Grammatical and Coherence-Based Factors in Pronoun Interpretation

Laura Kertz (kertz@ling.ucsd.edu)
Andrew Kehler (kehler@ling.ucsd.edu)
Department of Linguistics, 9500 Gilman Drive
La Jolla, CA 92093-0108 USA

Jeffrey L. Elman (elman@crl.ucsd.edu)
Center for Research in Language, 9500 Gilman Drive
La Jolla, CA 92093-0515 USA

Abstract

We describe a pronoun interpretation experiment in which a Coherence Hypothesis is tested against preference-based analyses of pronoun interpretation, including the Parallel Function Preference and the Subject Preference. We demonstrate that preferences can be systematically disrupted through the manipulation of coherence, and that only the Coherence Hypothesis can predict the full range of coreference patterns observed. We argue, following Kehler (2002), that apparent preferences follow from inferencing processes which support different types of coherence relation.

Introduction

Three decades of psycholinguistic research into pronoun interpretation have documented a broad range of biases and effects linked to syntax, semantics, pragmatics, and more. This has led some researchers to posit pronoun interpretation 'heuristics', that is, broadly-applicable strategies for matching pronouns to their antecedents. Of particular interest to this study are two proposals which gained prominence in the 1990's and which continue to play a role in theories of pronoun interpretation today: the Parallel Function Preference (Smyth 1994, Chambers & Smyth 1998, inter alia) and the Subject Preference (Crawley & Stevenson 1990, inter alia).

These preferences, however, are often in conflict, and each seems to surface only under certain conditions. Acknowledging this, proponents of preference-based models typically identify a basic antecedent selection preference, and then augment it with additional syntactic or semantic mechanisms. In this paper we offer an alternative analysis, following Kehler (2002), in which preference conflicts are best understood as side-effects of the establishment of different types of coherence.

Parallel Function Preference

The Parallel Function Preference holds that pronouns are preferentially resolved to antecedents that occupy a matching argument position, for example, subject pronouns prefer subject antecedents, and object pronouns prefer object antecedents, as demonstrated in (1-2) below.

Samuel threatened Justin with a knife, and
(1) … he blindfolded Erin with a scarf. [=Samuel]
(2) … Erin blindfolded him with a scarf. [=Justin]

However, the Parallel Function Preference fails in (3), where the object pronoun refers to a subject antecedent.
(3) Samuel threatened Justin with a knife, and
Erin stopped him. [=Samuel]
In light of examples like (3), Smyth (1994) suggests a modification to the Preference, which requires a perfect match between the two clauses in both argument structure (number and type of arguments) and sentence structure (presence/absence of adjunct modifiers).

In cases in which these conditions are not met, a subject preference kicks in, supporting resolution to a subject antecedent. This Qualified Parallel Preference makes the correct prediction in (3), because the first clause contains an adjunct modifier which is not matched in the second clause. It does not account for cases like (4), however, where the structures are fully parallel. The wrong prediction is made for (5) as well, where a lack of parallel structure does not trigger the subject preference. In both (4) and (5), the subject pronoun instead refers to an object antecedent.

Samuel threatened Justin with a knife, and
(4) he alerted security with a shout. [=Justin]
(5) he alerted security. [=Justin]
Note that neither the basic Parallel Function Preference nor its more restricted variant takes into account the semantic cues in (3-5) which support non-parallel reference.

Subject Preference

The Subject Preference, in its basic form, holds that pronouns are preferentially resolved to subject antecedents. The account assumes increased salience for arguments in subject position, which guides antecedent selection for both subject and object pronouns, as seen in (6-7).

Samuel threatened Justin with a knife, and
(6) … he blindfolded Erin with a scarf. [=Samuel]
(7) … Erin stopped him with pepper spray. [=Samuel]
The Subject Preference, as formulated by Crawley & Stevenson (1990), however, will not override explicit cues to antecedent reference, for example, pronoun gender. Nor is the Subject Preference predicted to override a resolution to a more plausible referent, as in (4-5) above, where the semantics of the sentence show a causal bias toward the non-subject antecedent. Note, however, that the Subject Preference does not account for cases like (2), where no such bias is present.

