, some of which were test items from an unrelated experiment.

### *Reading Time Results and Discussion*

The adjustment for word length was calculated as in the previous experiment, except that a modification in the presentation program allowed length to be recorded for every word in the experiment, rather than only for the experimental items. Each subject's length adjustment calculation was therefore based on words from 76 sentences, excluding only the practice items. Reading times for the ambiguous word position and all subsequent word positions in the sentence were trimmed, removing all reading times greater than four standard deviations above the condition mean at each word position; this affected 1.1% of the data.

The constraint-based model's prediction was that ambiguous and unambiguous reading times would differ only in the unsupportive bias condition, not in the supportive bias condition. The pattern of length-adjusted reading times in Fig. 4 supports this prediction.

Planned comparisons at each word at each level of bias revealed there was no reliable effect of ambiguity at any word in the supportive bias condition, shown in the left panel of Fig. 4 (all *p's* > .25). The unsupportive bias condition in the right panel did produce effects of ambiguity, however. Reading times at the ambiguous word and the next word (Words 9 and 10) were significantly shorter in the ambiguous condition than in the unambiguous condition (Word 9: $F_1(1,43) = 4.96$, $MS_e = 8162$, $p < .05$, $F_2(1,15) = 6.58$, $MS_e = 1784$, $p < .05$; Word 10: $F_1(1,43) = 4.70$, $MS_e = 21275$, $p < .05$, $F_2(1,15) = 12.48$, $MS_e = 4022$, $p < .01$). The direction of the ambiguous–unambiguous difference here appears to be consistent with the predictions of the delay model, but the delay model also predicts similar effects in the supportive bias condition, where they did not obtain. An analysis of the relative complexity of the alternative interpretations offers a plausible interpretation of these "reverse ambiguity" effects, in which ambiguous reading times are shorter than their unambiguous controls. When the ambiguous sequence is given the NN interpretation, the parser must represent the input in a noun phrase (e.g., [the warehouse fires]$_{NP}$). The NV interpretation, however, requires a more complex structure for the input, namely a noun phrase and an associated verb phrase, (e.g., [the warehouse]$_{NP}$ [fires]$_{VP}$). Because this more complex NV interpretation is required in the unambiguous conditions, their reading times were relatively long; reading times have previously been shown to be sensitive to syntactic complexity differences (Just & Carpenter, 1980; King & Just, 1991). The ambiguous reading times mimicked these unambiguous times when the supportive biasing information supported the correct NV interpretation. By contrast, ambiguous phrases with unsupportive bias were initially interpreted as part of a noun phrase, and subjects read noun phrases more quickly than the more complex sequence in the unambiguous condition. Thus reading times appear to reflect the most strongly active interpretation, as a function of available constraints; similar

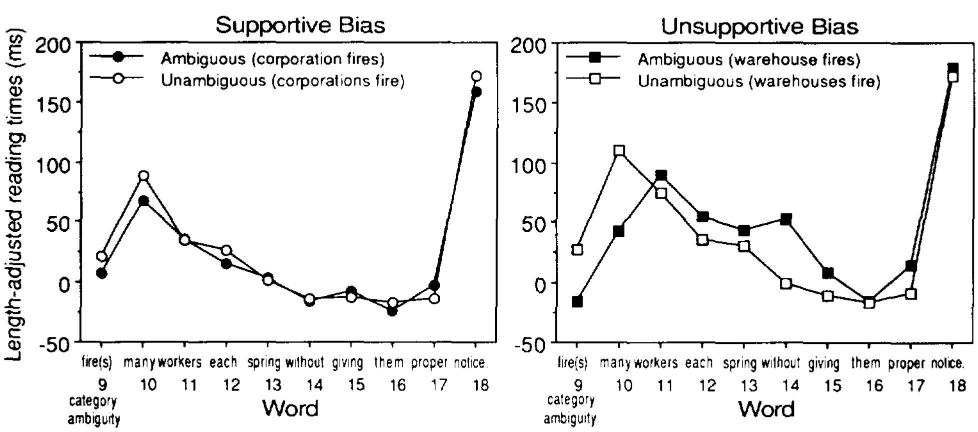


FIG. 4. Length-adjusted reading times in supportive and unsupportive bias conditions in Experiment 2.

patterns have been found with syntactic phrase structure ambiguities (MacDonald, 1993). The fact that this effect persisted to Word 10 is not surprising in light of the fact that this word did not provide a definitive disambiguation; such "spill-over" effects are common in self-paced reading experiments.

Beginning at Word 11, however, and throughout the rest of the sentence, ambiguous reading times were as long or longer than unambiguous times in the unsupportive bias condition. The difference across ambiguity was not reliable at Words 11–13 ($p$'s $> .15$), but at Word 14, ambiguous reading times were 52 ms longer than unambiguous times, $F_1(1,43) = 5.47$, $MS_e = 11273$, $p < .05$, $F_2(1,15) = 6.85$, $MS_e = 3289$, $p < .05$. The effect at Word 15 was marginal, $F_1(1,43) = 2.50$, $MS_e = 3078$, $p > .10$, $F_2(1,15) = 3.36$, $MS_e = 875$, $p < .10$. There was no difference at Word 16 ($F$'s $< 1$), but a marginal effect again appeared at Word 17($F_1(1,43) = 4.62$, $MS_e = 2528$, $p < .05$, $F_2(1,15) = 2.25$, $MS_e = 898$, $p < .10$). Differences at Word 18 were not reliable ($F$'s $< 1$).

