

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Coherence-driven effects in sentence and discourse processing

Permalink

<https://escholarship.org/uc/item/0v09m5zt>

Author

Rohde, Hannah

Publication Date

2008

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Coherence-Driven Effects in Sentence and Discourse Processing

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in
Linguistics

by

Hannah Rohde

Committee in charge:

Professor Andrew Kehler, Chair
Professor Ivano Caponigro
Professor Jeffrey L. Elman
Professor Victor Ferreira
Professor Roger Levy

2008

Copyright
Hannah Rohde, 2008
All rights reserved.

The Dissertation of Hannah Rohde is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2008

TABLE OF CONTENTS

Signature Page	iii
Table of Contents	iv
List of Abbreviations	viii
List of Figures	ix
List of Tables	xii
Acknowledgments	xiii
Vita	xvi
Abstract	xvii
1 Introduction	1
1.1 Discourse Factors in Comprehension	1
1.2 Phenomena Addressed in this Work	6
1.2.1 Pronouns and Coreference	7
1.2.2 Relative Clause Attachment	11
1.3 Outline of the Dissertation	13
2 Incorporating Discourse into Models of Processing	18
2.1 Predictability in Processing	19
2.2 The Role of Discourse Context	21
2.3 Intersentential Coherence Relations	24
3 Coreference: Event Structure & Causal Inference	29
3.1 Interpreting Ambiguous Pronouns	31
3.2 Experiment I: Event Structure	35
3.2.1 Methods	37
3.2.2 Results	40

3.2.3	Discussion of Event Structure Results	41
3.2.4	Coherence Analysis	42
3.2.5	Discussion of Coherence Results	48
3.2.6	Comparison with Previous Work	50
3.3	Experiment II: Object Type	52
3.3.1	Methods	53
3.3.2	Results	54
3.3.3	Discussion	58
3.4	Experiment III: Instruction Type	59
3.4.1	Methods	60
3.4.2	Results	61
3.4.3	Discussion	64
3.5	Signaling Intersentential Relationships	66
3.6	Implicit Causality	69
3.7	Experiment IV: Implicit Causality Biases	72
3.7.1	Methods	72
3.7.2	Results	74
3.7.3	Discussion	78
3.8	Conclusion	79
3.9	Acknowledgments	82
4	Bidirectional Influences in Coherence and Coreference	84
4.1	Experiment V: Unambiguous IC Contexts	86
4.1.1	Methods	90
4.1.2	Results	92
4.1.3	Discussion	97
4.1.4	Pronominalization in IC Contexts	99
4.2	Experiment VI: Ambiguous IC Contexts	102
4.2.1	Methods	103
4.2.2	Results	104
4.2.3	Discussion	109

4.3	Experiment VII: Pronouns in TOP Contexts	110
4.3.1	Methods	113
4.3.2	Results	114
4.3.3	Discussion	116
4.4	Experiment VIII: Ambiguous TOP Contexts	117
4.4.1	Methods	118
4.4.2	Results	119
4.4.3	Discussion	122
4.5	TOP Coreference Revisited	123
4.6	Conclusion	130
4.7	Acknowledgments	131
5	Syntactic Ambiguity	132
5.1	Relative Clause Attachment	133
5.1.1	Implicit Causality Biases	137
5.1.2	IC Biases in RC Attachment	138
5.2	Experiment IX: IC Biases	140
5.2.1	Methods	140
5.2.2	Results	144
5.2.3	Discussion	145
5.3	Experiment X: On-line IC Biases	148
5.3.1	Methods	149
5.3.2	Results	152
5.3.3	Discussion	154
5.4	General Discussion	155
5.5	Acknowledgments	157
6	Structuring a Discourse	158
6.1	Structuring a Discourse with Questions	161
6.2	Experiment XI: Monologues/Dialogs in IC Contexts	164
6.2.1	Methodology	165
6.2.2	Results	167

6.2.3 Discussion	168
6.3 Experiment XII: Monologues/Dialogs in Transfer Contexts	169
6.3.1 Methods	170
6.3.2 Results	171
6.3.3 Discussion	173
6.4 General Discussion	174
6.5 Acknowledgments	176
7 Conclusion	177
7.1 Contributions and Context	177
7.2 Overview of Findings	181
7.3 Future Research	183
7.4 Final Thoughts	185
A Experimental Stimuli	187
A.1 Experiment I	187
A.2 Experiment II	188
A.3 Experiment III	189
A.4 Experiment IV	189
A.5 Experiment V	192
A.6 Experiment VI	192
A.7 Experiment VII	192
A.8 Experiment VIII	195
A.9 Experiment IX	195
A.10 Experiment X	196
A.11 Experiment XI	198
A.12 Experiment XII	198
B Guidelines for Coherence Annotation	199
References	202

LIST OF ABBREVIATIONS

<i>CR</i>	Coherence Relation
<i>Elab</i>	Elaboration Coherence Relation
<i>Exp</i>	Explanation Coherence Relation
<i>GI</i>	Goal Interpretation
<i>IC</i>	Implicit Causality
<i>Imp</i>	Imperfective Verbal Aspect
<i>Occ</i>	Occasion Coherence Relation
<i>Par</i>	Parallel Coherence Relation
<i>Perf</i>	Perfective Verbal Aspect
<i>QUD</i>	Question Under Discussion
<i>RC</i>	Relative Clause
<i>Ref</i>	Referent
<i>Res</i>	Result Coherence Relation
<i>V-E</i>	Violated Expectation Coherence Relation
<i>VT</i>	Verb Type

LIST OF FIGURES

Figure 3.1: Effects of verbal aspect on pronoun interpretation in Experiment I	40
Figure 3.2: Effects of verb class and verbal aspect on pronoun interpretation in Experiment I	41
Figure 3.3: Pronoun interpretation across coherence relations for perfective passages in Experiment I	45
Figure 3.4: Pronoun interpretation across coherence relations for imperfective passages in Experiment I	47
Figure 3.5: Coherence relations and pronoun interpretation biases for perfective and imperfective passages in Experiment II	55
Figure 3.6: Coherence relations and pronoun interpretation biases for passages with normal and abnormal objects in Experiment II	57
Figure 3.7: (Lack of) effect of object on pronoun interpretation in Experiment II	58
Figure 3.8: Coherence relations following ‘Why’ instructions in Experiment III	62
Figure 3.9: Coherence relations following ‘What Next’ instructions in Experiment III	62
Figure 3.10: Effects of instruction type and verbal aspect on pronoun interpretation in Experiment III	64
Figure 3.11: Coherence and next-mention biases for IC-1 verbs in <i>because</i> and full-stop conditions in Experiment IV	75
Figure 3.12: Coherence and next-mention biases for IC-2 verbs in <i>because</i> and full-stop conditions in Experiment IV	76
Figure 3.13: Coherence and next-mention biases for non-IC verbs in <i>because</i> and full-stop conditions in Experiment IV	77
Figure 4.1: Effects of prompt type on proportions of NP1-biased Explanations and NP2-biased Results following IC-1 verbs in Experiment V	94
Figure 4.2: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Explanations following IC-2 verbs in Experiment V	95

Figure 4.3:	Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following non-IC verbs in Experiment V	96
Figure 4.4:	Effects of verb type and referent position on the the rate of pronominalization in Experiment V	100
Figure 4.5:	Effects of verb type and referent position on the rate of pronominalization in Experiment VI	106
Figure 4.6:	Effects of prompt type on proportions of NP1-biased Explanations and NP2-biased Results following IC-1 verbs in Experiments V & VI	107
Figure 4.7:	Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following IC-2 verbs in Experiments V & VI	108
Figure 4.8:	Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following IC-2 verbs in Experiments V & VI	109
Figure 4.9:	Effect of verbal aspect on choice of next mention across all data in Experiment VII and the pronoun-only subset	115
Figure 4.10:	Effect of referent position and verbal aspect on the rate of pronominalization in Experiment VII	115
Figure 4.11:	Effect of verb aspect and referent position on the rate of pronominalization in Experiment VIII	121
Figure 4.12:	Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Occasions following perfective verbs in Experiments VIII	121
Figure 4.13:	Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Occasions following imperfective verbs in Experiments VIII	122
Figure 5.1:	Proportions of high attachments and proportions of Explanation RCs	144
Figure 5.2:	Proportion of high attachments for Explanation and non-Explanation RCs	145
Figure 5.3:	Residual reading times	153
Figure 6.1:	Proportion of Explanation/Why responses in Experiment XI	168
Figure 6.2:	Effects of verbal aspect on Occasion/What-next responses in Experiment XII	172

Figure 6.3: Effects of verbal aspect on Elaboration/How responses in Experiment XII	172
Figure 6.4: Effects of verbal aspect on Explanation/Why responses in Experiment XII	173

LIST OF TABLES

Table 3.1:	Experiment I transfer-of-possession verb classes	38
Table 3.2:	Examples of coherence-annotated passages from Experiment I	44
Table 5.1:	Raw RTs and question accuracy (subject means \pm standard error)	152

ACKNOWLEDGEMENTS

I am hugely indebted to my advisor Andy Kehler whose clarity of thought and commitment to empiricism have shaped my education and my interest in linguistics. Andy bravely accompanied me down the psycholinguistics path, a new path for both of us, and guided the projects that were the foundation of this research. The sheer number of hours of his time that he has offered (or that I have stolen) allowed me to complete this work, and the results would not be what they are without Andy's advice and insights. It may sound like the most boring of compliments, but my favorite thing about Andy is that he is *reasonable*. I feel lucky to have worked with and learned from someone who keeps perspective, helps figure out the logical arc of an idea, asks the right questions, and does all this in a human and witty way.

I also want to thank Jeff Elman who played a critical role in helping kick-start this research. He sat with me as I ran my first ANOVA and has offered advice and support as I've taken this work on the road and faced new and different audiences. Roger Levy is of course irreplaceable; he has helped work out designs for experiments and analyses of the results and in general has been a fantastic resource for advice and collaboration. I owe particular thanks to Ivano Caponigro for his moral support and ready availability to dive in and discuss my work. Vic Ferreira, too, has been a great sounding board for ideas and has opened my eyes to different methodologies for my own projects and to the range of research in the psycholinguistics field.

Many thanks to those who provided sentence annotation and lab help: Jeff Derrenberger, Alex del Giudice, Erica Gold, Lara Taylor, Laura Kertz, and Brittany Young. Along with my committee, several people have had to put up with earlier versions of this material, and their suggestions and challenges were often just what I needed: Farrell Ackerman, Eric Bakovic, Chris Barker, Henry Beecher, Klinton Bicknell, Rebecca Colavin, Katie Cooke, Gabe Doyle, Laura Kertz, Cindy Kilpatrick, Robert Kluender, John Moore, and Sharon Rose. Elizabeth Smith read my equations

and used her degree in mathematics (well, probably her high school math) to help me with my analysis. danah boyd wrote a far longer and scarier dissertation that allowed me to take comfort in my own. Klinton Bicknell helped sort out statistics and scrawled on the whiteboard with Bayesian models. Josh Lewis, David Liu, Scott Richmond, Laura Staum, and Matt Tong wasted time with me playing online Scrabble. Additional thanks go to Philip Hofmeister for (unknowingly) letting me copy parts of the structure of his dissertation when I originally faced the blank screen of the first page of this thesis. I drew the line at plagiarizing his acknowledgments, not wanting to risk thanking someone else's mother.

An immeasurable debt of gratitude is owed to Laura Kertz and Cindy Kilpatrick who have been partners in crime through these years of attempting to do things we didn't know we could. It's an honor to be friends and colleagues with such decent and strong people, and it's hard to imagine facing linguistics without their daily support.

In the acknowledgments homestretch, I get to mention my parents, whose choice to instill in me an early love of language makes it easy to thank them. My mom is the best closet prescriptivist that I know, along with being my advocate and counsel. She is also a strong believer in the scientific method and the beauty of scientific pursuits. Although my father may deny being an academic, it was he who, sometime when I was in middle school, showed me the etymology of the word *posh*, and it's been downhill ever since.

Ronen Barzel provided advice and cheerleading at all hours and on all fronts from 500 miles away. Thanks for listening to me when the printer died. And thanks for taking none of this too seriously.

* * *

This work was supported by NIH grant T32-DC000041 to the Center for Research in Language at the University of California, San Diego and NSF IGERT grant DGE-0333451. In several chapters, the material discussed has been reported or submitted elsewhere in co-authored work. The dissertation author was the primary investigator for that material, and it is used here with permission from the co-authors. The material in Chapter 3 was reported, in part, in Rohde, Kehler, and Elman (2006) [Event Structure and Discourse Coherence Biases in Pronoun Interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*. Vancouver, July 26-29, 2006. Pages 697-702], in Rohde, Kehler, and Elman (2007) [Pronoun Interpretation as a Side Effect of Discourse Coherence. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Nashville, TN, August 1-4, 2007. Pages 617-622], and in Kehler, Kertz, Rohde, and Elman (2008) [Coherence and Coreference Revisited. *Journal of Semantics*. Volume 25: 1-44]. The material from Chapter 4 was reported, in part, in Rohde and Kehler (2008) [The Bidirectional Influence between Coherence Establishment and Pronoun Interpretation. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008] and is discussed, in part, in Ferretti, Rohde, Kehler, and Crutchley (Submitted) [Verb Aspect, Event Structure, and Coreferential Processing. *Journal of Memory and Language*. The material in Chapter 5 was reported, in part, in Rohde, Levy, and Kehler (2008) [Implicit Causality Biases Influence Relative Clause Attachment. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008]. The material in Chapter 6 was reported, in part, in Rohde and Kehler (2008) [Coherence-Driven Expectations in Discourse and Dialog. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008].

VITA

- 2003 B.A., Linguistics and Computer Science,
Brown University
- 2004 Research Assistant, Department of Linguistics
University of California, San Diego
- 2005 M.A., Linguistics,
University of California, San Diego
- 2006 Teaching Assistant, Department of Linguistics
University of California, San Diego
- 2008 Ph.D., Linguistics
University of California, San Diego.

ABSTRACT OF THE DISSERTATION

Coherence-Driven Effects
in Sentence & Discourse Processing

by

Hannah Rohde

Doctor of Philosophy in Linguistics

University of California, San Diego, 2008

Professor Andrew Kehler, Chair

This dissertation provides a psycholinguistic investigation of the influence of discourse on language comprehension. It examines factors that allow comprehenders to follow a discourse, to form representations of the events being described, and to make predictions about how subsequent utterances will relate to prior linguistic material. Previous work has recognized the importance of prediction in sentence-internal processing: transition probabilities at the phonemic level, semantic associations in lexical access, and structural frequencies at the syntactic level. The work presented here investigates whether learnable statistical regularities also exist at the discourse level, a topic that has remained largely unexplored in the psycholinguistics literature.

The dissertation presents a series of experiments testing the extent to which comprehenders use various pragmatic cues to make predictions about how a discourse will be continued. In order to quantify discourse-level information, the experiments use an inventory of coherence relations adopted from the theoretical linguistics and artificial intelligence literatures. The experimental results demonstrate that comprehenders do indeed make use of available pragmatic cues to generate expectations about upcoming coherence relations. Furthermore, the results show that the mechanisms for establishing coherence relations can inform our understanding of two well-

studied sentence-internal phenomena: coreference and syntactic ambiguity. The online results establish the importance of these pragmatic cues in comprehenders incremental sentence processing.

The coherence-based approach taken here provides a lens through which to view previous results in the domains of both coreference and syntactic ambiguity. The fact that phenomena in both these domains appear to be sensitive to coherence-driven biases suggests that these biases may be more pervasive than has been previously acknowledged. This work indicates that future processing models of sentence and discourse processing must take into account effects that emerge from discourse coherence.

1

Introduction

1.1 Discourse Factors in Comprehension

Most of the linguistic input that we encounter as comprehenders woefully underdetermines the crucial structural, semantic, and pragmatic relationships that make language meaningful. These invisible relationships must be inferred by the listener or reader, who manages to do this gracefully and automatically. From a stream of words, we extract a plausible grammatical structure, positing syntactic and semantic relations between words and phrases. At the same time, we also identify higher order relationships that hold between sentences and that allow us to link together a series of sentences to form a coherent discourse.

Almost all sentences that we encounter are embedded in larger multi-sentence discourse contexts, yet our current models of language comprehension tend to focus on the sentence-internal process of combining words to form local syntactic and semantic relationships. The aim of this dissertation is to show that comprehenders generate expectations about the direction the discourse is likely to take — that is, how upcoming sentences will relate to the current one — and that those expectations influence the interpretation of linguistic phenomena internal to the sentence. Take, for instance, the following two-sentence passage:

(1) Mary scolded Sue. She kicked her.

The sentences in (1) describe two events, a scolding event and a kicking event. Several different relations can be inferred to hold between the two events, and the different relations in turn yield different interpretations of the two pronouns. Under a reading in which the two events are part of a list of things that Mary did to Sue, the subject pronoun *She* will be preferentially interpreted to refer to Mary while the object pronoun refers to Sue (a reading like *Mary scolded Sue, and then, on top of that, she_{Mary} kicked her_{Sue}*). Under a reading in which the kicking event provides the reason for the scolding event, *She* will preferentially be interpreted to refer to Sue and *her* to Mary (*Mary scolded Sue because she_{Sue} kicked her_{Mary}*). The causal and temporal order can also be reversed such that the second sentence describes a resulting event in which the kicking event is a response to the scolding event (*Mary scolded Sue. As a result she_{Sue} kicked her_{Mary}*).

These different types of intersentential relationships have been formalized in work on discourse coherence (Hobbs, 1979; Kehler, 2002; see also Knott, 1996 for a review of other models of discourse structure and coherence). Formal definitions for a variety of relations will be given in Section 2.3. In the meantime, at an intuitive level, the three coherence relations sketched above for (1) can be defined as follows. The list reading represents an Occasion relation whereby two sentences are related because they describe a sequence of ordered events in which the end state of the first event acts as the start state of the second event. The reading in which the kicking event is the reason for the scolding event represents an Explanation relation: The event denoted by the second sentence explains the event denoted by the first. The reading in which the scolding brought about the kicking event represents a Result relation: The event denoted by the first sentence causes the event denoted by the second sentence. Kehler has argued that a variety of previously proposed heuristic strategies for phenomena such as pronoun interpretation are largely side effects of the process of establishing these and other coherence relations.

In this dissertation, I adopt Kehler’s model of coherence to investigate pragmatic biases in sentence and discourse processing. I use the term *coherence-driven biases* to refer to two types of discourse-level preferences. The first concerns the probability of certain coherence relations, conditioned on context, and the second concerns the probability of other linguistic outcomes, conditioned on the operative coherence relation. A coherence-driven model is therefore one that attempts to capture comprehenders’ behavior by considering the ways in which comprehension depends on inferences about the relationships that hold between sentences. This dissertation is a psycholinguistic investigation of a model of discourse coherence, so the units over which the probabilities are calculated comprise a discrete set of coherence relations. Throughout this dissertation, I will be looking at contextual cues that bias comprehenders regarding what the operative coherence relation is likely to be ($p(CR | Context)$) and at linguistic outcomes that depend on the operative coherence relation ($p(outcome | CR)$). I aim to show that the mechanisms for establishing coherence relations between sentences, a topic that has remained largely unexplored in the psycholinguistics literature, can inform our understanding of well-studied sentence-internal phenomena and can expand the types of questions we ask about what it takes for a comprehender to process language successfully.

The experiments described here test the effects of coherence-driven biases on language processing by considering two general phenomena that are crucial to language understanding and have received quite a bit of attention in the psycholinguistics literature: patterns of coreference and the resolution of local syntactic ambiguity. Coreference comprehension is concerned with tracking the relevant entities and individuals that are re-mentioned across sentences in a discourse; syntactic parsing is concerned with uncovering the intended structural relationships that hold between words and phrases in a sentence. In both cases, ambiguity must be resolved in order to understand the nature of the events being described and the roles of the relevant event participants. The coherence-based approach taken here provides a new lens through which to view previous results in both of these domains. For coreference,

what I find is that many of the surface-level heuristics that have been posited in the psycholinguistics literature can be better understood as components of more general discourse-level processes and representations. For syntactic ambiguity, I find that a widely-reported structural attachment bias can be reversed in contexts in which the pragmatic relationship between two clauses is relevant to their structural relationship.

The choice to use discourse coherence relations as a unit over which to measure expectations allows me to address the long-standing issue of incorporating discourse information (the nebulous pragmatic wastebasket, according to Kadmon, 2001) into probabilistic models of sentence processing. There is consensus about the importance of appealing to discourse factors in domains like pronoun interpretation (Grosz, Joshi, & Weinstein, 1995; Stevenson, Knott, Oberlander, & McDonald, 2000; Arnold & Griffin, 2007) and syntactic ambiguity (Gilboy, Sopena, Clifton, & Frazier, 1995; Frazier & Clifton, 1996; Traxler, Pickering, & Clifton, 1998; Hemforth, Konieczny, & Scheepers, 2000; Desmet, de Baecke, & Brysbaert, 2002; van Berkum, Brown, & Hagoort, 1999). In both areas, a host of morphosyntactic biases and heuristics have already been identified but the problems remain unsolved. Most researchers would not deny that complex inferencing affects language usage and comprehension, and there is a rich history of such analysis in the realms of formal semantics and artificial intelligence. But identifying a quantifiable unit over which to estimate predictions for processing has been hard to find. My claim is that new insights into language processing require models that appeal to novel discourse-level cues such as those that guide the establishment of intersentential coherence.

The approach presented here represents a departure from many previous approaches to these topics. Ambiguity resolution, both coreferential and structural, has often been cast as a problem to be solved when the situation arises rather than as a result of expectations: the presence of an ambiguous pronoun prompts a search for a salient referent (Gordon & Searce, 1995; Kaiser, 2003); a new word requires the selection of an appropriate attachment site within the syntactic structure (Frazier,

1978). The results presented here suggest that an understanding of pragmatic effects in linguistic *expectancy* can change the way we model ambiguity resolution. Instead of resolving ambiguity as it arises, comprehenders are seen as savvy anticipators. They generate forward-looking expectations about where the discourse is going, and these expectations about discourse direction in turn yield expectations about phenomena such as which referent is likely to be mentioned next or what role subsequent clauses will play within the sentence. If a pronoun or a syntactically ambiguous structure is encountered, the comprehender interprets it in part based on those expectations. In this light, the approach taken here is aligned more closely with recent expectation-based probabilistic models that have been proposed to account for the processing of sentence-internal levels of linguistic structure.

The existing work on expectation-based processing centers around the observation that statistical regularities are available at multiple levels of our linguistic input: transition probabilities at the phonetic level (Jusczyk, Luce, & Charles-Luce, 1994; Saffran, Aslin, & Newport, 1996), semantic associations in lexical access (Tabossi, 1988; Kamide, Altmann, & Haywood, 2003; DeLong, Urbach, & Kutas, 2005; Bicknell, Elman, Hare, McRae, & Kutas, 2008), and structural frequencies at the syntactic level (Elman, 1993; Hale, 2001; Levy, 2008). At each of these levels, comprehenders appear to track frequency patterns in their input in order to predict what is likely to occur next in specific contexts. It would be reasonable then to ask whether people are also aware of statistical regularities at the discourse level, tracking information about which types of contexts give rise to which types of discourse continuations and how those continuations affect sentence-internal phenomena.

Furthermore, if pragmatic cues come into play throughout sentence processing and interact with strategies for resolving local ambiguity, then no model that ignores such cues can fully account for these phenomena. Such a result, taken together with the existing evidence for expectation-based processing at other levels of linguistic representation, would endorse a model that combines cues of different representa-

tional granularity, simultaneously conditioning the probability of upcoming material on syntactic-, semantic-, and discourse-level factors.

The importance of combining information from different levels of linguistic representation has already been demonstrated in the interplay between syntactic parsing decisions and lexical semantic biases. To take one example, McRae, Spivey-Knowlton, and Tanenhaus (1998) show that comprehenders manage to avoid garden-path readings in the presence of lexical items that provide a better thematic fit for the non-garden-path reading. In sentences like *The cop arrested by the detective was found guilty*, the verb *arrested* is temporarily ambiguous between a main-verb reading (*The cop arrested the crook*) and a reduced-relative-clause reading (*The cop that was arrested was found guilty*). The plausibility of a cop being the one to arrest someone supports the garden-path main-verb reading. Substituting a more plausible patient for *the cop* makes the non-garden-path reduced-relative reading more likely: *The crook arrested by the detective was found guilty*. This work has been taken to show that lexical semantic cues from verb semantics (e.g., who is a likely agent of an arresting event?) can influence the resolution of temporary syntactic ambiguity (e.g., is *arrested* a main verb or part of a reduced relative clause?). The work described in this dissertation demonstrates a different type of contextual influence: the interplay between intersentential coherence and intrasentential ambiguity.

1.2 Phenomena Addressed in this Work

In order for psycholinguists to test processing models, one common strategy is to create contexts in which comprehenders are confronted with ambiguous constructions. Such experiments shed light both on the domain-specific aspects of processing such constructions as well as on more general mechanisms that comprehenders use throughout sentence comprehension.

A great deal of research has centered around the patterns observed in both

referential and syntactic ambiguity, and the experiments presented in this dissertation serve to clarify our understanding of these phenomena. At a more general level, though, these phenomena also provide a testing ground for the claim that a key part of language processing is the ability to track higher order relationships that hold between sentences. The evidence presented in this dissertation points to the conclusion that ambiguity resolution is sensitive to coherence-driven biases, and this conclusion is useful on two counts. On one hand, it adds to the body of knowledge about ambiguity resolution itself; on the other hand, it also establishes the first psycholinguistic evidence that comprehenders use contextual cues to generate expectations over units like coherence relations. Such a result establishes that the linguistic representations over which comprehenders generate expectations extend beyond sounds, words, and phrases to full sentences and discourse segments.

1.2.1 Pronouns and Coreference

Pronouns and other referring expressions represent prime targets for an investigation into coherence-driven expectations because their presence in discourse compels comprehenders to track material across multiple sentences. In order to follow a discourse, comprehenders must build a mental representation of the events being described. Pronouns and other referring expressions are part of what allow comprehenders to build these representations and track event participants from one sentence to the next.

Previous studies of coreference have generally treated pronouns as triggers for a backward-looking search, a search that uses morphosyntactic and thematic-role heuristics to locate the most likely referent in the preceding context (termed a ‘memory search’ by Clark and Sengal (1979), a process of ‘identifying the antecedent’ by Gordon and Searce (1995) and Almor (1999), or an implementation of ‘retrieval instructions’ according to Kaiser (2003)). The preferred referent for an ambiguous pronoun has been shown to correspond, in certain contexts and to varying degrees,

to the referent that occupies the subject position of the preceding clause (Crawley, Stevenson, & Kleinman, 1990), the referent in a parallel grammatical position (Smyth, 1994; Chambers & Smyth, 1998), the individual who was mentioned first in the preceding context (Gernsbacher & Hargreaves, 1988) or the one who appears in a salient thematic or event-level role (Stevenson, Crawley, & Kleinman, 1994).

Contrary to this retrieval-based approach, the coherence-driven model I am advocating posits that comprehenders generate forward-looking expectations about where the discourse is going. Expectations about who will be mentioned next are then conditioned on the comprehender's probabilistic belief about the coherence relation that is likely to ensue.

Other pronoun interpretation models have been posited which describe predictive mechanisms for tracking the shifting focus of a discourse and estimating which referent is likely to be mentioned next. The coherence-driven model proposed here, though different in many ways, reflects certain aspects of those previous models: namely, a notion of expectancy and a time course based on dynamic updates. Centering Theory (Grosz, Joshi, & Weinstein, 1995), for instance, distinguishes between a 'backward-looking center', which can be interpreted as the topic of the current sentence, and a 'forward-looking center', which can be interpreted as a predictor of which discourse entity is expected to be mentioned next. Centering Theory uses information about entities' structural prominence (e.g., subjecthood) to construct this 'forward-looking' ranked list, but the interpretation of a pronoun is more complicated (Brennan, Friedman, & Pollard, 1987). The coherence-driven model proposed here also posits that language processing is 'forward-looking', but it differs from Centering Theory in that comprehenders' predictions are modeled probabilistically and are presumed to update dynamically as a sentence unfolds, rather than clause-by-clause. Other models that rely on incremental updates include Stevenson, Knott, Oberlander, and McDonald's (2000) dynamic model of semantic focusing, which uses information associated both with entities' structural prominence and their semantic prominence.