The Coherence Hypothesis

An alternative account suggests that pronoun interpretation is not the sum of so many superficial cues, but rather the byproduct of a larger inferencing process, wherein hearers make sense of a discourse by inferring coherence relations between successive utterances (Hobbs 1979). Kehler (2002) extends Hobbs's proposal, arguing that interpretation 'preferences' are actually epiphenomena of the manner in which different types of coherence are established. Consider (8-9).

Dennis narrowly defeated Isaac, and
(8) … Lilly congratulated him.   RESULT
(9) … Lilly utterly trounced him.   PARALLEL

In (8), there exists a plausible causal relation between the two clauses. Interpreting the object pronoun as coreferent with *Dennis* supports this relation. Interpreting the pronoun instead as coreferent with *Isaac* requires the accommodation of extra inferences to explain why Lilly congratulated the loser and not the victor.

In (9), by contrast, a causal relation seems less likely. In this case, what stands out is the resemblance between the two clauses. Although Lilly's win was handier than Dennis', it is more likely that the passage coheres due to the parallelism between the two events and not due to some causal relationship. A PARALLEL coherence relation supports parallel coreference, and the pronoun corefers with *Isaac*.

Manipulating Coherence

Wolf et al. (2004) tested the Coherence Hypothesis against both the Subject Preference Hypothesis and the Parallel Function Hypothesis in a reading time experiment, finding evidence in support of the Coherence Hypothesis. In a 2x2 design, coherence frame (PARALLEL/RESULT) and antecedent position (subject/object) were varied. Coherence was manipulated through verb meaning and the use of the connectives *and similarly* and *and so* (for PARALLEL/RESULT relations, respectively). In PARALLEL frames, the verbs in both clauses were synonyms; in RESULT frames, verb semantics incorporated a bias toward the non-parallel referent. Antecedent position was signaled by pronoun gender (masculine/feminine), which matched a name in the preceding clause occurring in either subject or object position. In half of the stimuli, the coreference indicated by pronoun gender supported the coherence relation as indicated by verb semantics and the type of connective; in the remaining half it did not. Examples are given in (10-11).

(10) Fiona complimented Craig and similarly
James congratulated her/him after the match,
but nobody took any notice.   PARALLEL
(11) Fiona defeated Craig and so
James congratulated her/him after the match,
but nobody took any notice.   RESULT

Faster reading times were measured for the following coreference patterns: parallel coreference was preferred with a PARALLEL coherence frame; non-parallel reference with a RESULT frame. Wolf et al. thus confirmed that preferences for pronoun interpretation can be reversed by manipulating coherence, as predicted by Kehler (2002).

Open Questions

Although the Wolf et al. study provides crucial preliminary support for the Coherence Hypothesis, a number of important questions remain. Among these is whether the observed effect will transfer to a pronoun interpretation task with *ambiguous* pronouns. A second issue involves the range of possible resolution patterns. Wolf et al. restricted their study to object pronouns, demonstrating that object pronouns can be resolved to both subject and object antecedents. Subject pronouns are of particular interest, however, as both preferences predict a subject-pronoun-to-subject-antecedent resolution pattern, but neither predicts a subject-pronoun-to-object-antecedent resolution. This suggests the latter pattern is in fact dispreferred, but is such an assumption warranted?

Table 1: Possible Coreference Patterns

	subject antecedent	object antecedent
subject pronoun	✓	?
object pronoun	✓	✓

Finally, it is not clear, given the Wolf et al. design, whether the observed effects are necessarily due to coherence or whether an alternative explanation might apply. For example, all of the Wolf et al. stimuli include a prepositional phrase in the second clause. This introduces non-parallel structure, which, according to the Qualified Parallel Hypothesis, might disrupt the parallel coreference pattern. It has also been proposed that connectives can carry extra focusing properties which may disrupt other preferences (Stevenson et al. 1994, 2000). The Wolf et al. result does not rule out these possibilities.

Experiment

The present experiment asks whether the pronoun interpretation preferences reported in the literature can be interpreted as epiphenomena of the manner in which different types of coherence are established. As in the majority of the studies described above, we use an offline disambiguation task, focusing on the outcome of antecedent selection, as opposed to the time course of the processes supporting it. (Additional studies to address the latter are in development. See Discussion below.)

In a 2x2x2 design, we constructed stimulus sets with 8 variants, as shown in (12-15). Each stimulus contains two clauses: an introduction and a follow-on, both of which contain a transitive verb in active voice. The follow-on clause contains an ambiguous pronoun.