The diffuse appearance of ambiguity effects in the later words of the sentence presumably reflects the fact that these stimuli did not have one definitive point of disambiguation. Although the effects of ambiguity were not reliable at every word in the unsupportive bias condition, a comparison of the two panels in Fig. 4 clearly shows how bias modulated the ambiguity effects. An analysis of reading times over a broad disambiguation region, Words 11–17, reflects this pattern in a Bias × Ambiguity interaction, $F_1(1,43) = 3.80$, $MS_e = 1232$, $p < .06$, $F_2(1,15) = 8.44$, $MS_e = 286$, $p < .05$. When the bias supported the correct NV interpretation, ambiguous and unambiguous reading times differed by less than 1 ms/word in this region, $F$'s $< 1$, but in the unsupportive bias condition, ambiguous reading times were a mean of 20 ms/word longer than their unambiguous counterparts, $F_1(1,43) = 6.21$, $MS_e = 1440$, $p < .05$, $F_2(1,15) = 6.92$, $MS_e = 494$, $p < .05$.

Accuracy on comprehension questions for the experimental items with supportive bias was 90.3% for ambiguous and 89.8% for unambiguous sentences. For the unsupportive bias items, accuracy was at 87.5% and 92.5% for ambiguous and unambiguous items, respectively. These small differences yielded no reliable effects, all $p$'s $> .25$.

In sum, the reading time data showed that the manipulation of only one word in a noun phrase can create or remove effects of lexical category ambiguities, even though the frequency of the noun/verb ambiguous word itself always favored the incorrect NN interpretation. These results are not predicted by either a delay model or by a model in which only simple lexical frequency information dictated the interpretation that would be pursued. The results instead argue for the rapid use of combinatorial information during lexical category ambiguity resolution. It is reasonable to suppose that Frazier and Rayner would have found a similar pattern with stronger semantic biases.

Although these results support the constraint-based model, in which alternative interpretations are partially activated as a function of all available evidence, it is not yet clear what constitutes "available evidence" during ambiguity resolution. The next part of Experiment 2 investigates this question and explores the role of four stimulus factors in modulating the ambiguity effects observed in the reading time data. A better understanding of the nature of constraints is crucial to any constraint-based account of language comprehension.

### *Normative Data*

Two of the four stimulus factors investigated here are lexical, in that they reflect features of words in the ambiguous region independent of current context. The other two are combinatorial, reflecting information from combinations of words in the stimulus sentences. Data about each of these factors were collected from text corpora and from subjects' responses in a sentence completion task.

*Head vs modifying nouns.* The first lexical factor is the frequency with which the biasing nouns such as *warehouse* and *corporation* are used as heads of noun phrases in English, relative to the frequency with which they are modifiers in noun phrases. Most native speakers of English probably have the intuition that *warehouse* may often appear as a modifier of another noun, as in *warehouse inventory,* but noun phrases in which *corporation* is a modifier seem fairly rare. If this information is available on-line, it might be used to resolve head/modifier ambiguities in noun phrases, which could in turn affect interpretation of the subsequent lexical category ambiguity. That is, comprehenders might be quite confident that they have found the head of a noun phrase upon encountering the word *corporation,* and they would therefore be inclined to end the noun phrase at *corporation* and treat the next word, *fires,* as a verb (because if *corporation* is the head and *fires* is interpreted as another noun, an ungrammatical Head + Modifier sequence would result). As *warehouse* is frequently a modifier, however, comprehenders would be less sure that the noun phrase had ended at this word, and they might have difficulty determining whether *fires* is a verb or the head of the noun phrase, modified by *warehouse.*

Standard frequency counts do not record the percentage of uses in which a noun is the head of its noun phrase, so this information was coded from a corpus of text. The University of Pennsylvania's electronic *Wall Street Journal* (*WSJ*) corpus was chosen for its size; it contains over 40 million words from all articles published in the *Wall Street Journal* from 1987 through part of 1989. Although word frequency in the corpus correlates strongly with that of Francis and Kucera (1982), $r = .98$, there are some notable exceptions. The disadvantage of this corpus is that its technical financial articles use financial terms in ways that are unrepresentative of everyday language (e.g., the frequent occurrence of the noun phrase *sinking fund debentures* in the corpus inflates the modifier usage of *fund,* relative to everyday usage). Of the 32 biasing nouns from Experiment 2, 11 had business or financial connotations (*agency, bank, business, company, computer, corporation, employee, fund, office, official,* and *tax*). The *WSJ* corpus was not used for these words; the relative head vs modifier frequency was instead determined from the smaller but more representative Brown Corpus (containing one million words and forming the basis of the frequency information in Francis and Kucera, 1982). Frequency data for these 11 nouns were obtained entirely from the Brown corpus, with one exception: *computer* appeared only six times in the Brown corpus, and so a sample of sentences containing this word in the *WSJ* corpus were also coded, and the mean of the two codings was used.[3]

The coding proceeded as follows. First, a random sample of 100 sentences was selected from the appropriate corpus for each of the 32 biasing nouns; all sentences were included if less than 100 were available. One of two judges recorded the head/modifier status of the biasing noun in each sentence. Proper noun uses, hyphenated uses and verb uses (e.g., *to warehouse prisoners*) were excluded. For coding purposes, "noun phrase" was defined as the most local noun phrase that contained the biasing noun. For example, in the sequence [*the building near* [*the new warehouse*]$_{NP}$]$_{NP}$, *warehouse* was coded as the head of the noun phrase *the new warehouse,* even though that noun phrase is embedded in a larger noun phrase with *build-*

[3] A few additional words could perhaps be placed in the category of "business terms" and perhaps should be coded with the Brown corpus rather than the *WSJ* corpus. These words are *grocer, grocery, warehouse, client,* and *department.* In the case of the first four words, there are 0–14 instances of these words in the Brown corpus, and so the *WSJ* corpus was used. In the case of *department,* a sample was coded from the Brown corpus and compared to the *WSJ* coding; as the result differed by only a few percentage points, the *WSJ* coding was retained.