The dynamic focusing model posits a series of focus updates as new input becomes available. Under this model, an initial focus on the first mentioned individual gives way to focus on certain thematic roles which in turn may be strengthened or attenuated according to the attention-directing properties of various intersentential connectives. A subsequent pronoun is interpreted to refer to the most highly focused referent at the end of the sequence of focus updates. A final example of an expectation-driven model of coreference is Arnold's (2001) Expectancy Hypothesis. This model posits that pronoun interpretation follows from the accessibility of available referents and that referential accessibility varies as a function of how expected a referent is. Expectancy is calculated probabilistically and dynamically, based on multiple factors (recency, subjecthood, parallelism, thematic roles, etc) and is used as an index of the probability that a particular referent will be mentioned again with a pronoun.

In this dissertation, I revisit several of the results of previous pronoun interpretation models, but I appeal to more general discourse-level biases to explain the patterns observed. I consider whether the factors that have been claimed to contribute to a referent's expectancy are available and weighted equally across different discourse contexts. The idea is to move beyond models that simply list a set of factors that have been shown to matter and instead devise more linguistically motivated models that capture the discourse contexts in which those factors are or are not relevant.

The model that I set forth to account for the coreference results in Chapter 3 is one, like Stevenson et al.'s, that updates incrementally and, like Arnold's, tracks expectations about upcoming coreference. However, I question the direct link that has been posited to hold between salience/accessability/expectancy and pronoun interpretation. Instead, the model, as it is extended in Chapter 4, distinguishes between the primary (coherence-driven and lexical-semantically specified) factors that contribute to comprehenders' expectations about who will be mentioned next and those (structurally defined) factors that primarily determine whether a speaker will choose to use a pronoun or some other referring expression.

The two major coreference datasets that I consider consist of contexts with transfer-of-possession verbs and implicit-causality verbs. Both have received a fair bit of attention in the literature (Garvey & Caramazza, 1974; Caramazza, Grober, Garvey, & Yates, 1977; Au, 1986, McKoon, Greene, & Ratcliff, 1993; Stevenson, Crawley, & Kleinman, 1994; McDonald & MacWhinney, 1995; Stewart, Pickering, & Sanford, 2000; Stevenson, Knott, Oberlander, & McDonald, 2000; Arnold, 2001; Koornneef & van Berkum, 2006), but the patterns remain unclear. Below I give examples of these types of coreference contexts. In both cases, the aim is to determine whether coreferential processing is driven by surface-level cues (grammatical role, thematic role, first mention, recency) or rather is a byproduct of deeper discourse-level processes and representations (causal inference, event structure).

The first dataset consists of transfer-of-possession contexts as in (2). For contexts like these, comprehenders show a bias to interpret the ambiguous pronoun *He* as coreferential with the non-subject referent, Bill, as often as with the subject John. Bill is the individual associated with the end state of the transfer event and is said to occupy the Goal thematic role.

(2) John passed a comic to Bill. He ...

The model I present appeals to deeper discourse-level representations, allowing me to localize this previously reported Goal bias to contexts that denote events with salient end states and furthermore to coherence relations for which event structure is relevant.

The second dataset consists of so-called ‘implicit-causality’ contexts as in (3) and (4). For contexts like these, comprehenders have been reported to show a preference for re-mentioning the causally implicated referent (John in (3) and Bill in (4)).

(3) John infuriated Bill. ...

(4) John scolded Bill. ...

I localize the previously reported implicit-causality bias to Explanation coherence relations for which causal inferencing is relevant. The range of implicit-causality biases turns out to be much more complex once coherence is conditioned on.

For both datasets, I also address the apparent asymmetry that has been reported between pronoun production and pronoun comprehension. As will be shown, the coherence-driven model provides a unified account of diverse results in the pronoun literature and makes correct predictions for data that no prior model accounts for.

1.2.2 Relative Clause Attachment

The field of psycholinguistics has placed considerable emphasis on relative clause processing as a testing ground for a variety of phenomena, including ambiguity resolution (Cuetos & Mitchell, 1988), extraction (Kluender, 1992), optionality (V. S. Ferreira & Dell, 2000), expectation-based parsing (Levy, 2008) and thematic fit (McRae, Spivey-Knowlton, & Tanenhaus, 1998), among others. In this dissertation, I show how relative clause processing can provide answers to a different question: the time course over which comprehenders construct a discourse context and the ways in which coherence-driven biases influence syntactic disambiguation. Of particular interest is the question of whether or not sentence-internal coherence relations exist and, if so, how comprehenders generate expectations about the pragmatic relationship between a matrix and subordinate clause within the same sentence.

The phenomenon I consider is the relative clause attachment ambiguity that arises in sentences like (5).

- (5) Someone shot the servant of the actress who was on the balcony.

First discussed by Cuetos and Mitchell (1988), examples like (5) contain a relative clause (*who was on the balcony*) whose attachment site is ambiguous. The relative clause can be interpreted to modify one or the other of the two noun phrases

in the preceding complex noun phrase (*the servant of the actress*). In English, comprehenders show a default preference in cases like (5) for interpreting the relative clause to attach to the closer noun phrase, *the actress*, rather than the head noun, *the servant*. Certain other languages show the opposite pattern. The dispreference for the nearest attachment site in these other languages and the lack of a universal attachment preference has been problematic for theories of sentence processing that posit crosslinguistic syntactic constraints for the interpretation of words and clauses (Frazier's (1978) principles of Late Closure and Minimal Attachment). One could say more generally that investigations into 'default' linguistic preferences sidestep the deeper questions of what factors contribute to 'default' preferences and how these factors combine to yield the overall surface pattern. The notion of a default preference can have the effect of dismissing the subtlety of interacting factors in favor of an average observed value.

Whereas other analyses focus primarily on morphosyntactic cues to capture the observed patterns in relative clause attachment, the question I ask is whether pragmatic cues driven by coherence expectations can shift attachment preferences. By manipulating the expected contribution that the relative clause will make to the overall discourse structure, I show that the pragmatic relationship between the relative clause and the matrix clause (an *intrasentential* coherence relation) has an impact on relative clause attachment.

These two phenomena, coreferential processing and relative clause attachment, represent two fairly different domains. The discovery that discourse-level biases are relevant to both speaks to the generality of this coherence-driven approach and the importance of incorporating discourse-level information into our models of sentence processing. It also suggests that the success of existing models of expectancy will depend in part on their ability to include this type of information within their probabilistic characterization of language processing.

1.3 Outline of the Dissertation

In this section, I outline the remainder of the thesis. Chapter Two reviews research on prediction and expectancy in language processing, discusses examples of discourse-level biases that have been reported previously, and introduces the types of discourse coherence relations over which coherence-driven expectations will be calculated.

Together, Chapters Three and Four examine coreferential processing and show that a coherence-driven model can account for a variety of results by appealing to general mechanisms of discourse interpretation. The proposed model distinguishes three types of factors: factors that influence expectations about discourse direction (e.g., event structure and implicit causality of the prior context), factors that influence who will be mentioned next (e.g., the operative coherence relation and lexical semantic properties of the verb), and factors that determine whether a speaker will choose to use a pronoun instead of some other referring expression (e.g., subjecthood of the antecedent).

Chapter Three uses four story-continuation experiments to distinguish between two types of models of coreferential processing—those that rely primarily on surface-level heuristics and those that rely on deeper discourse-level biases and inferring. The elicited story continuations reveal expectations that comprehenders have about the direction the discourse is likely to take and about the individual who is most likely to be re-mentioned as the discourse proceeds. The first experiment shows that the way that an event is portrayed (as completed or ongoing) influences comprehenders' expectations about discourse direction which in turn influences pronoun interpretation. Specifically, transfer-of-possession contexts that describe completed events are found to yield continuations that tend to tell what happens next (sample context with underlined continuation and intended pronoun interpretation listed in brackets: *John handed a book to Bob. He took it and thanked John. [He = Bob_{GOAL}]). In these 'what-next' Occasion relations, an ambiguous pronoun prompt, i.e., *He*, is*

more likely to be interpreted to refer to the individual associated with the end state, namely the Goal. On the other hand, contexts that describe ongoing events yield continuations which are more likely to provide Elaborations and Explanations (*John was handing a book to Bob. He wanted Bob to read it.* [He = John_{SOURCE}]). In those continuations, it is the individual associated with the onset of the event, namely the Source, who is more likely to be mentioned next. The results for the passages that describe completed events replicate the Goal bias that has been reported in previous studies (Stevenson, Crawley, & Kleinman, 1994), but this bias disappears in the passages about ongoing events, suggesting that the previously reported Goal bias was a side effect of the prevalence of ‘what-next’ continuations in a stimuli set consisting of passages about completed transfer-of-possession events. The second and third story-continuation experiments confirm a prediction from the coherence-driven model — that a shift in the distribution of coherence relations should yield a shift in the pattern of pronoun interpretation. A final experiment investigates an additional surface/deep issue that emerges in passages with intersentential connectives. The question is whether next-mention preferences are driven by the connective itself or whether the connective simply signals the operative coherence relation. Strong next-mention biases have been found for passages with *because* connectives following so-called ‘implicit causality’ verbs: For story-continuation prompts like *Mary scolded John because ...*, the causally implicated referent John is most likely to be mentioned next. This same pattern emerges in the data presented here without connectives, but it is limited to the subset of the data in which an Explanation can be inferred to hold (*Mary scolded John. He cheated at Scrabble.*). The pattern disappears in continuations that embody other coherence relations (e.g., Elaborations favor the other referent, Mary: *Mary scolded John. She yelled at him in front of everyone.*). This final experiment localizes another widely reported result to a subset of a larger dataset, and again, the coherence-driven model provides a more general account of the results.

Chapter Four addresses the bidirectional nature of the dependency between coreference and coherence. Four story-continuation experiments explore the idea that

information about coreference can influence coherence, and the results complement the results from the preceding chapter showing that information about coherence can influence coreference. The chapter builds on the transfer-of-possession and implicit-causality results from Chapter Three, as well as on additional evidence that referents who were originally mentioned in subject position are more likely to be mentioned again with a pronoun. A generative model is introduced that combines the biases regarding coherence, next-mention, and pronominalization as a causal chain of influence: Contextual cues influence expectations about upcoming coherence relations; the operative coherence relation influences expectations about which referent will be mentioned next; and the grammatical role of that referent influences the probability that a subsequent mention of that individual will be pronominalized. Based on this model, a comprehender who encounters a pronoun, even an ambiguous one, can use that information to estimate the probability of a particular coherence relation. For example, in a sentence like *Bob scolded John. He _____*, the sentence-initial pronoun allows the comprehender to reverse engineer the coherence relation that likely generated that surface observation by inferring backwards up the chain: A reference to the subject Bob is more likely to have been pronominalized than a reference to the non-subject John, and Bob is a more likely referent for an Elaboration relation than an Explanation, even though Explanations are quite frequent following implicit-causality contexts. In other words, the mere presence of a pronoun, even a fully ambiguous one, can change comprehenders' expectations about where the discourse is going. These bidirectional influences also demonstrate the incremental nature of coherence and coreference processing: Comprehenders update their expectations about upcoming material using contextual cues such as the presence of a pronoun as that information becomes available.

Chapter Five explores the incrementality of coherence-driven biases in a different domain—syntactic processing—showing that expectations about upcoming discourse continuations can influence the resolution of local structural ambiguity. An off-line sentence-completion experiment and an on-line self-paced reading-time exper-

iment examine comprehenders' expectations for high/low relative clause attachments following implicit-causality and non-implicit-causality verbs (*John detests/babysits the children of the musician who...*). In the completion study, the widely reported low-attachment preference for English is observed in the non-implicit-causality condition, but this preference gives way to more high attachments in the implicit causality condition in cases in which (i) the verb's causally implicated referent occupies the high-attachment position and (ii) the relative clause provides an explanation for the event described by the matrix clause (e.g., *...who are arrogant and rude*). In the reading-time study, a similar preference for high attachment emerges in the implicit causality context, whereas the low attachment preference is consistent elsewhere. These results suggest that comprehenders construct discourse contexts dynamically during sentence processing, using available coherence-driven biases mid-sentence to generate expectations about the structural analysis of the remainder of the sentence.

Chapter Six considers whether the expectation-driven model of processing advocated in the previous chapters extends beyond coherence relations in single-speaker monologues to another domain, the domain of questions in dialogs. Using story and dialog continuations, I show that the contextual cues that bias participants in favor of certain coherence relations in monologue story continuations also bias them in favor of certain corresponding questions in dialog continuations — for example, contexts with implicit-causality verbs yield more Explanation coherence relations in single-speaker story continuations (SPEAKER1: *Mary scolded John. He cheated at Scrabble again.*), and they similarly yield more *Why*-type questions in dialog continuations (SPEAKER1: *Mary scolded John. SPEAKER2: Did he get in trouble because of his cheating?*). In this light, participants appear to be answering in a monologue context the questions that they are likely to pose in a dialog context, suggesting that a link can be drawn between models of discourse structure that center around inventories of intersentential coherence relations and models which posit that discourses are structured through a series of implicit and explicit questions.

The final chapter reviews the experimental evidence described in this dissertation and discusses some directions for future research. The fact that phenomena ranging from coreferential processing to syntactic disambiguation appear to be sensitive to coherence-driven biases suggests that these biases may be more pervasive than has been previously acknowledged. One would therefore expect that other linguistic constructions with simultaneous roles in the discourse structure and the syntactic or referential structure would be sensitive to coherence-driven biases. For future work, I consider contexts in which coherence relations are marked overtly, and I also discuss the implications for bridging models of discourse coherence and questions under discussion.

2

Incorporating Discourse into Models of Processing

This chapter situates an expectation-based coherence-driven model within the general landscape of previous work on discourse coherence and expectation-based processing. First, I review work on prediction and expectancy in language processing. This work shows that comprehenders do appear to track statistical regularities over sounds, words, and syntactic structures, but this previous work has not considered expectancy over higher order units which provide linguistic structure at the discourse level. The next section describes work showing that discourse factors do impact processing in a variety of ways. However, previous analyses of discourse factors have focused on the role of referential context and causal reasoning: who has been mentioned in the previous discourse and in the context of what other types of causal information. None of this work has provided a unit over which to calculate discourse-level expectancy. This dissertation aims to build an expectation-driven model of sentence and discourse processing by quantifying discourse-level predictions as probabilities over intersentential coherence relations. With that goal in mind, I close the chapter by describing the model of discourse coherence that I will rely on throughout the following chapters.

2.1 Predictability in Processing

Psycholinguistic research has long sought to capture the role of expectancy generation in language processing. The research has centered around comprehenders' ability to attend to statistical regularities in their input because learning statistical regularities is one of the factors that allows comprehenders to make predictions. In what follows, I point to (a subset of) the research showing that people are sensitive to regularities at the phonetic, word, and syntactic levels of linguistic input. This work left open questions regarding regularities at the discourse level and the role of expectancy in discourse processing.

As early as their first year, infants appear to be sensitive to distributional information in their input and are able to distinguish between predictable and unpredictable sequences of sounds and syllables. Jusczyk, Luce, and Charles-Luce (1994) have shown that infants prefer to listen to frequent phonotactic structures from their native language. Jusczyk et al. exposed infants to nonsense words containing legal phoneme sequences that were either frequent or infrequent in English. Infants distinguished between the two types of 'words', showing a preference for words that contained more frequent phonotactic properties (e.g., the sequence /ds/ can end but not begin syllables in English). Work by Saffran, Aslin, and Newport (1996) has further shown that infants can learn statistical regularities in an artificial language. Infants were exposed to a continuous sequence of syllables comprising a set of nonsense words (*bidakupadotigolabubidaku...*). Within a 'word' (e.g., *bidaku*), the probability of one syllable given the previous syllable was quite high ($p(da|bi)$), but at word boundaries, the predictability of the next syllable was low ($p(pa|ku)$). Infants distinguished between test items that adhered to or violated the word-internal conditional probabilities, showing that they were able to attend to and learn context-dependent frequencies in their input.

At the word level, comprehenders show sensitivity to context-driven contingencies and semantic associations (see Altmann (1998) for a review of earlier work).

Priming tasks reveal that verbs are read and pronounced more quickly when preceded by a noun that denotes a typical filler of one of the verb’s thematic roles (McRae, Hare, Elman, & Ferretti, 2005): the verb *praying*, for example, is read faster when preceded by *nun* than by *sniper*, suggesting that a verb’s agent can lead to increased activation of the verb itself. Eye-tracking studies show that a combination of an agent and a verb lead to expectations regarding likely patients (Kamide, Altmann, & Haywood, 2003). Kamide et al. show that comprehenders make anticipatory looks to images that correspond to likely next words before the words are actually uttered: In a visual world paradigm, comprehenders made more looks to a motorcycle following the words *The man will ride* than following *The girl will ride*, showing that they were generating predictions about plausible patients based on the context of an agent and a verb. Lastly, measures of brain ERPs (Event Related Potentials) show that comprehenders respond differently to semantically bleached items such as determiners in contexts in which the determiner does not match the noun that is expected to follow, as in *The day was breezy and the boy went outside to fly a/an...* (DeLong et al., 2005). Only if comprehenders are making predictions about the upcoming noun would encountering the determiner ‘an’ be a more surprising event than encountering the determiner ‘a’.

In keeping with the results from behavioral studies, computational models also make effective use of statistical regularities in language input in order to predict the next word or the next syntactic category in a sentence. The regularities in raw input allow a connectionist network to generate expectancies about the words that are likely to appear (Elman, 1990; Elman, 1993); regularities in syntactic structure allow a model built around a probabilistic context free grammar to predict the probability of upcoming words (Hale, 2001; Levy, 2008). Hale and Levy have both shown that a word in a clause may be unexpected because it requires a syntactic structure containing low-probability phrase-structure rules. Surprisal in this case is measured in information theoretic terms: a word is minimally surprising when it *must* appear (no new information, given an entirely deterministic context); a word is maximally sur-

prising when it is impossible given the preceding context (thereby providing a lot of new information about the contexts that license that word). Hale and Levy’s calculations of which syntactic outcomes generate increased surprisal correspond to reported points of difficulty in reading time tasks. For the well-known garden-path example *The horse raced past the barn fell*, the processing difficulty on ‘fell’ can be attributed to the low probability associated with the syntactic parse in which *raced past the barn* is a reduced relative clause modifying *the horse* and the verb *fell* acts as the main verb. Psycholinguistic results suggest that people make predictions about upcoming events so that likely words and structures are anticipated whenever possible, thereby minimizing the overall surprisal.¹

This prior work on predictability has focused on syntactic dependencies in tree structures and semantic contingencies between verbal arguments. The studies have not addressed contingencies between full clauses and sentences. The question I will be addressing in the following chapters is whether people encode similar probabilities about the direction that the discourse is likely to take based on a combination of the patterns they observe in their input, their own real-world knowledge, and their prior experience with the structure of an unfolding discourse.

2.2 The Role of Discourse Context

There is evidence that comprehenders do avail themselves of information beyond the bounds of the current sentence when they are processing sentence-internal words and phrases. As noted earlier, the term “discourse factor” has been used to refer to any contextual information that operates intersententially rather than intrasententially. The factors that have been studied include, among others, information about who has already been mentioned (unique/non-unique, given/new) as well as what event-level information has already been provided (causes/consequences). These fac-

¹The terms ‘surprisal’ and ‘prediction’ here are not meant to imply conscious anticipation on the part of the comprehender; rather, the unconscious pre-activation of certain words and constructions could allow for easier processing and integration.

tors have been shown to affect the processing of relative clauses (van Berkum, Brown, & Hagoort, 1999, Roland, O'Meara, Yun, & Mauner, 2007) and the choice of utterance type in speakers' story continuations, respectively (Simner & Pickering, 2005).

Work by van Berkum et al. shows that the status of a referent in the preceding discourse context can influence the interpretation of the word *that*, a word whose part of speech is ambiguous between a complementizer interpretation and a relative pronoun interpretation (as well as a determiner). Based on the observation that the primary function of a relative clause is usually to modify or restrict a referent (e.g., *the girl that was wearing a red hat*), van Berkum et al. conducted an experiment to show that relative clauses are processed more easily in contexts that contain two possible referents that might need to be distinguished. Comprehenders' brain potentials were measured as they read sentences containing a syntactic ambiguity between a complement clause and a relative clause (*David told the girl that...*). The ambiguous clause was presented in two different contexts: one containing only one possible referent for *the girl* and one in which two referents were available. The presence of the definite determiner *the* signals that the noun phrase *the girl* should be referentially unique in the preceding context. In a context with more than one referent for *the girl*, comprehenders were predicted to interpret the ambiguous *that* as marking a relative clause, one that would further restrict the identity of the referent, rather than as a complement clause, which would leave the referent of *the girl* as non-unique. The results from their ERP study showed that a target sentence that violated the expectations of the previous discourse context elicited a P600/SPS effect, an index of integration difficulty.

Discourse sensitivity in relative clauses has also been demonstrated by Roland et al. (2007). They found that the status of a referent as given or new in a discourse impacts the processing of object relative clauses. Object relative clauses in isolation are processed more slowly than subject relatives, but in a context in which the subject of the relative clause is given, the difference is neutralized. For example, the gap in

the relative clause in *The lady that the banker visited enjoyed the dinner very much* was processed with less difficulty after a context sentence that mentioned *the banker* than after a neutral context. The effect of discourse factors on relative clauses will be taken up again in Chapter 5. The experiments presented in that chapter go beyond the assumption that the function of a relative clause is to simply modify or restrict the identity of a referent. A different discourse function is considered, whereby a relative clause contributes propositional content to the overall discourse and the information it conveys enters into an intrasentential coherence relation with the matrix clause.

Beyond relative clauses, there is some evidence that discourse context also affects the probability of subsequent discourse continuations. Work by Simner and Pickering (2005) on the salience of causes and consequences in story continuations shows that an event that is preceded by information about its cause generates more continuations that describe possible consequences to the event; conversely, an event that is preceded by information about its consequence generates more continuations about the cause. This ‘fill in the gap’ strategy suggests that comprehenders keep track of utterance types in the discourse context and that they have preferences regarding upcoming utterance types (namely, that information which has already been provided need not be reiterated). Simner and Pickering leave open the question of what types of factors affect the selection of one of the remaining utterance types. The experiments in Chapter 3 identify factors related to event structure and causal biases that affect this selection.

The following section introduces the inventory of utterance types that will be used throughout the rest of this dissertation. Based on a model of discourse coherence, the coherence relations described below provide a set of utterance types over which comprehenders can be said to generate expectations.

2.3 Intersentential Coherence Relations

Current models of discourse aim to capture the relationships that are inferred to hold between adjacent discourse segments and that serve to make a discourse locally and globally coherent (Mann & Thompson, 1988; Webber & Joshi, 1998; Hobbs, 1990; Marcu, 2000; Kehler, 2002; Webber, 2006; Asher & Lascarides, 2003; with reviews in Knott, 1996 and Hutchinson, 2005). Most theories of discourse coherence take as their goal the identification and description of a set of licit discourse moves or relevant contributions that an utterance can make. The diversity of approaches and the variation in assumptions, terminology, and applications is an indication of the complexity of discourse analysis and the lack of consensus in the field. What the various models have in common is the assumption that discourse can be characterized as a complex structured linguistic representation. For the following chapters, I adopt Kehler's coherence model, which offers a concise inventory of philosophically motivated discourse coherence relations. Kehler's model is appealing because it draws on principles of general cognitive reasoning and logical inferencing that are not exclusive to language, and it applies those principles to the structure of a discourse. In what follows, I introduce Kehler's model and then I give definitions and examples for six of Kehler's coherence relations that I will use, with some adaptations, in all of the experiments described in this dissertation. Different properties of the coherence relations will become relevant in the experimental studies. In Chapter 5, I extend these definitions to intrasentential relationships (the relationship that is inferred to hold between a relative clause and matrix clause), and in Chapter 6, I further extend them to interrogatives (the relationship that can be inferred to hold between an utterance and an elicited question).

Kehler's model defines coherence relations as relationships that are inferred to hold between the propositions expressed by two sentences. The causal reasoning used in identifying causes and consequences (as in Simner & Pickering, 2005) represents just one of the inferential processes that comprehenders use to establish relationships

between sentences. The model establishes three key ways in which ideas can relate to each other, relying on propositional logic in order to define these relations. The system of relations is closely tied to that proposed by Hobbs (1990) but is organized into three classes: Cause-Effect, Resemblance, and Contiguity. These three categories were originally proposed by Hume (1748) as the three general ways in which ideas can be related.

The first three coherence relations that I present here, Explanation, Result, and Violated-Expectation, all fall within the *Cause-Effect* category originally posited by Hume. Explanations are characterized by the causal information that one sentence provides about the content in the preceding sentence. They may appear with the connective *because* as in (6).

EXPLANATION: infer propositions P and Q from the assertion of sentences S_1 and S_2 , where normally $Q \rightarrow P$.

- (6) Clinton and Obama both campaigned in Philadelphia because the Pennsylvania primary was the following Tuesday.

Understanding the passage in (6) requires both the real-world knowledge that Philadelphia is a city in Pennsylvania as well as the inference that Clinton and Obama were likely campaigning there for the presidential primary. To later learn that Clinton and Obama actually appeared in Philadelphia in order to support the local mayoral election would violate the inference that made the Explanation relation go through.

Result relations, sometimes appearing with connectives such as *therefore*, *so*, or *as a result*, are the counterpart to Explanations. They require that the comprehender draw a similar chain of cause-effect inferences to connect two eventualities, where, in this case, the cause appears in the first sentence and the effect in the second.

RESULT: infer proposition Q and P from the assertion of sentences S_1 and S_2 , where normally $Q \rightarrow P$.

- (7) Obama was criticized for being elitist, and therefore he started bowling with his constituents.

The passage in (7) is only coherent if the comprehender draws the necessary inference that going bowling, a pastime that is not typically considered elitist, might help Obama counter the criticism.

A last *Cause-Effect* relation, Violated-Expectation, is used to contrast an actual outcome with a desired or expected outcome. Violated Expectations may appear with connectives like *but*, *however*, or *nonetheless*.

VIOLATED-EXPECTATION: infer propositions P and Q from the assertion of sentence S_1 and S_2 , where normally $P \rightarrow \neg Q$.

- (8) Clinton was trailing Obama in the polls, the popular vote, and the number of pledged delegates and superdelegates. Nonetheless, she decided to stay in the race.

In the non-causal domain, a *Resemblance* relation like Parallel is used to describe a set of common or contrasting properties that hold between the entities a_1, a_2, \dots in one sentence and the entities b_1, b_2, \dots in a second sentence. Parallel relations are sometimes marked on the surface with connectives like *similarly* or *in contrast*.

PARALLEL: infer a common or contrasting relation $p(a_1, a_2, \dots)$ from sentence S_1 for entities a_i and infer a relation $p(b_1, b_2, \dots)$ from S_2 for b_i where some property q_i holds of a_i and b_i for all i .

- (9) Obama was endorsed by Governor Bill Richardson, and Clinton received support from Senator Dianne Feinstein.

Another *Resemblance* relation, Elaboration, is the limiting case of the Parallel relation, whereby the entities a_1, a_2, \dots in one sentence are identical to the entities b_1, b_2, \dots in the second sentence. Elaboration relations are generally restatements from a different perspective, requiring that the entities mentioned not only be similar, as in Parallel relations, but identical.

ELABORATION: Infer a relation $p(a_1, a_2, \dots)$ from the assertions of sentences S_1 and S_2 .

- (10) Clinton has been campaigning in the mid-Atlantic region with her family all week. Yesterday she appeared at rallies in Pennsylvania with Bill and Chelsea.

Lastly, the Occasion relation, adapted from Hobbs (1990), is the only relation in the *Contiguity* category. It is used to describe a sequence of events and can appear with connectives like *and then* or *next*. Establishing Occasion relations requires that a state of affairs be inferred as a point of connection between two eventualities, i.e., that the final state of the first eventuality is inferred to be the initial state of the second.

