Journal of Experimental Psychology:
Learning, Memory, and Cognition
2002, Vol. 28, No. 4, 748–769

0278-7393/02/$5.00 DOI: 10.1037//0278-7393.28.4.748

The Processing Role of Structural Constraints on the Interpretation of Pronouns and Anaphors

William Badecker and Kathleen Straub
Johns Hopkins University

The authors report 6 self-paced word-by-word reading studies of how morphosyntactic agreement, focus status, and the structural constraints of binding theory apply and interact during the online interpretation of pronouns (e.g., *him, her*) and anaphors (e.g., *himself, each other*). Previous studies held that structural conditions on coreference work as interpretive filters that impose exceptionless limits on which antecedent candidates can be evaluated by subsequent, content-based processes. These experiments instead support an interactive-parallel-constraint model, in which multiple weighted constraints (including constraints on binding) simultaneously influence the net activation of a candidate during preselection stages of antecedent evaluation. Accordingly, structurally inaccessible candidates can interfere with antecedent selection if they are both prominent in focus structure and gender–number compatible with the pronoun or anaphor.

Apprehending referential dependencies constitutes a central part of language understanding. Formal and experimental studies show that coreference processing is influenced by (a) the morphological and syntactic properties of the referentially dependent expression and its potential antecedents (Chomsky, 1981; Lasnik, 1989; Reinhart, 1983a, 1983b), (b) structural parallelism (Chambers & Smyth, 1998; Smyth, 1994), (c) the "causal" semantics of sentence predicates and connectives (Caramazza, Grober, Garvey, & Yates, 1977; Garnham, Traxler, Oakhill, & Gernsbacher, 1996; McDonald & MacWhinney, 1995), (d) the prominence or salience of the various entities referred to in the local discourse (Clifton & Ferreira, 1987; Gordon, Grosz, & Gilliom, 1993; Gordon, Hendrick, Ledoux, & Yang, 1999; Grosz, Joshi, & Weinstein, 1995; Gundel, Hedberg, & Zacharski, 1993; Prince, 1981), and (e) the real-world knowledge that speakers and listeners share about these entities. The studies reported here investigate coreference processing as a means of evaluating the role and influence of grammatical knowledge—specifically, knowledge about grammatical accessibility and the morphosyntactic features of pronouns. Our studies indicate that the accessibility constraints of the binding theory constitute one of several parallel, interactive determinants of pronoun (and reflexive) interpretation. The processes involved in assigning coreference and non-coreference provide a context in which we can examine what sorts of information are exploited and how they interact in language understanding.

Previous experimental research on pronoun interpretation has focused on two constraints on coreference processing: a pronoun's morphosyntactic features (person, number, and gender) and the status of a pronoun's potential antecedents in the local focus of attention. A pronoun's morphosyntactic features isolate a set of antecedent candidates (Arnold, Eisenband, Brown-Schmidt, & Trueswell, 2000; Erlich, 1980; Garnham & Oakhill, 1985; MacDonald & MacWhinney, 1990; Matthews & Chodorow, 1988; Shillcock, 1982; although see Greene, McKoon, & Ratcliff, 1992). Studies such as Caramazza et al.'s (1977) and Erlich's (1980) show that processing is facilitated if the local context contains only one referent that conforms to the pronoun's person, number, and gender. However, even if more than one referent fits these features, the potential ambiguity will not in itself complicate the process of identifying the pronoun's intended antecedent because the status of these feature-matching discourse entities will often be distinct in the local focus of attention. Relative prominence in the discourse also plays a large part in determining whether an entity is included in the initial candidate set (Grosz et al., 1995; Gundel et al., 1993). A pronoun that refers to the most prominent member of the current focus of attention is generally easier to interpret than repeated names or explicit descriptions that could otherwise refer successfully to the same entity (Gordon & Chan, 1995; Gordon et al., 1993, 1999). In addition to the morphosyntactic and focus-related constraints, principles of grammar that are based on structural relationships also govern the interpretations that can be assigned to referentially dependent expressions (Chomsky, 1981; Freidin, 1986; Lasnik, 1989; Reinhart, 1983a). To formulate these syntactic constraints, linguists distinguish two major classes of referentially dependent expressions: pronominals (*I* or *me*, *you*, *he* or *him*, *she* or *her*, etc.) and anaphors (reflexives [*myself*, *yourself*, *himself*,

William Badecker and Kathleen Straub, Department of Cognitive Science, Johns Hopkins University.

The research reported in this article was supported by National Institute of Mental Health Grant MH52675 to Johns Hopkins University. We thank Steve Anderson, Luigi Burzio, Chuck Clifton, Janet Fodor, Lyn Frazier, Mike McCloskey, Brenda Rapp, Paul Smolensky, Michael Tanenhaus, and Colin Wilson for many helpful comments, questions, and suggestions. We are also grateful to Fernanda Ferreira and two anonymous reviewers, whose comments and criticisms got us to clarify our discussion and claims. Finally we thank Angelique Budaya, Manas Mohapatra, Judy Huang, Samantha Levien, Courtney Burnette, Kaori Ozawa, Ben Kram, and Courtney McCollum for their assistance in scheduling and running participants for the experiments described in this article.

Correspondence concerning this article should be addressed to William Badecker, Department of Cognitive Science, Johns Hopkins University, 3400 North Charles Street, Baltimore, Maryland 21218-2685. E-mail: badecker@jhu.edu

herself, etc.] and reciprocals [*each other*]). Pronominals and anaphors are governed by complementary principles of binding theory (Chomsky, 1981) that rely on the structural relation of *c-command* to describe allowable coreference relations. A phrase α c-commands a phrase β if and only if every syntactic constituent that contains the phrase α also contains the phrase β. If β is coreferential with α and is c-commanded by α, then the phrase β is said to be bound by the phrase α. According to Principle A of the theory, an anaphor must be bound by an antecedent in its domain (where, for present purposes, the anaphor's domain is the minimal clause that contains both the anaphor and its subject). Principle B of binding theory stipulates that pronominals must not be bound by an antecedent in the pronoun's domain. Thus, pronouns and anaphors will appear in complementary distribution. A third principle (Principle C) prohibits the binding of proper names and other full referring expressions.

Binding principles account for the contrast between the sentences in Example 1, in which the pronouns and reflexives are bound by the italicized phrases preceding them, and those in Example 2, in which the binding relations indicated in the example are prohibited.

(1) a. *John* hurt *himself.*
b. *John* said that Mary hurt *him.*
c. *John* said that Bill hurt *him.*
(2) a. *John* said that Mary hurt *himself.* (violates Principle A)
b. (John said that) *Bill* hurt *him.* (violates Principle B)
c. *He* said that *Bill* fell. (violates Principle C)

Recent experimental work by Gordon and Hendrick (1997) confirmed that native English speakers' coreference judgments pattern along the predictions of the binding theory for contexts like those in Examples 1 and 2. However, that same study also reported cases in which speakers are reluctant to attribute coreference relations that are clearly licensed by the binding theory. Two such structures appear in Examples 3 and 4.

(3) [*Jeff* and Cindy] asked the bakery to make a cake for *him.*
(4) *His* brother thinks that *John* eloped.

In isolated contexts such as these, syntactic structure guides the introduction of new entities into the discourse model. Because *Jeff* and *Cindy* are introduced in a conjoined noun phrase (NP) in Example 3, they are initially represented in the discourse model as one (albeit plural) entity (Gordon & Hendrick, 1997; Gordon et al., 1999). Thus, resolving the antecedent for the subsequent pronoun *him* is relatively difficult because doing so requires unpacking the conjoined NP into its constituent elements. Not surprisingly, the acceptability of these coreference links improves substantially if the pronoun's intended referent is independently introduced and focused in the preceding discourse, as in Examples 5 and 6.

(5) Tomorrow is Jeff's birthday. Jeff and Cindy asked the bakery to make a cake for him.
(6) John went to visit his girlfriend yesterday and still hasn't returned. His brother thinks that John eloped.

This shift in felicity suggests that binding principles may not be independently sufficient to capture the factors that influence the process of coreference resolution: For Examples 5 and 6 the binding principles interact with discourse focus constraints. Altering the focus structure will not solve the problems with Example 2, though, because these sentences involve violations of explicit structural constraints on coreference relations.

Accordingly, an adequate description of the cognitive mechanisms that mediate the resolution of referential dependencies will need to identify (a) the various constraints that determine a unique interpretation for a pronominal or an anaphor, (b) how those constraints interact, and (c) how and when individual constraints will be sufficient to predict coreference judgments. A natural starting point for such an exploration is to focus on a specific constraint and explore its influence on antecedent identification when the effects of other constraints are systematically varied. For example, what happens when the focus status and morphosyntactic properties of a noun make it a strong antecedent candidate despite the fact that it is structurally inaccessible? That is, do the earliest stages of the interpretive process evaluate NPs that are ultimately excluded by purely structural criteria? Or is structural accessibility an immediate and primary filter for candidate status?

Surprisingly little experimental work addresses when and how binding theory mechanisms interact with other constraints on referential dependencies. In one notable exception, Nicol and Swinney (1989) presented a series of cross-modal priming studies that used sentences such as those shown in Examples 7a and 7b to probe whether and when the NPs preceding a pronoun or reflexive are evaluated as potential antecedents.

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

Using lexical recognition latencies for semantic associates of the to-be-evaluated NPs (relative to unrelated, matched controls) as a dependent measure, Nicol and Swinney reported facilitation for structurally accessible antecedents only. Specifically, at the offset of the reflexive in Example 7a, semantic associates of *doctor* are facilitated by approximately 100 ms, although associates for *boxer* or *skier* show no priming effect. Conversely, at the offset of the pronoun in Example 7b, semantic associates of both *boxer* and *skier* each realize approximately 50 ms of facilitative priming, whereas associates of *doctor* are not primed. Nicol and Swinney reasoned that the mechanism for anaphor resolution reactivates the antecedent and that activation then spreads to the antecedent's semantic associates. They maintained that the cross-modal, associate-priming paradigm provides a window to the candidate-evaluation process rather than the result of selection, because priming is observed for both potential antecedents in pronoun contexts like Example 7b. Because facilitative priming occurred only for semantic associates of grammatically possible antecedents, Nicol and Swinney adopted the position that a grammatical filter determines the candidate-referent set that content-based mechanisms subsequently evaluate. That is, "[t]he initial set of candidate antecedents contains *all and only* [italics added] those referents that bear the appropriate syntactic relation to the referentially dependent NP" (Nicol & Swinney, 1989, p. 7). If a particular NP fails the binding theory's syntactic relation test, then from the outset it is excluded from the set of referent candidates assessed by other evaluative mechanisms. In subsequent discussion, we refer to this as the *initial-filter model*.

The initial-filter model describes one of several alternative interpretations for the pattern of semantic-associate priming that Nicol and Swinney (1989) reported. An alternative interpretation is that the observed facilitation derives from the effects of selection rather than from candidate evaluation. The priming pattern for the reflexive context follows straightforwardly. Grammatical constraints stipulate that the antecedent of the reflexive in Example 7a is *doctor*. Therefore, one would expect (correctly) that facilitation would be observed for semantic associates of *doctor* but that there would be no facilitation for associates of *boxer* and *skier*. To interpret the results described for the pronoun resolution context (Example 7b), it is critical for one to note that the binding theory allows that either *boxer* or *skier* constitutes an equally plausible antecedent for *him*. Because other cues that might have influenced the process are carefully controlled in Nicol and Swinney's stimuli, it is reasonable to assume that participants ultimately selected each antecedent about half of the time. Therefore, the reported facilitation could represent priming for the associates of the selected antecedent *boxer* (and no priming for associates of *skier*) on half of the trials and priming for the associates of the selected antecedent *skier* (and no priming for associates of *boxer*) on the other half of the trials. In terms of the overall priming patterns, one might therefore expect that the outcome of averaging the facilitation effect across the bimodal distribution would decrease the absolute effect size by roughly one half. Although this alternative hypothesis effectively describes the reported priming pattern, it requires that pronoun resolution is fast enough to allow full antecedent selection—and its consequent activation of semantic associates—within the observed response window. This assumption regarding the required processing speed may undermine the plausibility of this interpretation.

A second, more credible alternative to the initial-filter model is that the observed priming pattern reflects the net activation for each antecedent candidate summed across parallel acting constraints during the preselection stage of antecedent evaluation. In contrast to Nicol and Swinney's (1989) assumption that initial processing of coreference relations entails exclusive, inhibitory filtering, work by MacDonald and MacWhinney (1990) and Gernsbacher (1989) converged on the claim that antecedent selection includes both excitatory and inhibitory mechanisms. We suggest that the various constraints acting during the antecedent-evaluation process independently assign either positive or negative activation to a candidate NP. During the preselection stage of candidate antecedent evaluation, the total activation level of an NP is the sum of the positive and negative activation apportioned to it by separate parallel acting constraints. Crucially, if a candidate NP receives positive support from one constraint (e.g., a content-based constraint such as morphosyntactic agreement) and inhibition from another (e.g., a structurally based constraint from the binding theory), then the excitatory activation for that candidate will be functionally canceled out. We refer to this as the *interactive-parallel-constraint model*. According to this model, cross-modal, associate priming may be useful to determine the net gain that a particular candidate antecedent received during initial processing, but the paradigm will not be sufficient to establish that a phrase has or has not been evaluated by content-based constraints in every circumstance. Therefore, additional evidence from alternative methodologies is required to converge on the set of candidates evaluated by the antecedent-resolution process. We present a series of word-by-word, self-paced reading experiments that vary the applicability of the major constraints that bear on the interpretation of referentially dependent expressions. In so doing, we provide data that further probe the membership of an initial candidate set.