*ing* as its head. When the rate of head and modifier uses differed by less than 25%, an additional 75 sentences (if available) were coded to check that the usage was stable across two samples of sentences, and a third sample was coded if the first two were not within 5% agreement. A total of over 3200 sentences were coded.

*Noun/verb ambiguous word usage*. The second lexical factor is the relative frequency of noun vs verb interpretations of the category ambiguous words. All of these words had more frequent noun interpretations than verb interpretations, but there was some variation across items. Those items with a somewhat higher percentage of verb interpretations of the ambiguous word might have been easier to comprehend than those with extremely rare verb interpretations. A random sample of 300 sentences containing the noun/verb ambiguous words were extracted from the *WSJ* corpus (or all sentences if less than 300 were available), and coders tallied noun and verb uses. Over 3300 sentences were coded.

*Word cooccurrence*. The third factor, the frequency with which the biasing noun and noun/verb ambiguous word appear together in noun phrases in English, is combinatorial. Some of the ambiguous phrases, such as *warehouse fires* seem to be fairly rare concepts, so that the cooccurrence of *warehouse* and *fires* is probably very rare in English noun phrases. Other phrases in the stimuli, e.g., *grocery stores*, seem much more frequent. If cooccurrence information influences ambiguity resolution, then those items with high cooccurrence between the biasing noun and noun/verb ambiguous word should be more likely to be interpreted with the (incorrect) NN interpretation, whereas low cooccurrence items will be more likely to promote the NV interpretation. The percentage head text corpora samples described above were used to assess cooccurrence; coders tallied the number of occurrences of modifier/head noun phrase sequences that were exact matches to ambiguous phrases in the Experiment 2 materials. The number of singular versions of these noun phrases (e.g., *warehouse fire*) were also counted.

*Plausibility*. The fourth factor is the semantic plausibility biases created by the combination of the biasing noun and the ambiguous noun/verb ambiguous word, such as the plausibility of a fire in a warehouse vs a warehouse firing someone. The bias factor was manipulated at two levels in Experiment 2, based on the author's intuitions; here it was assessed on a continuum via sentence completion norms. The three factors described above could all contribute to the plausibility factor, and it is impossible to assess the semantic plausibility effects without assessing the other effects in the same measure. However, the lexical factors do not capture the specific semantic combinatorial constraints of each phrase in the stimulus sentences, and thus the sentence completion norms measure something in addition to the other factors.

Data were collected from 96 native English-speaking subjects who participated as part of a course requirement in an introductory psychology class. Each subject received one of eight different versions of a questionnaire. The items were presented in the questionnaire as sentence fragments, each ending at the noun/verb ambiguous word. Subjects were instructed to write a completion to each fragment, and two examples were given. Subjects were encouraged to write down the first ending that occurred to them, and it was stressed that there were no right or wrong answers.

A large number of unambiguous fillers were included in the questionnaire to limit the chance that the subjects would become aware of the ambiguities; each questionnaire version contained 4 supportive biased and 4 unsupportive biased experimental fragments, intermixed with 23 practice and filler fragments of the same length as the experimental items.

### *Normative Data Results*

Appendix B gives the values obtained for each item on the following four measures.

*Percentage head*. The 32 biasing nouns

were used as head of noun phrases 72.1% of the time, $SD = 24.8$. For the nouns that created the supportive bias condition, there were 85.5% head uses, compared to 58.8% head usage for the unsupportive biased items, a highly reliable difference, $t(30) = 3.59$, $p < .005$ (two-tailed).

*N/V ambiguous word interpretations.* The mean percentage of verb interpretations for the 16 N/V ambiguous word was 6.5%, $SD = 6$. This small range (0–19% verb interpretations) is not surprising given that the words were selected to have rare verb interpretations. This result correlated reliably with the 6.8% mean verb usage in Francis and Kucera (1982), $r = .42$, $p < .05$. The major discrepancies between the two counts came primarily from items with very few observations in the Francis and Kucera count.

*Cooccurrence.* The biasing nouns and matched N/V ambiguous verbs from the stimuli appeared together as modifier/head sequences in only 1.1% of the biasing noun sentences in the corpora (considering exact matches and singular variants (e.g., *warehouse fire*) together, the total was 2.1%; this sum will be termed "combined cooccurrence"). The percentages of exact and combined cooccurrence were higher for the unsupportive bias items (2.1 and 42.%, respectively) than for the supportive bias items (0.1% for both exact and combined). These differences across bias were marginally reliable in $t$ tests, $p$'s $< .10$ (two-tailed).

*Sentence completion norms.* One judge who was blind to the experiment manipulations scored subjects' fragment completions for which interpretation was chosen; three illegible responses were discarded. The mean percentage of NV completions of all items was 30.2%, $SD = 30.4$. There were 50.9% NV completions for supportive bias items, which may seem small for items intended to favor the NV interpretation, but recall that the noun/verb ambiguous words are themselves biased to the NN interpretation. The NN biased items were given NV interpretations only 9.6% of the time. This difference across Bias was robust, $t(15) = 8.47$, $p < .001$.