OCCASION: infer a change of state for a system of entities in sentence S_2 , establishing the initial state for this system from the final state of sentence S_1 .

- (11) Obama flew from North Carolina to Pennsylvania. Later that day he gave a speech in Philadelphia about race relations.

The coherence of (11) requires that the end state of the event described in the first sentence (Obama landing in Pennsylvania) occur prior to the start state of the event in the second sentence (giving a speech in Philadelphia).


In the following chapters, I will return to these definitions because the logical inferencing associated with different coherence relations will be shown to be relevant to the phenomena being discussed. In each chapter, I will manipulate properties of the discourse context in an attempt to shift the probability of different types of coherence relations, and the properties described in these definitions will be the basis of the context manipulations. In Chapter 3, properties such as start states and end states in Occasion relations will become important, as well as the role of real-world knowledge in establishing cause-effect relations like Explanations. The coherence relations introduced here all act over propositions, but I will show in Chapter 5 that those propositions need not be independent clauses but may instead appear in the matrix and embedded clauses of the same sentence. As discussed at the beginning of this section, there are many different theories of discourse coherence, and there are also

other discourse models that rely on underlying ‘questions under discussion’ instead of underlying coherence relations. A link between these questions-under-discussion models and coherence will be drawn in Chapter 6. To refer back to these coherence relations, especially in figures, I will sometimes use the abbreviations *Elab*, *Exp*, *Occ*, *Par*, *Res*, and *V-E*, for Elaboration, Explanation, Occasion, Parallel, Result, and Violated-Expectation, respectively. These abbreviations can be found in the List of Abbreviations.

3

Coreference: Event Structure & Causal Inference

Experimental work on coreferential processing has left open questions as to the underlying mechanisms that guide pronoun interpretation strategies. The question of whether pronouns are interpreted based primarily on surface-level linguistic cues (grammatical role, thematic role, first mention, recency, parallelism) or as a byproduct of deeper discourse-level processes and representations (causal inference, event structure) remains unresolved in the literature. These two views come together in a story-continuation experiment by Stevenson, Crawley, and Kleinman (1994; see also Arnold, 2001) testing the influence of thematic roles on pronoun interpretation. One of the most striking results comes from passages like (12) with a transfer-of-possession context sentence followed by an ambiguous pronoun prompt.



The context sentence in these passages contains two possible referents for the pronoun, one that appears in subject position and fills the Source thematic role, and

one that appears as the object of a sentence-final prepositional phrase and fills the Goal thematic role. Stevenson et al. found that participants' continuations corresponded to a Goal interpretation for the pronoun almost as often as to a Source interpretation (henceforth 'Goal continuations' and 'Source continuations').

This result is unexpected in light of a variety of existing models of pronoun interpretation including: the subject preference (Crawley, Stevenson, & Kleinman, 1990), which predicts that an ambiguous pronoun will be coreferential with the subject of the previous clause, in this case the Source; the first-mention privilege (Gernsbacher & Hargreaves, 1988), which also predicts that the individual referred to as the subject Source will be the most salient referent; and lastly the grammatical parallelism preference (Smyth, 1994; Chambers & Smyth, 1998), which predicts that an ambiguous subject pronoun is resolved preferentially to a subject antecedent, again the Source for examples like (12).

To account for this and other pronoun interpretation results, Stevenson et al. posit a semantic focusing model in which comprehenders construct a mental model of the events being described and then locate the most highly focused individual in that representation when they encounter a pronoun. Under the semantic focusing model, comprehenders are posited to assign initial focus to the individual that has been mentioned in subject position; they then update their focus as the sentence continues, responding to cues in the sentence that either reinforce or diminish the salience of particular referents. In sentences like (12), the focus on the subject referent (a subject assignment strategy) together with the salience of the Goal (a thematic role bias) are posited to yield the roughly 50/50 distribution of Source/Goal continuations. Stevenson et al. offer two explanations for the unexpected salience of the non-subject Goal in examples like (12). The first is a thematic-role-level preference which amounts to a heuristic ranking Goals above Sources. The second explanation is an event-level bias for focusing on the end state of the previously described event. In this latter case, Stevenson et al. posit a bias to interpret a pronoun as co-referential with the

entity associated with the consequence of an event, where the Goal is considered to be more salient than the Source with respect to the end state in transfer-of-possession events.

This chapter aims to test whether the forces underlying coreferential processing reflect surface-level heuristics (i.e., the thematic role bias) or deeper event-level and discourse-level biases. Following a review in the next section of several existing coreference models and their predictions for data like (12), I present results from four passage completion experiments which address the surface/deep distinction. The first experiment aims to distinguish between the thematic-role and event-structure hypotheses. A more general approach based on discourse coherence is then presented, and that approach is tested in the second and third experiments. The last experiment investigates a further surface/deep distinction that emerges in passages with intersentential connectives.

3.1 Interpreting Ambiguous Pronouns

The Source-Goal sentences exemplified by example (12) were part of the larger study conducted by Stevenson et al. which was designed to explore the strategies that comprehenders use in coreferential processing. In the condition of interest, Stevenson et al. compared participants' interpretation of ambiguous pronouns following a variety of types of context sentences, as in (13).

- (13) a. John passed the comic book to Bill. He _____. [Source-Goal]
 b. Bill seized the comic book from John. He _____. [Goal-Source]
 c. Joseph hit Patrick. He _____. [Agent-Patient]
 d. Patrick was hit by Joseph. He _____. [Patient-Agent]
 e. Ken admired Geoff. He _____. [Experiencer-Stimulus]
 f. Geoff impressed Ken. He _____. [Stimulus-Experiencer]

The two context sentences in each pair describe similar events but use verb frames which vary the position of the two referents within the sentence. This manipulation allowed Stevenson et al. to assess the effects of grammatical role (subject/non-subject position) and thematic role (participant type within the evoked event) on comprehenders' pronoun interpretation preferences. As per Stevenson et al.'s definitions of the relevant thematic roles, the context sentences in (13a) and (13b) describe events with a Source referent, the individual away from which something moves, and a Goal referent, the individual towards which something moves. The Source appears in subject position and the Goal in the object of a prepositional phrase in (13a); these positions are reversed (13b). Sentences (13c) and (13d) describe events that contain an Agent referent, the instigator of an action, and a Patient referent, the individual affected by an action. The Agent is the subject in (13c) and appears as the object of a by-phrase in (13d); the Patient is the object in (13c) and then is promoted to be the subject in the passive construction in (13d). Sentences (13e) and (13f) describe states of affairs with an Experiencer referent, the individual having a given experience, and a Stimulus referent, the individual giving rise to a certain experience. The Experiencer is the subject and the Stimulus the object in (13e); these positions are reversed in (13f).

For each verb frame, Stevenson et al. assessed the number of continuations that corresponded to a first-mention interpretation of the pronoun. The experiment was used to probe participants' *interpretation* preferences, even though the task itself involves *producing* a continuation. As has been noted by Arnold (2001), writing a story continuation necessarily involves both comprehension and production: Participants must form a mental representation of the event described in the continuation prompt, up to and including the pronoun, before they produce their continuation. The elicited continuations are thus taken to reflect participants' expectations about where the story is likely to go next given the preceding context sentence and their interpretation of the ambiguous pronoun.

For half of the items, the pronoun interpretation results were compatible with a grammatical-role preference: Pronouns in passages like (13b,d,f) were more likely to be interpreted as coreferential with the subject of the preceding sentence than with the non-subject. Sample continuations are shown in (14), with continuations that correspond to the preferred pronoun interpretation highlighted in gray.

(14) Stevenson et al.'s items that exhibited a subject preference

b. John seized the comic book from Bill. He ...

... opened it and read the first page. [SUBJ]

... tried to grab it back. [NON-SUBJ]

d. Patrick was hit by Joseph. He ...

... promptly burst into tears. [SUBJ]

... apologized profusely for hitting his friend. [NON-SUBJ]

f. Ken impressed Geoff. He ...

... could sing like Pavarotti. [SUBJ]

... was easily impressed. [NON-SUBJ]

The pronoun interpretation preferences exemplified in (14) are in keeping with several earlier models that posit a subject bias or a first-mention privilege (Gernsbacher & Hargreaves, 1988; Crawley, Stevenson, & Kleinman, 1990). The preferences also are compatible with models of pronoun interpretation based primarily on grammatical parallelism. Under such models, a subject pronoun is resolved preferentially to a subject antecedent, and an object pronoun to an object antecedent (Chambers & Smyth, 1998; Smyth, 1994).

However, models relying primarily on morphosyntactic cues like subjecthood and structural parallelism fail to capture the coreference patterns in the other half of Stevenson et al.'s data. For context sentences like (13a,c,e), participants' continuations corresponded to a non-subject interpretation for the pronoun just as often as, or more often than, a subject interpretation. Sample continuations are shown in (15).

(15) Stevenson et al.'s items that did not exhibit a subject preference

a. Bill passed the comic book to John. He ...

... *told John to read the first page.* [SUBJ]

... *opened it and read the first page.* [NON-SUBJ]

c. Joseph hit Patrick.

... *then apologized profusely.* [SUBJ]

... *promptly burst into tears.* [NON-SUBJ]

e. Geoff admired Ken. He ...

... *was easily impressed.* [SUBJ]

... *could sing like Pavarotti.* [NON-SUBJ]

From the pattern of results in (14) and (15), Stevenson et al. concluded that an adequate model of pronoun interpretation must incorporate a set of thematic role biases alongside a subject assignment preference. Across the board, they found that the referents occupying Goal, Patient, and Stimulus roles were as likely, if not more likely, to be the topic of the continuation than referents occupying the Source, Agent, and Experiencer roles. When the preferred thematic role is located in subject position, the outcome appears to be a subject preference, but this pattern is reduced or even reversed in contexts in which the preferred thematic role is located in the non-subject position. Stevenson et al.'s semantic focusing model retains the subject assignment strategy alongside the thematic role bias because they found that the bias to the preferred thematic role was even stronger in passages in which the preferred referent appeared in subject position. For example, there were more Goal continuations following context sentences with Goal-Source verbs like *seize* as in (13b) (84.6% Goal continuations) than there were following context sentences with Source-Goal verbs like *pass* in (13a) (49.0% Goal continuations).

Stevenson et al. further consider whether the thematic-role biases can in turn be understood in terms of more general focusing strategies that arise from the structure of events and from causal inferences. For sentences that describe events (e.g., Source-Goal: *John passed the comic book to Bill*), their semantic focusing model

posits that the first-mentioned referent receives initial focus in the comprehender’s mental representation of the event, but this initial focus gives way to a focus on the consequence of the event as the representation is updated. The role of the end state is based on Moens and Steedman’s (1988) characterization of events as comprised of a preparatory process, a culmination, and a consequent state—a set of components that are represented in a comprehender’s mental model. Under the semantic focusing model, a subsequent connective can reinforce the end-state focus (e.g., *as a result*, *so*) or it can shift the focus back to the cause (e.g., *because*). Their proposal for sentences that describe ongoing states (e.g., Experiencer-Stimulus: *Geoff admired Ken*) is slightly different: The tripartite event structure does not apply; there is no consequence and hence no individual associated with the end state that can be focused for event-structural reasons. Instead, sentences that evoke states of affairs are said to create a bias to focus on the individual associated with the cause, with respect to which the Stimulus referent is considered to be more salient than the Experiencer.

An important question that Stevenson et al.’s model raises is whether factors like event structure and causality directly impact coreference or whether these factors are relevant for general discourse comprehension and only indirectly affect coreference. In this chapter, I present results showing that coreferential processing does indeed reflect event-structural biases and causal inference, but that these mechanisms can be better understood as side effects of deeper discourse-level processes, mediated by the operative intersentential coherence relation. Experiments I-III address the effects of event structure on pronoun interpretation, and Experiment IV addresses the effects of causal inference and intersentential connectives on comprehenders’ expectations about which referent will be mentioned next.

3.2 Experiment I: Event Structure

To recap, Stevenson et al. found that Source-Goal transfer-of-possession context sentences yielded almost as many Goal continuations (49%) as Source continu-

ations (51%). The Goal bias that they report is compatible with two explanations: a surface-level thematic role preference and a deeper event-structure bias. Experiment I aims to distinguish between these two explanations and to further investigate whether the Goal preference can receive deeper motivation from mechanisms that are used in establishing discourse coherence.

In order to manipulate event structure while holding thematic roles constant, minimal pairs were constructed with transfer-of-possession passages by varying only the verbal aspect of the Source-Goal verb, as in (16) and (17).

(16) COMPLETED EVENT (PERFECTIVE)

John_{SOURCE} passed a book to Bob_{GOAL}. He _____.

(17) INCOMPLETE EVENT (IMPERFECTIVE)

John_{SOURCE} was passing a book to Bob_{GOAL}. He _____.

The positions of the Source and Goal thematic roles are the same in (16) and (17), but the perfective verb in (16) is used to describe a completed event which is compatible with end-state focus, whereas the imperfective verb in (17) is used to describe an event as an ongoing process, rendering it incompatible with end-state focus (Moens & Steedman, 1988). The thematic role preference would predict a Goal bias in continuations elicited in both (16) and (17), whereas the event structure hypothesis predicts more Goal continuations for (16) than (17), because the Source is presumably still salient to an ongoing transfer-of-possession event.

The idea that salience within the end state affects pronoun interpretation suggests additional lexical semantic factors that could be examined. Verbs like *pass* and *throw* describe events with a co-located Source and Goal, meaning that the verb evokes an event in which the event participants are likely to share physical proximity. These events differ from those with an estranged Goal referent (ex. *send*, *mail*) where the Goal is not likely to be physically present with the Source. Similarly, verbs whose lexical semantics incorporate a default sense of successful transfer (ex. *pass*, *deliver*) differ from those that have no guarantee of successful transfer (ex. *throw*, *send*). The

event-structure hypothesis could be taken to predict that context sentences which describe events that lack a co-located Source and Goal or those that have a diminished sense of guaranteed transfer will yield fewer Goal continuations. The thematic role preference does not distinguish between different classes of transfer verbs, predicting a similar proportion of Goal continuations across the board.

3.2.1 Methods

All of the experiments described in this chapter employ story-continuation tasks. In these tasks, participants are presented with a context sentence and asked to write a natural continuation at the prompt provided. The elicited continuations for the first three experiments are interpreted as a measure of participants' expectations about where the discourse is likely to go next, given the event described in the context sentence, their interpretation of the pronoun, and their prior experience with the way that discourses naturally progress. These first three experiments mirror Stevenson et al.'s pronoun condition in which story continuations were elicited with an ambiguous pronoun prompt.

Participants

Forty-eight monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

The twenty-one experimental items consisted of a transfer-of-possession context sentence followed by an ambiguous pronoun prompt. The Source referent always appeared in subject position, and the Goal was the object of a *to*-phrase. All verbs described physical transfer events (e.g., *pass*, *throw*). Verbs were excluded that described abstract or conceptual transfer (e.g., *show*, *teach*) because they lacked a clear end state, even though they have been used in prior work on transfer of possession (Arnold, 2001). The actual stimuli from Stevenson et al.'s (1994) experiment are

not available for replication purposes, but the twenty-one verbs used here appear to be consistent with Stevenson et al.’s eight Source-Goal verbs in all relevant respects. Participants saw each verb only once in either perfective or imperfective form.

The twenty-one verbs in the stimuli were selected because they varied along two dimensions that were hypothesized to be relevant to end-state focus: co-location of event participants and guarantee of successful transfer, as shown in Table 3.1. Verbs in Class 1 are used to describe events in which the Source and Goal are co-located and the default assumption is successful transfer. Verbs in Class 2 lack this sense of guaranteed transfer, and verbs in Class 3 lack both a co-located Source and Goal and the sense of guaranteed transfer. Verbs in Class 2 also differ from those in Classes 1 and 3 in that they can take an *at*-complement (e.g., *throw a ball at someone*). The full stimulus set is given in the Appendix.

Table 3.1: Experiment I transfer-of-possession verb classes

Class 1	[co-located guaranteed transfer]	<i>bring, carry, deliver, give, hand, pass, serve</i>
Class 2	[co-located no guaranteed transfer]	<i>chuck, fling, lob, kick, roll, throw, toss</i>
Class 3	[not co-located no guaranteed transfer]	<i>fax, forward, mail, send, ship, transmit, wire</i>

Twenty-nine filler sentences were included that contained a non-transfer verb in the context sentence. These fillers consisted of transitive and intransitive verbs that appeared with perfective or imperfective aspect. Adverbs, proper names, or gender-unambiguous pronouns served as prompts.

Participants were given booklets containing fifty passages. They were instructed to imagine a natural continuation to each passage, writing the first continuation that came to mind and avoiding humor. They were also instructed to treat each item separately and to not try to tie the different passages together into a longer

story. There was no time limit, but an approximate time frame of 20 minutes was proposed in the instructions to encourage participants to move quickly through the items.

Evaluation and Analysis

Two trained judges, the author of this dissertation and one other graduate student, evaluated the participants' intended pronoun interpretations. They categorized the pronoun in each continuation as either a Source-referring or Goal-referring pronoun, in light of the story context and the elicited continuation. Judges were used instead of the participants themselves in order to avoid revealing the aim of the experiment.¹ To minimize the possible influence of personal pronoun-interpretation biases, the judges were instructed to substitute each referent in place of the pronoun and to consider the plausibility of the scenario. Judges were instructed to be cautious, erring on the side of categorizing a pronoun as ambiguous if the pronoun could be interpreted plausibly as coreferential with either referent. As such, not all responses could be disambiguated; analyses are reported for a conservative evaluation in which a pronoun was treated as ambiguous if at least one coder assessed it to be ambiguous.

Analyses of variance were conducted to measure the effects of two within-subjects factors (verbal aspect and verb class) on the pronoun interpretation outcome (Source/Goal). Verbal aspect varied within items and verb class varied between items. Because the pronoun interpretation outcome involves examining proportions, an arcsine transformation was applied to the percentages of pronoun interpretations before carrying out analyses of variance. All other analyses of variance over binary outcomes that are presented in this chapter use arcsine transformed percentages. For clarity of presentation in this section and throughout the chapter, I present means in the form of raw proportions.

¹The use of judges follows Arnold (2001). In Stevenson et al.'s experiments, on the other hand, participants were instructed to underline their intended referents after completing all the passages. However, Stevenson et al. ultimately relied on judges as well to remedy contradictions in the participants' underlining.

3.2.2 Results

From the set of all elicited continuations, 11.6% were excluded due to ambiguity, leaving a dataset of 891 continuations.

Verbal Aspect

As predicted by the event structure hypothesis, context sentences with perfective verbs yielded more Goal continuations (42.9%) than context sentences with imperfective verbs (20.9%; effect of verbal aspect: $F_1(1,47)=52.854$, $p<0.001$; $F_2(1,20)=30.079$, $p<0.001$). All means represent subject means, unless otherwise noted.² Figure 3.1 shows the effect of aspect on pronoun interpretation.

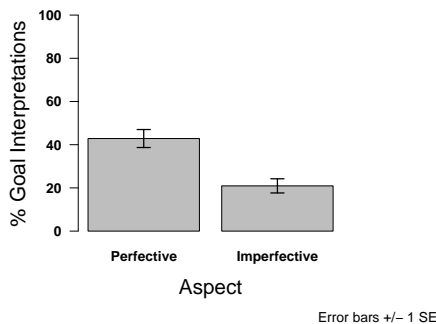


Figure 3.1: Effects of verbal aspect on pronoun interpretation in Experiment I

Verb Class

Also as predicted by the event structure hypothesis, pronoun interpretation differed across the three classes of verbs in the context sentences, as depicted in Figure 3.2. Context sentences with Class 3 verbs, those which describe events in which the Goal referent is not presumed to be co-located with the Source nor is there a sense of guaranteed transfer, yielded the fewest overall Goal interpretations (22.5%) compared to Class 1 (31.2%) and Class 2 (42.4%; effect of verb class on interpretation: $F_1(2,94)=14.973$, $p<0.001$; $F_2(2,18)=4.800$, $p<0.03$). There was also an aspect \times verb

²The results of this experiment were previously reported in Rohde, Kehler, and Elman (2006). However, the means presented here for the Goal bias differ slightly from those reported in Rohde et al. because the previously published means represented raw means (# Goal continuations/# total continuations) over the set of all continuations including ambiguous cases.

class interaction: Verbal aspect had a stronger effect in contexts sentences with Class 1 and Class 2 verbs than Class 3 verbs (interaction significant only by subjects: $F_1(2,94)=3.898$, $p<0.03$; $F_2<1$). As Figure 3.2 shows, all three classes showed a bias towards the Source following imperfective context sentences (Class 1: 17.0% Goal continuations; Class 2: 28.1%; Class 3: 18.2%), but interpretation patterns differed across verb classes following perfective context sentences. Class 2 verbs show the roughly 50/50 distribution following perfective context sentences (56.3% Goal interpretations); Class 1 verbs yielded fewer Goal continuations (41.7%), and Goal continuations were least frequent following Class 3 verbs (29.2%).

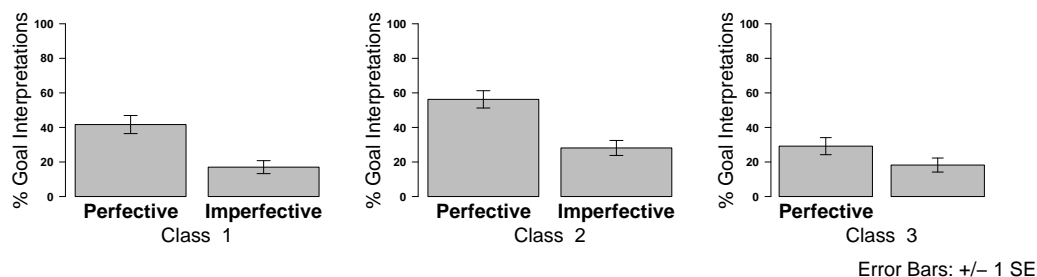


Figure 3.2: Effects of verb class and verbal aspect on pronoun interpretation in Experiment I

3.2.3 Discussion of Event Structure Results

The drop in the proportion of Goal continuations following imperfective context sentences is in keeping with the predictions of the event structure hypothesis, but not the thematic role preference: More Goal continuations emerge in environments that are compatible with end-state focus. The results for the perfective condition replicate Stevenson et al.'s results: The proportion of Goal continuations was roughly equal to the proportion of Source continuations. Whereas their items confounded thematic-role and event-structure biases, the verbal aspect manipulation described here shows that the bias to the Goal is reduced in imperfective contexts and that this previously reported bias can be recast as a side effect of event structure.

The different pronoun interpretation patterns across verb classes further support the notion that properties of the event, including those that are expressed by the lexical semantics of the verb, influence pronoun interpretation. For Class 3 verbs which describe events with non-co-located participants, the perfective condition revealed a strong Source preference not seen in the other two classes. In keeping with the event structure hypothesis, it appears that the lack of a co-located Goal diminishes the probability that the discourse will continue with references to the Goal even if the event is described as completed, whereas the non-subject Goal rivals the subject Source for coreference if the sentence describes a completed event with co-located event participants. The sense of guaranteed transfer did not appear to make a difference since the proportion of Goal continuations was actually higher for Class 2 verbs, which lack a sense of guaranteed transfer, than Class 1 verbs, which do guarantee successful transfer. The fact that the proportion of Goal continuations varies by verb class, however, suggests that properties of the event being described matter for coreference. As such, an account that relies on surface cues like thematic roles fails to predict this variation because the thematic roles were held constant across all items.

If the previously observed thematic role preference can be recast as a side effect of a deeper event structure bias, a further question is raised concerning whether or not the event structure bias itself can be understood as a side effect of another property of the passage. The next section considers this question, testing whether the event structure bias is merely epiphenomenal, a consequence of more general mechanisms used in establishing discourse coherence.

3.2.4 Coherence Analysis

Kehler has argued that a variety of heuristic pronoun interpretation strategies that have been proposed in the psycholinguistic literature are largely side effects of the process of establishing intersentential coherence relations. If the event structure bias is a side effect of coherence establishment, then the Goal bias is predicted to

be restricted to coherence relations that directly reference event-level information. Of the six coherence relations laid out in Section 2.3, the only relation that directly encodes information about the end state is the Occasion relation. The definition is repeated here:

OCCASION: infer a change of state for a system of entities in sentence S_2 , establishing the initial state for this system from the final state of sentence S_1 .

To see how coherence relations such as Occasions affect interpretation, consider the following ambiguous continuation that one participant in the experiment produced.

(18) Miriam sent a fruitcake to Rachel. She didn't want to eat it.

Both judges categorized this example as ambiguous, but closer inspection reveals that this ambiguity stems from the relationship that is inferred to hold between the two clauses. If the second clause is interpreted as an Explanation of the first — the ‘diswanting’ was the cause of the sending — then *She* will most likely be interpreted to refer to the Source, Miriam. If the second clause is taken to occasion the first — that is, the diswanting was connected to the sending but happened subsequently to it — *She* will most likely be interpreted to refer to the Goal, Rachel. Other coherence relations allow for different interpretations, but the point is that the interpretation of the pronoun varies with the intersentential coherence relation.

Just as the thematic role bias could be cast as a side effect of event structure, the question raised by examples like (18) is whether the event structure bias may be restricted to Occasion relations. If that is the case, one would expect to see stronger evidence for a Goal preference in contexts in which two clauses are related by an Occasion relation than in contexts in which other relations are operative.

To test this, the judges annotated all unambiguous responses for the coherence relation that held between the context sentence and the elicited continuation. The

coherence relations were annotated using the six relations laid out in Section 2.3, and judges resolved disagreements through discussion, following Stevenson et al. (1994). T-tests were conducted over raw means to measure the pronoun-interpretation bias for each coherence relation, as compared to a hypothetical mean of 0.5.

The coherence definitions laid out in Section 2.3 were applied here, with examples give in Table 3.2. Coherence relations are abbreviated *Elab*, *Exp*, *Occ*, *Par*, *Res*, and *V-E* for Elaboration, Explanation, Occasion, Parallel, Result, and Violated-Expectation, respectively. A more detailed description of the annotation guidelines is provided in the Appendix.

Table 3.2: Examples of coherence-annotated passages from Experiment I

<i>Relation</i>	<i>Transfer-of-possession example</i>
Exp	Matt passed a sandwich to David. He didn't want David to starve.
Res	Matt passed a sandwich to David. He said thanks.
V-E	Matt passed a sandwich to David. He wanted it for himself though.
Par	Matt passed a sandwich to David. He passed him an apple too.
Elab	Matt passed a sandwich to David. He did so carefully.
Occ	Matt passed a sandwich to David. He ate it up.

Pronoun Interpretation Biases across Coherence Relations

Figure 3.3 shows the interpretation breakdown by coherence for perfective passages — the passages in which the event structure bias emerged. Two distinct sets of biases are shown in Figure 3.3: First, the height of each bar signifies the percentage of continuations in which a particular coherence relation was inferred to hold (e.g., 37.2% of all continuations embodied Occasion relations); second, the shading within each bar signifies the pronoun interpretation bias associated with that coherence relation, where the dark gray represents the percentage of Goal continuations and the light gray represents the percentage of Source continuations. As predicted, Occasion relations following perfective context sentences were dominated by references to the Goal (78.0% Goal bias, significantly different from a hypothetical mean

of 0.5: $t_1(45)=5.354$, $p<0.001$; $t_2(20)=7.264$, $p<0.001$).




Figure 3.3: Pronoun interpretation across coherence relations for perfective passages in Experiment I

As Figure 3.3 also shows, Elaborations, Explanations, and Violated Expectations differed from Occasions in that they showed a preference for the Source (*Elab*: 3.2% Goal bias, $t_1(42)=19.66$, $p<0.001$; $t_2(19)=69.729$, $p<0.001$; *Exp*: 16.2% Goal bias, $t_1(30)=6.498$, $p<0.001$; $t_2(19)=9.112$, $p<0.001$; *V-E*: 25.8% Goal bias, marginal by items: $t_1(30)=3.165$, $p<0.004$; $t_2(12)=1.993$, $p<0.07$). There was also a Goal preference for Result relations (*Res*: 87.5% Goal bias: $t_1(15)=4.392$, $p<0.001$; $t_2(10)=9.238$, $p<0.001$), but the small set of Result continuations ($N=24$) was very homogeneous, more than half consisting of the form *X transfers Y to Z. Z thanks X.*³ There were too few Parallel cases ($N=8$) to support any conclusions.