Samuel threatened Justin with a knife, and
(12) … Erin blindfolded him (with a scarf).   PARALLEL
(13) … Erin stopped him (with pepper spray).   RESULT
(14) … he blindfolded Erin (with a scarf).   PARALLEL
(15) … he alerted security (with a shout).   RESULT

Notice that half of the variants contain an object pronoun, while half contain a subject pronoun. With this design, we are able to test the full range of possible coreference patterns in Table 1, including the subject-pronoun-to-object-antecedent resolution pattern that is not predicted by either preference account.

This design also allows us to test the Qualified Parallel Preference, as we have included variants in both a fully parallel condition (a modifier in the follow-on clause matches the modifier in the introductory clause) and a partially parallel condition (no modifier in the follow-on). Further, we can test the Qualified Subject Preference by excluding RESULT coherence frames, which, as in the Wolf et al. experiment, incorporate a semantic bias toward the non-parallel referent. Finally we rule out the possibility of a connective focusing effect by using the connective *and* across all stimuli.

Thus by manipulating coherence, structure, and pronoun position independently, we set up tests for each of the competing hypotheses and address the issues left open by Wolf et al (2004).

Predictions

The strongest form of the Subject Preference predicts an across-the-board preference for subject antecedents, while the Qualified Subject Preference predicts a subject effect only in non-biasing contexts (here, in PARALLEL coherence frames). The Parallel Function Hypothesis predicts across-the-board parallel coreference (main effect of pronoun position), while the Qualified Parallel Structure Hypothesis predicts parallel coreference for only those sentences which show full, as opposed to partial, parallel structure (interaction between pronoun position and sentence structure). The Coherence Hypothesis predicts parallel coreference in PARALLEL coherence frames, and non-parallel reference in RESULT frames (interaction between pronoun position and coherence relation).

Method

Participants

Participants were 32 undergraduates from the University of California, San Diego. All were self-reported monolingual native speakers of English. Participants received extra credit for participation.

Stimuli

16 stimulus sets were constructed, with 8 variants per set, for a total of 128 stimuli. Each set varied pronoun position (subject/object), sentence structure (fully/partially parallel), and coherence relation (PARALLEL/RESULT). An example set is provided above in (12-15).

Full/partial structural parallelism was based on the presence or absence of a modifier phrase in the second clause to match the modifier in the first clause. Modifiers were varied between pre-verbal adverbs and post-verbal prepositional phrases, balanced across sets. To allow for diversity in the stimulus set, verbs in the introductory clause were also varied across four types: physical action, social action, mental state verbs, and verbs of address.

Coherence frame was assessed in a prior norming phase, during which trained judges (Linguistics graduate students), who were blind to our hypothesis, categorized stimuli as instances of either a PARALLEL or RESULT coherence relation. For 119 of the 128 stimuli, 3 of 3 judges agreed on the coherence relation. For the remaining 9 stimuli, 2 of 3 judges agreed on the coherence relation with an average confidence score of 8 or more on an 11 point scale.

Design

A repeated measure design was employed, where each participant was tested on 2 stimuli from each of the 8 types, and where no two stimuli presented to the same participant were variants from the same set. The two replications were block randomized, and the 16 experimental stimuli were interleaved with 24 distracters; 16 of which also contained ambiguous pronouns. The resulting 16 lists were then reversed to rule out ordering effects, yielding 32 unique stimulus lists.

Participants were presented with a paper and pencil task, for which they read a two-clause passage and answered a question immediately after, as in (16).

(16) Samuel threatened Justin with a knife, and he
blindfolded Erin with a scarf.
Who blindfolded Erin?

The answer was taken to indicate the antecedent selected by the participant in interpreting the ambiguous pronoun. Where the answer matched the subject of the introductory clause, a score of 1 was assessed. Where the answer matched the object of the introductory clause, a score of 0 was assessed. This score was used as the dependent variable for data analysis.

Data Analysis

A 2x2x2 analysis of variance was conducted with the following factors: pronoun position (subject/object), sentence structure (fully/partially parallel), coherence relation (PARALLEL/RESULT). Separate analyses were conducted with participants (F_1) and items (F_2) as random variables. Two one-sample t tests were also conducted to compare the overall mean and the mean for a restricted set (PARALLEL coherence only) to chance.