To determine whether the sentence completion norms could be predicted from the percentage head, percentage verb, or cooccurrence measures, a stepwise regression was conducted with these factors. Only the percentage head measure entered the regression, $r^2 = .19$, $F(1,30) = 7.07$, $p = .01$. The correlation between these two measures was positive, indicating that as the percentage of head uses increased for the biasing noun, the percentage of NV completions of the sentence fragments also increased. The percentage verb interpretation factor and the cooccurrence factors (calculated as exact or combined cooccurrence) did not account for any additional variance, $F$'s $< 1$.

*Effects on reading times.* To determine the extent to which the four measures predict reading times, simple regression analyses were first performed with each predictor separately. When evaluating the relative contribution of the measures here, it is important to recognize that the stimuli were not designed to provide a broad range of values on all four measures, and that the restricted range of the percentage verb and cooccurrence measures in these particular stimuli may result in an underestimate of their power to account for reading times, relative to the percentage head and sentence-completion measures, both of which had a broader range of values in these stimuli. The cooccurrence measure is especially problematic in this regard; even considering the more lenient combined cooccurrence measure, fully two-thirds of the items had 0 cooccurrence values, and most nonzero values were very close to zero. It was therefore decided to recode this factor as a dichotomous factor reflecting whether or not any cooccurrence examples had been found in the corpora. All items which had a combined cooccurrence value greater than 0 were set to a value of 0, and all items with 0 combined cooccurrence values were set to 1. The assignment of 0 and 1 values here was designed to make the coding more consistent with the other factors. The percentage head, percentage verb interpretation of the N/V ambiguous word, and the percentage NV interpretations in sentence completion had all been scaled to have higher values in supportive bias conditions, leading to a prediction of a negative correlation with reading times—as supportiveness increases, reading times decrease. High cooccurrence values, however, are more unsupportive, as they promote the incorrect NN interpretation. Assignment of a larger (1) value for noncooccurrence and 0 for cooccurrence yields a negative correlation prediction for this factor as well, so that negative correlations with reading times any factor indicates that this factor was helpful.

The correlations for reading times in the ambiguous and unambiguous conditions are shown in Table 3, both for the individual words and for the multiword disambiguation region defined in the reading portion of the experiment. All four measures were correlated to varying degrees with reading times for the words in the ambiguous condition, and three measures showed marginal correlations with reading times at individual words in the unambiguous condition. Most correlations at and after Word 11 were negative, indicating that as support for the correct NV interpretation increased, reading times decreased at these positions. This pattern was stronger in the ambiguous reading times than in the unambiguous times, with more reliable correlations in ambiguous than unambiguous conditions at individual words and in the Word 11–17 Region (the last column of the table). The cor-

TABLE 3
CORRELATIONS BETWEEN FOUR MEASURES AND EXPERIMENT 2 READING TIMES

| | Word in sentence | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | 9 fire(s) | 10 many | 11 workers | 12 each | 13 spring | 14 without | 15 giving | 16 them | 17 proper | 18 notice. | Words 11–17 |
| | Ambiguous reading times | | | | | | | | | | |
| Biasing noun: percentage head uses | .19 | .22 | −.36** | −.32* | −.43** | −.34* | −.26 | .08 | −.27 | −.20 | −.56** |
| N/V ambiguous word: percentage verb uses | .02 | .09 | −.39** | −.34* | .12 | .13 | .00 | .08 | −.05 | .16 | −.20 |
| Cooccurrence | .02 | .31* | .05 | −.20 | −.38** | −.32* | −.36** | −.21 | −.24 | −.23 | −.43** |
| Sentence completion norms: percentage N/V interpretation | .01 | .06 | −.21 | −.34* | −.43** | −.26 | −.07 | −.06 | −.11 | −.28 | −.43** |
| | Unambiguous reading times | | | | | | | | | | |
| Biasing noun: percentage head uses | −.21 | −.06 | −.33* | −.32* | −.21 | −.13 | −.19 | .19 | −.26 | −.25 | −.39** |
| N/V ambiguous word: percentage verb uses | .01 | −.01 | −.18 | −.23 | .13 | 05 | −.15 | −.10 | −.09 | .23 | −.14 |
| Cooccurrence | −.36* | .05 | −.36* | .09 | −.08 | −.15 | .21 | .18 | .00 | −.22 | −.12 |
| Sentence completion norms: percentage N/V interpretation | .22 | −.21 | −.16 | −.07 | −.34* | .07 | .04 | .04 | .03 | .02 | −.16 |

* $p < .10$.
** $p < .05$.

relations with the percentage verb uses of the noun/verb ambiguous word appear to be the weakest of the factors, presumably as a result of the restricted range of verb interpretation frequencies in these stimuli.

The fact that some correlations appear for the unambiguous conditions suggests that plausibility, as measured by the various factors, affected reading times in the absence of lexical category ambiguity. It is not surprising that plausible sentences were read somewhat more quickly than their less plausible counterparts, especially considering that plausibility information might help comprehenders resolve temporary head/modifier ambiguities that are inherent in even "unambiguous" noun phrases. The primary question here, however, is whether the four factors specifically affect lexical category ambiguity resolution processes. To address this issue, the next set of regressions examines the relationship between the four factors and ambiguity effects.