Though all of the perfective passages summarized in Figure 3.3 describe a completed event with a salient endpoint, it is the coherence relation that dictates whether that endpoint is relevant. Occasion relations show a clear bias to the Goal, and they are precisely the relations that encode information about the end state. Recall that pronoun interpretation, when collapsed across coherence relations, showed a distribution that was much closer to 50/50. When the results are broken down by

³In any case, whereas Kehler's model of coherence would predict that causal inference plays a greater role in establishing Result relations than Occasion relations, the result component in a cause-effect sequence often arises from the end state of the first eventuality. It would perhaps not be surprising if the result in a cause-effect sequence exhibited a similar statistical bias toward focusing on the end state of the cause.

coherence, there is no evidence of a 50/50 split. Instead, each coherence relation carries its own stronger pronoun interpretation bias. This suggests that a different distribution of coherence relations could give rise to a different pattern of pronoun interpretation and that the previously reported Goal biases may have arisen due to a large number of Occasion relations elicited with a perfective-only stimuli set.

Effect of Verbal Aspect on Coherence Distribution

In considering what factors contribute to the distribution of coherence relations, notice in Figure 3.3 that the most frequent coherence relation following perfective context sentences is the Occasion relation. One interpretation of this is that the presence of a salient end state in the perfective context sentences licenses more Occasion relations. In contrast, the absence of a salient end state in the imperfective context sentences may give rise to a very different distribution of coherence relations. To test this prediction, analyses of variance were conducted to measure the effect of one within-subjects/within-items factor (verbal aspect) on the choice of coherence relation. For the analysis, coherence relations were coded as six binary outcomes: Explanation-or-not, Elaboration-or-not, etc., with an arcsine transformation applied to the percentages of each binary outcome.

As predicted, there were more Occasion relations following perfective context sentences (37.2%) than imperfectives (13.9%; $F_1(1,47)=85.082$, $p<0.001$; $F_2(1,20)=55.451$, $p<0.001$). In contrast, there were more Elaborations following imperfective context sentences (36.9%) than perfective context sentences (28.5%; $F_1(1,47)=14.17$, $p<0.001$; $F_2(1,20)=7.085$, $p<0.02$). There were also more Violated-Expectations following imperfectives (29.0%) than perfectives (8.4%; $F_1(1,47)=62.058$, $p<0.001$; $F_2(1,20)=30.818$, $p<0.001$). There was no significant difference between the proportion of Explanations following perfectives (18.5%) and imperfectives (16.4%; $F_s<1$), and Results and Parallel were very rare across the board. The distribution of coherence relations following perfective and imperfective context sentences can be seen by comparing the relative percentages of each coherence relation

in Figures 3.3 and 3.4.

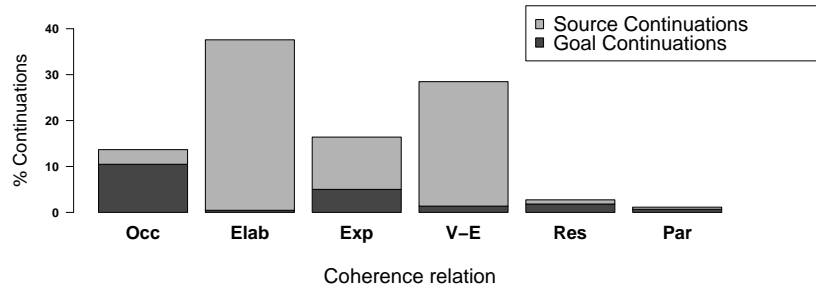


Figure 3.4: Pronoun interpretation across coherence relations for imperfective passages in Experiment I

Figure 3.4 also shows the pronoun interpretation breakdown for each coherence relation in imperfective passages — the passages in which the Goal bias was eliminated. Comparing Figure 3.4 to Figure 3.3, one can see that the pronoun interpretation biases observed across different coherence relations in the imperfective passages are quite consistent with those in the perfective passages: Following imperfective context sentences, Occasions were biased to the Goal whereas Elaborations, Explanations, and Violated-Expectations were biased to the Source (*Occ*: 70.7% Goal bias, significant only by subjects: $t_1(28)=2.72$, $p<0.02$; $t_2(19)=1.656$, $p<0.12$; *Elab*: 2.6% Goal bias, $t_1(47)=22.182$, $p<0.001$; $t_2(20)=53.851$, $p<0.001$; *Exp*: 36.0% Goal bias, marginal by subjects: $t_1(32)=1.837$, $p<0.08$; $t_2(18)=5.779$, $p<0.001$; *V-E*: 4.5% Goal bias, $t_1(42)=21.342$, $p<0.001$; $t_2(17)=22.303$, $p<0.001$; there were again very few Results ($N=12$) and Parallel relations ($N=5$)). To confirm that the pronoun interpretation bias associated with each coherence relation is independent of verbal aspect, one can test whether aspect is a significant factor in predicting the pattern of pronoun interpretation within a particular coherence relation (e.g., is the 78.0% Goal bias following perfective context sentences significantly different from the 70.7% Goal bias following imperfective context sentences?). The results show that, across all coherence types, there is no significant effect of aspect on the proportion of Goal

continuations (*Occ*: $F_s < 1$; *Elab*: $F_s < 1$; *Exp*: $F_1(1,26)=1.033$, $p < 0.32$; $F_2 < 1$; *V-E*: $F_1(1,27)=3.32$, $p < 0.08$; $F_2(1,12)=3.11$, $p < 0.11$; *Res*: $F_1(1,5)=1.0$, $p < 0.37$; *Par*: $F_s < 1$). The factor responsible for the pronoun interpretation differences between the perfective and imperfective conditions is thus the differing distributions of coherence relations.

3.2.5 Discussion of Coherence Results

Experiment I was designed to distinguish between two types of explanations that had been posited by Stevenson et al. for the salience of the Goal in transfer-of-possession contexts: a thematic role heuristic and a deeper event structure bias. The results of the verbal aspect and verb class manipulations offered support for the event structure hypothesis by showing that the Goal bias disappears in passages that describe events that lack a salient end state or in which successful transfer is not guaranteed. By breaking down the results by coherence, one can see that the bias to the end state is further limited to coherence relations are established by making inferences over start states and end states (i.e., Occasion and Result relations).

These results can be accounted for under a model that crucially separates the probability of a particular coherence relation from the probability of a particular pattern of pronoun interpretation, conditioned on coherence. The intuition is that the pattern of pronoun interpretation that is observed on the surface can be derived from two underlying biases about where the discourse is going and which referent is likely to be mentioned next, given the possible discourse directions. Following perfective context sentences, Occasion relations were quite frequent (37.2% Occasions) and were dominated by references to the Goal (78.0% Goal continuations), whereas following imperfective context sentences, Elaborations and Violated-Expectations were most common (28.5% and 29.0% of continuations, respectively), and these relations were dominated by references to the Source (97.4% and 95.5% Source continuations, respectively).

Equation (3.1) can be used to capture these biases and to calculate the probability that a pronoun will be interpreted as coreferential with the Goal ($p(Int=Goal)$) by summing over the set of possible coherence relations and multiplying the probability of each coherence relation ($p(CR)$) by the probability of a Goal interpretation given that coherence relation ($p(Int=Goal | CR)$). This type of model, one that conditions the probability of certain coherence relations on available contextual cues and that then conditions linguistic outcomes on the operative coherence relation, is what I will be referring to as a coherence-driven model throughout this dissertation.⁴

$$p(Int=Goal) = \sum_{CR} p(Int | CR) \times p(CR) \quad (3.1)$$

$$\begin{aligned} &= p(Int=Goal | Elab)p(Elab) + p(Int=Goal | Exp)p(Exp) \\ &\quad + p(Int=Goal | Occ)p(Occ) + p(Int=Goal | Par)p(Par) \\ &\quad + p(Int=Goal | Res)p(Res) + p(Int=Goal | V-E)p(V-E) \end{aligned} \quad (3.2)$$

The aim of presenting equation (3.1) here is to highlight the two distinct probabilities that contribute to the probability that a pronoun will be interpreted as co-referential with the Goal. The equation itself is true by definition if the two probabilities $p(Int | CR)$ and $p(CR)$ are conditionally independent of one another. Of course, both probabilities are conditioned on context, but the claim is that, for some terms, context may be more important than for others. This claim fits the result that $p(CR)$ reflects contextual manipulations like the verb aspect alternation, whereas $p(Int | CR)$ was consistent regardless of the aspect manipulation.

In Experiment I, perfective context sentences had a higher probability of an upcoming Occasion relation ($p(CR=Occ)$) than imperfective context sentences.

⁴Equation (3.1) makes the assumption that it is pronoun *interpretation* that is conditioned on coherence ($p(Int=Goal | CR)$). An alternative is that the probability of a Goal referent being re-mentioned, regardless of referring expression, is what is conditioned on coherence ($p(next\ mention=Goal | CR)$), and that the interpretation of an ambiguous pronoun is sensitive to other factors as well. Experiments II and III focus only on the effects of coherence on pronoun interpretation, whereas Experiment IV considers choice of next mention. The experiments in Chapter 4 will specifically address the interaction between the choice of next mention and the choice of referring expression.

From this, the large number of Goal continuations following perfective sentences can be attributed to a combination of the high probability of an Occasion relation ($p(CR=Occ)$) and the high probability of a Goal interpretation given an Occasion ($p(Int=Goal | CR=Occ)$). Following imperfectives, however, the preference for Occasions decreased even though the conditional probability of a Goal given an Occasion relation stayed high. The dispreference for Occasions following imperfective context sentences coincided with an increased preference for several Source-biased relations, yielding a larger number of Source continuations than Goal continuations overall in the imperfective condition.

3.2.6 Comparison with Previous Work

Arnold (2001) reports results from both a transfer-of-possession story-continuation experiment with perfective-only stimuli and a corpus analysis of naturally occurring transfer-of-possession passages. Her experiments replicate Stevenson et al.'s thematic role results, confirming the bias towards Goal continuations over Source continuations. To account for these results, she posits an Expectancy Hypothesis that uses probabilistic constraints to capture comprehenders' expectations about coreference. Under this model, comprehenders generate expectations about which referent is most likely to be mentioned next as a discourse continues, and basing their expectations on the frequency of particular references in their input. Her corpus and behavioral results are compatible with an event structure semantic focusing account, but she suggests that the salience of the Goal can be cast primarily in terms of frequency preferences at the thematic role level.

Arnold also investigates the relationship between the context sentence and the continuation. Instead of relying on connectives, Arnold coded participants' responses for the relationship that could be inferred to hold between the elicited continuation and the prior story context, distinguishing continuations that discussed the cause of the preceding event, the endpoint of the previous event, or some other aspect of the

event. In the paper, these categories were characterized with intuitive examples from participant responses. The categories were roughly equally distributed (36% *endpoint*, 25% *cause*, 38% *other*), but the proportion of Goal and Source continuations differed across these three types: Continuations that discussed the *endpoint* were more likely to re-mention the Goal, whereas continuations about the *cause* were more likely to re-mention the Source. Arnold admits that “the data do not reveal why [comprehenders] focus on one type of continuation rather than another (Arnold, (2001), p. 147). The results reported here for Experiment I show stronger biases when the continuations are annotated using a larger set of coherence relations, and the results also reveal that verbal aspect is one of the factors that appears to shift comprehenders’ focus towards particular coherence relations.

In comparing the coherence-driven model presented here with both Stevenson et al.’s semantic focusing model and Arnold’s Expectancy Hypothesis, the common themes that emerge are the notions of incremental updates to participants’ perceptions of referent salience and the generation of expectations about upcoming material based on contextual cues. Under the semantic focusing model, information is incorporated as it becomes available and contributes to comprehenders’ interpretation of a pronoun when and if a pronoun is encountered. The coherence-driven model also relies on incremental updates, but what is being updated differs. Whereas the semantic focusing model posits that comprehenders use available cues to update the salience of specific individuals within their mental model, the coherence-driven model posits that comprehenders use these same cues to update more general expectations about the direction the discourse is likely to take.

The Expectancy Hypothesis is appealing because it specifically models comprehenders’ *expectations* about upcoming material. It conditions the probability of next mention on factors such as grammatical role and thematic role. Like the semantic focusing model, though, the Expectancy Hypothesis posits that comprehenders specifically monitor referent salience, using cues in the context to update their

coreference-specific expectations. The coherence-driven model, on the other hand, does not require a domain-specific mechanism, and instead posits that contextual cues contribute to more general mechanisms of establishing discourse coherence and to predictions about the direction the discourse is likely to take. The probability of next mention is then conditioned on discourse direction, based on a comprehender’s prior experience with different coherence relations and the patterns of coreference observed for those relations. For the Goal bias, contexts that are compatible with end-state focus need not lead comprehenders to directly expect a subsequent mention of the Goal. Rather, the relationship between end-state salience and pronoun interpretation is more indirect: A salient end state makes certain coherence relations more likely—namely those relations that are established by linking the end state of one event to the start state of the next—and the coreference biases of those types of coherence relations then affect comprehenders’ expectations about which referent will be mentioned again as the discourse continues.

Given this alternative indirect mapping between event structure and coreference, a prediction emerges that a factor affecting only the distribution of coherence relations, but otherwise leaving the event structure untouched, could similarly induce a shift in the distribution of pronoun interpretation. In other words, within Equation (3.1), the term $p(CR)$ makes no reference to pronouns, yet changing that probability could have a strong impact on the overall percentage of Goal resolutions. The next two sections consider sentence-internal and sentence-external factors that influence the distribution of coherence relations, i.e., the $p(CR)$.

3.3 Experiment II: Object Type

In an attempt to shift the distribution of coherence relations in participant responses, the stimuli from Experiment I were modified to include different types of objects of transfer, as in (19) and (20), with the prediction being that certain objects might affect how participants choose to continue the passage.

(19) NORMAL TRANSFER OBJECT

John passed a book to Bob. He _____.

(20) ABNORMAL TRANSFER OBJECT

John passed a bloody meat cleaver to Bob. He _____.

Each normal object was paired with a bizarre, surprising, gross, or taboo object—an irrelevant manipulation for any surface-cue theory. Across all pairs of target items in this experiment, the two referents appear in the same grammatical position, they occupy the same thematic roles, and the event structure remains unchanged. The only difference is the ‘normality’ of the object of transfer. The prediction is that in contexts with abnormal objects, participants will be more likely to use the continuation to provide an explanation of the event in the context sentence. This prediction is based on the assumption that sentences which describe surprising events (i.e., events that are infrequent or implausible) require further clarification about the context that gave rise to such an unlikely situation. Given the results from Experiment I showing that Explanations are biased to the Source referent, an increase in the proportion of Explanation relations is predicted to yield more Source interpretations for the ambiguous pronoun following abnormal objects than normal objects.

3.3.1 Methods

Participants

Sixty-nine monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

The target and filler items were identical to Experiment I, except that each of the twenty-one transfer verbs was paired with a normal and abnormal object. The target items consisted of a transfer-of-possession context sentence followed by an ambiguous pronoun prompt. Participants saw each transfer verb only once, paired

with either a normal or abnormal object, in either perfective or imperfective aspect. The task was the same as in Experiment I. The full stimulus set is given in the Appendix.

Evaluation and Analysis

As in Experiment I, two judges, the author of this dissertation and one other graduate student (who was not a judge for Experiment I), assessed the pronoun interpretation and the coherence relation for each passage.

Analyses of variance were used to measure the effects of two within-subjects factors, object type and verbal aspect, on both pronoun interpretation and on the choice of coherence relation. Object type and verbal aspect varied within subjects and within items. Again, an arcsine transformation was applied to the percentages of pronoun interpretations before conducting analyses of variance. For this experiment and the remaining experiments in this chapter, the coherence relations were annotated using the six coherence relations laid out in Section 2.3. For the analysis, the coherence relations were coded as six binary outcomes with analyses of variance conducted on arcsine transformed percentages of each binary outcome. I present means in the form of raw proportions. For this experiment and the next, t-tests were also conducted over raw means of interpretation biases to measure the Goal bias for each coherence relation, as compared to a hypothetical mean of 0.5.

3.3.2 Results

The results reflect a conservative analysis in which a continuation was excluded if at least one judge assessed it as ambiguous. Out of all 1449 continuations, 23.5% were excluded due to ambiguity, leaving a dataset of 1109 continuations.

Replication of Previous Coherence-Driven Biases

The results replicate the original coherence biases and pronoun interpretation biases observed in Experiment I. Figures 3.5a and 3.5b show the coherence breakdown

for passages with perfective and imperfective objects, respectively. The figures also show the pronoun interpretation bias for each coherence relation.

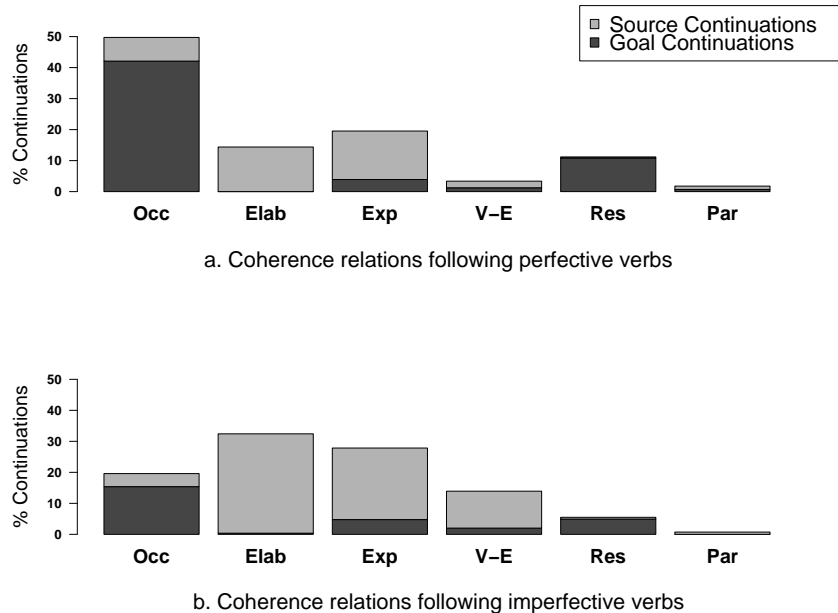


Figure 3.5: Coherence relations and pronoun interpretation biases for perfective and imperfective passages in Experiment II

As before, perfective and imperfective context sentences yielded different proportions of Occasions, Elaborations, and Violated-Expectations. Perfective context sentences yielded more Occasions (49.6%) than imperfectives did (19.7%; $F_1(1,67)=106.807$, $p<0.001$; $F_2(1,20)=62.563$, $p<0.001$), whereas Elaborations and Violated Expectations were more common following imperfectives than perfectives (*Elab*: 31.5% following imperfective compared to 13.8% following perfective; $F_1(1,67)=51.52$, $p<0.001$; $F_2(1,20)=30.188$, $p<0.001$; *V-E*: 12.6% following imperfective compared to 3.3% following perfective; $F_1(1,67)=43.994$, $p<0.001$; $F_2(1,20)=16.23$, $p<0.001$). For this dataset, there was a significant difference in the proportion of Explanations following perfective and imperfective context sentences (29.0% following imperfective compared to 19.9% following perfective; $F_1(1,67)=7.769$, $p<0.01$; $F_2(1,20)=6.733$, $p<0.02$). Result relations were

found to be more common following perfectives (11.5%) than imperfectives (6.5%; $F_1(1,67)=14.998$, $p<0.001$; $F_2(1,20)=13.213$, $p<0.002$).

As before, pronoun interpretation biases differed by coherence relation. As in Experiment I, Occasion and Result relations were biased to the Goal (*Occ*: 79.2% Goal bias, significantly different from a hypothetical mean of 0.5: $t_1(63)=7.665$, $p<0.001$; $t_2(20)=6.707$, $p<0.001$; *Res*: 93.1%; $t_1(47)=12.051$, $p<0.001$; $t_2(20)=7.787$, $p<0.001$), whereas Elaborations, Explanations, and Violated Expectations were biased to the Source (*Elab*: 1.9% Goal bias; $t_1(61)=29.506$, $p<0.001$; $t_2(20)=65.727$, $p<0.001$; *Exp*: 22.0%; $t_1(59)=6.703$, $p<0.001$; $t_2(20)=9.299$, $p<0.001$; *V-E*: 18.0%; $t_1(45)=6.911$, $p<0.001$; $t_2(20)=4.054$, $p<0.001$), and Parallel relations were quite rare ($N=14$). Given the pronoun interpretation biases for each coherence relation and the different distribution of coherence relations following the two verbal aspects, the overall rate of pronoun interpretation also differed by verbal aspect, as in Experiment I: Goal interpretations were more common following perfective context sentences (58.5%) than imperfectives (29.7%; $F_1(1,67)=75.72$, $p<0.001$; $F_2(1,20)=107.38$, $p<0.001$).

Effect of Object Type on Coherence and Interpretation

As predicted, abnormal objects did yield a significantly larger proportion of Explanation relations (30.0%) compared to normal objects (19.3%; $F_1(1,67)=14.261$, $p<0.001$; $F_2(1,20)=4.486$, $p<0.05$). Figures 3.6a and 3.6b show the coherence breakdown for passages with normal and abnormal objects, respectively. The figures also show the pronoun interpretation bias for each coherence relation.

This increase in Explanation relations following abnormal objects comes at the expense of three other relations: Elaborations (down to 19.9% compared to 25.4% following normal objects: $F_1(1,67)=4.379$, $p<0.05$; $F_2(1,20)=6.254$, $p<0.03$), Occasion relations (down to 31.1% compared to 38.1% following normal objects; significant only by subjects: $F_1(1,67)=8.310$, $p<0.006$; $F_2 < 1$), and Violated Expectations (down to 6.1% compared to 9.9% following normal objects; marginal by items: $F_1(1,67)=8.037$, $p<0.007$; $F_2(1,20)=3.181$, $p<0.09$). The increase in the proportion of Explanations

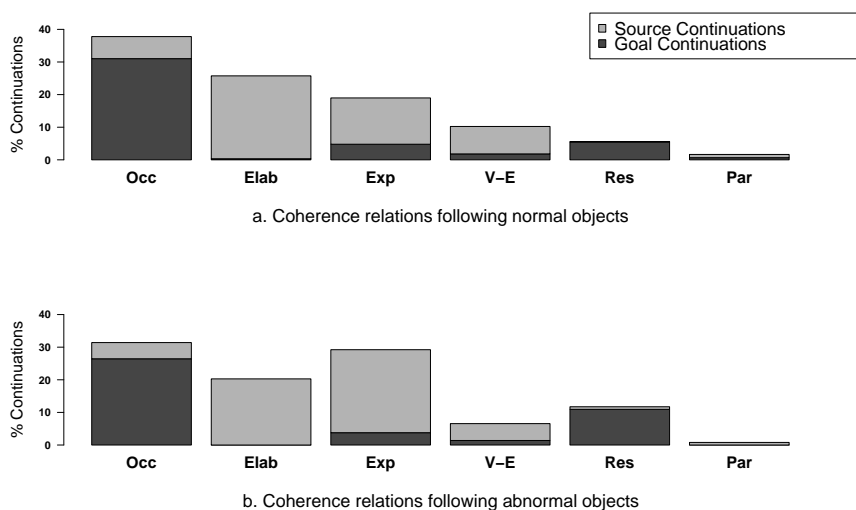


Figure 3.6: Coherence relations and pronoun interpretation biases for passages with normal and abnormal objects in Experiment II

following abnormal objects is accompanied by an increase in the proportion of Result relations (12.4% compared to 5.7% following normal objects; $F_1(1,67)=10.205$, $p<0.003$; $F_2(1,20)=11.948$, $p<0.003$).

Given the coherence distribution and the pronoun interpretation biases for each relation, one can see that the increase in Source-biased Explanation relations following abnormal objects comes at the expense of several other Source-biased relations, and the drop in Goal-biased Occasions is matched by an increase in Goal-biased Results. Although the manipulation succeeded in creating a different distribution of relations, the coherence-driven analysis predicts a canceling effect, and thus no significant difference can be expected in the overall pattern of pronoun interpretation when passages following normal and abnormal objects are compared.

Figure 3.7 shows that this is the case: Object type was not a significant factor for modeling pronoun interpretation (44.5% Goal bias following normal objects compared to 43.7% following abnormal objects; $F_s<1$). There was also no object type \times aspect interaction ($F_s<1$).

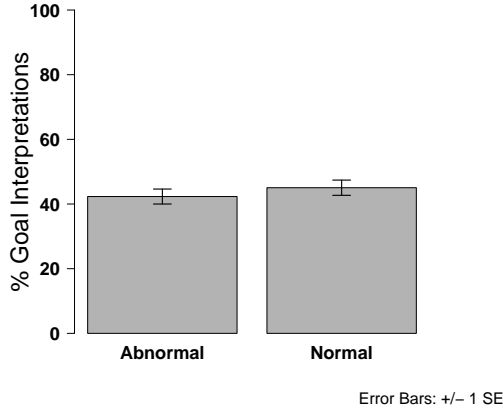


Figure 3.7: (Lack of) effect of object on pronoun interpretation in Experiment II

3.3.3 Discussion

Although the predicted increase in the proportion of Source continuations was not found, the results of Experiment II are in keeping with the model proposed in the previous section. Equation (3.1), repeated here as (3.3), predicts that the distribution of coherence relations in part dictates the pattern of pronoun interpretation.

$$p(Int) = \sum_{CR} p(Int | CR) \times p(CR) \quad (3.3)$$

Had the observed increase in Source-biased Explanations ($p(CR=Exp)$) come only at the expense of Goal-biased relations like Occasions and Results, then the overall pattern of pronoun interpretation would have been predicted to shift towards more Source continuations. What was not anticipated was that the increase in Explanations would be accompanied by a decrease in the proportion of Source-biased Elaborations and Violated-Expectations and an increase in the proportion of Goal-biased Results. The shift in the coherence distribution between passages with normal and abnormal objects had a canceling effect. In this sense, the lack of a difference in the pattern of pronoun interpretation does fit the proposed model, but the effect that abnormal objects had on the distribution of coherence relations differed from the initial predictions. The actual outcome should not have been altogether surprising if

one assumes that events involving abnormal objects have interesting causes *and* interesting consequences. Since the referent associated with the cause is the Source and the referent associated with the consequence is the Goal, increasing the proportion of continuations about both had a negligible effect on the aggregate pattern of pronoun interpretation.

Importantly, the event participant associated with the cause and the consequence remained consistent across conditions in Experiment II. Note that Equation (3.3) depends on the consistency across contexts of the pronoun interpretation bias associated with each coherence relation ($p(Int=Goal | CR)$). As described earlier, the pronoun interpretation biases for each coherence relation were consistent with those observed in Experiment I. If pronoun interpretation biases do stay consistent, then a manipulation that successfully shifts the distribution in favor of only Source-biased relations or only Goal-biased relations should yield the predicted shift in the overall pattern of pronoun interpretation. Experiment III was designed to induce such a shift by using a factor external to the context sentence.

3.4 Experiment III: Instruction Type

The set-up for Experiment III was identical to Experiment I with one exception: different participants were given different instructions. The instructions enforced a bias toward particular coherence relations by giving an explicit question that participants were instructed to answer in their story continuations. Depending which group they were assigned to, participants were instructed to write a continuation that answered either the question ‘Why’ or ‘What happened next’ — the aim being to elicit Explanation and Occasion relations, respectively. Given the interpretation biases from Experiments I and II, the prediction was that these two types of relations would correspond to very different patterns of pronoun interpretation, making instruction type a good predictor of interpretation. The change in instructions is a discourse-level manipulation that is irrelevant to surface-cue theories: The two

potential referents still occupy the same grammatical and thematic roles, and the event structure remains the same.

The materials were taken directly from Experiment I, so, depending on the group they were assigned to, participants wrote either ‘What next’ or ‘Why’ continuations for a set of passages that varied by verbal aspect. The main prediction is that instruction type will shift the distribution of coherence relations, and that the shift in coherence will in turn shift the pattern of pronoun interpretation, even though the stimuli themselves remain identical. For the effects of verbal aspect on the coherence distribution and on pronoun interpretation, the prediction is that these effects will be overwhelmed by the preference for Goal-biased Occasions in the ‘What next’ condition and the preference for Source-biased Explanations in the ‘Why’ condition.