Previous self-paced reading studies reveal that reading times for material immediately following a pronoun varies according to the antecedent options available for the pronoun (Badecker & Straub, 1994). In contexts like Example 8, the longest reading times occurred when no gender-matched antecedent was available for the pronoun (as shown in Example 8d). The fastest reading times were observed for Example 8a, in which the only NP matching the pronoun in gender was the matrix subject. Intermediate reading times were observed for sentences like Examples 8b and 8c.

(8) a. Kenny assured Lucy that he was prepared for the new job.
b. Julie assured Harry that he was prepared for the new job.
c. Kenny assured Harry that he was prepared for the new job.
d. Julie assured Lucy that he was prepared for the new job.

There are two important observations embedded in this finding. First, the processing load observed in the context of Example 8d—in which there is no suitable antecedent candidate—is significantly greater than that observed in contexts with at least one suitable antecedent (Examples 8a and 8b). We refer to this as the *no-antecedent effect*. Second, additional processing is required to identify a unique antecedent when, as in Example 8c, more than one candidate exists. We refer to the processing-load difference between sentences like Examples 8a and 8c as the *multiple-candidate effect*. Whereas the interactive-constraint model holds that this effect will be observed whenever more than one salient candidate NP is available in the local discourse, the initial-filter model predicts that the presence of additional matching NPs will increase processing load only when the additional candidates are grammatically accessible to the pronoun or reflexive.

The experiments presented here offer self-paced reading evidence that supplements Nicol and Swinney's (1989) cross-modal, semantic-associate priming results to provide a clearer picture of the candidate identification and selection process. Taken together, the data suggest that the initial-filter model may overstate the role of accessibility constraints in the early stages of candidate identification. We argue that although the effects of structural constraints on coreference may emerge quickly in the comprehension process, they nonetheless exert their influence over time and in parallel with other interpretive mechanisms. Whereas the initial-filter model asserts that referential processing is a cascaded process in which structure based principles act as a preemptive filter on the input to subsequent semantic and pragmatic processing mechanisms, we propose that all feature-matching referents that are members of the local focus of attention are evaluated by content-based mechanisms in the antecedent selection process—independent of their structural relation to the to-be-interpreted pronoun or reflexive. That is, the interactive-constraint model suggests that the initial stages of interpreting referentially dependent expressions are governed not only by structural constraints imposed by the binding theory but also by the focus status of candidate antecedents and by the morphosyntactic properties of the pronoun or reflexive that is in need of an interpretation.

Experiment 1: Evaluating Structurally Inaccessible Candidates

If membership in the initial candidate antecedent set is limited by the binding principles, then the content of structurally inaccessible NPs should not influence the processing load associated with resolving pronominal relations. Experiment 1 tested this prediction by contrasting self-paced reading performance for sentences that vary the morphosyntactic compatibility of potential antecedent NPs with the to-be-resolved pronoun (henceforth referred to as *gender compatibility*) and their grammatical accessibility. In these sentences, exemplified in Examples 9–12, gender compatibility was varied by alternating names normed for gender recognizability (along with other factors) within a given sentence position (e.g., *John–Jane* for the matrix subject position). Grammatical accessibility is varied according to Binding Principle B.

(9) *multiple match*
John thought that Bill owed him another chance to solve the problem.

(10) *accessible match*
John thought that Beth owed him another chance to solve the problem.

(11) *inaccessible match*
Jane thought that Bill owed him another chance to solve the problem.

(12) *no match*
Jane thought that Beth owed him another chance to solve the problem.

Both the initial-filter model and the interactive-constraint model predict that the fastest reading times will be observed for the accessible-match condition and that an observable no-antecedent effect will be observed for the no-match cases. However, the models differ with respect to the relative processing load anticipated for the multiple-match condition.

According to Principle B, the matrix subject *John* is a structurally acceptable antecedent for the pronoun in these sentences, whereas the subject of the embedded clause, *Bill,* is not. If binding principles preemptively eliminate all grammatically inaccessible nouns from the antecedent-candidate set, as the initial-filter model maintains, then Examples 9 and 10 should appear identical to subsequent coreference-processing mechanisms. That is, the gender compatibility of inaccessible names should have no observable effect on anaphor resolution. In contrast, the interactive-constraint model posits that the binding principle constitutes only one of several parallel-acting mechanisms for evaluating candidate NP. For the accessible match versus multiple match contrast, if the morphosyntactic mechanism computing gender compatibility for the pronoun and preceding NPs can identify two potential candidates in Example 9, then additional processing will be needed to select a unique antecedent. This should result in an increased processing load following the pronoun. Therefore, we predicted that reading times for the multiple-match condition would be longer than those for the accessible-match condition. That is, the interactive-constraint model predicts a multiple-candidate effect.

The no-antecedent conditions in the experiment also allowed us to gauge how early the constraints derived from binding theory exerted their influence. On the basis of the fact that Nicol and Swinney (1989) observed selective associate priming for targets presented immediately following the pronoun in sentences like Example 5 above, we also expected to see an effect for structural accessibility of gender-matched candidates. We hypothesized that when only one NP matches the pronoun in gender, the effect of structural position for that matching NP should be evident in the form of longer reading times in the region following the pronoun in the inaccessible-match condition (Example 11) in comparison with the corresponding material in the accessible-match condition (Example 10). A simple no-antecedent effect (longer reading times in the no-match condition in comparison with the accessible-match condition) was also expected on both models.

In summary, the initial-filter model and the interactive-constraint model made materially different predictions for processing times associated with the accessible- and multiple-match items. The initial-filter model predicted equivalent reading times for the multiple-match and accessible-match conditions because both contain a single accessible antecedent. In contrast, the interactive-constraint model predicted faster reading times for the accessible-match condition, in which all constraints converge positively on a single antecedent candidate, than for the multiple-match condition.

Method

Participants. Twenty undergraduates from the University of Rochester subject pool participated in this experiment for class credit. Participants were right handed, native English speakers with normal or corrected-to-normal vision, reporting a Verbal SAT score of 500 or higher and no history of reading or language disorders.

Materials and design. All of the experiments presented here use a word-by-word, self-paced reading task coupled with a sentence-final probe-recognition task and intermittent yes–no comprehension questions. The secondary probe-recognition task, which accompanied each trial, and the comprehension questions, which occurred on approximately one quarter of the trials, were included to encourage participants to attend to the content of the stimulus sentences. With the exception of the target stimulus design, the materials-development procedure for all of the subsequent experiments was identical to that of Experiment 1. For this reason, it is presented in detail here only.

Target stimuli reflecting the minimal contrasts for Experiment 1 were constructed on the model of the sentences in Examples 9–12. In all items in this and subsequent experiments (with the exception of Experiment 4), gender-unambiguous proper names were used for these manipulations rather than definite descriptions because names have been shown to be more effective in establishing discourse focus than descriptive referring expressions (Sanford, Moar, & Garrod, 1988). Each sentence set was associated with a single probe word. The probe word for half of the target sets was selected from among the content words consistent across the set: Critically, these "yes" probes were never the pronoun or either of the proper names in the test or filler trials. This was to ensure that the probe-recognition task did not induce reading strategies that would not otherwise be adopted to comprehend the sentence stimuli, as well as to avoid highlighting pronoun interpretation as a salient interest of the experiment. Apart from this constraint, "yes" probes were selected from all sentence regions (initial, medial, and final) to avoid cuing likely probe locations. For the remaining sets, words that did not occur in the sentence(s) were selected. Among these "no" probes, one third were semantic associates to some content word occurring in the trial set, one third were morphological neighbors to some word in the set, and one third were both semantically and morphologically unrelated to any of the content words contained in the target-sentence set. Comprehension questions were developed for one quarter of the target-sentence sets. Responses to comprehension questions were designed to be independent of coreference resolution,

just as with the probe stimuli. The stimulus sentence, probe and, when appropriate, comprehension-question sets were inserted into four presentation lists, each of which contained one variant of each of the target stimulus sentences, with equal numbers of tokens of each treatment condition occurring in each list. The target stimulus lists were then interleaved into a 77-item filler set to derive a fixed, pseudorandomized stimulus set. Filler sentences varied in length and syntactic complexity and on the types of referring expressions they contained. As with the target items, filler items were associated with probe words selected to reflect these constraints: Filler probes were balanced on yes–no responses, and "no" probes were selected so that one third were semantically related, one third were morphologically related, and one third were neither semantically nor morphologically related to any content word in the trial. Yes–no comprehension questions were developed for one quarter of the filler trials. As with the target trials, correct response to the questions following filler sentences did not require resolving any pronoun-antecedent relations that occur in the sentence. Finally, five complete trials were constructed to serve as practice trials.

Experiment 1 consisted of a 2 × 2 design in which the factors accessible-NP fit (match vs. mismatch) and inaccessible-NP fit (match vs. mismatch) were crossed to provide the four treatment conditions. In the example, the sentences in Examples 9 and 10 constituted the accessible-NP-match conditions, and the items in Examples 11 and 12 made up the accessible-NP-mismatch conditions; the inaccessible-NP-match conditions were Examples 9 and 11, and the inaccessible-NP-mismatch conditions were Examples 10 and 12. The treatment conditions for each sentence set were created by alternating the gender of the proper names while holding the gender of the pronoun constant. Thus, in one quarter of the target sentences the morphosyntactic characteristics of the pronoun matched those of both preceding NPs, although only one antecedent was grammatically licensed (multiple match); in one quarter the accessible NP matched the morphosyntactic gender of the pronoun (accessible match); in one quarter the inaccessible NP matched the following pronoun in gender (inaccessible match); and in one quarter of the items there was no morphosyntactically matched antecedent, accessible or otherwise, for the pronoun (no match). Pronoun gender was balanced across the 24 stimulus sets, with half of the items being developed around *her* and half around *him*. Finally, the name pairs used to create the target sentences (i.e., *John–Jane* and *Bill–Beth*, in the example) were matched on letter and syllable length. The names in each pair were also matched on gender typicality to eliminate gender-ambiguous proper names and to ensure that the names used in the match and mismatch conditions were comparable in their gender identifiability. This was accomplished by using gender-categorization norms that we collected in a preliminary name-recognition study: Participants (N = 20) were asked to identify the gender of individually presented proper names as quickly as possible in a forced-choice reaction-time task. Names with a gender consensus of less than 80% were discarded. The remaining names were distributed into male–female pairs on the basis of their match on gender-consensus rate, speed of categorization, length (in syllables and phonemes), and initial consonant (or consonant cluster).

Procedure. Participants were tested using a PC running the Micro Experimental Lab protocols for collecting self-paced word-by-word reading times. The primary task on each trial was to read the sentence that appeared, one word at a time; the secondary tasks were probe recognition (every trial) and question answering (intermittent trials). Individual trials began with a "Ready?" prompt. When a participant pressed the spacebar in response to the prompt, word-by-word presentation of the stimulus sentence began, with subsequent words replacing previous words in the center of the video screen at the rate of presentation determined by the participant's spacebar presses. At the end of each sentence, a probe word appeared in a specified location and the participant used the *d* and *k* keys (for "no" and "yes," respectively) to indicate whether that probe word had occurred in the just-read sentence. Similarly, on trials including a comprehension question, participants used the keyboard to indicate their responses. Auditory feedback was provided to indicate accuracy on both of the secondary tasks. Participants were instructed to read the text at a quick, comfortable pace, but carefully enough to respond accurately on the probe and comprehension tasks. Testing sessions lasted approximately 25 min.

Analysis

As in the *Method* section, details of data analysis that are consistent over all of the experiments are presented only for Experiment 1. Several types of data were collected for each participant. First, probe and comprehension data were collected to create an accuracy profile through which we might identify participants who failed to attend to the text. Only data from participants scoring greater than 80% accuracy on both the probe identification and the comprehension questions were used in the analyses for this and the following experiments. However, because the probe items and comprehension questions following target items included both "yes" and "no" responses (and because only one quarter of the target trials were followed by comprehension questions), the data for these secondary tasks were not included in the analyses of the manipulations relating to coreference resolution.

Reading times provided the central data set for these studies. Our analyses of processing load focused on word-by-word reading times summed over designated text regions, following the approach taken by Trueswell (1996) and others. The text regions for these analyses were chosen on the basis of previous findings in pronoun contexts that showed that reading-time effects for manipulations of antecedent names typically appeared on the first two words following a pronoun. The summed reading times were contrasted across the various treatment conditions. There were several critical text regions that were considered in this and the following experiments: (a) The region preceding the pronoun/anaphor/referring expression; (b) the pronoun/anaphor/referring expression, itself; and (c) the first two (usually phrasal) pairs of contiguous words following the pronoun (region_1 and region_2, respectively). Our tests for reading-time differences in the region preceding the referentially dependent expression were comprehensive in that we compared both individual words and groupings of adjacent words across match conditions. It was predicted that reading times would not differ across treatment conditions in this region, because this is the region in which the lexical material was identical across conditions (modulo the manipulation of proper names). However, if the task was sensitive enough to capture the processing-load differences in antecedent-candidate evaluation and selection created by manipulating the possible candidate set, then reading-time differences should have emerged on or after the to-be-resolved pronoun. On the basis of our findings in earlier studies using this self-paced word-by-word protocol (Badecker & Straub, 1994), it was expected that reading-time differences would be observed in region_1—the paired adjacent words that immediately followed the referentially dependent expression—and not on the pronoun itself. It must be noted, however, that there was no a priori reason to assume that this particular sentence region constituted a theoretically privileged grouping—or even to assume that the specific regions in which processing complexity differences derive from anaphor resolution would be identical across different anaphor types. There could instead have been cases in which important differences emerged as early as the referring expression

itself, or others in which meaningful differences appeared later or were distributed across a larger region of text (e.g., $region_1$ and $region_2$ combined). Such variation could have arisen, for example, as a function of how apparent the match between candidate and pronoun was. When the system is provided with clear gender assignments for two entities, then one might expect to see evidence of competition between two matching entities during the course of selecting a referent for a pronoun (even when one of these entities corresponds to a phrase that is structurally inaccessible to the pronoun). If the match between pronoun and candidate is more difficult to detect, though (e.g., because the system has failed to commit to a particular gender assignment or because it registers some degree of uncertainty regarding the assignments that it makes), then interference from a matching nonantecedent may take longer to register (i.e., reading-time differences may occur further downstream of the pronoun–anaphor) or may be more variable with respect to the locus of the effect. Alternatively, the interference itself may be successfully sidestepped because the competition between the candidates is so slow to develop under such circumstances that the interpretive constraints imposed by binding principles may head off competition from the structurally inaccessible candidate. We discuss other potential sources of variability below. Although they are not reported in the body of the text, comparisons of both adjacent-word groups and single words have been computed over the collected reading times. We report mean reading times, standard errors, and reading-time differences as calculated on grand means. In the interest of readability, our exposition highlights only theoretically critical contrasts and focuses on differences (or on the failure to observe differences) in $region_1$ unless otherwise noted. Interested readers are invited to examine the table of word-reading latencies provided in the Appendix to extract a more explicit picture of the processing-load profile for the various experiments. On the basis of the counterbalanced design of the experiments, the analyses for this and subsequent studies were computed using participants as a random factor (Raaijmakers, Schrijnemakers, & Gremmen, 1999). All reported comparisons are reliable at the $p < .05$ level.