Ambiguous minus unambiguous reading time differences were calculated for each word, and a step-wise regression (with $F$ = 2.0 to enter) was conducted for the four factors at each word position. Although none of the factors had reliably predicted reading times for the category ambiguity itself (Word 9) in Table 3, two of the factors combined to predict the difference in reading times between ambiguous and unambiguous conditions at this position. The percentage head factor entered the regression first, accounting for a marginally reliable portion of the variance, $r^2 = .09, p < .10$. The sentence completion norms next entered the equation, yielding $r^2 = .20$, $F(2,29) = 3.68, p < .05$.[4] At this final step in the regression, the partial correlation between reading time differences and the percentage head measure was .47, while the partial correlation with the sentence completion norms was −.37. To understand this pattern, recall that ambiguous word reading times were shorter in the ambiguous condition than in the unambiguous times with unsupportive bias. It was hypothesized that after encountering an unsupportive bias noun, subjects were incorrectly forming a simple NN interpretation and were reading quickly, whereas they were forming the more complex NV interpretation and reading slowly in the unambiguous and supportive bias ambiguous conditions. The positive correlation between difference scores and the percentage head measure indicates that as the probability of a biasing noun's being the head noun increased, ambiguous reading times approached unambiguous reading times, producing a difference near zero, but as head probability decreased and a modifier interpretation became more likely, the ambiguous–unambiguous difference became negative as subjects read the more unsupportive biased items more quickly. Given this effect, there was then an effect of semantic plausibility in the regression: The items with a high percentage of sentence completions with the NV interpretation—those in which the biasing noun and ambiguous word were plausible as a noun + verb sequence—produced shorter reading times than those with less plausible noun + verb sequences. The other two factors did not account for any additional variance at this word, $F$'s < 1.

At later word positions, the sentence completion and cooccurrence factors accounted for reading time differences; there were no other cases in which multiple factors entered the stepwise regression. No factors entered the regression at Words 10, 11, and 13, but at Words 12 and 14, the sentence completions accounted for a marginal amount of variance, in both cases $r^2 = .07, p < .15$. At Word 15, the cooccurrence

[4] Because the $r^2$ in the first step was .09 and in the second step was .20, it may appear that the second factor, the sentence completion norms, is accounting for more variance. This is not the case, however. The sentence completion norms were able to account for variance in reading times only *after* the percentage head measure had partialed out some of the variance; the simple regression of the sentence completion norms on the difference scores at this word yielded only $r^2 = .03, F < 1$.

factor accounted for a significant portion of the variance in reading time differences, $r^2 = .26$, $p < .01$. Cooccurrence also accounted for a marginal portion of variance at Word 16, $r^2 = .10$, $p < .10$ and Word 17, $r^2 = .08$, $p < .15$. Finally, the sentence completion norms accounted for a small portion of variance at Word 18, the last word of the sentence, $r^2 = .07$, $p < .15$. In all of these cases, correlations with difference scores were negative, indicating that as these factors more strongly favored the NV interpretation, ambiguous–unambiguous difference scores decreased. In no case did the percentage verb interpretation factor enter any equation.

### *Discussion*

The important result in this experiment was that several lexical and combinatorial factors accounted for a significant portion of the variance in reading times of ambiguous sentences and in the difference between ambiguous and unambiguous reading times at the category ambiguity (Word 9) and in the disambiguation region. These results indicate that both lexical and combinatorial information influenced the ambiguity resolution process as soon as the ambiguity was introduced, and that the effects were continuous: The more the constraints promoted the NV interpretation, the smaller the effect of ambiguity. Models which posit the delayed use of this information in favor of using only the relative interpretation frequency of the noun/verb ambiguous word or in favor of waiting for a clear disambiguation are not supported by these data.

Given the interesting effects of the percentage head measure in accounting for ambiguity effects at the ambiguous word, it is important to explore the nature of this factor. Frequency effects in psycholinguistics are hardly controversial, but the notion that comprehenders keep track of the proportion of times that a noun is used as a head of a noun phrase may seem to be an unusual example of frequency information. Perhaps head frequency is merely a correlate of some other information that comprehenders access. Two such hypotheses were explored for the 32 biasing nouns in this experiment.

The first alternative questions whether animate nouns were more likely to be considered heads; perhaps comprehenders access and use animacy information rather than head frequency. Animacy could be coded in several ways, depending on whether words like *community* are considered animate, but no matter how animacy was defined, animate biasing nouns were not treated as heads more than inanimate nouns in these stimuli, all $t$'s $< 1$. The second hypothesis was that the percentage head effects are not frequency effects, but that a morphological analysis influences ambiguity resolution. A few of the biasing nouns have a morphologically related adjective or noun that could be used as a modifier and which has virtually the same meaning as the biasing noun used as a modifier (e.g., *corporation/corporate*). Perhaps morphological processes during word recognition provide information about related forms, which influences decisions about head location. Again, it is possible to code the stimulus items in several different ways (e.g., *technical* and *technician* do not mean exactly the same thing when used as modifiers, but they are clearly morphologically related), but no division on morphological criteria produced any reliable differences in the percentage head measure, $t$'s $< 1$. These results suggest that there is no reason to consider the percentage head measure as an artifact of these other factors.

Indeed, perhaps it is not so unlikely that the percentage head information is part of the subtle lexical information that is available for nouns. When a comprehender has encountered a noun in English, there is always an ambiguity of whether that noun is the head of the noun phrase or whether it is modifying some not-yet-encountered head noun. If we view decisions about noun phrase heads as another kind of ambiguity resolution, then we would expect that probabilistic information such as the percentage head information would affect this ambiguity resolution process, in the same way that relative frequency of alternatives affects lexical ambiguity resolution (Simpson, 1981).