Results

The data showed strong support for the Coherence Hypothesis, and were not consistent with any of the other hypotheses. There was no evidence of a Subject Preference or of a Parallel Function Preference, in either their basic or qualified formulations.

Overview

Examining gross percentages across participant responses, no evidence of a subject preference is found. Across all

stimuli, the percentage of resolutions to subject and to object antecedents are very close at 52:48. (See Table 2.) If we omit the stimuli which exhibit semantic bias and consider only those in PARALLEL coherence frames, the distribution is 54:46, and as such there is no evidence of a Qualified Subject Preference.

Table 2: Resolutions (%) to Subject/Object Antecedent

	Antecedent Subj	Obj	n
Subject Preference			
all pronouns	0.52	0.48	512
Qualified Subject Preference			
non-biasing context	0.54	0.46	256
Parallel Structure Preference			
subject pronouns	0.51	0.49	256
object pronouns	0.52	0.48	256
Qualified Parallel Preference			
subject pronouns: fully parallel structure	0.52	0.48	128
object pronouns: fully parallel structure	0.50	0.50	128
Coherence Hypothesis			
subject pronouns: PARALLEL coherence	**0.98**	0.02	128
subject pronouns: RESULT coherence	0.05	**0.95**	128
object pronouns: PARALLEL coherence	0.10	**0.90**	128
object pronouns: RESULT coherence	**0.94**	0.06	128

Similarly, the percentages do not support a Parallel Function Preference. We find that subject pronouns are resolved to subject antecedents about as often as to object antecedents (51:49). Object pronoun resolutions are similarly even at 52:48 for subject/object antecedents respectively. The Qualified Parallel Preference fares no better, with subject pronouns resolving to subject antecedents 52% of the time when structure is fully parallel. Object pronouns in the fully parallel condition resolve to object antecedents 50% of the time.

Given these near 50:50 splits, the coherence data are particularly dramatic. We find parallel coreference for subject pronouns in a PARALLEL frame 98% of the time. Subject pronouns show non-parallel coreference in the RESULT frame at a rate of 95%. Object pronouns show parallel coreference in the PARALLEL frame 90% of the time, and object pronouns show non-parallel coreference (resolve to subject) in the RESULT frame at a rate of 94%.

Statistical Analysis

Statistical analysis of these data confirms that the interaction between coherence frame and pronoun position, predicted by the Coherence Hypothesis, is significant [$F_1(1,31) = 1379.23$, p < .0001; $F_2(1,15) = 2016.158$, p < .0001]. A second, smaller effect, which we did not predict, was found for coherence alone [$F_1(1,31) = 4.429$, p = .044; $F_2(1,15) = 7.105$, p = .018].

The confidence interval for the overall mean is .516 +/- .043. (Recall that a score of 1 indicates resolution to a subject antecedent, 0 to an object.) Based on a one-sample t test comparing the overall sample mean to a hypothetical mean of .5, we conclude that it is not significantly different from chance [$t(511) = .707$, p = .240], contra the Subject Preference Hypothesis. Comparing the mean for PARALLEL relations only (.539 +/- .061) to the hypothetical mean of .5, we find the mean for this restricted set is not significantly different from chance [$t(255) = 1.251$, p = .106], ruling out a Qualified Subject Preference as well.

The main effect of pronoun position, predicted by the Parallel Function Preference, is not statistically significant [$F_1(1,31) = .088$, p = .768; $F_2(1,15) = .105$, p = .751], nor is the interaction between sentence structure and pronoun position, predicted by the Qualified Parallel Preference [$F_1(1,31) = 1.130$, p = .300; $F_2(1,15) = 1.552$, p = .232].

Discussion

These results strongly support the Coherence Hypothesis, confirming our prediction that pronoun interpretation preferences can be reversed through the manipulation of coherence relations. Furthermore, the Coherence Hypothesis makes correct predictions across the full range of interpretation patterns, as described below.

Against Preference-Based Accounts

Recall that both the Subject Preference Hypothesis and the Parallel Function Hypothesis predict only a subset of possible resolution patterns. The Coherence Hypothesis, on the other hand, makes correct predictions for all four of the patterns under investigation.