Results

All participants in Experiment 1 scored greater than 80% accuracy on both the probe recognition and the comprehension questions. Accuracy analyses of Probe Recognition × Treatment Condition were unremarkable. Analyses of variance (ANOVAs) are reported for the summed reading times for $region_1$ for the independent variables accessible-NP fit (match vs. mismatch) and inaccessible-NP fit (match vs. mismatch).

As anticipated, no reading-time differences were observed for the region leading up to the pronoun (nor for the pronoun itself) on either the planned-text groupings or the single-word reading-time analyses.

Reading-time differences were observed after the pronoun. Within $region_1$, there emerged both a main effect of accessible-NP fit, $F_1(1, 19) = 6.69$, and an interaction between accessible-NP fit and inaccessible-NP fit, $F_1(1, 19) = 7.18$. There was no main effect for inaccessible-NP fit. Three planned comparisons were carried out on the reading times for $region_1$ to evaluate the predictions regarding a no-antecedent effect and a multiple-candidate effect and to determine how soon an effect emerges for candidate accessibility in contexts with only one matching candidate.

The first of these comparisons probed for the predicted no-antecedent effect by contrasting the reading times for the accessible match and no-match sentences exemplified in Examples 10 and 12. As shown in Figure 1, reading times for the material in $region_1$ (the boxed area in the figure) was 111 ms longer in the no-match condition than in the accessible-match condition (accessible match = 706 ms, $SE = 24$; no match = 817 ms, $SE = 36$; $F_1[1, 19] = 12.61$). This result replicates the no-antecedent effect

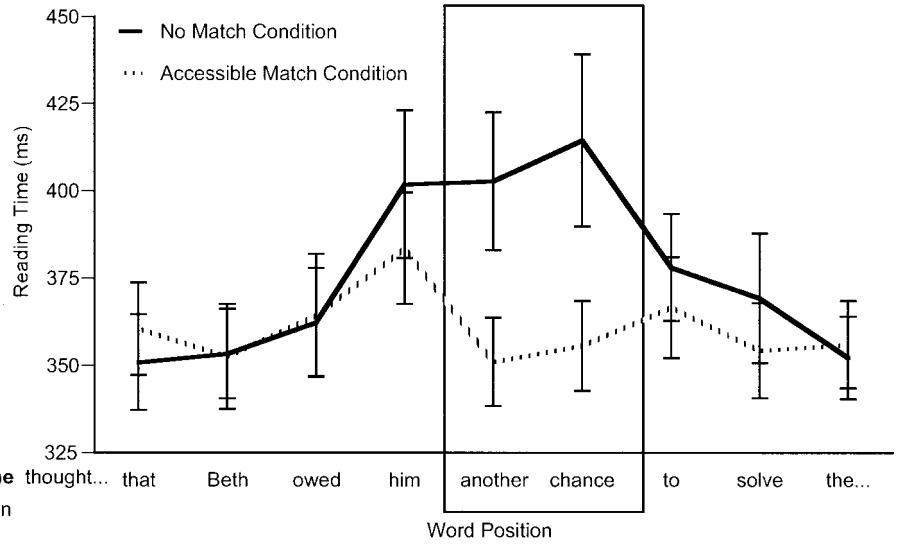

Figure 1. Word-by-word reading times for the no-match and accessible-match conditions in Experiment 1. Error bars indicate standard error of the participant mean. The name in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Reading-time differences represent a no-antecedent effect.)

observed (in comparison with gender-match conditions) when none of the accessible names that precede a pronoun are of the appropriate gender (Badecker & Straub, 1994).

Figure 2 provides evidence that accessibility constraints can be quick to influence candidate evaluation. For contexts in which only one of the two preceding names fit with the gender of the pronoun, reading times for the material in $region_1$ were 94 ms faster when the match was structurally accessible than when it was not (accessible match = 706 ms, SE = 24; inaccessible match = 801 ms, SE = 34; $F_1[1, 19] = 11.51$). This differs from the no-antecedent effect in that the increase in processing load in the inaccessible-match condition derives from the cost of trying to interpret the pronoun when the only matching candidate is structurally inaccessible to it (in contrast to the optimally unambiguous and well-formed accessible-match condition). To understand this result, one must first determine whether there was also an effect for the presence of multiple candidates when only one of these candidates was structurally accessible.

The final comparison focused on how structurally inaccessible NPs affect pronoun processing. As seen in Figure 3, reading times in the critical region were 74 ms longer in the multiple-match items than in the accessible-match items (multiple match = 780 ms, SE = 33; accessible match = 706 ms, SE = 24; $F_1[1, 19] = 8.17$). Given that the initial-filter model asserts that inaccessible candidates will be invisible to the coreference processor, this finding is clearly problematic for that model.

Discussion

There are several results from this experiment that deserve mention. The first relates to the presence versus absence of a viable antecedent candidate for a pronoun: Reading times in the postpronominal region were significantly longer in sentences in which (structurally accessible) candidates failed to match the pronoun's gender than they were in sentences in which the corresponding candidates fit the pronoun's gender. From this finding we drew two simple but important conclusions. First, from a purely methodological perspective, the no-antecedent effect demonstrates that (a) pronoun interpretation has a measurable impact on processing complexity and (b) the self-paced-reading protocol is sensitive enough to capture that effect. Second, from a more theoretical perspective, it indicates that gender is used to identify the referent of a pronoun more or less automatically as part of the comprehension process. Greene et al. (1992) have argued that participants must use strategically controlled interpretive mechanisms to differentiate gender matching from nonmatching candidates and that these mechanisms are not applied in experimental settings unless required by the task (e.g., comprehension questions that depend on disambiguating pronominal reference). But if the use of gender were so easily suspended, then one should not observe a no-antecedent effect in the task used here; apart from the self-imposed comprehension requirements, our participants were not obliged by any aspect of the task to resolve pronoun–antecedent relations in the target sentences. The fact that there is clear evidence that participants did forge such links supports the automaticity of such processing (see also Arnold et al., 2000; Gordon et al., 1999).

The reading-time measures of Experiment 1 also provide evidence that the structural accessibility of an antecedent candidate is assessed rapidly (at least in the contexts examined here). However, the inaccessibility of a candidate does not block its membership in the initial candidate set. The processing-load differences observed between the accessible-match and inaccessible-match conditions suggest that the structural status of the inaccessible candidate begins to influence the evaluation process very soon after the pronoun is encountered, perhaps as soon as gender information comes into play. The initial-filter model of coreference processing would hold that the inaccessible candidate is rejected outright and

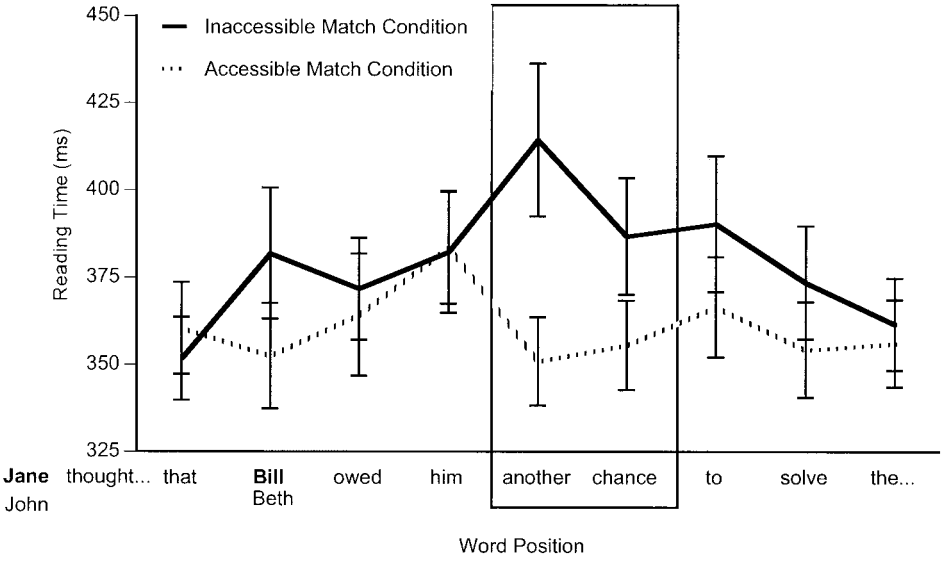

Figure 2. Word-by-word reading times for the inaccessible-match and accessible-match conditions in Experiment 1. Error bars indicate standard error of the participant mean. The names in bold correspond to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Reading-time differences represent effect of structural accessibility for the single proper name that matches the pronoun in gender.)

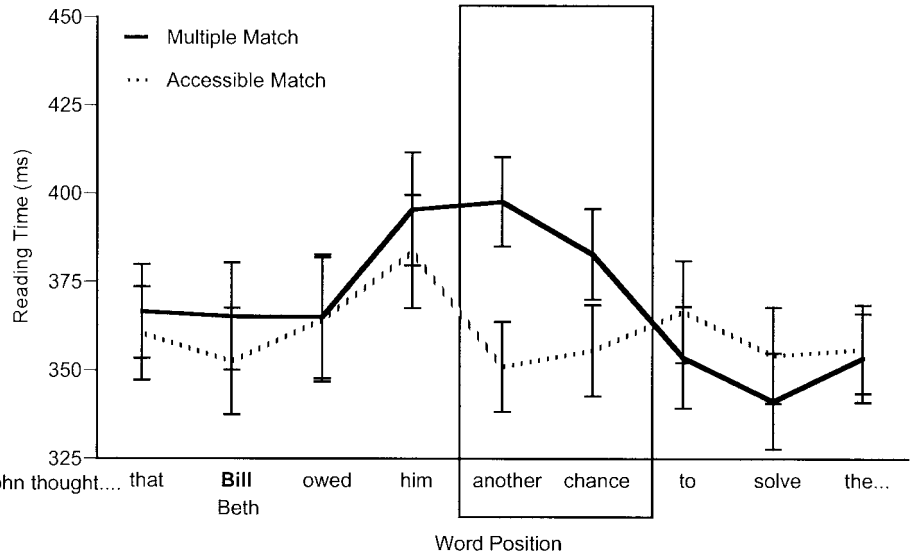

Figure 3. Word-by-word reading times for the accessible-match and multiple-match conditions in Experiment 1. Error bars indicate standard error of the participant mean. The name in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Reading-time differences represent interference from structurally inaccessible proper name in multiple-match condition.)

that what one observes in this contrast is a simple no-antecedent effect. However that interpretation is undermined by the third and most important finding of the experiment: that postpronominal reading times were slowed in the multiple-match condition in comparison with the accessible-match condition. If the structural inaccessibility of a phrase renders the inaccessible candidate completely invisible to the coreference processor (i.e., if it can exclude the inaccessible candidate entirely from the initial candidate set), then no effect for the gender of the inaccessible name should be observed. That we did observe such an effect is evidence that the candidate is evaluated on the basis of its fit with the pronoun. In other words, the mechanisms that impose interpretive constraints on the basis of binding principles do not render the inaccessible candidate completely invisible to other components of the evaluation process.

There are two methodological considerations that must be addressed before the results from Experiment 1 can be taken at face value. The first concern is that the multiple-candidate effect observed in the comparison of sentences like Examples 9 and 10 might not derive from the interpretation of the pronoun. Instead, it might be argued, sentences containing two proper names of the same gender are simply more difficult to encode than analogous sentences in which the proper names differ in gender.[1] Although the locus of the reading-time differences suggests that the processing-load increase reflects the cost of coreference resolution, it is nevertheless conceivable that the differences reflect a processing confound incurred prior to the pronoun. In this view, the asynchrony between the locus of the complexity and its reading-time consequence must derive from task constraints: Participants in self-paced, word-by-word reading tasks will sometimes advance text at a rate that at least partially outpaces their interpretation of that text (Magliano, Graesser, Eymard, Haberlandt, & Gholson, 1993). If the added complexity of encoding two proper names of the same gender is the source of the observed reading-time differences, then the same effects should emerge even when there is no referentially dependent NP to process. This alternative explanation for the ambiguity effect was evaluated (and eliminated) in Experiment 2.

The second methodological concern is that the multiple-candidate effect might only reflect processing strategies that participants adopted to accommodate the fact that the experimental list included a number of sentences with pronouns that lacked an antecedent—a violation of the general expectation that pronouns will be meaningful. To rule this explanation out, we showed (again, in Experiment 2) that the effect is also observed with stimulus sets in which the prevalence of no-antecedent sentences is greatly reduced.