The dichotomous cooccurrence factor also accounted for reading times. This result is interesting because it suggests that comprehenders keep track of the cooccurrence of individual words (in this case modifier/head pairs) in the input, and that this effect is distinct from lexical information about the relative frequency of a noun's being a head or a modifier, or from semantic plausibility, as measured by sentence completion norms. This result is compatible with other recent demonstrations of cooccurrence effects during sentence processing (Juliano & Tanenhaus, 1993), but it seems to contradict previous findings that the predictability of individual words in text is in general extremely low and rarely useful during comprehension (Fischler & Bloom, 1979). In fact the present results may be quite consistent with these earlier findings, in that most of the Experiment 2 modifier/head pairs had cooccurrence percentages at or near zero. The effects of cooccurrence might therefore be carried by those rare items with high cooccurrence values, whereas other information like percentage head information and plausibility (as was measured by sentence completion norms) might be important to a broader range of items. Additional research including explicit manipulations of the various factors is needed to test this and any other hypothesis about the relative contribution of these different constraints during ambiguity resolution.

## General Discussion

Several different accounts of lexical category ambiguity resolution have been considered here. One alternative is the delay model (Frazier & Rayner, 1987), in which the parser delays processing of the ambiguous phrase until a definitive disambiguation is reached. A second "local analysis" hypothesis posits that one alternative interpretation is chosen rapidly on the basis of some metric such as simplicity or the relative frequencies of the alternative interpretations. The third hypothesis, which has been advocated here, is that alternative interpretations are partially activated and rapidly constrained by probabilistic information such as the relative frequency of the alternative interpretations, combinatorial semantic information, and word cooccurrence information. The results of Experiments 1 and 2 argue against the first two alternatives. Experiment 1 showed that a pattern of reading times that had been taken to support the delay model appear to be due to added differences across ambiguous and unambiguous conditions. Experiment 2 showed that the effects of ambiguity depended on semantic biases. Crucially, the effects of bias appeared at the introduction of the ambiguity; this result argues against models in which semantic information is delayed in its influences on ambiguity resolution.

The argument for a constraint-based account of lexical category ambiguities offers the possibility that theories of ambiguity resolution at different levels of linguistic representation could be unified under one processing architecture. That is, the nature of the ambiguity and the available evidence may differ across different levels of linguistic representation, but the processes that resolve the ambiguities at each level may all operate via the partial activation of alternative interpretations in the light of available evidence. This account is clearly compatible with models of lexical ambiguity resolution (Kawamoto, 1988; Seidenberg et al., 1982; Tanenhaus et al., 1979) and versions have been suggested for syntactic ambiguity resolution (Altmann & Steedman, 1988; Burgess & Hollbach, 1988; MacDonald, 1993; MacDonald et al., 1993; McClelland et al., 1989; Trueswell et al., 1993a) and for quantifier scope ambiguity resolution

(Kurtzman & MacDonald, 1993). This view contrasts with one in which ambiguities at different levels of representation require qualitatively different processing mechanisms (Frazier, 1989). The continued investigation of ambiguities that span several levels of representation should inform this debate.

At its current stage of development, the constraint-based accounts have a weakness in the specification of the kinds of information that can affect the ambiguity resolution process and how this information is weighed. The second half of Experiment 2 was designed to address this concern, and it provides some important information about the kinds of constraints that are used during ambiguity resolution, as well as information concerning the time course of constraint use. These results have clearly not exhausted the possible factors that could modulate ambiguity resolution, however. For example, it seems likely that the noun/verb ambiguous word's bias to the noun or verb interpretation should be extremely important, even though it did not matter within the very small range of variation that was allowed here. Another possibly important factor is the length of a noun phrase. Noun phrases can in principle be infinitely long, but in practice, those with very many words are rarer than their shorter counterparts. Thus as the length of the noun phrase grows, it becomes increasingly likely that the end of the noun phrase has been reached, so that the next word will be more likely to be considered as a verb rather than part of a still longer noun phrase. This sort of information is combinatorial rather than lexical, but it differs from the contextual information in Experiment 2 in that it relies on comprehenders' knowledge of typical noun phrase length in the abstract, rather than semantic information in the current input. It is currently unknown whether such factors influence ambiguity resolution, and if so what the strength and time course of such factors are.

The notion of a "syntactic frequency" factor such as noun phrase length is particularly incompatible with parsing models in which syntactic structure is constructed anew each time, because graded effects such as word frequency have traditionally been considered part of the "access" mechanism—the more often a word is accessed, the faster that process becomes (e.g., through thresholds in the Logogen model, Morton, 1969). On this view, there would be no mechanism for keeping track of noun phrase length, because syntactic structures like noun phrases are not stored and accessed.

An alternative view that does not maintain the strict dichotomy between accessed words and constructed syntactic representations is receiving increasing support from a number of independent sources. First, syntactic theories have recently placed increasing importance on lexical representations (Levin & Pinker, 1991), so that lexical representations do much of the "work" of constraining the form of the syntactic structure. Second, recent psycholinguistic research has shown that comprehenders are sensitive to extremely subtle probabilistic syntactic and plausibility information during syntactic processing (e.g., Altmann & Steedman, 1988; MacDonald, 1993, Pearlmutter & MacDonald, 1992; Taraban & McClelland, 1988; Trueswell et al., 1993a, 1993b); the rapid use of probabilistic syntactic information is not consistent with a purely "constructed" view of syntactic processing. Third, there is evidence that syntactic structures, or at least the processes that build them, do have activation levels (Bock, 1986; Bock & Loebell, 1990; see also Berg, 1991), suggesting that partial activation of syntactic structures is a coherent concept. Finally, several recent connectionist simulations of language processing have demonstrated that sensitivities to cooccurrence and other probabilistic information at many different levels of linguistic representation are a natural consequence of the learning algorithms that are employed in these models (e.g., Elman, 1990; McClelland et al., 1989; Seidenberg, 1992). On this view, frequency effects such as head noun frequency, word cooccurrence, etc., are entirely expected, and it would be the absence of effects such as these that would constitute a surprising result, according to this perspective.