3.4.1 Methods

Participants

Forty-two monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

The twenty-one experimental stimuli and twenty-nine fillers were taken directly from Experiment I, without modification. As before, participants wrote continuations for fifty passages. Depending on the instruction condition, they were asked to imagine a natural continuation that answered either the question ‘Why’ or ‘What happened next’. Participants produced several practice continuations in order to demonstrate that they understood what constituted an answer to the question ‘Why’ or ‘What happened next’.

Evaluation and Analysis

Again, two trained judges (the author of this dissertation and another graduate student who had not been a judge for Experiment I or II) assessed the pronoun

interpretation and the coherence relation for each passage. Analyses of variance were conducted to measure the effects of one within-subjects factor (verbal aspect) and one between-subjects factor (instruction type) on both pronoun interpretation and on the choice of coherence relation. Both factors varied within items.

3.4.2 Results

The results reflect a conservative analysis in which a continuation was excluded if at least one judge assessed it as ambiguous. Out of all 882 continuations, 11.2% were excluded due to ambiguity, leaving a dataset of 783 continuations.

Replication of Previous Coherence-Driven Biases

The results replicate the original pronoun interpretation biases conditioned on coherence that were observed in Experiment I. Again, Occasion and Result relations were biased to the Goal (*Occ*: 69.0%, significantly different from a hypothetical mean of 0.5: $t_1(21)=3.467$, $p<0.003$; $t_2(20)=4.000$, $p<0.001$; *Res*: 87.5%; $t_1(11)=4.18$, $p<0.002$; $t_2(12)=8.124$, $p<0.001$), whereas Elaborations, Explanations, and Violated Expectations were biased to the Source (*Elab*: 0% Goal bias; *Exp*: 18.4% Goal bias; $t_1(26)=6.929$, $p<0.001$; $t_2(20)=9.133$, $p<0.001$; *V-E*: 21.0% Goal bias; $t_1(20)=5.262$, $p<0.001$; $t_2(20)=5.493$, $p<0.001$), and Parallel relations were rare ($N=28$). The probability of each coherence relation and the pattern of pronoun interpretation following context sentences with perfective and imperfective aspect are compared below.

Effect of Instruction Type and Verbal Aspect

As predicted, pronoun interpretation and coherence differed by instruction type, and the pattern of interpretation corresponded directly to the distribution of coherence relations. The effects of instruction type and verbal aspect can be seen by comparing Figures 3.8 and 3.9. Each figure contains three graphs: on the top, the distribution of coherence relations collapsed across perfective and imperfective aspect and on the bottom, the distributions broken down by verbal aspect. The figures also

show the pronoun interpretation biases for each coherence relation, which are highly consistent across conditions.

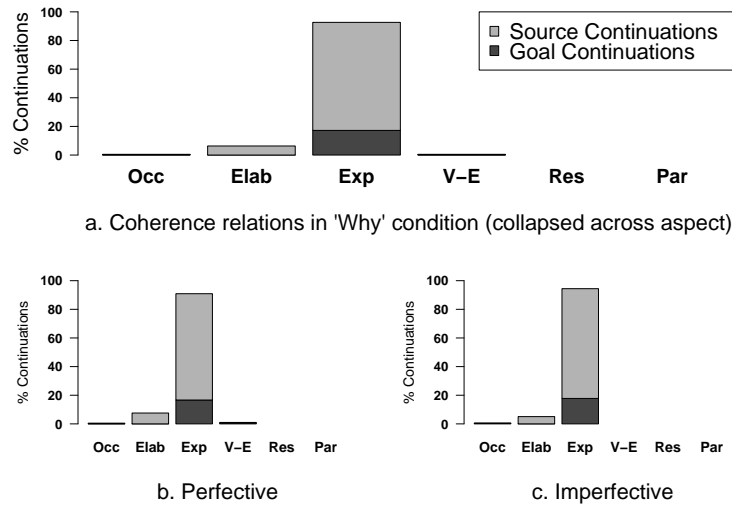


Figure 3.8: Coherence relations following ‘Why’ instructions in Experiment III

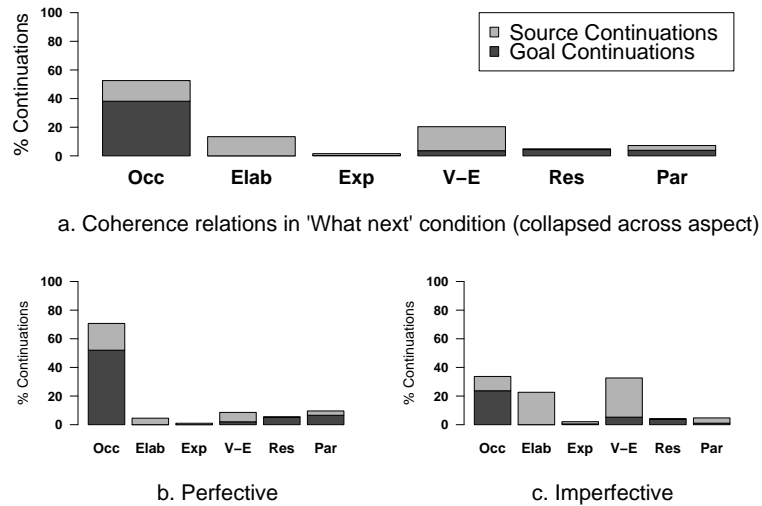


Figure 3.9: Coherence relations following ‘What Next’ instructions in Experiment III

Considering the proportion of Explanation relations, there were, as predicted, significantly more Explanations elicited in the ‘Why’ condition (92.5%) than the ‘What next’ condition (1.5%; $F_1(1,40)=1077.9$, $p<0.001$; $F_2(1,20)=1017.0$, $p<0.001$). The proportion of Explanations also varied marginally by verbal aspect: There

were slightly more Explanation continuations following imperfective context sentences (48.3%) compared to perfective context sentences (45.9%; significant only by items: $F_1(1,40)=2.167$, $p<0.15$; $F_2(1,20)=4.527$, $p<0.05$), a result which is in keeping with the verbal aspect results reported in Experiments I and II. There was no instruction \times aspect interaction on the proportion of Explanation relations ($F_s<1$).

Considering the proportion of Occasion relations, there were, as predicted, significantly more Occasions elicited in the ‘What next’ condition (52.2%) than the ‘Why’ condition (0.5%; $F_1(1,40)=315.23$, $p<0.001$; $F_2(1,20)=263.34$, $p<0.001$). The proportion of Occasions also varied by verbal aspect: As in Experiments I and II, there were more Occasion relations following perfective context sentences (35.5%) compared to imperfective context sentences (17.3%; $F_1(1,40)=71.329$, $p<0.001$; $F_2(1,20)=42.244$, $p<0.001$). There was also an instruction \times aspect interaction: Aspect had a greater effect on the proportion of Occasions in the ‘What next’ condition than the ‘Why’ condition ($F_1(1,40)=72.627$, $p<0.001$; $F_2(1,20)=44.714$, $p<0.001$).

The pattern of pronoun interpretation can be seen in Figure 3.10. There was, as predicted, a significant main effect of instruction type: more Goal interpretations in the ‘What next’ condition (49.8%) compared with the ‘Why’ condition (18.1%; $F_1(1,40)=16.705$, $p<0.001$; $F_2(1,20)=33.213$, $p<0.001$). There was also a significant main effect of aspect, in keeping with the results from Experiment I and II: Goal interpretations were more common following perfective context sentences (41.8%) than imperfective context sentences (26.3%; $F_1(1,40)=19.893$, $p<0.001$; $F_2(1,20)=25.868$, $p<0.001$). There was also an instruction \times aspect interaction: Aspect had a greater effect on pattern of pronoun interpretation in the ‘What next’ condition than the ‘Why’ condition ($F_1(1,40)=25.342$, $p<0.001$; $F_2(1,20)=36.227$, $p<0.001$).

Notice in Figure 3.8 that the bias to Explanations for the ‘Why’ condition is strong in both the perfective and imperfective conditions, which accounts for the low Goal bias in pronoun interpretation for the ‘Why’ condition, regardless of verbal aspect, as shown in Figure 3.10. On the other hand, notice in Figure 3.9 that the

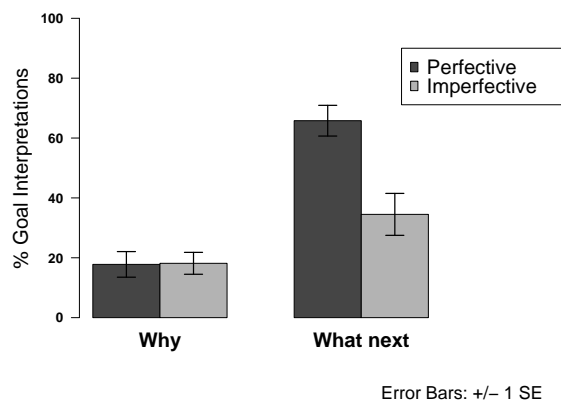


Figure 3.10: Effects of instruction type and verbal aspect on pronoun interpretation in Experiment III

bias to Occasions for the ‘What next’ condition is stronger in the perfective condition than in the imperfective condition, which accounts for the clear Goal bias seen in Figure 3.10 for the ‘What next’ perfective condition and the reduced Goal bias for the ‘What next’ imperfective condition. These patterns are reflected in the presence of the instruction \times aspect interactions for both the the proportion of Occasion relations and the proportion of Goal continuations and the lack of an instruction \times aspect interaction for the proportion of Explanation relations.

3.4.3 Discussion

In accordance with the original predictions, a shift in the distribution of coherence relations yields a shift in the pattern of pronoun interpretation: In contexts in which participants were instructed to write story continuations that answered the question ‘Why’, they wrote more Source continuations; in contexts with instructions to answer the question ‘What next’, they wrote more Goal continuations. This experiment provides strong evidence in favor of a model of pronoun interpretation that incorporates coherence-driven biases. The stimuli for the ‘Why’ and ‘What next’ conditions were identical, and it was only the information that the participants had about the discourse direction that influenced their pronoun interpretation.

The effect of verbal aspect on pronoun interpretation was predicted to be overwhelmed by the instruction-type manipulation. This was the case for in the ‘Why’ condition, but there was nonetheless a significant effect of verbal aspect that was driven by the effects of aspect in the ‘What next’ condition: Goal continuations were more frequent following perfective context sentences compared to imperfective context sentences, but only in the condition in which participants were instructed to answer the question ‘What next’. This difference between the ‘Why’ and ‘What next’ conditions is represented in the instruction type×verbal aspect interaction on the proportion of Occasion relations. The interaction reinforces the claim that contextual cues influence comprehenders’ expectations about the upcoming coherence relation. Whereas the question ‘Why’ consistently yielded Explanations regardless of aspect, the question ‘What happened next’ yielded mostly Occasions following perfective context sentences but a larger variety of relations following imperfective context sentences. The variation in the ‘What next’ imperfective condition may stem in part from the lack of a salient endpoint in events described with imperfective aspect, making Occasion relations less likely and yielding a wider range of alternative coherence relations instead. The prevalence of Violated Expectations may reflect the fact that a transfer-of-possession event described with imperfective aspect fails to establish a natural endpoint for an event which usually has one.

Crucially, the interaction is significant both for the proportion of Occasion relations and for the pattern of pronoun interpretation: In other words, the pattern of pronoun interpretation reflects the underlying distribution of coherence relations, which reinforces the claim that variation in the distribution of coherence relations yields direct effects on the pattern of pronoun interpretation. The question ‘Why’ yielded a large proportion of Source-biased Explanations following both perfective and imperfective context sentences, and that consistency gave rise to the strong Source bias in the ‘Why’ condition. The question ‘What happened next’ yielded a large proportion of Goal-biased Occasions following perfective context sentences, but there was a mix of Goal-biased Occasions and Source-biased Elaborations and Violated-

Expectations following imperfective context sentences. In other words, contexts in which the coherence distribution was more mixed (‘What next’ with imperfective context sentences) are precisely the contexts in which the interpretation pattern is also more mixed.

These subpatterns within the data show that different contextual environments yield different distributions of coherence relations. No model of interpretation that relies only on surface cues can account for these results.

The next section considers contexts in which the intersentential coherence relation need not be inferred because it is marked on the surface with an intersentential connective. The question is whether the coreference biases observed in connective-free contexts can offer an explanation of effects that have been attributed to the connective itself in cases in which the relationship between two clauses is marked overtly.

3.5 Signaling Intersentential Relationships

As was noted in the review of Stevenson et al.’s semantic focusing model in Section 3.1, the presence of a connective has been posited to directly influence which referent is the most highly focused within a comprehender’s mental model of the events being described. Recall that under this model, verbs that describe events (e.g., *pass*, *hit*) are posited to focus the referent associated with the end state of the event, whereas verbs that describe states (e.g., *impress*, *admire*) are posited to focus the referent associated with the cause. These biases in turn are said to be strengthened or weakened depending on the presence of a connective like *because* or *so*. Under such a model, the connective *because* is said to highlight the individual associated with the cause of the event, whereas the connective *so* highlights the individual associated with the consequence or end state. An alternative, coherence-driven explanation of such effects would posit that the connective simply constrains the operative coherence relation and the resulting patterns of coreference correspond to the coreference biases

of the relevant coherence relations.

Evidence for apparent connective-driven biases comes from an experiment that Stevenson et al. conducted in which they elicited story continuations using prompts like those in (21). In evaluating the elicited continuations, Stevenson et al. compared which referent was most likely to be mentioned next depending on the verb and the type of connective in each prompt.

- (21) a. John passed the comic book to Bill because/so _____. [Source-Goal]
 b. Bill seized the comic book from John because/so _____. [Goal-Source]
 c. Joseph hit Patrick because/so _____. [Agent-Patient]
 d. Patrick was hit by Joseph because/so _____. [Patient-Agent]
 e. Ken admired Geoff because/so _____. [Experiencer-Stimulus]
 f. Geoff impressed Ken because/so _____. [Stimulus-Experiencer]

Stevenson et al. found that for the Source-Goal and Goal-Source prompts, the preference to re-mention the Goal was limited to the *so* condition, whereas the next-mention biases were more mixed in the *because* condition. The results for Agent-Patient and Patient-Agent prompts were similar — the preference to re-mention the Patient was also limited to the *so* condition, whereas the next-mention biases were more mixed in the *because* condition. For the Experiencer-Stimulus and Stimulus-Experiencer prompts, however, the Experiencer was consistently preferred in the *so* condition, and the Stimulus was consistently preferred in the *because* condition. Stevenson et al. interpret these results as evidence that the connective interacts with the thematic role biases to bring different referents into focus. As comprehenders read the context sentence, the available referents are said to become more or less focused based on biases from the verb and from the presence of the connective. At the end of the sequence of focus updates, whichever referent is most highly focused is the one that is most likely to be mentioned next.

The analysis presented here offers an alternative interpretation of these results, one that is in line with Hobbs’s (1979) and Kehler’s (2002) discourse coherence models. In this account, the connective simply signals the operative coherence relation: The connective *because* enforces an Explanation relation and thus essentially raises the $p(CR=Exp)$ to 1; the connective *so* signals a Result relation and essentially raises the $p(CR=Res)$ to 1.⁵ As was shown in Equation (3.1), varying the $p(CR)$ can influence the overall pattern of coreference, even though that term makes no reference to pronouns or next mentions. Given the presence of an intersentential connective, the pattern of coreference is then simply a reflection of the coreference bias of the relevant coherence relation on this analysis ($p(Ref | CR)$).

This section and the next focus on the biases that emerge in contexts with *because* connectives in order to determine whether the surface-level focusing properties posited by Stevenson et al. can be better understood in terms of deeper discourse-level effects related to establishing intersentential coherence. The prompts in (21e-f) are of particular interest because examples like these with Stimulus-Experiencer and Experiencer-Stimulus verbs exhibited a clear bias to one referent in the *because* condition. Furthermore, examples like these have received a great deal of attention in the coreference literature, and the role of the connective has led to particular confusion. In Section 3.6, I review the relevant literature on these verbs, which are members of the class of so-called ‘implicit causality’ verbs. Then in Section 3.7, I present Experiment IV which is designed to test whether a coherence-driven model can clarify the effect of connectives in contexts with implicit causality verbs.

⁵Stevenson et al. (1994) note that *so* is actually ambiguous between a result interpretation (as in *where-upon*) and a purpose interpretation (as in *so that*). They return to this ambiguity in Stevenson, Knott, Oberlander, and McDonald (2000) in order to confirm the original results in an experiment by using judges to annotate the elicited continuations for the participants’ interpretation of *so*. They interpret the results as support for the semantic focusing model: The presence of the connective is said to focus entities in the discourse separately from the way it establishes the relationship between events.

3.6 Implicit Causality

Previous researchers have noted that certain verbs, including ones like the Experiencer-Stimulus/Stimulus-Experiencer verbs used in Stevenson et al.'s (1994) experiments, contain within their meaning additional information about certain aspects of the event that they evoke (Chafe, 1972). This information includes, among other things, concepts pertaining to the cause and consequences associated with the event. Chafe points out that a verb like 'burn down' licenses a subsequent mention of concepts such as 'the damage' or 'the insurance'. Furthermore these subsequent mentions can be uttered with definite determiners (*the* damage, *the* insurance), referring expressions that are usually reserved for entities that have already been mentioned in the discourse. The fact that 'the cause' is likewise available for subsequent mention with a definite determiner has been taken as evidence that this verb contains within its meaning an inherent feature of causation that is accessible for subsequent mention without prior introduction.

Garvey and Caramazza (1974) refer to this causal information as part of people's "knowledge of the world", knowledge that comprehenders have and that they share with other discourse participants. This knowledge of the world is what allows comprehenders to interpret the ambiguous pronouns in examples like (22) and (23) (examples adapted from Caramazza, Grober, Garvey, & Yates, 1977) in which the only difference is the verb in the first clause.

(22) Mary infuriated Jane because she had stolen a tennis racket.

(23) Mary scolded Jane because she had stolen a tennis racket.

In (22) and (23), the ambiguous pronoun is preferentially interpreted to refer to the causally implicated individual in the event evoked by the first clause (Mary in (22) and Jane in (23)). Understanding the sentence in (22) requires inferring that an event in which Mary stole a tennis racket could be the cause behind an event in which Jane became infuriated. Likewise, understanding the sentence in (23) requires inferring

that an event in which it was Jane who stole a tennis racket could be the cause of an event in which Mary scolded Jane. This type of inferencing has been shown to affect coreference in disambiguation and reading time tasks using passages like (22-23), as well as in sentence-completion experiments using prompts like (24-25) (Garvey & Caramazza, 1974; Brown & Fish, 1983; Au, 1986; McKoon, Greene, & Ratcliff, 1993; Koornneef & van Berkum, 2006).

(24) Mary infuriated Jane because _____.

(25) Mary scolded Jane because _____.

Sentence completion frames like (24) and (25) yield strong biases regarding which referent is likely to be mentioned next: Mary for (24) and Jane for (25). Both the next-mention biases for passages like (24-25) and the pronoun interpretation biases for passages like (22-23) have been attributed to comprehenders' knowledge about the *implicit causality* of the verbs. Implicit causality (IC) verbs are classified as biased either to the first referent (IC-1, as in *infuriate*, *impress*, *disappoint*) or the second referent (IC-2, as in *scold*, *admire*, *detest*). Both IC-1 and IC-2 verbs have been reported to differ from other verbs in that they define a "locus of the underlying cause of the action or attitude" (Garvey & Caramazza, 1974). They form a class because they systematically attribute cause to one individual, namely the one that initiates the action or is implicated as the cause. This individual is taken to be the most likely referent for an ambiguous pronoun in contexts like (22-23) and is also likely to be mentioned again in contexts like (24-25). However, this next-mention bias is reduced or eliminated with a full-stop prompt as in (26) and (27) (see Au, 1986).

(26) Mary infuriated Jane. _____

(27) Mary scolded Jane. _____

The results from experiments with stimuli like (24-27) have led to confusion in the literature about the role of the connective *because*. If Stevenson et al. (1994) are

correct, the connective constrains patterns of coreference by directly modifying the salience of entities, in a role distinct from that of constraining coherence relations. If discourse coherence models like those of Hobbs (1979) and Kehler (2002) are correct, the connective *because* simply enforces an Explanation coherence relation and any subsequent patterns of coreference reflect the coreference biases of Explanation relations.

Given the coherence-driven model proposed in Equation (3.1) which distinguishes between biases for upcoming coherence relations ($p(CR)$) and upcoming coreference conditioned on coherence ($p(\textit{referent} | CR)$), the differences in next-mention biases that have been reported for (24-25) and (26-27) may be attributable to the fact that the elicited continuations in (26-27) represent a broader range of coherence relations, each with its own next-mention bias, compared to the Explanation-only continuations elicited with *because* prompts. This would suggest that the previously reported IC coreference patterns may represent only part of a larger and more complex system of biases that will become apparent once coreference is conditioned on the operative coherence relation.

The experiment in the next section tests the prediction that the IC biases found in sentence completions using a *because* prompt are statistically indistinguishable from those found in a full-stop condition when only freely generated Explanation continuations are analyzed—in other words, whether $p(\textit{referent} | CR=Explanation)$ is consistent across prompt conditions even if $p(CR=Explanation)$ varies depending on the presence or absence of a *because* prompt. In categorizing responses by coherence relation, the aim is to test whether the previously reported IC bias can be localized to Explanation relations and whether a richer picture emerges when choice of next mention is conditioned on a wider range of coherence relations.

3.7 Experiment IV: Implicit Causality Biases

Experiment IV crosses verb type (IC verb vs. non-IC verb) with prompt type (*because* vs. full-stop vs. dialog prompt; see Section 6.2 for discussion of the dialog condition) in order to test whether the next-mention biases that emerge in a *because* condition match the next-mention biases observed in freely generated Explanations. Sample story continuation prompts are shown in (28-30).

- (28) Mary infuriated John(./because) _____. [IC-1 VERB]
- (29) Mary scolded John(./because) _____. [IC-2 VERB]
- (30) Mary chatted with John(./because) _____. [NON-IC VERB]

Items with *because* prompts are predicted to yield a set of next-mention biases that replicate the biases reported in earlier studies: In the *because* condition, IC-1 verbs are predicted to create a bias towards the individual mentioned in the first noun phrase (NP1), IC-2 verbs are predicted to create a bias towards the second noun phrase (NP2), and non-IC verbs in aggregate are predicted to create no strong next-mention bias. These next-mention biases are expected to parallel the biases observed in full-stop continuations when the analysis is restricted to full-stop continuations that embody Explanation relations. This parallel is predicted to disappear, on the other hand, when the *because* continuations are compared with all full-stop continuations, regardless of coherence relation. The distribution of coherence relations following IC verbs is also compared with the distribution following non-IC verbs in order to see what effects verb type has on comprehenders' expectations about where the discourse is going and what types of coherence relations are likely to ensue.

3.7.1 Methods

The task described here uses an adaptation of the procedure for Experiments I-III. Participants were instructed to imagine that they were talking on the phone

with a friend. They were told that they would see statements that the friend made and they would be asked to either continue what the friend was saying (monologue condition) or pose a question to the friend (dialog condition). I discuss here the results for the monologue condition. (See Section 6.2 for discussion of the items in the dialog condition.) The monologue continuations are interpreted as an index of the participants' expectations about where the discourse is going and who is likely to be mentioned next.

Participants

Seventy-five monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

Forty IC verbs and forty non-IC verbs were taken from McKoon et al. (1993) with three replacements. The verbs *cheat*, *jeer*, and *dread* were felt to sound awkward in these sentence frames, and were replaced with *offend*, *mock*, and *fear*, respectively. The forty IC verbs consisted of twenty IC-1 verbs and twenty IC-2 verbs. All context sentences mentioned two possible referents, one male and one female, with gender balanced across items. The full stimulus set is given in the Appendix. Twenty filler sentences used non-IC verbs and were followed by various interclausal connectives (monologue continuation) or a dialog response that contained the beginning of a question (dialog continuation), for a total of 100 items.

Story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their continuation.

Evaluation and Analysis

The annotation and analysis follows that described in the previous sections. Two trained judges, the author of this dissertation and a UCSD Linguistics undergraduate, evaluated the participants' continuations. For continuations elicited with

the *because* prompt, the judges annotated the participants' choice of next mention (NP1 or NP2). Continuations following a 'because' prompt were all classified as Explanations. There were some continuations that, without the 'because' connective, would have been ambiguous between an Explanation and another coherence relation, but the annotators classified them as Explanations because they assumed that, in that context, the participant intended a causal reading.⁶

For the full-stop prompt, they annotated both choice of next mention and the coherence relation that could be inferred to hold between the context sentence and the continuation. Pronouns that appeared in participants' continuations were disambiguated given the genders of the two possible referents. Analyses of variance were conducted to measure the effect of one within-subjects factor (verb type) on both the choice of next mention and on the choice of coherence relation. Verb type varied between items.

3.7.2 Results

After setting aside continuations elicited for the dialog condition as well as those continuations in which the participant wrote only a fragment or for which the coherence relation was not clear, the remaining dataset contained 3899 continuations. Of those, only continuations that contained a mention of at least one of the possible referents from the context sentence were included in the analysis (N=3208). It is customary to present results for IC data separately by verb type, so in what follows, I describe the next-mention results for the IC-1, IC-2, and non-IC verbs separately and then compare across verb types for the coherence analysis.

Effect of Verb Type on Next-Mention Biases

The next-mention results for IC-1 verbs (e.g., *infuriate*, *impress*, *disappoint*)

⁶For example, the following continuation written by a participant in the experiment was classified as an Explanation even though, without the 'because' prompt, a Violated-Expectation relation can be inferred to hold:

(i) Greg corrected Sally because she kept on making grammatical mistakes.

exhibit the predicted pattern: The NP1 bias in the *because* condition (85.0%) is statistically indistinguishable from the NP1 bias in the full-stop condition (81.2%) when the analysis is restricted to Explanations ($F_s < 1$). When the analysis includes all full-stop continuations, regardless of coherence relation, the full-stop NP1 bias is reduced (59.8%) and is significantly different from the NP1 bias in the *because* condition (effect of prompt type on choice of next mention: $F_1(1,74)=48.315$, $p < 0.001$; $F_2(1,19)=49.308$, $p < 0.001$).⁷

As was reported for the transfer-of-possession contexts in the previous sections, individual coherence relations are associated with different next-mention biases in the full-stop condition. The biases for IC-1 verbs are shown in Figure 3.11.




Figure 3.11: Coherence and next-mention biases for IC-1 verbs in *because* and full-stop conditions in Experiment IV

Following IC-1 verbs, Explanation relations are biased to NP1 (84.2%; significantly different from a hypothetical mean of 0.5: $t_1(74)=24.315$, $p < 0.001$; $t_2(19)=13.213$, $p < 0.001$), Results are biased to NP2 (84.8%; $t_1(51)=7.134$, $p < 0.001$; $t_2(17)=13.652$, $p < 0.001$), and Elaborations are biased to NP1 (62.6%; significant only by items: $t_1(32)=1.609$, $p < 0.12$; $t_2(17)=2.231$, $p < 0.04$). Other relations represented less than 5% of the continuations and are not reported here.

The next-mention results for IC-2 verbs (e.g., *scold*, *admire*, *detest*) also show

⁷As in the rest of this dissertation, the means listed in this section all represent subject means. The results of this experiment were previously reported in Kehler, Kertz, Rohde, and Elman (2008). However, the means presented here differ from those reported by Kehler et al. because the previously published means were raw means (# NP1 continuations/# total continuations).

the predicted pattern: The NP2 bias in the *because* condition (91.1%) is statistically indistinguishable from the NP2 bias in the full-stop condition (87.7%) when the analysis is restricted to Explanations ($F_1 < 1$; $F_2(1,19)=1.224$, $p < 0.29$). When the analysis includes all full-stop continuations, regardless of coherence relation, the full-stop NP2 bias is reduced (83.5%) and is significantly different from the NP2 bias in the *because* condition (effect of prompt type on choice of next mention: $F_1(1,74)=12.568$, $p < 0.001$; $F_2(1,19)=6.024$, $p < 0.03$).