Experiment 2: Disambiguating the Effect of Multiple Matching Names

One might argue that the longer reading times in the multiple-antecedent condition of Experiment 1 may instead arise from the basic processing requirements of developing and maintaining a discourse representation. In such an account, the added processing cost associated with the multiple-match condition merely indicates that it is easier to maintain a discourse representation when the participants in an event are easier to distinguish. This alternative account was easy to test, because it predicted that if we manipulated the number of gender-matching names, then we would observe the same processing-load effect for sentences with and without pronouns. However, if the reading-time differences observed in Experiment 1 reflected the cost of resolving coreference relations, then we should have observed processing differences only for sentences that require pronoun resolution. Experiment 2

[1] We thank Janet Fodor for bringing this concern to our attention.

evaluated these contrasting predictions. This experiment also addressed the concern that the multiple-candidate effect observed in Experiment 1 could be an artifact induced by the prevalence of target sentences that lacked viable pronoun antecedents. The stimulus set for Experiment 2 contains a minimal number of filler sentences containing pronouns with no available antecedent (less than 10%) and no target items of that sort. If the effect for a gender-matching name that is grammatically inaccessible to a pronoun only arises when a high proportion of target sentences lack pronoun antecedents, then it should not have been observed here.

Method

Participants. Forty-four undergraduates from the Johns Hopkins University participated in return for a payment of $5. All participants were right handed, native English speakers who had normal or corrected-to-normal vision and no reported history of reading or language disorders.

Materials and design. The materials used in this experiment, derived from the multiple-match and accessible-match items used in Experiment 1,[2] were developed on the model shown in Examples 13 and 14 to produce a 2 × 2 design with sentence type (pronoun vs. name) and match condition (single vs. multiple match) as factors. Sentences with pronouns (single and multiple-pronoun match) were paired with control items (single and multiple-name match) by replacing the pronouns from the pronoun items with proper names.

(13) a. *pronoun condition, single match*
John thought that Beth owed him another opportunity to solve the problem.
b. *pronoun condition, multiple match*
John thought that Bill owed him another opportunity to solve the problem.

(14) a. *name condition, single match*
John thought that Beth owed Jim another opportunity to solve the problem.
b. *name condition, multiple match*
John thought that Bill owed Jim another opportunity to solve the problem.

Twenty-four stimulus sets based on the model were developed using name substitutions based on the name-pairing procedures described in Experiment 1. All other details of experiment design and stimulus development, including the selection of probe items and the type and distribution of comprehension questions, were identical to those of Experiment 1.

Procedure. The procedure used in this experiment was identical to that of Experiment 1. Participants took approximately 15 min to complete the experiment.

Results

All participants scored above 80% on both the probe-identification and comprehension questions. ANOVAs of reading times summed over the planned text regions were carried out for the factors sentence type (pronoun vs. name) and match condition (multiple match vs. single match). No reading-time differences were observed in positions prior to and including the pronoun or reflexive. As in Experiment 1, only critical significant differences for planned comparisons over the two-word, postpronominal region are discussed in the text below. Word-by-word comparisons appear in the Appendix. The results are summarized in Figure 4.

Full model analyses revealed a main effect for sentence type, $F_1(1, 43) = 8.34$, and a Sentence Type × Match Condition interaction, $F_1(1, 43) = 8.05$. Within the pronoun items, participants were 52 ms slower reading region$_1$ when the inaccessible name matched than when it did not match the pronoun gender (multiple match = 787 ms, $SE = 20$; single match = 735 ms, $SE = 17$; $F_1[1, 43] = 14.73$). No differences were observed for match condition in the name sentences.

Discussion

The results of this experiment bolster the claim that content-based properties of an NP that is structurally inaccessible to a pronoun can affect the processing complexity associated with interpreting that pronoun. The effect of manipulating the gender of the embedded-clause subject in the pronoun condition of Experiment 2 replicates Experiment 1: Postpronoun reading times were longer in the multiple-match condition than in the single-match condition. However, there were no effects for the manipulation of the embedded-clause subject in the name condition in any sentence position. These data also speak to the concern that the processing cost observed in the multiple-match condition in Experiment 1 could be an artifact of a strategy that participants adopted in response to the high proportion of no-antecedent sentences in that study. Although we nearly eliminated the number of sentences containing pronouns with no available antecedent in this experiment, we still observed the multiple-match effect. Thus, we can conclude that this effect of a grammatically inaccessible gender-matched NP arises because mechanisms that evaluate candidate antecedents for pronouns consider the full set of prominent, gender-matched NPs and that this set is not initially filtered by the constraints imposed by binding theory.

The results of Experiments 1 and 2 provide evidence against the initial-filter theory of processing if there exist explicit, structure-based constraints on pronoun interpretation. However, there are variants of binding theory. One might hold that the initial-filter model is correct but that the grammar does not contain an explicit statement corresponding to Principle B (the constraint that stipulates that a pronoun must not be bound by a local c-commanding NP). Some authors have proposed that Principle A (which governs anaphors) is explicitly stated as part of the grammar but Principle B (which governs pronominals) is not. Instead, the generalization corresponding to Principle B is derived pragmatically (Levinson, 1987) or by appeal to an "elsewhere" principle (e.g., Burzio, 1989, 1996, 1998). For the sake of economy, here we outline only the generic, pragmatics-based account of the so-called Principle-B effects. When a pronominal appears in a context that would have permitted linking between a reflexive and a structurally local NP, the use of a pronoun instead of the reflexive is interpreted via the inference that this particular local NP is not the intended referent: If the intended antecedent of the referentially dependent expression were the NP in its binding domain, then in this context the speaker would have chosen to use the type of dependent expression that is specifically required to be linked to that NP (by

[2] One may notice a slight change in the available-match example from Experiment 1 to Experiment 2. This change, from "another chance to solve the problem" to "another opportunity to solve the problem," reflects an effort to use "heavy" words in the region immediately following the pronoun. This strategy was invoked to reduce the likelihood of floor effects on reading times in the critical region.

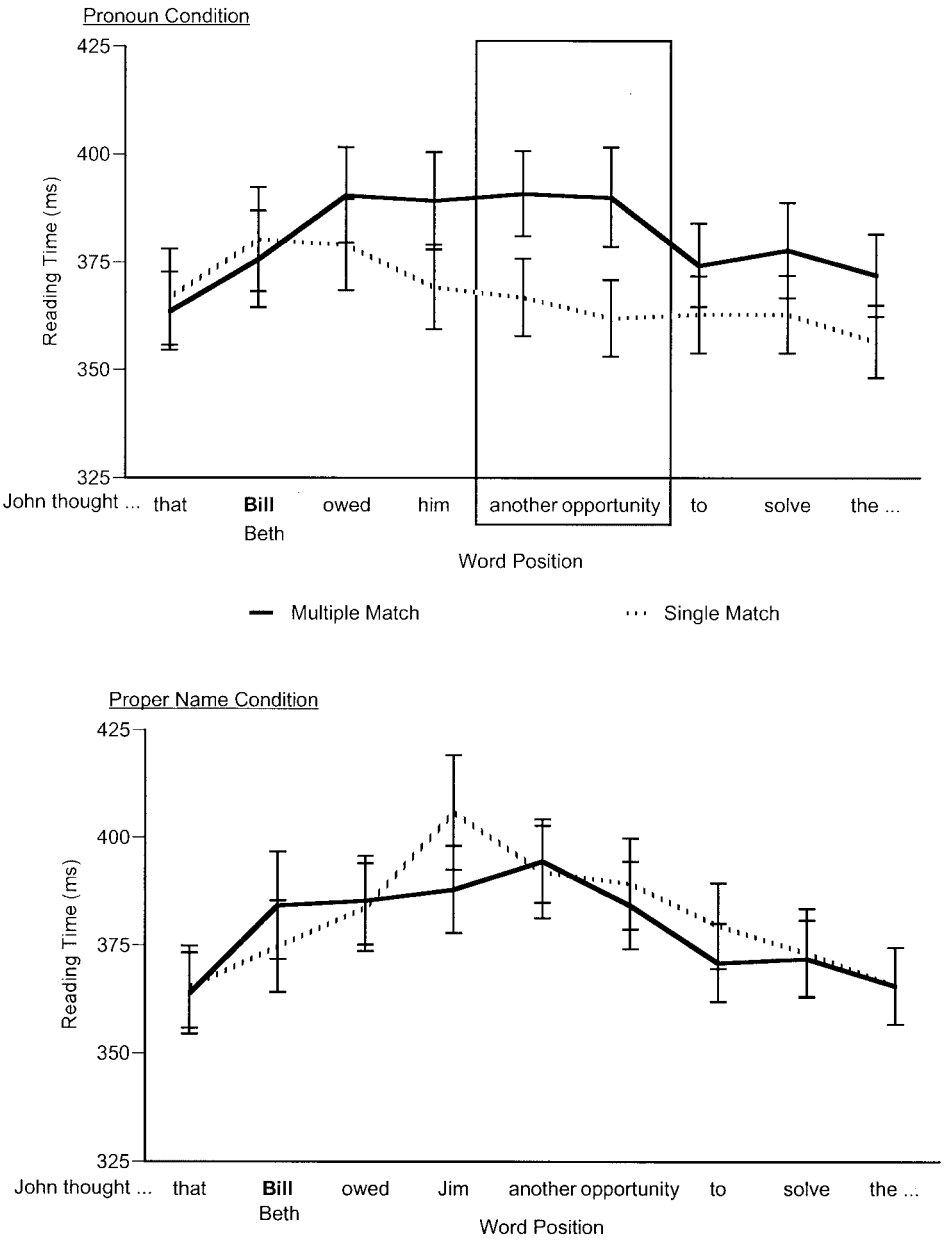

Figure 4. Word-by-word reading times for pronoun versus name conditions in the single-match and multiple-match conditions in Experiment 2. Error bars indicate standard error of the participant mean. The name in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Reading-time differences are limited to pronoun sentences, showing that effect for structurally inaccessible name relates to pronoun interpretation.)

Principle A). Because the speaker did not encode the intended relations by using a reflexive, the interpretation must be other than what is appropriate to the reflexive. That is, the interpretive options for the pronoun do not derive from an explicit principle of grammar but from the implications of not using an expression with interpretive options that are explicitly governed by the grammar. In this interpretation of Principle-B effects, the results of Experiments 1 and 2 do not challenge the initial-filter hypothesis, because no explicit binding principle would have applied. That is, one might expect to see an effect for NP arguments that are inaccessible to a pronoun but not one for arguments that are inaccessible to a reflexive if the structure-based binding principle functions as an initial filter. Experiment 3 was designed to address this issue.

Experiment 3: Comparing Pronouns and Reflexives

The initial-filter hypothesis cannot be evaluated independently of the assumptions the we make about the grammatical knowledge

that is being exploited. If the grammar contains only Principle A (governing the distribution and interpretation of anaphors), and Principle B merely represents a generalization by which the complementary distribution of pronouns and reflexives is derived, then the results of Experiments 1 and 2 would not directly challenge models of anaphor resolution in which binding principles function as initial filters. To test the initial-filter account, one would have to observe an effect for the content of an inaccessible NP on the processing of a grammatical anaphor. Experiment 3 examined whether participants' performance on sentences with reflexives differs from their performance on sentences with pronouns. The single-match and multiple-match conditions for the pronoun and reflexive conditions in this experiment are indicated in Examples 15 and 16.

(15) *pronoun conditions*
   a. John thought that Beth owed him another opportunity to solve the problem.
   b. John thought that Bill owed him another opportunity to solve the problem.

(16) *reflexive conditions*
   a. Jane thought that Bill owed himself another opportunity to solve the problem.
   b. John thought that Bill owed himself another opportunity to solve the problem.

If the anaphor-based description of the initial-filter account is correct, then no effect for the single-versus-multiple-match manipulation should be observed for reflexive items like those in Example 16. The content of the main clause subject in these sentences should have no effect on reading times because the interpretive options corresponding to these NPs should be excluded from the candidate set from the outset by Principle A. When the initial-filter model is construed in terms of the implication-based account of Principle B, though, it makes no prediction regarding the manipulation in the pronoun sentences of the inaccessible NPs (i.e., the embedded clause subject). Therefore, on the basis of the results of the preceding experiments, the initial-filter model equipped only with Principle A predicts that one should observe a Sentence Type $\times$ Match Condition interaction. In contrast, the interactive-parallel-constraint model predicted that no such interaction would be observed. We hypothesize that if the set of candidates that content-sensitive processing mechanisms can have access to is not determined solely by the constraints of binding theory, then there should be a multiple-candidate effect for the inaccessible NPs in both the pronoun and reflexive sentences.

Method

Participants. Thirty-two undergraduates from the Johns Hopkins University participated in return for a payment of $5. All participants were right handed, native English speakers who had normal or corrected-to-normal vision and no reported history of reading or language disorders.

Materials and design. Although some modifications to the sentence materials were necessary to render the content of the sentence context compatible with both pronoun and reflexive direct objects, the form of the sentence stimuli used in Experiment 3 was identical to the form of the items used in the previous experiments. It is important to note that in these items the accessible NP always matched the subsequent referentially dependent expression in gender. The experimental manipulation consisted of a 2 $\times$ 2 design in which Sentence Type (pronoun vs. reflexive) was fully crossed with Match Condition for the referentially dependent expression (single match vs. multiple match) as indicated by Examples 15 and 16. Twenty-four stimulus sets were constructed and pseudorandomly distributed across four presentation lists such that each list contained one item from each stimulus set and an equal number of items from the four treatment conditions. The presentation lists were distributed into the four copies of the same experimental filler list, which consisted of 92 additional sentences of varying length, syntactic complexity, and anaphor presence. Probe words, comprehension questions, and practice trials were created following the procedures for Experiments 1 and 2. As in the previous experiments, correct responses to comprehension questions did not require resolving the interpretation of the pronoun or reflexive in either target or filler sentences. Five items were devised as practice trials.