In sum, this combination of work from linguistics, psycholinguistics, and simulation modeling suggests that psycholinguists may have been substantially underestimating the probabilistic information that is available during comprehension and the speed with which probabilistic information can be used. Constraint-based models of language comprehension pose an opportunity to unify accounts of lexical and syntactic processing within one general processing architecture and to explore a number of subtle interacting constraints. At the same time, such models challenge researchers to develop strong predictions about the nature of interactions between different levels of linguistic representation and the limits on the use of probabilistic information during comprehension.

## Appendix A

### Raw and Length-Adjusted Reading Times per Word in Experiment 1 for the Two-Word Ambiguous Region (Amb 1 and Amb 2) and the Four-Word End Region (End 1–4).

| | Raw reading times (ms) | | | | | |
|---|---|---|---|---|---|---|
| | Word | | | | | |
| | Amb 1 | Amb 2 | End 1 | End 2 | End 3 | End 4 |
| NN interpretation | | | | | | |
| 1.* Inherent Amb, the | 334 | 364 | 396 | 361 | 371 | 533 |
| 2.* Inherent Amb, this/these | 340 | 375 | 373 | 361 | 363 | 482 |
| 3. Inherent Unamb, the | 342 | 369 | 374 | 370 | 355 | 486 |
| 4. Inherent Unamb, this/these | 356 | 391 | 378 | 351 | 361 | 523 |
| NV interpretation | | | | | | |
| 5.* Inherent Amb, the | 329 | 352 | 383 | 390 | 377 | 519 |
| 6.* Inherent Amb, this/these | 347 | 369 | 433 | 370 | 377 | 511 |
| 7. Inherent Unamb, the | 356 | 379 | 368 | 361 | 366 | 493 |
| 8. Inherent Unamb, this/these | 351 | 399 | 384 | 382 | 362 | 488 |

| | Length-adjusted reading times (ms) (raw minus length-predicted) | | | | | |
|---|---|---|---|---|---|---|
| | Word | | | | | |
| | Amb 1 | Amb 2 | End 1 | End 2 | End 3 | End 4 |
| NN interpretation | | | | | | |
| 1.* Inherent Amb, the | −67.22 | −29.09 | 24.32 | −28.52 | −11.78 | 112.17 |
| 2.* Inherent Amb, this/these | −60.25 | −18.37 | 0.47 | −26.50 | −25.31 | 62.92 |
| 3. Inherent Unamb, the | −77.27 | −21.97 | 0.51 | −16.59 | −30.85 | 62.94 |
| 4. Inherent Unamb, this/these | −61.67 | −1.17 | 4.20 | −29.73 | −27.96 | 102.49 |
| NV interpretation | | | | | | |
| 5.* Inherent Amb, the | −70.50 | −40.70 | 11.84 | −0.29 | −13.21 | 112.06 |
| 6.* Inherent Amb, this/these | −56.51 | −25.78 | 66.70 | −15.46 | −8.70 | 101.83 |
| 7. Inherent Unamb, the | −45.98 | −18.28 | −0.75 | −35.00 | −18.53 | 86.40 |
| 8. Inherent Unamb, this/these | −52.33 | −0.21 | 14.98 | −11.32 | −18.14 | 78.61 |

*Note*. Conditions that are replications of Frazier and Rayner (1987) are marked with an asterisk.

## Appendix B

Stimulus sentences and normative data from Experiment 2. Biasing nouns are given in the order supportive/unsupportive bias in the sentence, each followed by the percentage head value. The first number following the N/V ambiguous word indicates the percentage verb usage in the corpora. The next pair of numbers indicate "combined" cooccurrence percentages (exact matches plus singular variants) for the previous two words, in the order (supportive bias + category ambiguity cooccurrence/unsupportive bias + category ambiguity cooccurrence). The sentence-final numbers reflect percentage of NV interpretations in the sentence completion norms in the supportive bias/unsupportive bias conditions. Unambiguous conditions were created by changing the plurality of the biasing noun and agreeing verb.

1. The townspeople were pleased that the new sheriff (99)/prison (64) guards (29) (0/0) the community from the dangerous criminals in the area. 33/0
2. The doctor refused to believe that the shrine (100)/miracle (85.5) cures (19) (0/2.4) people of many fatal diseases like cancer and AIDS. 90/41.5
3. The prospective students were informed that the university (50)/fraternity (64) houses (3) (0/5.6) a handful of promising undergraduates in luxurious new rooms. 33.5/9
4. The government auditors were surprised that the agency (82)/bank (88) accounts (15) (0/0) for all expenditures at the end of the year. 26/0
5. The efficiency expert recently reported that the company (89)/office (92) supplies (13) (0/0.5) more than their share of manpower in getting clients. 54/4;
6. It says in the manual that the technician (100)/computer (47) programs (0) (0/0) the printer to write the text with wide margins. 46/0
7. Bob explained to the police that the report (100)/official (29) documents (4) (0/0) the effects of radiation on the citizens of Smithville. 37.5/0
8. Ellen was pleased to see how the curtain (82)/window (81) frames (6) (0/0) the scene outdoors, bringing the dull room to life. 8/0
9. Since the company's policy was revised, the community (60)/employee (48) (0/14.3) benefits (8) from increased awareness of the need for adult education. 54.5/0
10. At the beginning of every summer, the coach (86)/college (26) loans (0) (0/0.8) equipment for the scouts to use in their camp. 100/29
11. Fred wanted to be sure that the client (78)/business (78) contacts (1) (0.6/0) him directly if the merchandise does not arrive promptly. 54.5/0
12. The seamstress may have problems, because the ruler (99)/tape(69) (0/0.6) measures (11) accurately enough to make only the simplest of dresses. 83/58
13. In developing the new promotion policy, the department (69)/army (77) bases (5) (0/1.7) their decisions on some technical reports from top analysts. 78/0
14. The young inspector soon learned that the grocer (97)/grocery (4) stores (1) (0/34.9) the fresh beef in huge, locked refrigerators at night. 24.5/4
15. The union told the reporters that the corporation (98.5)/warehouse (54) fires (8) (0/2.8) many workers each spring without giving them proper notice. 62.5/4
16. The residents were all pleased that the fund (80)/tax (36) returns (8) (0/3.7) some of the money to the community each year. 29/4