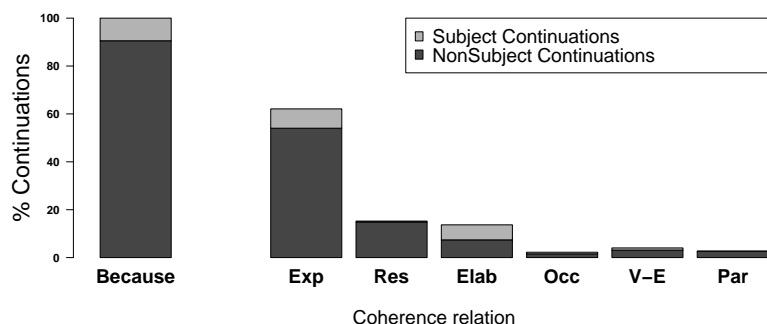


Figure 3.12: Coherence and next-mention biases for IC-2 verbs in *because* and full-stop conditions in Experiment IV

This difference is again due to variation in next-mention biases across different coherence relations, as can be seen in Figure 3.12. Following IC-2 verbs, Explanation relations are biased to NP2 (90.0%; $t_1(74)=30.086$, $p < 0.001$; $t_2(19)=17.487$, $p < 0.001$); the NP2 bias is even stronger with Result relations (96.9%; $t_1(42)=19.27$, $p < 0.001$; $t_2(14)=4.036$, $p < 0.002$), and almost disappears with Elaboration relations (56.9%; $t_s < 1$). Other relations represented less than 5% of the continuations and are not listed here.

Lastly, the next-mention results for non-IC verbs (e.g., *chat with*, *see*, *work with*) show the same pattern: The NP1 bias in the *because* condition (55.4%) is statistically indistinguishable from the NP1 bias in the full-stop condition (54.0%) when the analysis is restricted to Explanations ($F_1 < 1$; $F_2(1,36)=1.460$, $p < 0.24$). When the analysis includes all full-stop continuations, regardless of coherence relation, the full-

stop NP1 bias is reduced (48.5%) and is significantly different from the NP1 bias in the *because* condition (effect of prompt type on choice of next mention: $F_1(1,74)=4.187$, $p<0.05$; $F_2(1,19)=7.404$, $p<0.01$).

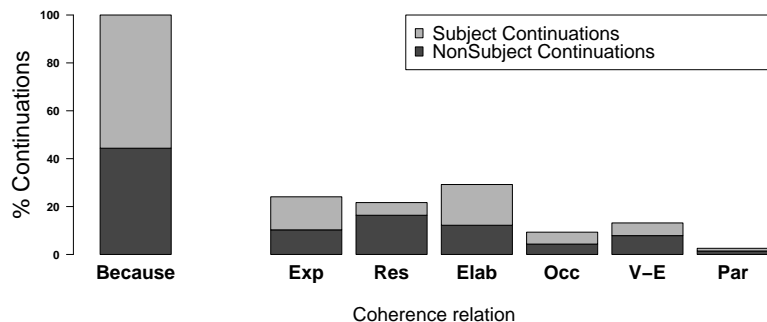


Figure 3.13: Coherence and next-mention biases for non-IC verbs in *because* and full-stop conditions in Experiment IV

As was the case for the IC verbs, this difference is due to the variation in next-mention biases across different coherence relations in the full-stop condition. The patterns can be seen in Figure 3.13. Following non-IC verbs, Explanation relations are biased to NP1 (55.8%; $t_1(74)=3.180$, $p<0.003$; $t_2(39)=2.373$, $p<0.03$), Results are biased to NP2 (68.2%; $t_1(56)=3.433$, $p<0.002$; $t_2(30)=2.453$, $p<0.03$), and Elaborations are biased to NP1 (59.3%; significant only by subjects: $t_1(68)=2.137$, $p<0.04$; $t_2(36)=1.655$, $p<0.11$). Other common coherence relations included Occasions, which showed no strong next-mention bias (51.0% NP2; not significant by subjects or by items: $t_s<1$), and Violated Expectations, which also showed no strong next-mention bias (55.1% NP2; not significant by subjects or by items: $t_1<1$; $t_2(25)=1.423$, $p<0.17$).

Effect of Verb Type on Coherence Distribution

Not only do next-mention biases differ by verb type, but the distribution of coherence relations differs by verb type as well. Comparing the probability of each coherence relation following IC verbs and non-IC verbs, the most striking difference is the bias regarding upcoming Explanations: As can be seen by comparing

the percentages of each coherence relation in Figures 3.11, 3.12, and 3.13, IC verbs yield significantly more Explanations (60.0%) compared to non-IC verbs (24.9%; $F_1(1,74)=145.65$, $p<0.001$; $F_2(1,78)=71.658$, $p<0.001$). IC and non-IC verbs yield roughly the same proportions of Result relations (18.4% and 21.2%; $F_s<1$). Non-IC verbs, however, exceed IC verbs in the proportions of Elaborations (28.9% following non-IC vs. 12.3% following IC; $F_1(1,74)=54.86$, $p<0.001$; $F_2(1,78)=24.403$, $p<0.001$), Occasions (9.4% vs. 3.1%; $F_1(1,74)=16.57$, $p<0.001$; $F_2(1,78)=12.027$, $p<0.001$), and Violated Expectations (12.9% vs. 3.6%; $F_1(1,74)=41.516$, $p<0.001$; $F_2(1,78)=9.748$, $p<0.003$).

3.7.3 Discussion

The aim of this experiment was to test whether IC next-mention biases depend on the presence of a *because* connective or whether the biases are coherence-driven, predicting that they can be replicated without the connective if coherence is controlled. The predictions were borne out: For all three verb types (IC-1, IC-2, non-IC), the next-mention biases were (i) significantly different between the two prompt conditions, but were nonetheless (ii) statistically indistinguishable when only the full-stop continuations that embodied Explanation relations were analyzed. Whereas (ii) can only be coincidental under the previous analyses that attribute IC biases to a combination of the verb and the connective, the results pattern with the data from the previous sections in that the overall statistics conceal a consistent system of stronger biases once coherence relations are conditioned on. Recall that the full-stop prompt in Stevenson et al.'s story continuation experiment yielded an overall NP1 bias following IC-1 context sentences (63.1% to NP1) compared to an overall NP2 bias following IC-2 context sentences (71.5% to NP2). The biases observed here are similar to Stevenson et al.'s when the results are collapsed across coherence relations, but this experiment reveals that the aggregate proportions dilute the biases and conceal a much more complex pattern underlying coreference.

The results also reveal an additional IC bias with respect to the probability of certain upcoming coherence relations: IC verbs are significantly more likely to elicit Explanation continuations than non-IC verbs in the full-stop condition. Whereas the strong next-mention bias for IC verbs has long been the bedrock observation in the psycholinguistics IC literature, this latter bias has not been demonstrated before because previous studies have not categorized their data by coherence.

Whereas the earlier experiments described in this chapter tested a coherence-driven model of *pronoun interpretation*, this experiment suggests that coherence-driven biases may contribute more generally to comprehenders' expectations about next mention, as distinct from the form of reference chosen. If that is the case, the model that was advocated in the earlier sections may be better understood as a model of next-mention biases, as in (3.4).

$$p(\text{NextMention}=\text{Ref}) = \sum_{CR} p(\text{NextMention}=\text{Ref} \mid CR) \times p(CR) \quad (3.4)$$

In order to relate the next-mention biases in the model in (3.4) to the pronoun interpretation biases in the original model proposed in (3.1) for Experiment I, one would need a term that links the choice of next mention with the choice to use a pronoun. The relationship between next-mention biases and pronoun interpretation biases will be taken up in the next chapter.

3.8 Conclusion

The approach to coreferential processing taken in this chapter aims to clarify the source of surface-level biases by reanalyzing them as side effects of deeper discourse-level processes. The results from the four story-continuation experiments support the conclusion that the focusing effects posited by researchers like Stevenson et al.—the thematic-role biases of certain verbs and the attention-shifting properties

of intersentential connectives—are in fact side effects of deeper event-level and causal biases that emerge from general coherence establishment processes.

Experiment I showed that pronoun interpretation is sensitive to changes in event structure (verbal aspect) and to differences in the locations of event participants (verbal lexical semantics). Furthermore, the event structure bias itself is associated with only certain types of coherence relations in the discourse. Experiments II and III showed that other contextual cues (normal and abnormal objects or a question posed in the instructions) can shift the distribution of coherence relations, and that the observed pattern of pronoun interpretation corresponds to the underlying distribution of coherence relations. Lastly, Experiment IV showed that the coherence-driven model can be successfully applied to a different set of biases, those that emerge in contexts with implicit-causality verbs. The previously reported implicit-causality biases were shown to be limited to passages in which the operative coherence relation is an Explanation. In fact, next-mention biases observed in passages with an intersentential connective that enforces an Explanation coherence relation were found to be indistinguishable from next-mention biases in passages in which the Explanation relation must be inferred. Taken together, these results suggest that factors which previous researchers had modeled as contextual cues that directly modify the salience of particular referents can instead be understood as cues that indirectly shift comprehenders' expectations about where the discourse is going — which in turn has an effect on comprehenders' expectations about coreference.

The specific manipulations of verbal aspect, verb type, object type, and the question posed in the experimental instructions are less important than the fact that different discourse environments yield different distributions of coherence relations. Given a distribution of coherence relations combined with the lexical semantics of the verbs, it appears to be the coreference biases of those individual relations that play a crucial role in determining the overall pattern of coreference. These results suggest that even where pockets of data show evidence of grammatical-role, thematic-role,

event-level, or causal-inferencing biases, these biases can and should be understood within a much richer model of discourse comprehension. These experiments also point to the importance of controlling for coherence factors when investigating surface-level or information-structural preferences in pronoun interpretation (see also Wolf, Gibson, & Desmet, 2004; Kertz, Kehler, & Elman, 2006).

The model of coreferential processing that was sketched in Section 3.2 and revised in Section 3.7 contained two types of biases: a bias regarding the upcoming coherence relation ($p(CR)$) and a conditional probability representing the probability of a specific coreference outcome given the operative coherence relation ($p(Ref | CR)$). The experiments presented in this chapter demonstrate that different contextual cues can shift the probability of the upcoming coherence relation ($p(CR)$), but that the coreference biases for each coherence relation remain largely consistent across experiments. Experiment I showed that aspectual changes can shift the first bias: perfective aspect increases the number of Occasion relations. Experiment II showed that an abnormal direct object increases the number of Explanations and Result relations. Experiment III showed that a ‘Why’ or ‘What next’ question posed in the instructions increases the number of Explanations or Occasions, respectively. Experiment IV showed that IC verbs create a bias towards upcoming Explanation relations. Despite the shifts in the distribution of coherence relations, the coreference biases were in general quite consistent. Across all of the transfer-of-possession contexts, Occasions and Results were consistently biased to the Goal, whereas Elaborations, Explanations, and Violated Expectations were consistently biased to the Source. The IC contexts were a bit different: Although Results were consistently biased to NP2 and Elaborations contained more NP1 references, the next-mention biases for Explanations depended both on the coherence relation and the verb lexical semantics — Explanations following IC-1 verbs were biased to NP1, whereas Explanations following IC-2 verbs were biased to NP2, and Explanations following non-IC verbs showed no strong next-mention bias.

Aside from shedding light on a set of results in the coreference literature, this chapter also suggests that expectation-driven biases apply at the discourse level as well as at other levels of linguistic structure. It shows that comprehenders are aware of contextual cues that provide signals about where a discourse is going (either overt cues like intersentential connectives or more indirect cues like the aspect of the verb or surprising objects of transfer). This fits with a series of results discussed in the introduction that show that comprehenders are sensitive to contextual cues at other non-discourse levels of language processing and that they use these cues to generate expectations about upcoming material. The results in this chapter show that there is nothing unique about discourse-level expectations; they can be thought of as simply operating over different linguistic units—coherence relations rather than sounds or words or syntactic structures.

The next chapter considers whether one of the factors that can influence comprehenders' expectations about upcoming coherence relations may be referring expressions themselves. The story-continuation experiments in the next chapter vary the continuation prompts between a full-stop prompt and a pronoun prompt. This manipulation provides an opportunity to sort out the relationship between next-mention biases and pronoun interpretation, and, furthermore, to test whether or not information about which referent has been mentioned next in the discourse can affect comprehenders' expectations about where the discourse is going and what the operative coherence relation is likely to be.

3.9 Acknowledgments

The experiments described in this chapter have been reported in earlier work and are discussed here with permission from the coauthors. The dissertation author was the primary investigator for the material presented here. Experiment I was reported in Rohde, Kehler, and Elman (2006) [Event Structure and Discourse Coherence Biases in Pronoun Interpretation. In *Proceedings of the 28th Annual Conference*

of the Cognitive Science Society. Vancouver, July 26-29, 2006. Pages 697-702]. Experiments II and III were reported in Rohde, Kehler, and Elman (2007) [Pronoun Interpretation as a Side Effect of Discourse Coherence. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Nashville, TN, August 1-4, 2007. Pages 617-622]. Experiments I and IV appeared in Kehler, Kertz, Rohde, and Elman (2008) [Coherence and Coreference Revisited. *Journal of Semantics*. Volume 25: 1-44].

4

Bidirectional Influences in Coherence and Coreference

The experiments in this chapter build on the results of the previous chapter by considering another side of the interaction between coherence and coreference. Whereas the last chapter addressed the question of how comprehenders' expectations regarding upcoming coherence relations influence patterns of coreference, this chapter considers the opposite causal order by testing whether the distribution of coherence relations is itself influenced by cues about which referents have been mentioned again as the discourse progresses. If comprehenders are aware that particular coherence relations are associated with particular coreference biases, then the hypothesis is that simply acquiring information about who has been mentioned next could bias comprehenders about the coherence relation that is most likely to ensue. Consider the story-continuation prompts in (31).

- (31) a. John infuriated Mary. _____. [no-pronoun prompt]
b. John infuriated Mary. He _____. [NP1-referring prompt]
c. John infuriated Mary. She _____. [NP2-referring prompt]

The verb *infuriate* in (31) is a member of the class of so-called 'implicit causality' (IC) verbs that were analyzed in the final story-continuation experiment in Chap-

ter 3. These verbs were shown to yield strong biases concerning which referent was likely to be mentioned next, and more importantly, those next-mention biases were strongly conditioned on the operative coherence relation. For example, continuations elicited with no-pronoun prompts like (31a) were shown to yield more references to the subject (NP1) than the non-subject (NP2), just in case the operative coherence relation was Explanation, whereas the same context sentence was shown to yield more references to NP2 if the operative coherence relation was Result.

The prediction that emerges from these coherence-driven biases is that the mere presence of a referring expression that signals who has been mentioned next could shift comprehenders' expectations about the upcoming coherence relation—biasing them in favor of relations that are associated with that referent. The NP1-referring pronoun in (31b) is thus predicted to yield more continuations that provide an explanation of the event in the context sentence because Explanations are NP1-biased (e.g., *John infuriated Mary. He cheated at Scrabble*); likewise, the NP2-referring pronoun in (31b) is predicted to yield more continuations that describe the result of the event in the context sentence because Results are NP2-biased (e.g., *John infuriated Mary. She told him to take a hike*).

These predictions are tested in the first experiment using story-continuation prompts like those in (31) with no-pronoun and unambiguous-pronoun prompts. The no-pronoun prompt replicates the next-mention biases reported in Chapter 3: Different coherence relations are associated with different patterns of next-mention (coherence influences coreference). The NP1-referring and NP2-referring pronoun prompts show that, as predicted, comprehenders' coherence expectations are affected by a subsequent referring expression (coreference influences coherence). The continuations elicited with the no-pronoun prompt (31a) reveal additional information about participants' use of referring expressions—how often they choose to use pronouns and who they use pronouns to refer back to. The results reveal a bias towards producing pronouns when the referent occupies the subject position of the context sentence, re-

ardless of whether the subject referent is preferred for next mention (see also Stevenson, Crawley, & Kleinman, 1994 and Arnold, 2001). The next experiment builds on these pronoun production biases and investigates whether evidence of an effect of coreference on coherence can also be found in contexts with ambiguous pronouns.

The last two experiments address the question of whether the bidirectional effects extend beyond IC passages to other types of contexts. The additional context tested here is one that is familiar from the pronoun interpretation experiments in Chapter 3: passages with transfer-of-possession verbs. The results for the transfer-of-possession passages suggest that the conclusions from Chapter 3 can be better understood within a more general model of pronoun interpretation and production—one that crucially distinguishes between factors that influence choice of next mention (verb type and the operative coherence relation) and factors that influence choice of referring expression (grammatical role of the referent). A Bayesian approach to this data is laid out in the discussion in Section 4.5

4.1 Experiment V: Unambiguous IC Contexts

The interaction between coherence and coreference is tested first using story continuations with IC and non-IC verbs. Sample passages are shown in (32) with the IC-1 verb *infuriate*, the IC-2 verb *scold*, and the non-IC verb *chat with*.¹

- (32) a. NO-PRONOUN PROMPT
 John infuriated_{IC1} Mary. ____.
 John scolded_{IC2} Mary. ____.
 John chatted_{with_{nonIC}} Mary. ____.
- b. NP1-REFERRING PROMPT
 John infuriated_{IC1} Mary. He ____.
 John scolded_{IC2} Mary. He ____.
 John chatted_{with_{nonIC}} Mary. He ____.

¹This chapter will continue to use the IC-1 and IC-2 labels for verbs even though Chapter 3 showed that these biases are restricted primarily to Explanation relations.

c. NP2-REFERRING PROMPT

John infuriated_{IC1} Mary. She ___.

John scolded_{IC2} Mary. She ___.

John chatted_with_{nonIC} Mary. She ___.

In this experiment, context sentences are followed by either a no-pronoun prompt or an unambiguous pronoun prompt. When participants write a continuation following a no-pronoun prompt (32a), they are free to complete it however they wish, and they are assumed to do so based on their expectations about where the story is going and who is likely to be mentioned next. The responses in the no-pronoun condition are therefore interpreted as a reflection of comprehenders' expectations about upcoming coherence relations, conditioned only on the preceding context sentence without any next-mention information.

In the unambiguous-pronoun-prompt conditions (32b-c), participants again generate expectations based on the context sentence but their continuations are constrained by the unambiguous pronoun referring to one of the two individuals in the context sentence. The elicited continuations are therefore interpreted as a reflection of comprehenders' updated expectations about the operative coherence relation, conditioned on the preceding context sentence and the information about which referent has been mentioned next.

If comprehenders use cues about who has been mentioned next to determine which coherence relation is likely to be operative, then an NP1-referring pronoun is predicted to shift comprehenders' expectations in favor of NP1-biased coherence relations, whereas an NP2-referring pronoun is predicted to shift expectations in favor of NP2-biased coherence relations. In order to test the effects of an NP1-referring or NP2-referring pronoun, pairs of coherence relations are needed with opposite next-mention biases. In what follows, I review the biases for IC-1, IC-2, and non-IC verbs and identify a pair of relations for each verb type that will make it possible to test the effects of coreference on coherence. Then I lay out the predictions for the different

verb types in story-continuation contexts with no-pronoun, NP1-referring, and NP2-referring prompts.

Implicit Causality Biases

For contexts with IC-1 verbs like *infuriate*, Experiment IV showed that no-pronoun prompts yield an overall bias to re-mention NP1, e.g. John in (33). The NP1 bias is apparent when the data is collapsed across coherence relations. However, a system of stronger next-mention biases emerges when the biases are conditioned on coherence: Explanation relations following IC-1 verbs show a strong NP1 bias, whereas Results show a strong NP2 bias. Experiment IV showed that the overall NP1 bias can be attributed to the prevalence of NP1-biased Explanations.

(33) John infuriated Mary. _____. [IC-1]

For context sentences with IC-2 verbs like *scold*, on the other hand, no-pronoun prompts reveal an overall bias to re-mention NP2, e.g. Mary in (34). The NP2 bias is apparent when the data is collapsed across coherence relations. Again, however, the biases differ by coherence relation: Explanations following IC-2 verbs show a strong NP2 bias, whereas Elaborations contain more NP1 references. Experiment IV attributed the overall NP2 bias to the prevalence of NP2-biased Explanations.

(34) John scolded Mary. _____. [IC-2]

Finally, Experiment IV showed that context sentences with non-IC verbs, like that in (35), also yield different next-mention biases for different coherence relations. With the set of non-IC verbs tested in Experiment IV, the two most common coherence relations were found to be Elaborations and Results, and these two relations were shown to have opposite next-mention biases: Elaboration relations are biased to NP1, whereas Result relations are biased to NP2.

(35) John chatted with Mary. _____. [IC-2]

Given these different next-mention biases across verb types and coherence relations, the predictions for the current experiment are as follows (stimuli repeated from (32)). For context sentences with IC-1 verbs, the no-pronoun condition (36a) is predicted to replicate the previously reported NP1 bias for Explanations and NP2 bias for Results. The NP1-referring prompt (36b) is predicted to yield more Explanation continuations than the NP2-referring prompt (36c). The NP2-referring prompt (36c), on the other hand, is predicted to yield more Results than (36b).

- (36) a. John infuriated Mary. _____. [IC-1, no-pronoun prompt]
 b. John infuriated Mary. He _____. [IC-1, NP1 prompt]
 c. John infuriated Mary. She _____. [IC-1, NP2 prompt]

For context sentences with IC-2 verbs, the condition with the no-pronoun prompt (37a) is predicted to replicate the previously reported large number of NP1 references for Elaborations and the NP2 bias for Explanations. The NP1-referring prompt (37b) is predicted to yield more Elaboration continuations than the NP2-referring prompt (37c). The NP2-referring prompt in (37c), on the other hand, is predicted to yield more Explanations than (37b).

- (37) a. John scolded Mary. _____. [IC-1, no-pronoun prompt]
 b. John scolded Mary. He _____. [IC-1, NP1 prompt]
 c. John scolded Mary. She _____. [IC-1, NP2 prompt]

For context sentences with non-IC verbs, the no-pronoun prompt (38a) predicted to replicate the previously reported NP1 bias for Elaborations and NP2 bias for Results. The NP1-referring prompt (38b) is predicted to yield more Elaboration continuations than the NP2-referring prompt (38c). The NP2-referring prompt in (38c), on the other hand, is predicted to yield more Results than (38b).

- (38) a. John chatted with Mary. _____. [non-IC, no-pronoun prompt]
 b. John chatted with Mary. He _____. [non-IC, NP1 prompt]
 c. John chatted with Mary. She _____. [non-IC, NP2 prompt]

Across all three verb types, the continuations elicited with the no-pronoun prompt are predicted to replicate the previous next-mention biases, but the main contrast of interest is the difference in the proportion of coherence relations following NP1-referring and NP2-referring pronoun prompts. For this reason, any effect of prompt type on coherence is further analyzed using pairwise comparisons in order to localize the source of the effect — checking for differences among the no-pronoun, NP1-referring, and NP2-referring prompts. After the discussion in Section 4.1.3 of the bidirectional results, Section 4.1.4 reports an additional analysis concerning the rate of pronominalization. The results of that analysis lead to a prediction that is tested in Experiment VI.

4.1.1 Methods

The story continuation task described here follows the procedure used for Experiment IV (Section 3.7.1). Participants were asked to imagine that they were having a phone conversation with a friend. For the target items, they were presented with a sentence that the friend had supposedly uttered, and they were asked to write a natural continuation at the prompt provided. Other non-target items required the participant to pose or answer a question as part of the simulated dialog. The continuations are interpreted as an index of participants' expectations about where the discourse is going and, for items in the no-pronoun condition, also about which referent is likely to be mentioned next.

Participants

Thirty-two monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

For the experimental items, eighteen IC verbs and eighteen non-IC verbs were taken from Experiment IV (originally adapted from McKoon et al., 1993). The verbs

were selected for the strength of their next-mention biases: The eighteen IC verbs consisted of nine canonical (i.e., strongly NP1-biased) IC-1 verbs and nine canonical (strongly NP2-biased) IC-2 verbs; the eighteen non-IC verbs had an average next-mention bias close to 50/50 for NP1/NP2. Context sentences mentioned two possible referents, one male and one female, with gender balanced across items. Participants saw each verb only once, appearing in one of three conditions: no-pronoun prompt, NP1-pronoun prompt, or NP2-pronoun prompt. The number of items reflects the three conditions: Each participant saw six IC verbs, and six non-IC verbs across the three prompt conditions. The full stimulus set is provided in the Appendix. The experimental stimuli consisted of thirty-six experimental items intermixed with forty-eight stimuli for an unrelated experiment with transfer-of-possession contexts (Experiment XII described in Section 6.3) as well as forty-eight additional fillers, for a total of one hundred and thirty-two items. The additional fillers consisted of monologue and dialog prompts with context sentences containing non-IC verbs followed by intersentential connectives, full stops, or pronoun prompts.

Story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their continuation. The entire experiment took roughly forty-five minutes, but participants were encouraged to have an hour available so that the experiment could be completed in one session. Participants could leave the website and return at a later time by identifying themselves with an ID number. They were instructed to imagine a natural story continuation for each prompt, writing the first continuation that came to mind and avoiding humor.

Evaluation and Analysis

Two trained judges, the author of this dissertation and a UCSD Linguistics undergraduate, annotated the participants' continuations. For both the no-pronoun and pronoun-prompt conditions, they assessed the coherence relation that could be

inferred to hold between the context sentence and the continuation. For the no-pronoun condition, they also annotated the participants' choice of next mention. Pronouns that appeared in participants' continuations were disambiguated by gender.

For the analysis in this experiment and in the remaining experiments in this chapter, I consider each verb type separately and use analyses of variance to test the effect of one factor (prompt type) on the choice of coherence relation. Prompt type varied within subjects and within items. For each verb type, the analysis focuses on one pair of coherence relations with opposing biases: Explanation/Result for IC-1; Elaboration/Explanation for IC-2; Elaboration/Result for non-IC. Pair-wise tests were also conducted to identify the source of the effect of prompt type, using t-tests on raw means to compare the proportion of each relevant coherence relation between no-pronoun, NP1-pronoun, and NP2-pronoun conditions. Again, the main prediction is that the NP1-prompts and NP2-prompts will yield significantly different proportions of the relevant oppositely-biased coherence relations.

As in Experiments I-IV, coherence relations and next mentions were coded as binary outcomes. Because the analyses involve examining proportions, the percentages of coherence relations and next mentions were first arcsine-transformed before carrying out analyses of variance. For clarity of presentation in this section and throughout the chapter, I present means in the form of raw proportions.

4.1.2 Results

After setting aside continuations in which the participant wrote only a fragment or for which the coherence relation was not clear (4.0% of the total), the remaining dataset contained 1106 continuations. Of those, only continuations that mentioned at least one referent from the context sentence were included in the analysis (N=1024).

Replication of Previous Coherence-Driven Biases

The next-mention biases observed in the no-pronoun condition replicate the biases reported in Experiment IV. I report here only the biases for the pair of relevant oppositely-biased coherence relations for each verb type. As before, context sentences with IC-1 verbs yield an overall NP1 bias (56.3% NP1 continuations), and Explanation relations in these contexts have a stronger NP1 bias (77.8%) than Result relations (16.7%; $t_1(37)=4.639$, $p<0.001$; $t_2(14)=4.176$, $p<0.001$). Context sentences with IC-2 verbs yield an overall NP2 bias (82.0% NP2 continuations), and Explanation relations in these contexts have a stronger NP2 bias (86.4%) than Elaborations (55.6%; significant only by subjects: $t_1(37)=2.243$, $p<0.04$; $t_2(11)=0.85$, $p<0.42$). Finally, context sentences with non-IC verbs yield an overall NP2 bias (74.0% NP2 continuations which is stronger than that reported for non-IC verbs Experiment IV), and Result relations in these contexts have a stronger NP2 bias (85.4%) than Elaboration relations (44.7%; $t_1(36)=3.105$, $p<0.01$; $t_2(21)=2.661$, $p<0.02$).²

The results from the no-pronoun condition similarly replicate the original biases observed in Experiment IV regarding upcoming coherence relations: Explanation continuations are more frequent following IC verbs (58.2%) than non-IC verbs (16.4%; $F_1(1,31)=61.978$, $p<0.001$; $F_2(1,33)=30.002$, $p<0.001$).