Procedure. The procedure was identical to those of Experiments 1 and 2. Testing sessions lasted approximately 25 min.

Results

All participants scored above 85% on both the probe-identification and comprehension questions. Reading times are plotted in Figure 5. An ANOVA was carried out examining sentence type (pronoun vs. reflexive) and match condition (single match vs. multiple match). No reading-time differences were observed in positions prior to and including the pronoun or reflexive.

There was no effect for sentence type, nor was there an interaction between sentence type and match condition. However, there was a main effect for match condition, $F_1(1, 31) = 6.81$. Planned comparisons of reading times for region$_1$ within sentence type revealed a reliable effect for match condition for the reflexive items (multiple match = 897, $SE = 40$; single match = 842, $SE = 26$; $F_1[1, 31] = 6.39$) but not for the pronoun items. However, as indicated in Figure 5, the main effect for match condition was carried largely by differences within region$_1$ (in particular, by the second word following the referentially dependent expression). When reading times on this word were examined alone, match condition had a reliable effect for both sentence types: Reading times were 35 ms longer in the multiple-match condition than in the single-match condition for the pronoun items (multiple match = 433 ms, $SE = 16$; single match = 398 ms, $SE = 12$; $F_1[1, 31] = 10.33$) and 42 ms longer for the reflexive items (multiple match = 454 ms, $SE = 28$; single match = 412 ms, $SE = 13$; $F_1[1, 31] = 4.63$).

Discussion

This experiment demonstrated that the match effect for a structurally inaccessible NP does not depend on the type of referentially dependent expression (pronoun vs. reflexive). Therefore, the longer reading times in the multiple-match conditions observed in Experiments 1 and 2 cannot be discounted by adopting a variant of binding theory that posits explicit constraints governing possible interpretations for anaphors but not for pronouns. Instead, it appears that content-based mechanisms for interpreting referentially dependent expressions can have access to candidates that are eventually excluded in virtue of their syntactic inaccessibility. It is indisputable that grammatical constraints play a dominant role in the interpretation process and that the antecedent that is ultimately selected for the reflexives or pronouns in these experiments must allow the binding relations that these expressions enter into to satisfy Principles A and B (cf. Gordon & Hendrick, 1997). The reading-time results from these experiments reflect the relative

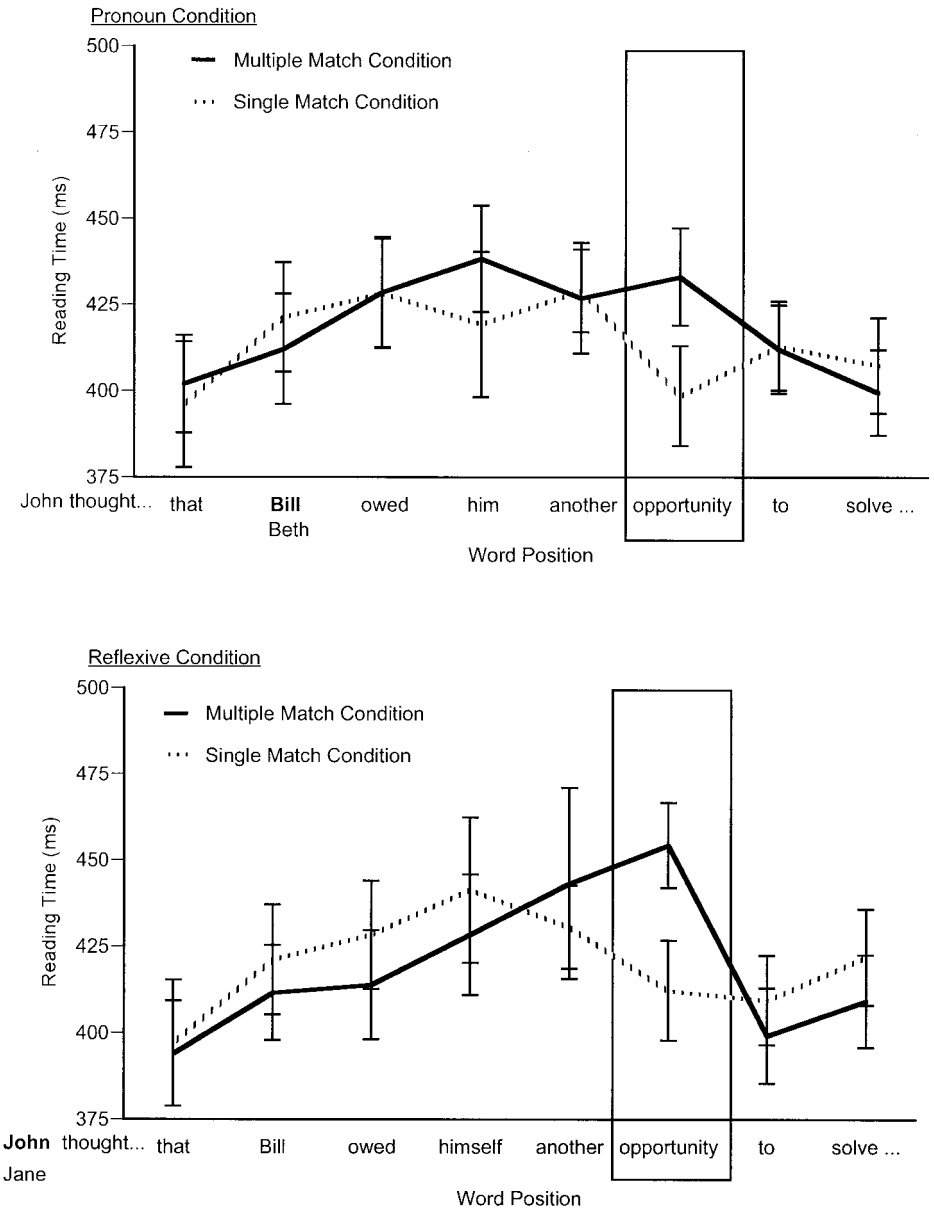

Figure 5. Word-by-word reading times for pronoun and reflexive sentences in the single-match and multiple-match conditions in Experiment 3. Error bars indicate standard error of the participant mean. The name in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Effect for structurally inaccessible name is observed for both pronoun and reflexive constructions.)

complexity of identifying the antecedent of a pronoun or reflexive when more than one preceding NP matches the content of that dependent expression. It is clear that the grammatical principles do not exert their control by rendering content-based mechanisms blind to the presence of grammatically inaccessible antecedent candidates.

Experiment 4: Reciprocal Anaphors

In each of the preceding experiments, the early stages of interpreting a referentially dependent expression were examined by manipulating the gender of a structurally accessible or inaccessible proper name. This is because (a) proper names have been shown to be more effective in establishing discourse referents than NPs headed by relational or occupational predicates like *teacher* or *dentist* (Sanford et al., 1988), and (b) they are less variable than common nouns in terms of their gender bias. We expected these properties of proper names to maximize the opportunities to observe signs that readers attempt to link pronouns (and reflexives) to prominent discourse entities. However, gender is not the only feature that a referentially dependent expression and its antecedent must share. Experiment 4 investigated whether the number specification of an anaphor is also effective in shaping the initial

candidate set. This was done by varying the grammatical number of the inaccessible main-clause subject in sentences containing reciprocal anaphors (e.g., *each other*) to derive the single-match and multiple-match conditions on the model of Examples 17a and 17b, respectively.

(17) a. The attorney thought that the judges were telling each other which defendants had appeared as witnesses before.
b. The attorneys thought that the judges were telling each other which defendants had appeared as witnesses before.

Like reflexives, reciprocals are grammatical anaphors, and therefore their interpretation is governed by Principle A of the binding theory. On the basis of this constraint, *each other* can only be bound by *the judges* in these sentences. If the number specification of an NP functions like gender in formulating the initial set of candidate antecedents, then one might predict a reading-time effect for the manipulation of the inaccessible subject *the attorneys* following a reciprocal anaphor.

Method

Participants. Thirty-six Johns Hopkins undergraduates participated in return for a payment of $5. All participants were right handed, native English speakers who had normal or corrected-to-normal vision and no reported history of reading or language disorders.

Materials and design. Twelve stimulus sets based on the model in Example 17 were developed. In each set, the number of the grammatical subject of the embedded clause conformed with the reciprocal direct object in that clause. The manipulation consisted of a one-way design in which the number of the inaccessible, matrix-subject NP varied with respect to its consistency with the reciprocal anaphor. Composition of the stimulus sets and the secondary probe and comprehension tasks followed the model of the preceding experiments. There were 96 trials in all (12 critical items and 84 filler sets).

Procedure. The experiment used the same procedure as the previous experiments and lasted approximately 20 min.

Results

All participants scored above 85% on both the probe-identification and comprehension questions. Reading times are plotted in Figure 6. An ANOVA of reading times for the independent variable match condition (single match vs. multiple match) was carried out over one-, two-, and four-word-reading-intervals regions preceding and following the reciprocal. No significant differences emerged for any individual words or two-word regions. However when reading times are collapsed across the four positions following the reciprocal, reading times for this region were 48 ms longer in the multiple-match than in the single-match condition (multiple match = 1,565, SE = 37; single match = 1,517, SE = 36; $F_1[1, 35] = 4.37$).

Discussion

The data here are weakened by the fact that one must look beyond region$_1$ (the first two words following the anaphor) to find a measurable difference between the two conditions. However, when reading times are collapsed over the four positions following the reciprocal, one does see evidence that morphological number contributes to identifying the initial set of antecedent candidates. The anaphor was more difficult to process when both the main-clause and embedded-clause subjects satisfied its number requirement in comparison with when this requirement was satisfied by

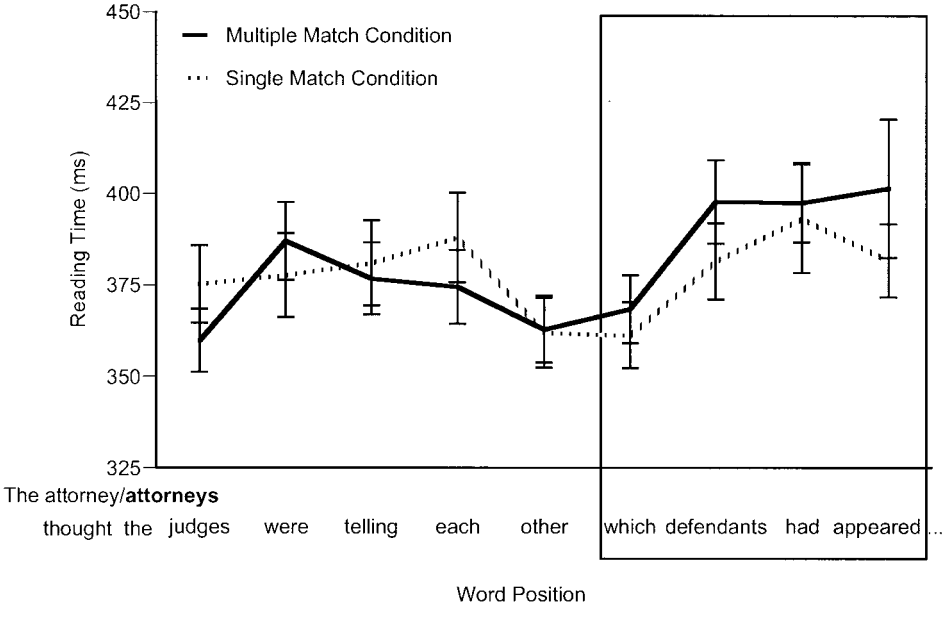

Figure 6. Word-by-word reading times for reciprocal sentences in the single-match and multiple-match conditions in Experiment 4. Error bars indicate standard error of the participant mean. The noun in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box. (Attenuated effect for structurally inaccessible noun phrase is observed with manipulation of grammatical number.)

the structurally accessible embedded-clause subject only. The effect for the content of an inaccessible NP on the reciprocal expression was smaller and more variable in this experiment than what had been observed in Experiments 1–3 when the full interpretation of target sentence depended on linking the referent of proper names with pronouns or reflexive anaphors. One explanation for this attenuation is that participants found the interpretation of reciprocals to be less transparent than the simple coreference relation that exists between a pronoun and its antecedent, and this weakened the temporal link between identifying the reciprocal phrase and the multiple-candidate effect. The fact that referring expressions like *the lawyers* may not be as effective as proper names in establishing discourse entities (cf. Sanford et al., 1988) also may have exacerbated this difference in semantic complexity. It may simply have taken more computational effort for the composition of the candidate set to become apparent to the relevant processing mechanisms. That is, the amount of processing needed to register the inaccessible number-matching phrase as a candidate was simply more variable.

It is also possible that participants do not work out interpretations for reciprocal anaphors as consistently as they do for pronouns or reflexives. Nevertheless, the results of this experiment pattern the results of earlier ones. The interpretation of a reciprocal is dependent on its being linked to a plural antecedent, and when there are multiple prominent plural expressions, the resolution of the linking relation is slowed. We defer further discussion of the differences in effect size and location to the General Discussion.

Experiment 5: Genitive NPs as Antecedent Candidates

In Experiments 1–4, the grammatically inaccessible NPs that slowed the interpretation of a pronoun, reflexive, or reciprocal, always functioned as the grammatical subject of the clause containing the semantically dependent expression. It was tempting to suggest, on the basis of these findings, that local salience constitutes a factor that elevates discourse entities into the set of potential antecedent candidates. However, such a conclusion would have been premature. For instance, it was also possible that every referent introduced by an NP in the local syntactic context would be treated as a member of the initial candidate set. Experiment 5 addressed this by exploring the impact of grammatically nonsalient referential anchors, such as genitives, on the resolution of antecedent relations. If the initial candidate set for a reflexive included all gender-matching discourse entities, then we should have seen an effect for the choice of genitive.