Table 3: Comparison of Data Coverage (Pronoun:Antecedent)

	subj:subj	subj:obj	obj:obj	obj:subj
Subject Preference:	✓	✗	✗	✓
Parallel Function:	✓	✗	✓	✗
Coherence:	✓	✓	✓	✓

Further, we find that the qualifications introduced to extend both the Parallel Function Hypothesis and the Subject Preference Hypothesis are untenable. We show, for example, that parallel coreference does not depend on perfectly parallel structure; parallel coreference was observed in PARALLEL coherence frames, whether the structures were fully or partially parallel. Similarly, even when we exclude the RESULT coherence relations, there is no evidence of a subject preference; antecedent selection is

split between subjects and objects. And finally, we rule out a focusing role for connectives (Stevenson et al. 1994, 2000).

Table 4: Summary of Results

Hypothesis		Prediction
Subject Preference:	✖	Across-the-board preference for subject antecedents
Qualified Subject Preference	✖	Preference for subject antecedent in non-biasing contexts (i.e. PARALLEL relations)
Parallel Function Preference:	✖	Across-the-board preference for parallel coreference
Qualified Parallel Preference:	✖	Preference for parallel coreference in fully-parallel structure condition
Coherence:	✓	Parallel coreference with PARALLEL coherence; Non-parallel coreference with RESULT coherence

Perhaps more compellingly, however, we show that the *basic* preferences supported by the Subject Preference Hypothesis and the Parallel Function Hypothesis simply disappear when the relevant factors are balanced. By manipulating pronoun position and coherence frame, we were able to systematically 'break' these preferences, supporting our claim that coherence plays a crucial role in the interpretation of pronouns.

Parallelism

In fact, coherence appears to be at the heart of the parallelism account. The kinds of data that meet Smyth (1994)'s requirements for strict parallelism are often instances of PARALLEL coherence relations. This suggests that the relevant similarity between two clauses is not simply structural, but more importantly propositional. Recognizing a PARALLEL coherence relation, we propose, depends upon being able to match up meaningful units in distinct propositions. And constructing sentence pairs that have similar propositional content is very likely facilitated by similarity in structure. For example, one would expect that verbs with similar meaning would have similar argument structures, taking the same number of arguments, in the same order. The exact nature of the relationship between parallel structure and PARALLEL coherence, however, remains an important issue for future research.

The Subject Effect

The Subject Preference has been documented by a variety of researchers, and is supported by a plausible model of discourse processing which takes into account issues of salience and focus of attention (c.f. related proposals identifying a topic preference and first-mention effects). So

why do our results show a subject preference occurring at close to chance levels? As suggested above, the subject preference disappears *when the relevant factors are balanced.*

In this case, the relevant factor is coherence frame. As described in the immediately preceding section, there is no default subject preference within a PARALLEL coherence frame. Subject pronouns tend to corefer with subject antecedents and object pronouns with object antecedents. Similarly, there is no default antecedent position in a RESULT coherence frame. In this experiment, we examined cases where causal bias supports non-parallel reference, as in (17) where a subject pronoun refers to an object antecedent, but causal inferencing might also support a parallel subject antecedent, as in (18).

Peter snapped at Ethan, and
(17) … he sulked the rest of the afternoon. [=Ethan]
(18) … he felt guilty the rest of the afternoon. [=Peter]

Thus there appears to be little support for a subject preference in PARALLEL and RESULT coherence frames.

But what about other types of coherence? There may be reason to believe that the subject preference is more robust within the 'OCCASION' coherence relation, where multiple discourse segments combine to form a narrative. (See Kehler 1997 for discussion.) Consider the passage in (19):

(19) Rubens passed Lopez in the final lap, and
he went on to win the race. [=Rubens]

The flow of the narrative seems to be following Rubens, and the passage as a whole makes the most sense if the driver in front is the one who wins the race.

Whether and how salience interacts with coherence is an open question, and we are currently developing studies to address the issue. While the current study followed the previous literature in using an offline task, experiments currently in preparation will utilize online measures such as self-paced reading time, which we hope will provide a fuller picture of the processes supporting pronoun interpretation. For example, although the ultimate resolution for (17) is to a non-subject antecedent, one might ask, in light of the subject and parallelism preferences, whether reading times are delayed in comparison to cases like (18).