Effects of Prompt Type on Coherence Distribution

In contexts with IC-1 verbs (e.g., *John infuriated Mary. He/She/∅...*), the effect of prompt type on the distribution of coherence relations exhibits the predicted pattern. The results are shown in Figure 4.1. The largest proportion of Explanation relations was observed in the condition with NP1-referring pronoun prompts (75.5%), followed by no-pronoun prompts (46.6%), and then NP2 prompts (10.9%; effect of prompt type on proportion of Explanations: $F_1(2,62)=47.661$, $p<0.001$; $F_2(2,16)=26.202$, $p<0.001$). In the pairwise comparisons, the proportion of Explanations following NP1 prompts is greater than the proportion following no-pronoun

²For this experiment and all other experiments in this chapter, all means represent subject means, unless otherwise noted.

prompts ($t_1(62)=4.652$, $p<0.001$; $t_2(16)=2.252$, $p<0.04$), which is in turn greater than the proportion following NP2 prompts ($t_1(62)=5.466$, $p<0.001$; $t_2(16)=3.722$, $p<0.002$). As predicted, the NP1-referring prompts yield significantly more Explanations than the NP2-referring prompts ($t_1(62)=11.844$, $p<0.001$; $t_2(16)=6.650$, $p<0.001$).




Figure 4.1: Effects of prompt type on proportions of NP1-biased Explanations and NP2-biased Results following IC-1 verbs in Experiment V

Also as predicted for IC-1 verbs, the proportion of NP2-biased Result relations is found to increase in precisely the conditions in which the proportion of NP1-biased Explanation relations decreases: NP2-referring pronoun prompts yield the largest proportion of Results (73.2%), followed by no-pronoun prompts (22.1%), and then NP1 prompts (4.7%; effect of prompt type: $F_1(2,62)=51.351$, $p<0.001$; $F_2(2,16)=68.263$, $p<0.001$). In the pairwise comparisons, the proportion of Results following NP2 prompts is greater than the proportion following no-pronoun prompts ($t_1(62)=6.393$, $p<0.001$; $t_2(16)=5.004$, $p<0.001$), and the proportion of Results following no-pronoun prompts is in turn greater than the proportion following NP1 prompts ($t_1(62)=2.676$, $p<0.01$; $t_2(16)=2.711$, $p<0.02$). As predicted, the NP2-referring prompts yield more Results than the NP1-referring prompts ($t_1(62)=11.584$, $p<0.001$; $t_2(16)=8.515$, $p<0.001$).

In contexts with IC-2 verbs (e.g., *John scolded Mary. He/She/Ø...*), the effect of prompt type on the distribution of coherence relations also exhibits the

predicted pattern. The results are shown in Figure 4.2. NP1-referring pronoun prompts yielded the largest proportion of Elaboration relations (30.7%), followed by no-pronoun prompts (13.5%), and then NP2 prompts (4.7%; effect of prompt type: $F_1(2,62)=10.976$, $p<0.001$; $F_2(2,15)=7.114$, $p<0.007$). In the pairwise comparisons, the proportion of Elaborations following NP1 prompts is numerically greater than the proportion following no-pronoun prompts (only significant by subjects: $t_1(62)=2.805$, $p<0.007$; $t_2(15)=1.563$, $p<0.14$), which is in turn greater than the proportion following NP2 prompts (marginal by subjects and non-significant by items: $t_1(62)=1.954$, $p<0.06$; $t_2(15)=1.285$, $p<0.22$). As predicted, the proportion of Elaborations following NP1 prompts differs significantly from the proportion following NP2 prompts ($t_1(62)=4.979$, $p<0.001$; $t_2(16)=2.846$, $p<0.02$).




Figure 4.2: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Explanations following IC-2 verbs in Experiment V

Also as predicted for IC-2 verbs, the proportion of NP2-biased Explanations is found to increase in precisely the conditions in which the proportion of NP1-biased Elaboration relations decreases: NP2-referring pronoun prompts yield the largest proportion of Explanation relations (78.1%), followed by no-pronoun prompts (71.6%), and then NP1 prompts (41.7%; effect of prompt type: $F_1(2,62)=11.643$, $p<0.001$; $F_2(2,15)=6.13$, $p<0.02$). In the pairwise comparisons, the proportion of Explanations following NP2 prompts is numerically greater than the proportion following no-pronoun prompts (not significant by subjects or items: $t_1(62)=0.839$, $p<0.41$;

$t_2(15)=1.340$, $p<0.21$), and the proportion of Explanations following no-pronoun prompts is in turn greater than the proportion following NP1 prompts (marginal by items: $t_1(62)=3.500$, $p<0.001$; $t_2(15)=1.813$, $p<0.09$). As predicted, the proportion of Explanations following NP2 prompts differs significantly from the proportion following NP1 prompts ($t_1(62)=4.505$, $p<0.001$; $t_2(16)=3.233$, $p<0.006$).

In contexts with non-IC verbs (e.g., *John chatted with Mary. He/She/∅...*), the effect of prompt type on the distribution of coherence relations also exhibits the predicted pattern. The results are shown in Figure 4.3.

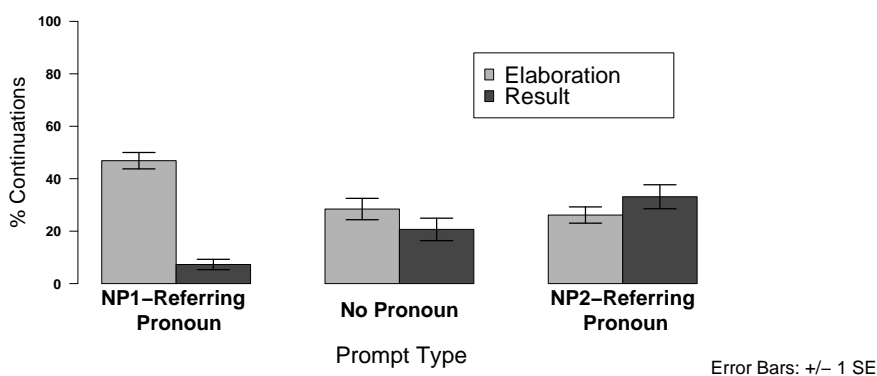


Figure 4.3: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following non-IC verbs in Experiment V

NP1-referring pronoun prompts are found to yield the largest proportion of Elaboration relations (46.9%), followed by no-pronoun prompts (28.4%), and then NP2 prompts (26.1%; effect of prompt type: marginal by items: $F_1(2,62)=11.304$, $p<0.001$; $F_2(2,34)=2.898$, $p<0.07$). In the pairwise comparisons, the proportion of Elaborations following NP1 prompts is greater than the proportion following no-pronoun prompts (marginal by items: $t_1(62)=3.591$, $p<0.001$; $t_2(34)=1.915$, $p<0.07$), which is in turn numerically greater than the proportion following NP2 prompts (not significant by subjects or items: $t_1(62)=0.448$, $p<0.66$; $t_2(34)=0.452$, $p<0.66$). As predicted, the proportion of Elaborations following NP1 prompts differs significantly from the proportion following NP2 prompts ($t_1(62)=4.706$, $p<0.001$; $t_2(34)=2.463$, $p<0.02$).

Also as predicted for non-IC verbs, the proportion of NP2-biased Result relations is found to increase in conditions in which the proportion of NP1-biased Elaborations decreases: NP2-referring pronoun prompts yield the largest proportion of Result relations (33.1%), followed by no-pronoun prompts (20.7%), and then NP1 prompts (7.3%; effect of prompt type: $F_1(2,62)=10.720$, $p<0.001$; $F_2(2,34)=7.368$, $p<0.003$). In the pairwise comparisons, the proportion of Results following NP2 prompts is numerically greater than the proportion following no-pronoun prompts (marginal by items and non-significant by subjects: $t_1(62)=1.989$, $p<0.06$; $t_2(34)=1.240$, $p<0.23$) which is in turn greater than the proportion following NP1 prompts ($t_1(62)=2.843$, $p<0.007$; $t_2(34)=2.218$, $p<0.04$). As predicted, the proportion of Results following NP2 prompts differs significantly from the proportion following NP1 prompts ($t_1(62)=5.191$, $p<0.001$; $t_2(34)=3.146$, $p<0.004$).

4.1.3 Discussion

The results presented here, taken together with the results from Chapter 3, support the development of a model in which coherence and coreference exert bidirectional influences on each other. Whereas the experiments in Chapter 3 show that coherence impacts coreference, Experiment V shows that comprehenders use information about which referent has been mentioned next to update their expectations about the operative coherence relation. In passages with IC-1, IC-2, and non-IC verbs, the presence of an NP1-referring pronoun prompt shifted the distribution of elicited coherence relations in favor of NP1-biased relations; the presence of an NP2-referring pronoun prompt shifted the distribution in favor of NP2-biased relations.

Across the three verb types, the same coherence relations are not necessarily biased to the same referents, but in all cases the presence of an unambiguous pronoun shifts the distribution of coherence relations to accommodate the enforced coreference. Explanation relations are NP1-biased following IC-1 verbs and NP2-biased following IC-2 verbs, whereas Elaborations are consistently NP1-biased and

Results are consistently NP2-biased. For each verb type, the analysis focused on one pair of coherence relations with opposing biases. As predicted, the NP1-referring prompts yielded more NP1-biased coherence relations, whereas the NP2-referring prompts yielded more NP2-biased relations.

Across verb types, the distribution of coherence relations in the no-pronoun condition generally differed from that in the pronoun-prompt conditions. This effect lends support to a model of processing in which comprehenders update their expectations incrementally. An incremental model would account for the fact that comprehenders appear to use the additional information provided by the presence of an unambiguous pronoun to immediately revise their earlier expectations that were based only on the information available in the context sentence. Rather than waiting until the full propositions expressed by the two sentences are available, comprehenders appear, at least in the off-line results presented here, to start generating expectations and revising those expectations as soon as they have information about where the story may be going next. In this light, the unambiguous pronoun prompt is informative because it makes certain coherence outcomes more likely while reducing the probability of other outcomes.

At first blush, it may seem that only unambiguous pronouns would affect the coherence distribution. This is because unambiguous pronouns crucially provide new information about who has been mentioned next, whereas ambiguous pronouns appear to provide no new information. In contexts with fully ambiguous pronouns, it seems that comprehenders would simply retain the expectations they had from the context sentence and accommodate the pronominal referring expression as they write their continuation. However, results from experiments investigating speakers' choice of referring expression and their choice of next mention in no-pronoun and pronoun-prompt conditions suggest that pronouns are not completely neutral in this respect (Stevenson et al., 1994; Arnold, 2001). This previous work suggests that speakers are more likely to use a pronoun than a proper name to re-mention a referent that

occupied the subject position of the previous sentence. If ambiguous pronouns carry their own coreference biases, thereby distinguishing themselves from more general next-mention biases, then the presence of an ambiguous pronoun could also be ‘informative’ to a comprehender.

The continuations elicited here for Experiment V conveniently offer a glimpse into participants’ use of pronouns. I describe those results below and briefly note the implications that will be explored in later sections of this chapter.

4.1.4 Pronominalization in IC Contexts

The continuations elicited with no-pronoun prompts as in (36a), (37a), and (38a), repeated as (39)-(41), reveal which referring expressions participants prefer to use for re-mentioning certain individuals.

(39) John infuriated Mary. _____. [IC-1, no-pronoun prompt]

(40) John scolded Mary. _____. [IC-2, no-pronoun prompt]

(41) John chatted with Mary. _____. [non-IC, no-pronoun prompt]

The continuations allow us to address the question of whether or not choice of referring expression is tied to the same factors that govern choice of next mention: Are participants more likely to use a pronoun when they re-mention the preferred referent for a particular verb type or is pronominalization sensitive to other factors, such as referent position, as has been suggested in previous work (see Section 4.5 for review of this work in the context of the results of this chapter and Chapter 3)?

To test this, the judges annotated the no-pronoun continuations for the participants’ choice of referring expression. Analyses of variance were conducted on the choice of referring expression in the no-pronoun condition in order to test whether referent position (subject/non-subject) and verb type (IC-1, IC-2, non-IC) affect the choice of referring expression. The choice of referring expression was analyzed as a

binary outcome (pronoun or name). Because these analyses involve examining proportions, an arcsine transformation was first applied to the percentages of referring expressions before carrying out analyses of variance.




Figure 4.4: Effects of verb type and referent position on the the rate of pronominalization in Experiment V

As can be seen in Figure 4.4, references to the preceding subject were pronominalized more often (98.4% of the time) compared to references to the non-subject (77.3%; effect of referent position on proportion of pronouns produced: $F(1,27)=13.721$, $p<0.001$; $F_2(1,26)=18.474$, $p<0.001$). The rate of pronominalization also differed by verb type: IC-1 verbs yielded a higher rate of pronominalization overall (91.7%) than IC-2 verbs (85.4%) and non-IC verbs (75.9%; effect of verb type in a model in which referent position and verb type are factors, marginal by subjects: $F_1(2,57)=3.144$, $p<0.06$; $F_2(2,30)=3.816$, $p<0.04$). There was no referent-position \times verb-type interaction ($F_1<1$, $F_2(2,26)=1.884$, $p<0.18$).³

As the lack of interaction shows, the rate at which subject and non-subject referents are pronominalized does not depend on the next-mention bias of the verb

³For the effect of verb type, note that the non-subject argument of the non-IC items often appeared as the object of a non-argument prepositional phrase. By restricting the analysis to only IC-1 and IC-2 verbs, for which the subject and non-subject referents both appear as main arguments of the verb, the effect of verb type disappears ($F_1(1,22)=2.587$, $p<0.13$; $F_2(1,14)=3.245$, $p<0.10$), but the effect of referent position is still significant ($F_1(1,25)=5.742$, $p<0.03$; $F_2(1,14)=6.528$, $p<0.03$), and there is still no referent position \times verb type interaction ($F_s<1$).

type: IC-1 verbs, which are biased to the subject for next mention, create a slight numerical bias to pronominalize the subject compared to IC-2 verbs (rate of pronominalization of subjects following IC-1: 100%; IC-2: 97.0%), but IC-1 verbs also create a slightly stronger bias to pronominalize the non-subject compared to IC-2 verbs (rate of pronominalization of non-subjects following IC-1: 84.8%; IC-2: 82.5%).

To summarize, the overall rate of pronominalization was quite high, but references to the subject were nonetheless pronominalized more often than references to the non-subject. This result is interesting given that the subject referent was not always the preferred referent for next mention. For example, following context sentences with IC-2 verbs, participants showed a bias to re-mention the non-subject NP2 referent; however, participants used far more pronouns to refer back to NP1, the dispreferred subject, whereas more proper names were used to refer back to NP2 than to NP1.

The implications, which will be explored more fully in Section 4.5, are two fold. First, these results point to an asymmetry: Although pronoun interpretation is believed to reflect comprehenders' expectations about which referent will be mentioned next, the pronominalization results shown here indicate that the referent which is most likely to be mentioned next is not necessarily the one that will be mentioned with a pronoun. This would suggest that biases for pronoun interpretation are sensitive to other factors beyond biases for next mention. This idea will be taken up again in Section 4.5 when I revisit the pronoun interpretation results from the previous chapter in light of the findings from this chapter. Related work that provides a similar conclusion will be reviewed then (Stevenson et al., 1994; Arnold, 2001), along with speculation as to why referent position might matter for pronominalization.

Secondly, for the bidirectional influences that are the focus of this chapter, the fact that pronouns are linked to the subject in the form of a production bias gives rise to a novel prediction: If comprehenders are aware of the production bias in pronominalization, then a pronoun prompt, even a fully ambiguous one, will bias a

comprehender toward the subject of the preceding sentence and thus the continuation is therefore more likely to embody a subject-biased coherence relation. In other words, even the presence of a fully ambiguous pronoun prompt could influence the distribution of coherence relations. This prediction is tested in Experiment VI.

4.2 Experiment VI: Ambiguous IC Contexts

The stimuli used in Experiment VI were adapted from the stimuli from Experiment V, changing the context sentences so that they contain two same-gender referents and varying the prompt type between a no-pronoun prompt and an ambiguous-pronoun prompt, as in (42)-(44). As in Experiment V, the analysis for each verb type is restricted to a pair of coherence relations with opposing next-mention biases: Explanations/Results for IC-1 verbs, Elaborations/Explanations for IC-2 verbs, and Elaborations/Results for non-IC verbs. Compared with the results from Experiment V, the prediction is that the ambiguous pronoun prompt here will yield similar (but reduced) coherence biases: More NP1-biased coherence relations will be elicited with the NP1-biased ambiguous pronoun prompt (akin to the NP1-referring unambiguous pronoun prompt in Experiment V) than with the no-pronoun prompt.

(42) John infuriated Bob. (He) _____. [IC-1]

(43) John scolded Bob. (He) _____. [IC-2]

(44) John chatted with Bob. (He) _____. [non-IC]

As in Experiment V, pairs of oppositely-biased coherence relations will be analyzed here to test whether the presence of a pronoun changes the distribution of coherence relations. For IC-1 contexts like (42), the pronoun prompt is predicted to yield more Explanations and fewer Results than the no-pronoun prompt. For IC-2 contexts like (43), the pronoun prompt is predicted to yield more Elaborations and

fewer Explanations. For non-IC contexts like (44), the pronoun prompt is predicted to yield more Elaborations and fewer Results. These predictions only hold if, even in gender-ambiguous contexts, references to the subject are pronominalized at a higher rate than references to the non-subject. If this is the case, participants' no-pronoun continuations are predicted to show evidence of this bias—more pronominal referring expressions when the participant re-mentions NP1 than NP2.

4.2.1 Methods

The story continuation task used here follows the procedure from Experiment V. The only difference is that the continuation prompts vary between no-pronoun prompts and ambiguous-pronoun prompts.

Participants

Twenty-eight monolingual English-speaking undergraduates from UC San Diego participated in the experiment either for extra credit in Linguistics courses or for the chance to be entered in a raffle to win a gift certificate.

Materials and Procedures

For the experimental items, the IC verbs and non-IC verbs were taken from the original IC experiment in Chapter 3, Experiment IV. Context sentences for the experimental items mentioned two referents, either both male or both female, with gender balanced across items. Participants saw each verb only once, in one of two conditions: a no-pronoun prompt or a pronoun prompt. The stimuli is provided in the Appendix. The experiment consisted of one hundred items: eighty experimental items (40 IC, 40 non-IC) intermixed with twenty non-IC fillers. The additional fillers consisted of prompts with context sentences containing non-IC verbs followed by intersentential connectives, no-pronoun or pronoun prompts. Story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which

participants were instructed to write their continuation.

Evaluation and Analysis

As in Experiment V, two judges, the author of this dissertation and a UCSD Linguistics undergraduate, annotated the participants' continuations. For all conditions, they assessed the coherence relation that could be inferred to hold between the context sentence and the continuation. They also assessed the participant's choice of next mention. Pronouns were disambiguated as NP1-referring or NP2-referring in light of the context sentence and the elicited continuation. For the no-pronoun condition, they also annotated the participants' choice of referring expression. Each verb type is tested separately for an effect of one within-subjects/within-items factor, prompt type, on the relevant coherence outcome. The no-pronoun continuations are also tested for effects of referent position and verb type on the rate of pronominalization.

4.2.2 Results

The analysis reported here represents a conservative analysis in which a continuation was excluded if at least one coder assessed the coherence relation or the intended referent of the pronoun to be ambiguous (15.7% of the total number of continuations). For continuations elicited in the no-pronoun condition, the analysis also excludes continuations in which the participant failed to mention one of the referents in the context sentence or used a conjoined noun phrase (*John and Bob*), a plural pronoun (*they*), or a possessive (*his*) (16.6% of the total). The remaining dataset contained 1516 continuations.

Replication of Previous Coherence-Driven Biases

The results for the no-pronoun condition replicate the original coherence and next-mention biases observed in Experiment IV and V. Context sentences with IC-1 verbs yield an overall NP1 bias (73.0% NP1 continuations), but a set of stronger biases emerges once coherence is conditioned on: Explanation relations following IC-1 verbs contain significantly more NP1 references (94.8%) compared with Result relations (9.4%: $t_1(40)=14.638$, $p<0.001$; $t_2(31)=16.422$, $p<0.001$). In contrast, context sentences with IC-2 verbs yield an overall NP2 bias (73.6% NP2 continuations), but Explanation relations contain more NP2 references (80.1%) than Elaborations (39.8%: $t_1(44)=3.605$, $p<0.001$; $t_2(30)=4.270$, $p<0.001$). Finally, context sentences with non-IC verbs yield a small overall NP1 bias (54.3% NP1 continuations), but Result relations following non-IC verbs contain more NP2 references than Elaboration relations (79.2% NP2 continuations vs. 47.8%: marginal by items: $t_1(45)=3.304$, $p<0.002$; $t_2(47)=1.826$, $p<0.08$).

As before, Explanation continuations were more frequent following IC-1 and IC-2 verbs (54.3%) compared to non-IC verbs (20.8%; $F_1(1,27)=141.12$, $p<0.001$; $F_2(1,79)=35.078$, $p<0.001$).

Replication of Pronominalization Biases

The pattern of pronominalization replicates that observed in Experiment V. As shown in Figure 4.5, references to the preceding subject were pronominalized significantly more often (79.4% of the time) than references to the non-subject (23.7%; effect of referent position on proportion of pronouns produced: $F(1,23)=82.864$, $p<0.001$; $F_2(1,66)=188.620$, $p<0.001$). The rate of pronominalization differed by verb type: IC-1 verbs yielded more pronouns (62.3%) than IC-2 verbs (37.8%) and non-IC verbs (52.1%, but this effect is not significant in an analysis in which referent position is already a factor because the effects of verb type and referent position are highly correlated: $F_1<1$; $F_2(2,74)=1.661$, $p<0.20$). There was no referent-position \times verb-type

interaction ($F_1(2,34)=1.521$, $p<0.24$; $F_2<1$).⁴

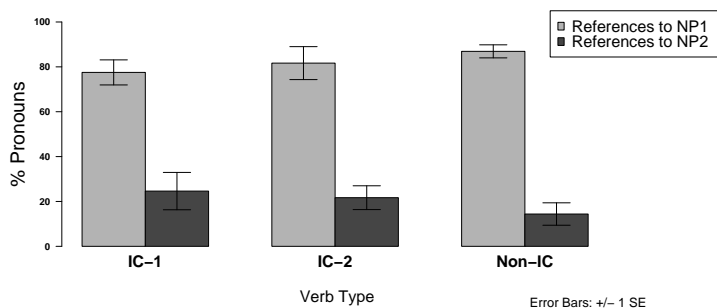


Figure 4.5: Effects of verb type and referent position on the rate of pronominalization in Experiment VI

As the lack of interaction shows, the rate at which subject and non-subject referents are pronominalized does not depend on the next-mention bias of the verb type. For example, IC-1 verbs, which are biased to the subject for next mention, do not create a stronger bias to pronominalize the subject than IC-2 verbs (rates of subject pronominalization: IC-1: 77.5%; IC-2: 80.8%; non-IC: 85.0%). Similarly, IC-2 verbs, which are biased to the non-subject, do not create a stronger bias to pronominalize the non-subject than IC-1 verbs (rate of non-subject pronominalization: IC-1: 26.6%; IC-2: 21.7%; 16.5%).

Effects of Prompt Type on Coherence Distribution

In contexts with IC-1 verbs (*John infuriated Bob. (He) ...*), the effect of prompt type on the distribution of coherence relations exhibits the predicted pattern. The results can be seen in Figure 4.6a. The proportion of Explanation relations is greater following pronoun prompts (48.7%) than no-pronoun prompts (36.4%; $F_1(1,27)=7.405$, $p<0.02$; $F_2(1,19)=13.120$, $p<0.002$). The proportion of Explanations goes up at the expense of Result relations: The proportion of Result relations is smaller following pronoun prompts (4.4%) than no-pronoun prompts (18.0%;

⁴In an analysis of just the IC contexts, references to the preceding subject were still pronominalized significantly more often (76.9% of the time) than references to the non-subject (25.4%; effect of referent position: $F(1,25)=62.278$, $p<0.001$; $F_2(1,35)=86.228$, $p<0.001$). The effect of verb type was still not significant ($F_1<1$; $F_2(1,36)=2.406$, $p<0.13$), nor was the referent-position \times verb-type interaction ($F_s<1$).

$F_1(1,27)=16.322$, $p<0.001$; $F_2(1,19)=17.995$, $p<0001$). These patterns are consistent with the distribution of coherence relations following NP1-referring and no-pronoun prompts in the IC-1 condition of Experiment V; the similarity can be seen by comparing Figures 4.6a and 4.6b.

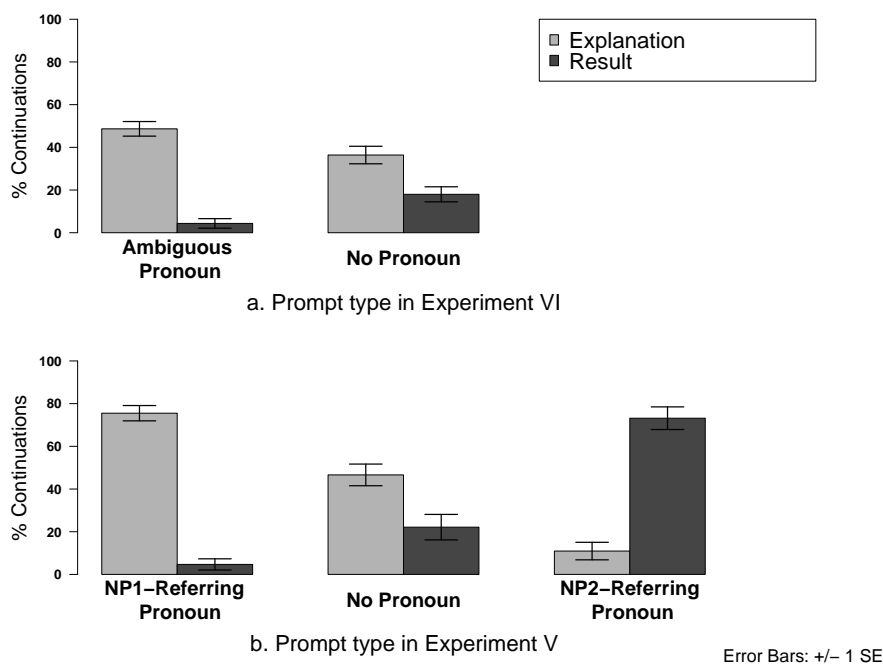


Figure 4.6: Effects of prompt type on proportions of NP1-biased Explanations and NP2-biased Results following IC-1 verbs in Experiments V & VI

In contexts with IC-2 verbs (*John scolded Bob. (He) ...*), the effect of prompt type on the distribution of coherence relations also exhibits the predicted pattern. The results are shown in Figure 4.7a. The proportion of Elaboration relations is greater following pronoun prompts (33.3%) than no-pronoun prompts (14.6%; $F_1(1,27)=13.670$, $p<0.001$; $F_2(1,19)=24.407$, $p<0.001$). The proportion of Elaborations goes up at the expense of Explanation relations: The proportion of Explanation relations is smaller following pronoun prompts (58.2%) than no-pronoun prompts (72.6%; $F_1(1,27)=6.202$, $p<0.02$; $F_2(1,19)=16.649$, $p<0.001$). These patterns are consistent with the distribution of coherence relations following NP1-referring and no-pronoun prompts in the IC-2 condition of Experiment V. The similarity can be

seen by comparing Figures 4.7a and 4.7b.




Figure 4.7: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following IC-2 verbs in Experiments V & VI

Finally, in contexts with non-IC verbs (e.g., *John chatted with Bob. (He) ...*), the distribution of coherence relations exhibits the predicted pattern numerically, but the results fail to reach significance. The results are shown in Figure 4.8a. The proportion of Elaboration relations is greater following pronoun prompts (48.0%) than no-pronoun prompts (42.5%, not significant by subjects or by items: $F_1(1,27)=1.805$, $p<0.20$; $F_2<1$). The proportion of Elaborations goes up at the expense of Result relations: The proportion of Result relations is smaller following pronoun prompts (12.1%) than no-pronoun prompts (17.0%, not significant: $F_1(1,27)=1.227$, $p<0.28$; $F_2<1$). The numerical patterns are consistent with the distribution of coherence relations for NP1-referring and no-pronoun prompts in the non-IC condition of Experiment V, and the similarity can be seen in Figure 4.8.