Alternatively, we hypothesized, it could be that nonsubject referents will be projected onto the initial candidate set only if they play a central role in the local context (e.g., if they are semantic arguments of the verb). Experiment 6 explored this by examining whether nonsubject argument NPs affect the initial stages of coreference processing. If semantic arguments of the main or embedded verb were automatically considered as initial candidates as long as they match on the appropriate morphosemantic–morphosyntactic dimensions, then we should have seen reading-time differences between the single and multiple-match conditions for this experiment. Given these options, we expected some degree of consistency in terms of the processing-load effects we would observe in Experiments 5 and 6. That is, we believed that we might observe an effect for compatibility of the reflexive and the inaccessible NP in both experiments, in neither experiment, or in Experiment 6—in which the inaccessible NP is a thematic argument of the main clause predicate—but not in Experiment 5—in which the inaccessible NP is not assigned a thematic role by either the main-clause or embedded-clause predicate. However, we did not expect to observe an effect for the inaccessible-NP manipulation in Experiment 5 unless it was also observed in Experiment 6.

In Experiment 5, the role of local salience in determining the composition of the initial candidate set was examined by manipulating the gender of a genitive noun within a subject NP in sentences containing a reflexive. The multiple-match and accessible-match conditions in the reflexive stimuli are shown in Examples 18a and 18b, respectively.

(18) *reflexive sentences*
- a. Jane thought that Bill's brother owed himself another opportunity to solve the problem.
- b. Jane thought that Beth's brother owed himself another opportunity to solve the problem.

The genitive is not assigned a pivotal thematic role within the sentence. Instead, it functions as a referential anchor for the discourse entity denoted by the NP that contains the genitive. Therefore, the genitive NP differs from the phrases manipulated in the first four experiments (a main or embedded-clause subject) in that it is less prominent and does not receive its thematic role from a verb. Using similar sentences, Gordon et al. (1999) found that a pronoun is more easily interpreted if it refers to a full NP (e.g., *Bill's aunt*) than if it refers to an NP contained in the complex NP (*Bill*). This leads us to predict that the referent of the genitive NP will not be sufficiently prominent to make it a member of the initial candidate set for the reflexive.

The prospect of a null effect for the manipulation in Example 18 introduced the need to verify that participants would attempt to assign interpretations to the referentially dependent expressions and that the content of a genitive NP would at some point affect the outcome of these assignments. Even if one sees no effect for the genitive-NP manipulation in the reflexive conditions, the absence of a viable antecedent in pronoun sentences should increase reading times following the pronoun in comparison with when there is one suitable antecedent. To test this, we included in the design pronoun sentences in multiple-match and inaccessible-match conditions like Examples 19a and 19b.

(19) *pronoun sentences*
- a. Jane thought that Bill's brother owed him another opportunity to solve the problem.
- b. Jane thought that Beth's brother owed him another opportunity to solve the problem.

On the basis of the previous experiments, we predicted that if every name in the sentence projects onto the set of potential antecedent candidates, then we should observe an ambiguity effect for the reflexive items and a no-antecedent effect for the pronoun items. That is, in reflexive sentences, material following the anaphor in Example 18a should have been read more slowly than the corresponding material in Example 18b. In the pronoun sentences, in which the genitive specifier is the only entity that can serve as an antecedent in the match condition, reading times should have been longer when the genitive failed to match the pronoun, as in Example 19a. If mention alone was not sufficient to project a

discourse element into the initial set of antecedent candidates, though, then a different set of expectations would emerge. On the basis of the no-antecedent effect observed in Experiment 1, we still expected to see significantly shorter reading times for pronoun sentences like Example 19a (because the genitive NP is an available antecedent candidate) relative to the no-antecedent context in sentences like Example 19b. Here the effect arises from the processing costs incurred by (possibly second-pass) mechanisms that must search the discourse representation exhaustively for an antecedent when the initial set fails to provide any semantically–morphosyntactically consistent candidates. Thus, we predicted that we should observe a no-antecedent effect regardless of whether the entity corresponding to the accessible genitive is included in the initial candidate set. In contrast, we predicted that if the entity denoted by the genitive NP was not projected into the initial candidate set, then there would be no differences between the single and multiple-match conditions within the reflexive sentences like Example 18.

Method

Participants. Twenty-eight Johns Hopkins undergraduates participated in return for a payment of $5. All participants were right handed, native English speakers with normal or corrected-to-normal vision.

Materials and design. Twenty-four sets of materials based on the model in Examples 18 and 19 were constructed. The sentences were derived from the materials from Experiment 3 by marking the proper name in the embedded clause with the genitive morpheme and introducing a common (relational) noun that was appropriate to the sentence context as the head of the embedded subject. The common noun (e.g., *brother, teacher, lawyer*) was selected on the basis of two criteria. The genitive NP plus noun combination had to make sense (cf. *Mary's teacher* vs. ?*Mary's policeman*); and it had to be biased toward the gender of the referentially dependent expressions for the sentence set (so that it would serve as a good antecedent in the reflexive contexts). The sentence materials reflected a 2 × 2 manipulation crossing sentence type (reflexive vs. pronoun) and match condition (match vs. mismatch) for the genitive proper name. Apart from these details, the materials and design were as in Experiments 1–3.

Procedure. The procedure used in this experiment was identical to those of the previous experiments. Testing sessions lasted approximately 20 min.

Results

All participants scored above 80% on both the probe-identification and comprehension questions. Reading times are plotted in Figure 7. An ANOVA of reading times for region$_1$ was carried out for the factors sentence type (pronoun vs. reflexive) and match condition (match vs. mismatch).

Although there was no main effect for sentence type, the ANOVA revealed a main effect for match, with significantly longer reading times in region$_1$ in the mismatch condition, $F_1(1, 27) = 7.30$, and a Sentence Type × Match Condition interaction, $F_1(1, 27) = 4.13$.

It is important to recall that the match and mismatch conditions had a very different status in the pronoun and reflexive sentences in this experiment. For reflexives, mismatch was a single antecedent condition, whereas for the pronoun sentences mismatch was a no-antecedent condition (with respect to the accessible NPs only). The match condition constituted a multiple-antecedent condition (ignoring accessibility) for both the reflexive and pronoun conditions—but the only potential matching NP for the pronoun (apart from the genitive *Bill's*) was the structurally inaccessible subject NP that contains it (*Bill's brother*). No reading-time differences were observed between the match and mismatch conditions within the reflexive sentences. However, as indicated in Figure 7, an effect for match condition was seen within the pronoun items: Reading times for region$_1$ were 50 ms faster when the genitive NP and the pronoun match (a multiple-match condition) than when they mismatch (an inaccessible-match condition; match = 720, SE = 21 ms; mismatch = 771, SE = 26; $F_1[1, 27] = 6.53$). This difference replicated the no-antecedent effect observed for the pronoun sentences in Experiment 1 and in Badecker and Straub (1994).

Discussion

The manipulation of gender for the grammatically inaccessible nonsubject NP did not have the same effect the manipulation of subject NPs had in Experiment 3. Here, there was no multiple-candidate effect in the reflexive conditions. This supports Gordon et al.'s (1999) finding that a prenominal possessive NP is not as prominent as the full NP that contains it. However, there was a no-antecedent effect in the pronoun conditions. We consider these effects further in the discussion section for Experiment 6.

Experiment 6: Other Nonsubject Antecedent Candidates

The manipulation of the inaccessible NP in the reflexive sentences in Experiment 5 varied the compatibility of a potential referent that functions as a referential anchor in the sentence context. The fact that this position was not assigned a thematic role by any of the predicates in the sentence context invites concern that this thematic status alone excluded the NP from the initial candidate set. Therefore, in Experiment 6, we manipulated the content of an inaccessible nonsubject that nevertheless functions as a thematic argument of the main-clause predicate. Examples of the multiple-match and accessible-match conditions are given in Examples 20a and 20b, respectively.

(20) a. It appeared to John that Bill owed himself another opportunity to solve the problem.
b. It appeared to Jane that Bill owed himself another opportunity to solve the problem.

If thematic argument status is sufficient to procure membership in the initial candidate set, then one might expect longer reading times following the reflexive when the inaccessible argument name matches the reflexive in gender, as in Example 20a, in comparison with when it does not, as in Example 20b. Note, however, that it is also possible that the inaccessible indirect object–goal expression is not sufficiently prominent in the local focus of attention to be included in the initial candidate set projected by the reflexive. This might be the case, for example, because only subject NPs are sufficiently focused for the purpose of acquiring membership in the initial candidate set (cf. Gordon & Scearce, 1995). In that event, manipulating the gender of the proper name should have no effect on reading times.

Method

Participants. Forty-eight Johns Hopkins undergraduates participated in return for a payment of $5. All participants were right handed, native

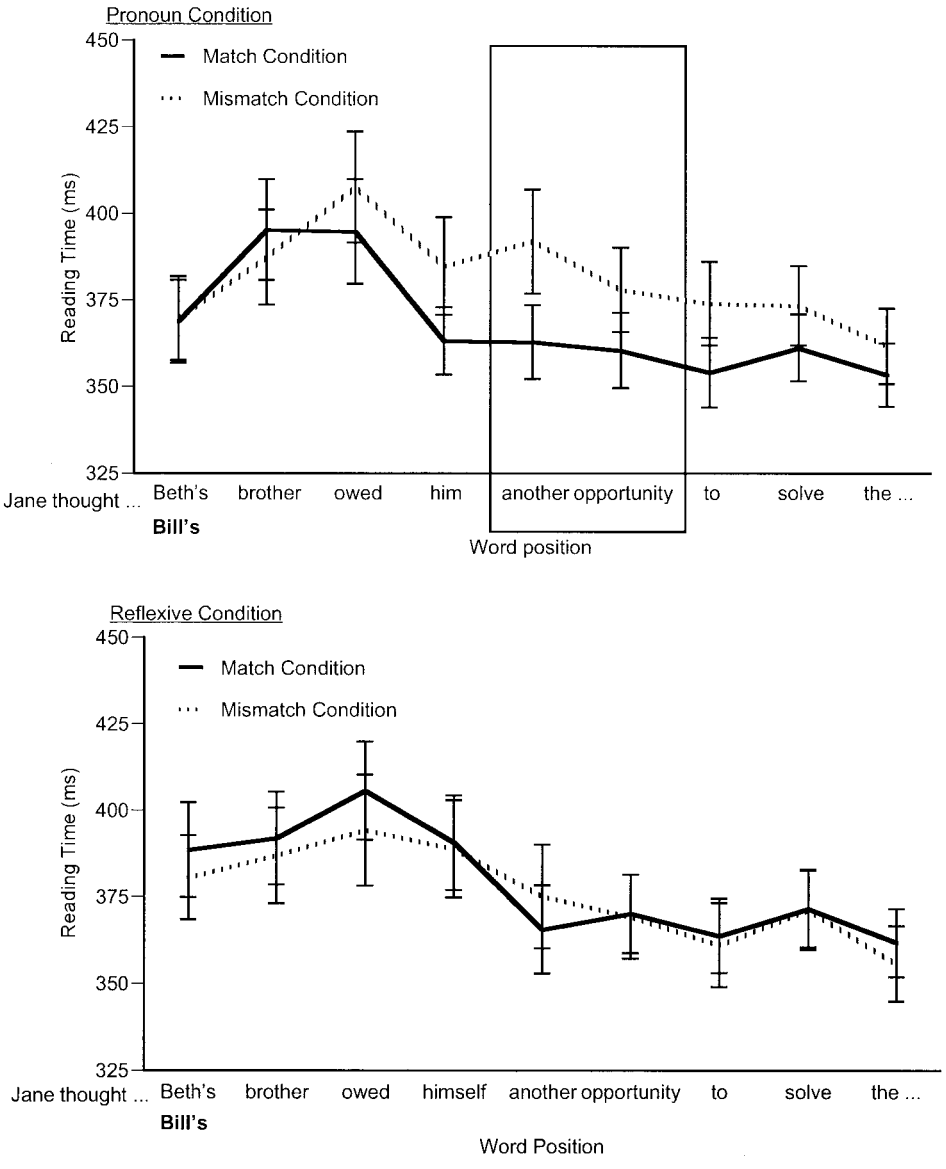

Figure 7. Word-by-word reading times in Experiment 5 for pronoun sentences in the match and mismatch conditions (multiple match and inaccessible match, respectively) and for the reflexive sentences in the match and mismatch conditions (multiple match and accessible match, respectively). Error bars indicate standard error of the participant mean. The name in bold corresponds to the condition plotted by a solid line. Critical differences occur in the region indicated by the box.

English speakers with normal or corrected-to-normal vision and no reported history of reading or language disorders.

Materials and design. Twelve stimulus sets were developed on the basis of the model in Example 20. In each set, the gender of the reflexive expression always matched the gender of the grammatical subject of the clause containing the anaphor. The manipulation in this experiment consisted of a one-way design in which the gender of an inaccessible, nonsubject NP was varied with respect to its compatibility with the reflexive (match vs. mismatch). In other respects, the details of the experimental trials paralleled the previous experiments.

Procedure. The procedure was identical to those of the previous experiments. Sessions lasted approximately 20 min.

Results

All participants scored above 85% on both the probe-identification and comprehension questions. No significant reading-time differences were observed for individual words or adjacent-word positions in the region leading up to or including the reflexive nor in any of the regions following the reflexive.

Discussion

The absence of an effect for the gender of the (syntactically inaccessible) indirect object (Experiment 6) and for the gender of

the genitive NP in reflexive sentences (Experiment 5) supports the notion that the initial candidate set is composed of only highly salient or focused entities in the local discourse context. One might be tempted to hypothesize that the absence of an effect for the inaccessible NP in these two experiments suggests that for some reason the participants in these studies did not resolve the referential relations in the sentence stimuli. This is highly unlikely. Although there is no explicit requirement that participants assign antecedents to referentially dependent expressions, it is clear from the no-antecedent effect on pronouns that participants did register whether sentences like Example 19 contained an accessible antecedent. However, the grammar does not oblige a pronoun to be assigned an antecedent in these sentences: Pronouns may refer to entities mentioned outside the sentence that contains the pronoun. If participants were attending to the assignment of a referential antecedent of an expression that is not required to have an antecedent in the same sentence, then it seems implausible that they would fail to process the grammatically required coreference (for reflexive anaphors) in the same experiment. Taken together with the results of the previous studies, these findings point to limitations on the initial set of referents that are considered as candidate antecedents for pronouns and reflexives.