Semantic Focusing

Stevenson et al. (2000) make reference to coherence in formulating their Semantic Focusing Model of pronoun interpretation. They propose a focusing mechanism associated with discourse connectives and then contrast this revised focusing account with a 'relational' analysis. Stevenson et al. find that the revised Semantic Focusing Model makes better predictions than the Relational Model in a sentence continuation task.

We suggest that the reliance on discourse connectives to signal coherence in their work is a critical flaw. Under such an analysis, antecedent preferences should be stable unless a focusing connective is introduced. The present experiment shows, however, that antecedent preferences can be shifted through the manipulation of coherence in the absence of

discourse connectives. Furthermore, the Relational Model assumed by Stevenson et al. requires a one-to-one mapping between coherence relation and preferred referent (for example, the patient argument is always preferred in a result coherence relation). That assumption is not a prediction of the coherence model they cite (Hobbs 1979), and the data reported here refute the assumption, showing that pronouns can refer back to either a subject (agent) or an object (patient) in either the PARALLEL or the RESULT coherence relation.

The Coherence Model

The Coherence Model is conceptually distinct from preference-based accounts of pronoun interpretation. Each of the preference models holds that morpho-syntactic cues have a direct bearing on pronoun interpretation. The Coherence Model, on the other hand, holds that such effects are mediated by the interpretation of coherence.

Table 5: Two Models of Pronoun Interpretation

Model	Mechanism
Preference Models:	morpho-syntax $\Rightarrow$ pronoun interpretation
Coherence Model:	coherence morpho-syntax $\nearrow$ $\nwarrow$ pronoun interpretation

Evidence for the mediating effect of coherence was provided in the present experiment, where morpho-syntactic structure (including transitivity, tense, aspect, voice, and syntactic structure) was matched across stimulus pairs, but coherence was manipulated. Further evidence comes from ambiguous passages like (20), taken from Kehler (2002), where two competing interpretations of the object pronoun are possible.

(20) Colin Powell defied Dick Cheney, and
George Bush punished him.

The object pronoun *him* can be interpreted as coreferent with either the subject of the preceding clause, *Colin Powell,* or the object, *Dick Cheney*. Crucially, these competing interpretations are each aligned with a distinct coherence frame. Under one interpretation, Cheney is having a bad day: first Powell defies him, and then Bush punishes him. A PARALLEL relation supports parallel coreference. Under the alternative, Powell got what he deserved: he defied Cheney, and so Bush punished him. In this case, a RESULT relation supports non-parallel reference.

As this example involves a single utterance with two potential interpretations, there is clearly no morpho-syntactic trigger which can select between the two. The only difference between the two interpretations is the supporting coherence frame. We argue that any successful account of the data presented here, including ambiguities like (20), must appeal to coherence.

Summary

We described an off-line pronoun interpretation experiment in which a Coherence Hypothesis was tested against the Subject Preference Hypothesis and the Parallel Function Hypothesis. We demonstrated that preferences can be systematically disrupted through the manipulation of coherence relations, and that when the relevant factors are balanced, preferences disappear. We addressed issues left open by previous work (Wolf et al. 2004) and showed that only the Coherence Hypothesis makes correct predictions for the full range of coreference patterns under investigation. We suggested further investigation be carried out using online measures and outlined areas for future research, including possible interactions between salience and coherence.

References

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

Crawley, R. A., & Stevenson, R. J. (1990). Reference in single sentences and in texts. *Journal of Psycholinguistic Research*, *19*, 191-210.

Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science, 3*, 67-90.

Kehler, A. (1997). Current theories of centering for pronoun interpretation: a critical evaluation. *Computational Linguistics*, *23(3)*, 467-475.

Kehler, A. (2002). *Coherence, reference, and the theory of grammar.* Stanford, CA: CSLI Publications.

Stevenson, R. J., Crawley, R. A., & Kleinman, D. (1994). Thematic roles, focus, and the representation of events. *Language and Cognitive Processes*, *9,* 519–548.

Stevenson, R. J., Knott, A., Oberlander, J., & McDonald, S. (2000). Interpreting pronouns and connectives: Interactions between focusing, thematic roles and coherence relations. *Language and Cognitive Processes, 15*, 225–262.

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

Wolf, F., Gibson, E. & Desmet, T. (2004). Discourse coherence and pronoun resolution. *Language and Cognitive Processes, 19(6)*, 665–675.