## References

Altmann, G. T. M. (Ed.) (1990). *Cognitive models of speech processing.* Cambridge, MA: MIT Press.

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

Berg, G. (1991). *Learning recursive phrase structure: Combining the strengths of PDP and x-bar syntax.* Department of Linguistics and Cognitive Science Tech Report, SUNY Albany.

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, **18**, 355–387.

Bock, J. K., & Loebell, H. (1990). Framing sentences. *Cognition*, **35**, 1–40.

Burgess C., & Hollbach, S. C. (1988). A computational model of syntactic ambiguity as a lexical process. *Proceedings of the 10th Annual Conference of the Cognitive Science Society.* (pp. 263–269).

Elman, J. L. (1990). Representation and structure in connectionist models. In G. T. M. Altmann (Ed.)., *Cognitive Models of Speech Processing* Cambridge, MA: MIT Press.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *The Journal of Memory and Language*, **25**, 348–368.

Francis, W. N., & Kucera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar.* Boston: Houghton Mifflin.

Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading.* Hillsdale, NJ: Erlbaum.

Frazier, L. (1989). Against lexical generation of syntax. In W. Marslen-Wilson, (Ed.) *Lexical representation and process.* Cambridge, MA: MIT Press.

Frazier, L., & Rayner, K. (1987). Resolution of syntactic category ambiguities: Eye movements in parsing lexically ambiguous sentences. *Journal of Memory and Language*, **26**, 505–526.

Juliano, C., & Tanenhaus, M. K. (1993). Contingent frequency effects in syntactic ambiguity resolution. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society.*

JUST, M. A., & CARPENTER, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, 87, 329–354.

JUST, M. A., CARPENTER, P. A., & WOOLLEY, J.D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 3, 228–238.

KAWAMOTO, A. (1988). Interactive processes in the resolution of lexical ambiguity. In S. Small, G. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Computational, linguistic, and psychological perspectives*. NY: Morgan Kauffman.

KING, J., & JUST, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, 30, 580–602.

KURTZMAN, H., & MACDONALD, M. C. (1993). Resolution of quantifier scope ambiguities. *Cognition*, in press.

LEVIN, B., & PINKER, S. (1991). Introduction to special issue of *Cognition* on lexical and conceptual semantics. *Cognition*, 41, 1–8.

MACDONALD, M. C. (1993). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes*, in press.

MACDONALD, M. C., PEARLMUTTER, N. J., & SEIDENBERG, M. S. (1993). The lexical nature of syntactic ambiguity resolution. Manuscript in preparation.

MCCLELLAND, J. L., ST. JOHN, M., & TARABAN, R. (1989). Sentence comprehension: A parallel distributed processing approach. *Language and Cognitive Processes*, 4, SI287–335.

MORTON, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165–178.

PEARLMUTTER, N. J., & MACDONALD, M. C. (1992). Plausibility and syntactic ambiguity resolution. *Proceedings of the Fourteenth Annual Conference of the Cognitive Society*.

RAYNER, K., CARLSON, M., & FRAZIER, L. (1983). The interaction of syntax and semantics during sentence processing. *Journal of Verbal Learning and Verbal Behavior*, 22, 358–374.

SEIDENBERG, M. S. (1992). Connectionism without tears. In S. Davis (Ed.), *Connectionism: Theory and Practice*. London/New York: Cambridge Univ. Press.

SEIDENBERG, M. S., TANENHAUS, M. K., LEIMAN, J. M., & BIENKOWSKI, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, 14, 489–537.

SIMPSON, G. B. (1981). Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior*, 20, 120–136.

SMALL, S., COTTRELL, G., & TANENHAUS, M. (Eds.). (1988). *Lexical ambiguity resolution*. San Mateo, CA: Morgan Kaufmann.

SPIVEY-KNOWLTON, M., TRUESWELL, J., & TANENHAUS, M. (1993). Context effects in syntactic ambiguity resolution: Discourse and semantic influences in parsing reduced relative clauses. *Canadian Journal of Psychology*, in press.

SWINNEY, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645–660.

TANENHAUS, M. K., LEIMAN, J. M., & SEIDENBERG, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, 18, 427–440.

TARABAN, R., & MCCLELLAND, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language*, 27, 597–632.

TRUESWELL, J. C., TANENHAUS, M. K., & GARNSEY, S. M. (1993a). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language*, in press.

TRUESWELL, J. C., TANENHAUS, M. K., & KELLO, C. (1993b). Verb specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory & Cognition*, in press.

TYLER, L. K., & MARSLEN-WILSON, W. D. (1977). The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behavior*, 16, 683–692.

(Received October 22, 1992)
(Revision received March 5, 1993)