A more serious concern with regard to the central thesis of this study is that the contrast between the effects we observed for structurally inaccessible subjects in Experiments 1–4 and the absence of effects for nonsubjects in Experiments 5 and 6 could indicate that the structural status of a phrase is considered in formulating the initial candidate set. Should this effect of subjecthood be taken as support for the initial-filter model? To see why it should not, one needs to consider how the structural contrast between subject and nonsubject is exploited in pronoun interpretation. Although it is not possible to enumerate every way in which the structural status of an NP can enter into the interpretive process, there are at least three critical means by which the grammatical role of an NP can have an effect on pronoun–anaphor interpretation. How these different aspects of grammatical role assignment relate to the time course of interpretation is what matters. The parser assigns a structural role to a phrase almost as quickly as the categorical status of the phrase is established (Rayner, Carlson, & Frazier, 1983; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, l995). Immediate word-by-word analysis of structural relations is one feature of human sentence processing that makes rapid incremental interpretation possible. In addition to the role that they play in the recognition of thematic relations, these rapid structural assignments also have consequences for the dynamic tracking of NP discourse status. To the extent that the grammatical role of an NP determines the local discourse prominence of the entity it refers to, parsing an NP as a subject, direct object, or indirect object can either alter, reaffirm, or newly establish the referent's degree of local salience. Like thematic role assignment, the prominence-related consequences of parsing an NP as a subject or nonsubject are immediate. When a pronoun occurs later in the sentence, the viability of an established referent as a potential antecedent will depend on whether that entity remains a highly salient member of the local focus of attention. Therefore, when one considers only prominence, the effect of a candidate's structural role is determined prior to the point at which the pronoun occurs. If the discourse representation of an entity serves as the primary basis for coreference processing, then the effect of the grammatical role of an antecedent candidate will be both immediate and indirect when the pronoun is encountered. The contrast between Experiments 1–4, on the one hand, and Experiments 5 and 6, on the other, reflects the role that structure plays in setting up the prominence relations among the established discourse entities.

The results from Experiments 1–4 suggest that whereas the current prominence of a previously interpreted NP may be immediately recoverable for the coreference processor, the NP's structural position may not. When initiated by a pronoun (or anaphor), coreference processing must extract information about the structural relation that exists between a candidate antecedent and the pronoun (anaphor). Unless one assumes that these structural relations are tracked and recorded for every possible phrasal pair in an incremental and cumulative fashion, the structural accessibility of a particular candidate must be evaluated when the pronoun is encountered. The accessibility of a candidate may not be immediately discernable to the coreference processor because the system might need to reestablish what its structural position had been. Thus, the processing asymmetry between referents of subject versus nonsubject antecedent candidates provides no support for the initial-filter model. Instead, the results suggest that the language processor accesses a set of discourse referents that are treated as candidate antecedents in the earliest phase of reference resolution. Membership in the initial candidate set is limited by the content of the referring expressions that introduced these candidates but also by the role that the referring expressions play in the local context. Discourse entities introduced as functionally salient participants (e.g., subject phrases) are members of this candidate set, whereas those introduced in less prominent roles do not appear to have automatic membership. Therefore, whereas we have indications from these and other studies that structural constraints on the interpretation of pronouns or reflexives are quickly brought to bear on the evaluation of the candidate set, the composition of the initial set appears to originate on the basis of the discourse status and the content of the referentially dependent phrase.

General Discussion

There is no disputing that grammatical principles influence the interpretations that can be assigned to a given utterance. However, our understanding of the interaction between the components of one's grammatical knowledge and other sources of knowledge remains incomplete. In the domain of coreference processing, findings from cross-modal priming studies have been used to support the claim that the structural constraints on binding relations function as an initial filter on antecedent evaluation and selection (Nicol & Swinney, 1989). On this view, mechanisms that subsequently nominate or evaluate potential antecedents on the basis of the content of a referring expression assess only syntactically accessible candidates. The present study suggests a different story. In four experiments, varying the gender of a syntactically inaccessible phrase altered the processing load associated with interpreting a pronoun or anaphor. In each instance, the reading times for positions following the dependent expression were longer when a grammatically inaccessible subject phrase agreed in gender (and number) with the pronoun or anaphor. These results indicate that the binding-theory principles do not function as initial filters on the input to all stages of coreference processing. Instead,

the data presented here support the interactive-parallel-constraint model. The initial candidate set is composed of the focused discourse entities (or sentence constituents) that are compatible with the lexical properties of the referentially dependent expression: If a phrase (or corresponding discourse entity) is prominent in the requisite sense, the mechanisms for interpreting referentially dependent expressions will have access to it even in instances in which the phrase fails to meet the structural requirements for coreference and will ultimately be excluded by these grammatical constraints. These focused entities include (at least) the grammatical subjects of the local clauses. It is also evident that the initial candidate set excludes entities that are introduced in the local context in less prominent roles (such as genitive specifiers or other nonsubject roles).

It is important to note that this interpretation of our data is consistent with the view that there may be processing costs associated with applying binding principles. Such a proposal was made by Stowe (1986) to explain reading-time effects for false syntactic-gap positions inside of *wh*-islands. However, even if the longer reading times observed following referentially dependent expressions in Experiments 1–4 arise because there is a cost associated with excluding a particular referent on the basis of Principle A or B, it is a cost that varies according to whether the inaccessible phrase conforms to the gender–number constraints of the pronoun or anaphor. To be apprised of this compatibility, the parser must first evaluate the properties of the preceding referential expressions and identify some as satisfying the requirements for an antecedent that are imposed by the referentially dependent phrase (i.e., its number and gender features). Clearly, membership in this candidate set cannot have been first determined by mechanisms that care about structure but not about content. In other words, the initial candidate set is not composed of just those entities that are referred to by phrases that satisfy the principles of binding theory. This observation holds regardless of the specific nature of the content in question.

One might suppose that the cost observed for matching inaccessible phrases reflects effects for the gender and number specification of an antecedent candidate that is part of the discourse representation. That is, the content in question is the content of the discourse representation of a referent and not of the surface grammatical forms that encoded these details of the mental model. Evidence that referential dependencies are identified by linking representations in a discourse structure would support this interpretation (Cloitre & Bever, 1988). On the other hand, one might wish to argue that the gender and number manipulations that induced these effects are morphosyntactic in nature, and therefore the content under scrutiny remains grammatical in nature. Given the requirements imposed by pronouns and anaphors on their antecedents, and given that discourse focus status is not easily distinguished from the syntactic role of an antecedent candidate, these considerations may make the specific representational level at which these effects occur difficult or impossible to sort out. In either case, though, the structural relationship alone between a referring phrase and a referentially dependent expression will not serve as the basis for an initial filter on the earliest composition of the set of candidate antecedents for the pronoun or anaphor in the strict sense of filter being considered here.

Another aspect of our results that deserves discussion is the variability in the onset of the multiple-candidate effect following a pronoun or anaphor. Although our analysis of Experiments 1–3 found that the effect of a feature-matching inaccessible NP can be detected in the reading times for region$_1$ (the first two words following the referentially dependent expression), the effect cannot be further localized to the same lexical positions across experiments. For example, in Experiment 3 the differences that exist in region$_1$ are largely attributable to differences in reading times for the second member of this phrasal pair. One could make a case that the effect of the gender manipulation in the pronoun items of Experiment 2 begins to emerge as early as the pronoun itself. There are two points to be made regarding this concern. First, we can at least attribute any differences that emerge to coreference processing. One normally interprets longer reading times on a specific word (or phrase) in a sentence as an indication that there is a processing difficulty associated with processing that very word (or phrase). However, the initial recognition of an ambiguity or processing difficulty in a self-paced reading task can also induce a participant to speed up the response to advance the text at a rate that more closely resembles that of normal reading. Because individual participants may adopt this approach intermittently, the exact point at which a reader slows the responses may be less than perfectly coupled with the lexical material that initiated this increased processing load. The greatest immediate concern is that reading-time differences that relate directly to material that occurs before the pronoun might show up on or after the pronoun. We examined this possibility in Experiment 2. If the processing-load effects were caused directly by the comprehension of material preceding the pronoun, then we would have seen evidence of this in both the pronoun and the no-pronoun contexts. The absence of an effect for the proper-name manipulation in the no-pronoun sentences satisfies one important prerequisite of our argument.

The second point is that once the role of the pronoun's morphosyntactic properties is established there is a fundamental problem for the initial-filter model of coreference processing no matter how soon or late the gender of a structurally inaccessible phrase affects reading performance following the pronoun (at least when the variability is within reason). As one reviewer observed, the logic of the situation differs in an important respect from component interactions in other sentence-processing domains—as when one tries to argue that an early effect for content should be taken to show whether other, structure-based parsing decisions had been made. In the present context, evidence that a structurally inaccessible phrase alters the process of interpreting a pronoun or reflexive is prima facie evidence that the structural constraint does not insulate the interpretive process from structurally inappropriate candidates. As Nicol, Fodor, and Swinney (1994) noted in their defense of the associate-priming paradigm for filler-gap dependencies in *wh*-constructions, the exact locus of an effect relating to an anaphoric link may be influenced by subtle differences in processing load that arise from (sometimes small) differences in the stimuli across experiments. Differences such as those relating to where reading-time effects emerge in Experiments 1–3 are not as important as the consistency with which structurally inaccessible proper names alter reading times. Therefore, although we acknowledge the more variable performance in Experiment 4, we must also identify the most consistent interpretation of our findings. What is noteworthy is that the findings of Experiment 4 conform to those of Experiments 1–3: Interpreting the reciprocal anaphor *each other* appears to be affected by the content of a main

clause subject, despite the fact that this subject NP is grammatically inaccessible (on the basis of Principle A).

Although our interpretation is incompatible with that of Nicol and Swinney (1989), our account can accommodate their priming data. Their results impose important constraints on the time course of the processing that is initiated by a referentially dependent expression. The priming results, along with the reading-time data presented here, indicate that the structure-based principles of binding theory are enlisted quickly to evaluate and select viable antecedent candidates. The critical issue, though, is that the priming results cannot provide a complete picture of the initial candidate set. The absence of associate priming cannot be taken as unambiguous evidence regarding the candidate status of an NP. Still, Nicol and Swinney's results provide compelling evidence that binding information is brought to bear quickly in candidate evaluation.[3]

Before we leave behind Nicol & Swinney's (1989) associate-priming results, though, we should discuss the role presentation rate may play in deriving the effects that we have observed in our experiments. One reviewer expressed a concern that the speed with which a sentence unfolds in the different experimental paradigms may account for why Nicol and Swinney appear to have evidence for the initial-filter model, whereas we appear to have evidence against that model. One might propose, for example, that morphological factors may be brought into play at slower presentation rates in ways that they could not (or simply are not) with faster presentation rates. We agree with the general point that the speed of presentation may alter the relative availability of different types of information associated with the constituents of the sentence stimulus, and this might affect the way that processing components interact. Because the timing of the input may allow one type of information to diminish or head off effects evoked by another type of information, one could suppose that effects relating to a particular source of information might come or go as a function of the experimental task. However, these observations do not undermine our arguments for the interactive-parallel-constraint model of coreference processing. First, the hypothetical differences, if they do exist, relate to when information of a particular type is used under specific presentation conditions (and not to what information is available under these presentation conditions). There is clear evidence from both cross-modal associate-priming studies (e.g., Nicol, 1988; Shillcock, 1982) and from paradigms using eye-tracking measures of gender effects (e.g., Arnold et al., 2000; Garrod, Freudenthal, & Boyle, 1994) that coreference processing makes early and immediate use of gender information. This evidence leaves little room for the prospect that in the self-paced reading task, the gender (or number) of a referentially dependent expression could initiate specific coreference processes that are not initiated at faster presentation rates (or are initiated in some order that departs from the norm). The observed effects suggest that these hypothetical processes must be focus dependent (because we observed an effect for subject but not nonsubject candidates). In addition, they are also sensitive to the gender properties of candidates that are, by hypothesis, excluded by the principles of binding theory. The availability of this inaccessible-candidate information requires that the outcome of structure-based processes are ignored in self-paced reading but not in normal listening or reading, because the proposal that gender information can be used more quickly in the self-paced reading task is undermined by the evidence that gender (and focus) information about candidates is available at the outset of the interpretive process in all presentation modes.

Even so, this view is still not compatible with the initial-filter model. If the reading-time results were obtained because the outcome of early acting binding principles was subsequently ignored by later, exceptional processes triggered by gender information, the hypothetical shift in timing would also alter the manner in which binding principles are hypothesized to influence coreference processing. That is, setting aside the issue of whether the presentation rate changes when information is available and used, the reading-time results still falsify the claim that the binding principles blind other interpretive components to structurally inaccessible phrases. These considerations suggest that one productive strategy for uncovering the role of structure-based principles in coreference processing will be one that integrates the associative-priming and reading-time results.

Our interpretation of the preceding experiments might be taken to suggest a model in which a somewhat unconstrained set of sentence-processing mechanisms is placed in conflict with an underlyingly rigid set of grammatical principles. If the interpretation of referentially dependent phrases must be true to the principles of binding theory, then how did a processing system in which the initial candidate set need not be fully consistent with these principles ever arise to begin with? One possibility is that the presence of grammatically inaccessible entities in the candidate set might arise as the result of modularity. If the rapid interpretation of referentially dependent expressions like pronouns relies on the ability of a pronoun to nominate salient entities in the current discourse representation as potential antecedents, it may be more efficient to do so on the basis of content relations alone. Because some of the candidates may not be discourse entities contained in the same sentence as the pronoun, and because the set of focused entities may be sufficiently small, the cost of taking additional structural properties of candidates into account may add unnecessarily to the computational costs of evoking this initial collection of candidates. Such an architecture might accelerate the processing of referentially dependent expressions by allowing the procedures for identifying potential antecedents to commence before the parser has identified all of the details of structure that are relevant to computing accessibility for the candidates at hand. In this view, a pronoun identifies the members of the initial candidate set without regard to the structure of the preceding sentence fragment for the same reason that a lexical form accesses associated mean-

[3] We do not endorse the view that data from cross-modal associate priming of the sort presented in Nicol and Swinney (1989) can be ignored. McKoon, Ratcliff, and Ward (1994) found that their attempts to replicate contingent priming effects in *wh*-constructions raised the concern that the appearance of associate priming could be derived from integration effects of the sort observed by Wright and Garrett (1984). If the semantic associate in some conditions is more easily integrated with the context than with the matched control, then the relative ease of integrating the probe with the sentence context could be the true source of the apparent priming effect. Although this is true in principle, the fact that the nature of the priming effect in Nicol and Swinney's study showed complementary effects as a function of the type of referentially dependent expression suggests that the task truly does tap into the process of interpreting the pronoun or reflexive. (See Nicol et al., 1994, and McKoon & Ratcliff, 1994, for further discussion.)

ings without regard to the compatibility of those associated meanings with the grammatical and semantic requirements of the preceding fragment (Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979).[4] Thus, the initial candidate set can include grammatically inaccessible nominees because identifying candidates on the basis of content and evaluating candidates on the basis of structure are computational tasks carried out over different types of information and are therefore invoked by different, albeit parallel, processing mechanisms.

In summary, we have processing-load evidence that reveals how structural constraints interact with other types of information about antecedent candidates in the process of interpreting referentially dependent expressions. This evidence, along with findings from cross-modal associate priming (Nicol & Swinney, 1989), indicates that grammatical constraints exert their effect on the selection of candidates early in the course of interpretation by evaluating members of a set that is initially composed of all currently prominent discourse entities. This initial candidate set is projected from the set of focused discourse entities by the grammatical features of the pronoun or anaphor and may include entities that will fail the test of structural accessibility. In other words, the early influence of structural constraints does not completely exclude inaccessible entities from the candidate set at the outset. Instead, the initial candidate set is created in a manner analogous to the process of activating the set of interpretations associated with an ambiguous lexical form: The salient entities of the local discourse (the focus of attention, in the terminology of Grosz et al., 1995) that are consistent with the number and gender specification of the pronoun or anaphor make up the set of available interpretations for the referentially dependent expression, and the grammatical constraints on interpretation operate quickly and effectively in the process of selecting from among these options.

[4] The resemblance to other lexical ambiguities may be amplified if coreference processing is initiated before lexical recognition mechanisms have time to distinguish pronouns and reflexives. In Experiments 1–3, both the pronoun *him* and anaphor *himself* require a masculine singular antecedent, but the morphophonological relation between *him* and *himself* may have some impact on the course of lexical access and on the interpretive processes that are initiated by activating a particular lexical item. However, the effects observed in Experiment 4 (with reciprocal anaphors) suggest that the initial candidate set is driven in large part by the morphological properties of the actual pronoun or anaphor in the sentence.

References

Arnold, J., Eisenband, J., Brown-Schmidt, S., & Trueswell, J. (2000). The rapid use of gender information: Evidence of the time course of pronoun resolution from eyetracking. *Cognition*, *76*, B13–B26.

Badecker, W., & Straub, K. (1994, March). *Evidence that binding principles participate in a constraint satisfaction process*. Poster session presented to the Seventh Annual CUNY Sentence Processing Conference, New York.

Burzio, L. (1989). On the non-existence of disjoint reference principles. *Revista di Grammatica Generativa*, *14*, 3–27.

Burzio, L. (1996). The role of the antecedent in anaphoric relations. In R. Freidin (Ed.), *Current issues in comparative grammar* (pp. 1–45). Dordrecht, Netherlands: Kluwer.

Burzio, L. (1998). Anaphora and soft constraints. In P. Barbosa, D. Fox, P. Hagstrom, M. McGinnis, & D. Pesetsky (Eds.), *Is the best good enough? Optimality and competition in syntax* (pp. 93–113). Cambridge, MA: MIT Press/MIT Working Papers in Linguistics.

Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behavior*, *16*, 601–609.

Chambers, C., & Smyth, R. (1998). Structural parallelism and discourse coherence: A test of centering theory. *Journal of Memory and Language*, *39*, 593–608.

Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht, Netherlands: Foris.

Clifton, C., & Ferreira, F. (1987). Discourse structure and anaphora: Some experimental results. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 635–654). London: Erlbaum.

Cloitre, M., & Bever, T. (1988). Linguistic anaphors, levels of representation, and discourse. *Language and Cognitive Processes*, *3*, 293–322.

Erlich, K. (1980). Comprehension of pronouns. *Quarterly Journal of Experimental Psychology*, *32*, 247–255.

Freidin, R. (1986). Fundamental issues in the theory of binding. In B. Lust (Ed.), *Studies in the acquisition of anaphora* (pp. 151–188). Dordrecht, Netherlands: D. Reidel.

Garnham, A., & Oakhill, J. (1985). On-line resolution of anaphoric pronouns: Effects of inference making and verb semantics. *British Journal of Psychology*, *76*, 385–393.

Garnham, A., Traxler, M., Oakhill, J., & Gernsbacher, M. (1996). The locus of implicit causality effects in comprehension. *Journal of Memory and Language*, *35*, 517–543.

Garrod, S., Freudenthal, D., & Boyle, E. (1994). The role of different types of anaphor in the on-line resolution of sentences in discourse. *Journal of Memory and Language*, *33*, 39–68.

Gernsbacher, M. (1989). Mechanisms that improve referential access. *Cognition*, *32*, 99–156.

Gordon, P., & Chan, D. (1995). Pronouns, passives, and discourse coherence. *Journal of Memory and Language*, *34*, 216–231.

Gordon, P., Grosz, B., & Gilliom, L. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, *17*, 331–347.

Gordon, P., & Hendrick, R. (1997). Intuitive knowledge of linguistic co-reference. *Cognition*, *62*, 325–370.

Gordon, P., Hendrick, R., Ledoux,. K., & Yang, C. L. (1999). Processing of reference and the structure of language: An analysis of complex noun phrases. *Language and Cognitive Processes*, *14*, 353–379.

Gordon, P., & Scearce, K. (1995). Pronominalization and discourse coherence, discourse structure and pronoun interpretation. *Memory & Cognition*, *23*, 313–323.

Greene, S., McKoon, G., & Ratcliff, R. (1992). Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 266–283.

Grosz, B., Joshi, A., & Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, *21*, 203–225.

Gundel, J., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, *69*, 274–307.

Lasnik, H. (1989). *Essays on anaphora*. Dordrecht, Netherlands: Kluwer.

Levinson, S. (1987). Pragmatics and the grammar of anaphora: A partial pragmatic reduction of binding and control phenomena. *Journal of Linguistics*, *23*, 379–434.

MacDonald, M., & MacWhinney, B. (1990). Measuring inhibition and facilitation in pronouns. *Journal of Memory and Language*, *29*, 469–492.

Magliano, J. P., Graesser, A., Eymard, L., Haberlandt, K., & Gholson, B. (1993). Locus of interpretive and inference processes during text comprehension: A comparison of gaze durations and word reading times. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 704–709.

Matthews, A., & Chodorow, M. (1988). Pronoun resolution in two-clause sentences: Effects of ambiguity, antecedent location, and depth of embedding. *Journal of Memory and Language*, *27*, 245–260.

McDonald, J., & MacWhinney, B. (1995). The time course of anaphor resolution: Effects of implicit causality and gender. *Journal of Memory and Language*, *34*, 543–566.

McKoon, G., & Ratcliff, R. (1994). Sentential context and on-line lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1239–1243.

McKoon, G., Ratcliff, R., & Ward, G. (1994). Testing theories of language processing: An empirical investigation of the on-line lexical decision task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1219–1228.

Nicol, J. (1988). *Coreference processing during sentence comprehension*. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Nicol, J., Fodor, J., & Swinney, D. (1994). Using cross-modal lexical decision tasks to investigate sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1229–1238.

Nicol, J., & Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research*, *18*, 5–19.

Prince, E. (1981). Toward a taxonomy of given–new information. In P. Cole (Ed.), *Radical pragmatics* (pp. 223–255). New York: Academic Press.

Raaijmakers, J., Schrijnemakers J., & Gremmen, F. (1999). How to deal with "The language-as-fixed-effect fallacy": Common misconceptions and alternative solutions. *Journal of Memory and Language*, *41*, 416–426.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior*, *22*, 358–374.

Reinhart, T. (1983a). *Anaphora and semantic interpretation*. Chicago: University of Chicago Press.

Reinhart, T. (1983b). Coreference and bound anaphora: A restatement of the anaphora questions. *Linguistics and Philosophy*, *6*, 47–88.

Sanford, A., Moar, K., & Garrod, S. (1988). Proper names as controllers of discourse focus. *Language and Speech*, *31*, 43–56.

Seidenberg, M., Tanenhaus, M., Leiman, J., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, *14*, 489–537.

Shillcock, R. (1982). The on-line resolution of pronominal anaphora. *Language and Speech*, *25*, 385–401.

Smyth, R. (1994). Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, *23*, 197–229.

Stowe, L. (1986). Parsing *wh*-constructions: Evidence for on-line gap location. *Language and Cognitive Processes*, *1*, 227–245.

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

Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., & Sedivy, J. (l995, June 16). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632–1634.

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

Wright, B., & Garrett, M. (1984). Lexical decision in sentences. *Memory & Cognition*, *12*, 31–45.

Appendix

Statistics for Word-by-Word Reading-Time Differences

Although the analyses presented in the body of this article focus on reading-time differences observed across preselected text regions (i.e., the phrase consisting of the first two words following the pronoun or reflexive), the data-collection paradigm that was used allows for analyses of individual word reading times as well. The following table is presented so that readers can explore the time course of the reading-time differences across the experiments.

In the interest of space, reading-time differences for word positions occurring before the critical referring expressions, which neither differed nor reflected consistent trends toward differing, are omitted.

Table A1

Reading Times (and Standard Errors) in ms by Position for Single Word Comparisons Calculated as Grand Means

Condition	Lead-in	him	another*	chance*	to	solve
Experiment 1 (Figure 1)						
No match	*Jane* thought that Beth owed	402 (16)	403 (13)	414 (13)	378 (14)	369 (14)
Accessible match	*John*	383 (14)	351 (13)	356 (16)	367 (21)	354 (25)
Experiment 1 (Figure 2)		him	another*	chance*	to	solve
Inaccessible match	*Jane* thought that *Bill* owed	382 (17)	414 (22)	387 (17)	390 (19)	373 (16)
Accessible match	*John* *Beth*	383 (16)	351 (13)	356 (13)	367 (14)	354 (14)
Experiment 1 (Figure 3)		him	another*	chance	to	solve
Multiple match	John thought that *Bill* owed	396 (16)	398 (13)	383 (17)	354 (14)	341 (14)
Accessible match	*Beth*	383 (15)	351 (12)	355 (12)	366 (14)	354 (13)

Table A1 (*continued*)

Condition	Lead-in					
		Experiment 2 (Figure 4)				
		him*	another*	opportunity*	to	solve
Pronoun						
Multiple match	John thought that *Bill* owed	392 (11)	394 (10)	393 (11)	377 (10)	382 (11)
Single match	*Beth*	373 (10)	370 (9)	365 (9)	366 (9)	366 (9)
		Jim	another	opportunity	to	solve
Name						
Multiple match	John thought that *Bill*	391 (10)	398 (10)	387 (10)	373 (9)	374 (9)
Single match	*Beth*	410 (13)	396 (11)	393 (11)	382 (10)	377 (10)
		Experiment 3 (Figure 5)				
		him	another	opportunity*	to	solve
Pronoun						
Multiple match	John thought that *Bill* owed	438 (15)	427 (15)	433 (16)	412 (14)	399 (13)
Single match	*Beth*	419 (16)	429 (21)	398 (12)	413 (14)	407 (13)
		himself	another	opportunity*	to	solve
Reflexive						
Multiple match	*John* thought that Bill	428 (16)	443 (17)	454 (28)	399 (12)	409 (14)
Single match	*Jane*	441 (15)	430 (14)	412 (13)	409 (13)	422 (14)
		Experiment 4 (Figure 6)				
		other	which	defendants	had	appeared
Multiple match	The *lawyers* thought the judges told each	363 (9)	368 (9)	398 (11)	398 (19)	402 (10)
Single match	*lawyer*	362 (10)	361 (9)	382 (10)	393 (15)	382 (10)
		Experiment 5 (Figure 7)				
		him*	another*	opportunity	to*	solve
Pronoun						
Mismatch	Jane thought that *Beth's* brother owed	385 (15)	392 (12)	378 (12)	374 (12)	373 (12)
Match	*Bill's*	363 (10)	363 (11)	360 (11)	354 (10)	361 (10)
		himself	another	opportunity	to	solve
Reflexive						
Mismatch	Jane thought that *Beth's* brother owed	389 (13)	375 (11)	369 (11)	361 (11)	371 (11)
Match	*Bill's*	391 (11)	366 (13)	370 (11)	364 (11)	372 (11)

*$p < .05$, on by-subject analysis for the single word.

Received May 5, 2000
Revision received November 26, 2001
Accepted December 5, 2001 ■