Figure 4.8: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Results following IC-2 verbs in Experiments V & VI

4.2.3 Discussion

The results for choice of referring expression in the no-pronoun condition replicate the no-pronoun results from Experiment V: Participants were more likely to pronominalize a reference to the subject than a reference to the non-subject. They were also slightly less likely to use a pronoun at all following an IC-2 verb than an IC-1 verb, but crucially the rate of pronominalization of subjects and non-subjects did not differ by verb type. This lack of an interaction between referent position and verb type reaffirms the claim from the previous section that the choice of referring expression can be disassociated from the choice of next mention.

As the coherence biases across the no-pronoun and pronoun-prompt conditions show, participants appear to incorporate these pronominalization biases into their expectations regarding upcoming coherence relations. As predicted, the mere presence of an ambiguous pronoun shifts the distribution of elicited coherence relations in favor

of subject-biased relations. The effects were stronger for the IC verbs than the non-IC verbs, and this may be due to the variety of biases introduced by non-IC verbs — only in aggregate do they comprise a “neutral” non-IC category.

Whereas the unambiguous pronouns in Experiment V provided a deterministic cue regarding which referent had been mentioned, the ambiguous pronouns in this experiment show that probabilistic cues to next mention can also influence expectations about upcoming coherence relations in much the same way that deterministic information can. The next two sections consider probabilistic cues to next mention in a dataset that was the focus of several experiments in Chapter 3.

4.3 Experiment VII: Pronouns in TOP Contexts

Transfer-of-possession (TOP) passages were the focus of the first three experiments in Chapter 3, and the analysis that was proposed in that chapter centered around coherence-driven effects for pronoun interpretation. The experiments used story continuation prompts like those in (45).

(45) John handed a book to Bob. He _____.

The context sentence in (45) is used to describe a transfer event; the Source referent appears in subject position, and the Goal referent appears as the object of a prepositional phrase. The set of elicited continuations showed that different patterns of pronoun interpretation were associated with different coherence relations, and further experimentation showed that shifting the distribution of coherence yielded a shift in the overall pattern of pronoun interpretation.

In order for bidirectional effects to emerge in contexts like (45) with a TOP verb and an ambiguous pronoun prompt, two conditions must be met: (i) There must be coherence relations with opposing next-mention biases, and (ii) the preference that was observed in IC contexts for pronominalizing references to the subject must hold for

TOP contexts as well. Condition (i) was established in Chapter 3: Elaborations and Explanations were shown to be biased to the subject Source whereas Occasions and Results were shown to be biased to the non-subject Goal. For condition (ii), previous work has found that references to the subject in TOP contexts are pronominalized at a higher rate than references to the non-subject (Stevenson et al., 1994; Arnold, 2001). This previous work on pronominalization biases in TOP contexts will be discussed more fully in Section 4.5 when the results of this chapter are compared with the results of Chapter 3.

To confirm the pronominalization bias, this section presents Experiment VII, a story continuation experiment with no-pronoun prompts following TOP context sentences, as in (46).

(46) John handed a book to Mary. _____.

The story continuation experiment was conducted as a norming study for an ERP (event-related brain potentials) experiment (Ferretti, Rohde, Kehler, & Crutchley, submitted). The ERP experiment tested the time course of the event-structure biases that had been observed in Experiment I, but the norming data constitutes a set of continuations that were elicited with no-pronoun prompts—continuations which can be analyzed here for participants' pronominalization preferences. Like Experiment I, the norming study and the ERP experiment used stimuli that described either completed events (perfective aspect) or ongoing events (imperfective aspect), as in (47)-(49).

(47) Experiment I stimuli

- a. John passed_{PERF} a book to Bob. He _____
- b. John was passing_{IMP} a book to Bob. He _____

(48) ERP norming stimuli

- a. Sue handed_{PERF} a timecard to Fred. _____
- b. Sue was handing_{IMP} a timecard to Fred. _____

(49) ERP stimuli

- a. Sue handed_{PERF} a timecard to Fred. She asked about the upcoming meeting.
- b. Sue was handing_{IMP} a timecard to Fred. She asked about the upcoming meeting.

Experiment I (stimuli as in (47)) showed that comprehenders exhibit a bias toward the referent associated with the end state, the Goal, in cases in which the context sentence describes a completed transfer event rather than an ongoing transfer event. These event-structural biases were the focus of the ERP work. Whereas Experiment I used ambiguous pronoun prompts as in (47) to probe participants' pronoun interpretation preferences, the ERP norming experiment used no-pronoun prompts as in (48), and the ERP experiment used passages with unambiguous pronouns as in (49). The unambiguous pronouns in the ERP experiment allowed Ferretti et al. to measure participants' response when they encountered a mention of a particular referent.⁵ The continuations elicited as part of the norming study are useful for the purposes of this chapter because they show which referring expression participants prefer to use to re-mention particular referents. If the pattern from Experiment V and VI with IC verbs holds up in other contexts, the prediction is that participants will be more likely to use a pronoun to re-mention the subject referent than the non-subject referent. If the event structure results of Experiment I are replicated here, the prediction is that participants will produce more Goal continuations following context sentences about completed events than following ongoing events, but crucially this bias to the Goal is not predicted to impact participants' preference to use pronouns to refer back to the subject Source.

⁵Briefly, the results confirmed the off-line effects from Chapter 3: Comprehenders showed less difficulty integrating a subsequent reference to the non-subject Goal following context sentences that described completed events (perfective aspect) compared to context sentences that described ongoing events (imperfective aspect). The results are noted here in part to clarify the context in which Experiment VII was conducted and in part to show that the off-line coreference factors described in Chapter 3 have an impact on comprehenders' on-line processing. The ERP results showed that, for both verbal aspects, Source-referring pronouns yielded greater anterior negativity than Goal-referring pronouns. However, the effect came earlier with perfective context sentences and was followed by a P600 not seen with imperfective context sentences. The ERP experiment did not explicitly manipulate coherence, but the results are consistent with a coherence-driven model.

4.3.1 Methods

Participants

Fifty-four monolingual English-speaking undergraduates from Wilfrid Laurier University participated in the experiment for credit in Psychology courses.

Materials and Procedures

The stimuli consisted of 72 target items and 72 fillers like those in (48). Target items contained a context sentence with a Source-Goal TOP verb. The Source and Goal referents were both proper names that differed in gender. Each participant saw half the verbs with perfective aspect and half with imperfective, and no participant saw any verb more than once. Gender was balanced across stimuli. The 72 target items were randomly mixed with 72 fillers which described non-transfer-of-possession events involving one or two individuals that were mentioned with proper names. Half of the filler verbs were perfective and half were imperfective. The stimulus set is given in the Appendix.

Story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their continuation.

Evaluation and Analysis

The elicited story continuations were coded for several factors: choice of next mention (Source or Goal), referring expression of next mention (name or pronoun), and the position of next mention (first word or not). The choice of next mention and the coherence relation were assessed by one annotator (the author of this dissertation). Pronouns were disambiguated given the genders of the two possible referents. Analyses of variance were conducted to measure the effects of one within-subjects/within-items factor (verbal aspect) on the choice of next mention. Analyses of variance were also conducted to measure the effects of referent position and verbal aspect on the choice of referring expression.

4.3.2 Results

The continuations from 14 participants were eliminated because the participants misunderstood the task (writing only sentence fragments or questions) or did not complete the entire experiment. From the remaining continuations (N=2880), 13% were excluded because the continuation did not mention either the Source or the Goal (or else referred to the Source and Goal together as *they*), along with 3.6% that consisted only of a sentence fragment (e.g., a prepositional phrase or a relative clause), 3.4% that referred to the Source or Goal only with a possessive (*his* or *hers*), and less than 1% that were nonsensical or contained mistakes concerning the interpreted gender of the male/female names. The analysis is restricted to no-pronoun continuations in which either the Source or Goal was mentioned as the first word of the continuation (N=1859) and the subset of those in which the reference to the Source or Goal was pronominalized (N=912).

Replication of Next-Mention Biases

The effect of aspect on next mention can be seen in Figure 4.9. Considering the full dataset, regardless of referring expression (N=1859), Goal references made up a large majority of the responses following both perfective context sentences (85.7%; different from a hypothetical mean of 0.5: $t_1(39)=13.10$, $p<0.001$; $t_2(71)=17.71$, $p<0.001$) and imperfective context sentences (77.0%; $t_1(38)=9.29$, $p<0.001$; $t_2(71)=11.88$, $p<0.001$). However, Goal references were significantly more common following context sentences with perfective verbs than imperfective verbs, which replicates the bias shown in Experiment I (effect of aspect: $F_1(1,38)=18.76$, $p<0.001$; $F_2(1,71)=9.97$, $p<0.003$). The results were similar for the data with pronominal references (N=912): the Goal bias was stronger following perfective sentences (75.7%; $t_1(35)=6.19$, $p<0.001$; $t_2(71)=7.90$, $p<0.001$) than imperfective sentences (57.0%; significant only by items: $t_1(33)=1.38$, $p<0.18$; $t_2(71)=4.64$, $p<0.001$). Goal references were significantly more common following perfective context sentences

than imperfective sentences ($F_1(1,32)=16.51$, $p<0.001$; $F_2(1,71)=6.86$, $p<0.02$).⁶

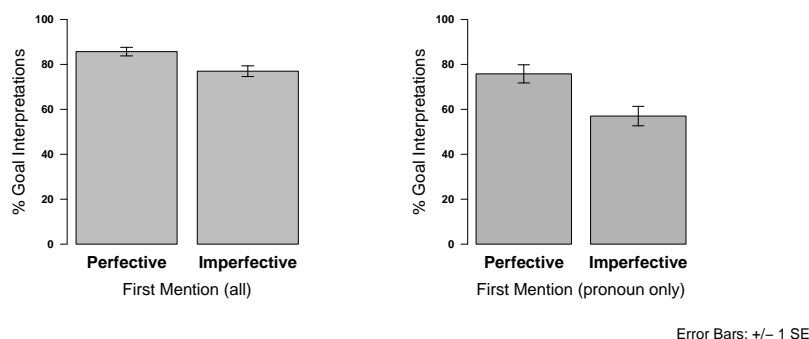


Figure 4.9: Effect of verbal aspect on choice of next mention across all data in Experiment VII and the pronoun-only subset

Effect of Referent Position on Choice of Referring Expression

As predicted, subject mentions were pronominalized at a higher rate (76.2%) than non-subject mentions (52.1%; $F_1(1,35)=52.746$, $p<0.001$; $F_2(1,66)=87.242$, $p<0.001$). The rate of pronominalization did not differ by verbal aspect ($F_1<1$; $F_2(1,70)=1.626$, $p<0.21$) nor was there a referent-position \times aspect interaction ($F_1(1,28)=1.290$, $p<0.27$; $F_2<1$). The results are shown in Figure 4.10.

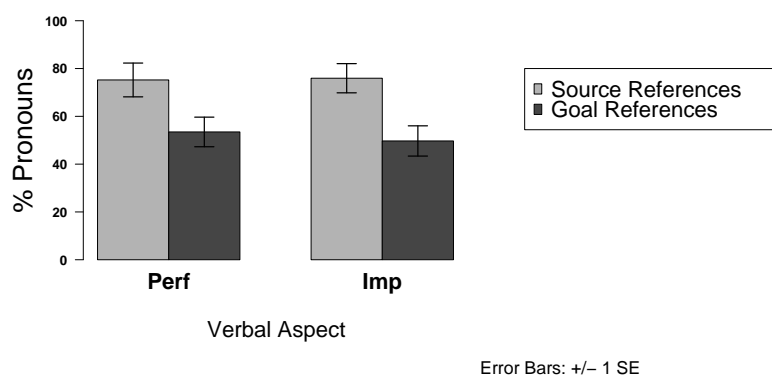


Figure 4.10: Effect of referent position and verbal aspect on the rate of pronominalization in Experiment VII

⁶These numbers differ slightly from those reported in Ferretti et al. (submitted) due to differences in participant inclusion and exclusion.

4.3.3 Discussion

The results for Experiment VII confirm the pattern of pronominalization that was noted for IC contexts and for TOP contexts in earlier work: Subsequent mentions of the subject referent (in this case, the Source) are pronominalized at a higher rate than are mentions of the non-subject referent (in this case, the Goal). In both perfective and imperfective contexts, participants were more likely to re-mention the Goal referent, while at the same time producing more pronouns in the cases in which they re-mentioned the Source. The difference in strength of the Goal next-mention bias between perfective and imperfective conditions did not result in a difference in the percentages of subject-referring pronouns.

The next-mention results are both similar to and different from the results in Experiment I with TOP verbs and ambiguous pronoun prompts. On one hand, the event structure effects align well (Goal continuations are more frequent following perfective context sentences than imperfective context sentences), but on the other hand, there were overall far more Goal continuations in Experiment VII than in Experiment I. Whereas the Goal bias in Experiment I was 42.9% following perfective context sentences and 20.9% following imperfective context sentences, the Goal bias here, when all continuations are considered, was 85.7% following perfectives and 77.0% following imperfectives. This difference may stem in fact from the connection between pronouns and subject referents—the hypothesis would be that the ambiguous pronoun prompt in Experiment I may actually have functioned as a ‘subject-biased’ pronoun prompt. In that case, comprehenders who encounter a pronoun, even a fully ambiguous one, may be more likely to interpret it as coreferential with the subject of the preceding sentence precisely because they are aware that references to the subject are more likely to be produced with a pronoun. For the coherence~coreference bidirectional effects that are the focus of this chapter, the presence of an ambiguous pronoun in a TOP context is hypothesized to bias a comprehender toward a subject-biased coherence relation. Experiment VIII tests this hypothesis.

4.4 Experiment VIII: Ambiguous TOP Contexts

The stimuli used in Experiment VIII were taken directly from Experiment I, with the only addition being the prompt manipulation. Sample items are shown in (50) and (51), with the ambiguous pronoun listed in parentheses to distinguish the pronoun-prompt and no-pronoun conditions.⁷

(50) John handed_{PERF} a book to Bob. (He) ____.

(51) John was handing_{IMP} a book to Bob. (He) ____.

The sentences in (50) and (51) contain two referents, a subject Source (NP1) and a non-subject Goal (NP2). (50) describes a completed event (perfective aspect), whereas (51) describes an ongoing event (imperfective aspect). In the original experiment, context sentences that contained perfective verbs yielded more NP2-biased Occasion and Result relations than context sentences that contained imperfective verbs. Imperfective context sentences, on the other hand, yielded more NP1-biased Elaborations, Explanations, and Violated-Expectations. Across all continuations in Experiment I, Occasions and Elaborations were the two most common coherence relations.

The frequency and opposing next-mention biases of Elaborations and Occasions provide an opportunity to test whether the mere presence of a pronoun, even an ambiguous one, can influence the distribution of coherence relations. NP1-biased Elaborations are predicted to be more frequent in the pronoun prompt condition compared to the no-pronoun condition. NP2-biased Occasions, on the other hand, are predicted to be disfavored following the pronoun prompt compared to the no-pronoun prompt. This prediction only holds if the pronoun production bias favoring the subject, as was observed in unambiguous contexts, is also apparent in gender ambiguous

⁷This prompt manipulation with TOP contexts is in part a replication of Stevenson et al., with the addition of a verb aspect manipulation. Stevenson et al. analyzed the elicited continuations for patterns of next mention. For our purposes, the continuations elicited with the different prompt types allow us to test whether prompt type affects the distribution of coherence relations.

contexts. If this is the case, participants' no-pronoun continuations are predicted to provide evidence of this bias—more pronominal referring expressions when the participant re-mentions the subject than the non-subject.

In considering the factors that influence upcoming coherence relations, note that no interaction is predicted between the verbal aspect and prompt type manipulations. Perfective aspect is predicted to bias participants towards upcoming Occasion relations, whereas imperfective aspect is predicted to bias participants towards upcoming Elaboration relations. No-pronoun prompts are predicted to also yield more Occasions, whereas pronoun prompts are predicted to yield more Elaborations. There is, however, no reason to expect that the two factors would conspire to create an effect that is more than additive.

4.4.1 Methods

Participants

Fifty-two monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

For the experimental items, the context sentences came directly from Experiment I, containing two same-gender referents, a Source and Goal, with gender balanced across items. The full stimulus set is provided in the Appendix. The experiment consisted of sixty-three items: twenty-one experimental items interleaved with twenty-one relative-clause stimuli for an unrelated experiment (Experiment IX) and twenty-one additional fillers. The stimuli for the interleaved experiment are described in Section 5.2 and consisted of a matrix clause followed by the beginning of a relative clause with an ambiguous relative pronoun as the story continuation prompt. The additional fillers consisted of sentences with non-transfer verbs and a variety of prompts (pronouns, proper names, relative pronouns, adverbs).

Story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their continuation.

Evaluation and Analysis

Two judges, the author of this dissertation and a UCSD Linguistics undergraduate, annotated the participants' continuations. For both the no-pronoun and pronoun-prompt conditions, they assessed the coherence relation that could be inferred to hold between the context sentence and the continuation. For pronouns in the prompt condition and the no-pronoun condition, they assessed a Source or Goal interpretation for each pronoun in light of the context sentence and the elicited continuation. For the no-pronoun condition, they annotated the participant's choice of referring expression. Analyses of variance were conducted on the proportions of Elaborations and Occasions to test for effects of the within-subjects/within-items factors of verbal aspect and prompt type.

4.4.2 Results

The analysis reported here represents a conservative analysis in which a continuation was excluded if at least one coder assessed the coherence relation or the intended referent of the pronoun to be ambiguous (10.8% of the total number of continuations). For continuations elicited in the no-pronoun condition, the analysis also excludes continuations in which the participant failed to mention one of the referents in the context sentence or used a conjoined noun phrase (*John and Bob*), a plural pronoun (*they*), or a possessive (*his*) (8.5% of the total number of continuations). The remaining dataset contained 881 continuations.

Replication of Previous Coherence-Driven Biases

I limit the discussion here to the results for the relevant coherence relations: Occasions and Elaborations. The results for the pronoun-prompt condition repli-

cate the original coherence biases observed in Experiment I. For the effect of aspect on the probability of upcoming coherence relations, Occasion relations were more frequent following perfective context sentences (31.2%) compared to imperfective context sentences (21.8%; $F_1(1,50)=10.002$; $p<0.003$; $F_2(1,20)=7.581$; $p<0.02$). Elaborations, on the other hand, were more frequent following imperfective context sentences (19.1%) compared to perfective context sentences (13.2%; $F_1(1,50)=6.469$; $p<0.02$; $F_2(1,20)=5.749$; $p<0.03$). In keeping with the next-mention biases observed in the original experiment, Occasion relations were biased to NP2 (82.0%; significantly different from a hypothetical mean of 0.5: $t_1(37)=7.055$, $p<0.001$; $t_2(19)=5.545$, $p<0.001$), whereas Elaborations were biased to NP1 (98.5%; $t_1(38)=45.985$, $p<0.001$; $t_2(18)=59.578$, $p<0.001$).

Replication of Pronominalization Biases

The pattern of pronominalization replicates that observed in Experiment VII and can be seen in Figure 4.11. References to the preceding subject were pronominalized significantly more often (77.8% of the time) than references to the non-subject (11.1%; effect of referent position: $F(1,36)=99.877$, $p<0.001$; $F_2(1,19)=163.091$, $p<0.001$). For the effect of aspect, there were numerically more pronominal references following imperfective context sentences (26.7%) compared to perfective context sentences (24.3%; marginal by items and not significant by subjects: $F_1(1,49)=2.062$, $p<0.16$; $F_2(1,19)=3.457$, $p<0.08$). There was no referent position \times verb type interaction ($F_s<1$).

Effects of Prompt Type on Coherence Distribution

As predicted, both the proportion of Occasion relations and the proportion of Elaboration relations differed by prompt type: Elaborations were more frequent in the pronoun-prompt condition (21.7%) than the no-pronoun condition (9.1%; $F_1(1,50)=23.889$; $p<0.001$; $F_2(1,20)=30.412$; $p<0.001$), whereas Occasions were more frequent in the no-pronoun condition (31.9%) than the pronoun-prompt condition (22.4%; $F_1(1,50)=8.411$; $p<0.006$; $F_2(1,20)=11.007$; $p<0.004$). The effect of aspect

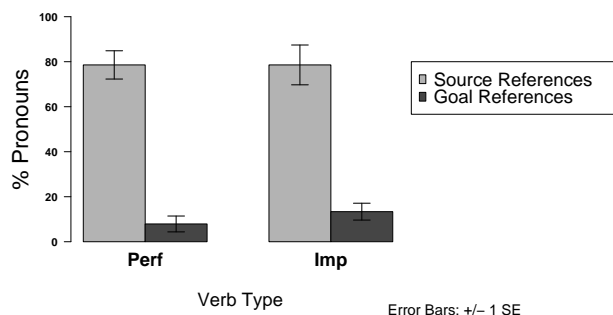


Figure 4.11: Effect of verb aspect and referent position on the rate of pronominalization in Experiment VIII

on the distribution of coherence relations was noted above, and there were no prompt-type \times aspect interactions for the proportion of Elaborations or the proportion of Occasions ($F_s < 1$).

Figure 4.12 shows the perfective results. In contexts with perfective transfer verbs, NP2-biased Occasions were more frequent than NP1-biased Elaborations. As predicted, though, the bias to Occasions is stronger in the no-pronoun condition (36.0% Occasion vs. 6.0% Elaboration) than in the pronoun-prompt condition (27.8% Occasion vs. 20.4% Elaboration).




Figure 4.12: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Occasions following perfective verbs in Experiments VIII

Figure 4.13 shows the imperfective results: In contexts with imperfective transfer verbs, the frequency of NP1-biased Elaborations and NP2-biased Occasions varies

by prompt type: Elaborations are favored over Occasions in the pronoun condition (24.0% Elaboration vs. 17.4% Occasions), whereas Occasions are favored over Elaborations in the no-pronoun condition (26.8% Occasion vs. 11.9% Elaboration), as predicted.

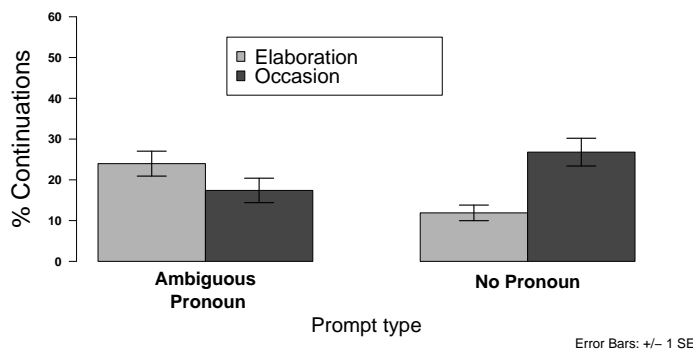


Figure 4.13: Effects of prompt type on proportions of NP1-biased Elaborations and NP2-biased Occasions following imperfective verbs in Experiments VIII

4.4.3 Discussion

The results for choice of referring expression in the no-pronoun condition are in keeping with the no-pronoun results in Experiments V and VI with IC verbs and VII with TOP verbs: Participants were more likely to pronominalize a reference to the subject than a reference to the non-subject. The rate of pronominalization for the non-subject, however, was lower than the rates reported for the IC contexts. This may be due in part to the grammatical position of the non-subject. In these transfer-of-possession contexts, NP2 appears as the object of a prepositional phrase, whereas in the implicit-causality contexts in Experiment VII, NP2 usually appears in the direct object position. It is possible that the rate of pronominalization differs depending on the grammatical position of the verbal argument, yielding fewer pronominal references even for a required argument like a transfer verb’s Goal thematic role. What is crucial, though, is that the NP2 next-mention bias does not create a corresponding bias to pronominalize NP2, even though NP2 is the most likely (or, in Arnold’s (2001)

terms, the most ‘accessible’ or ‘expected’) referent. The rate of pronominalization observed here is even lower than the rate reported in Experiment VII for the other TOP contexts. This may be due to the ambiguity that a pronoun creates in these contexts: The contexts here contain two same-gendered referents so that a pronoun is necessarily ambiguous, whereas in Experiment VII, the referents differed in gender.

The effects of prompt type on coherence show that participants appear to incorporate these pronominalization biases as they generate expectations regarding upcoming coherence relations. The result is in keeping with the results for ambiguous IC contexts that probabilistic information about which individual has been mentioned again can influence coherence expectations. It extends the previous result by establishing that these coherence-coreference influences are not limited to implicit causality contexts.

4.5 TOP Coreference Revisited

Recall that Experiment VII, a story-continuation experiment with no-pronoun prompts, confirmed that the bias to pronominalize subject next mentions persists in TOP contexts. The experimental set-up was similar to that for Experiment I in the previous chapter: Both experiments used TOP verbs and manipulated the verbal aspect of the TOP verb; however, the stimuli for Experiment I contained same-gendered referents and pronoun prompts (as in (52)), whereas the stimuli for Experiment VII contained opposite-gendered Source and Goal referents and no-pronoun prompts (as in (53)).

(52) John passed_{PERF}/was passing_{IMP} a book to Bob. He _____
[Experiment I stimuli]

(53) Sue handed_{PERF}/was handing_{IMP} a timecard to Fred. _____
[Experiment VII stimuli]

In the discussion of Experiment VII, similarities and differences were noted between the results for the two experiments. On one hand, the verbal aspect manipulation had the same effect in both experiments: Goal continuations were more common following perfective context sentences than imperfective context sentences. On the other hand, there was an overall bias to the non-subject Goal in Experiment VII that was not seen in Experiment I with the pronoun prompt. This section provides an explanation for the different coreference patterns in conditions with no-pronoun and pronoun prompts. The explanation is based on the biases that are introduced by the mere presence of a pronoun, even a fully ambiguous one like that in the Experiment I stimuli.

Previous Work on (Lack of) Pronoun Neutrality

Differences between no-pronoun and pronoun-prompt conditions have been noted in previous work. Stevenson et al. (1994) conducted a series of story-continuation experiments, one of which was the basis for the pronoun interpretation experiments in Chapter 3. Stevenson et al.'s original experiment also included conditions with no-pronoun prompts and context sentences with other verb types. Sample items from their experiments are shown in (54), with the ambiguous pronoun listed in parentheses to distinguish the pronoun-prompt and no-pronoun conditions.⁸

- (54) a. John passed the comic book to Bill. (He) _____. [Source-Goal]
 b. Bill seized the comic book from John. (He) _____. [Goal-Source]
 c. Joseph hit Patrick. (He) _____. [Agent-Patient]
 d. Patrick was hit by Joseph. (He) _____. [Patient-Agent]
 e. Ken admired Geoff. (He) _____. [Experiencer-Stimulus]
 f. Geoff impressed Ken. (He) _____. [Stimulus-Experiencer]

⁸A fourth condition served as a control, with stimuli involving Agent-Source/Goal verbs (*Simon ran towards/away from Richard*).

For the no-pronoun condition, Stevenson et al. report strong next-mention biases: Participants were more likely to write continuations that mentioned the referents appearing in the Goal, Patient, and Stimulus thematic roles in the context sentence than they were to write continuations that mentioned the Source, Agent, or Experiencer referents. The Goal bias in no-pronoun TOP contexts has been replicated by Arnold (2001) and was confirmed in Experiment VII.

Stevenson et al. also report that, in the no-pronoun condition, references to the subject of the preceding sentence were pronominalized more often than references to the non-subject. This was the case across all verb types, regardless of which referent was the preferred referent for next mention. These pronominalization effects have been replicated for the Source-Goal and Goal-Source contexts (Arnold 2001, and Experiments VII and VIII reported here) and for the Experiencer-Stimulus and Stimulus-Experiencer contexts (see Experiments V and VI).

As further evidence of the lack of neutrality of pronominal referring expressions, Stevenson et al. report that the strength of the Goal/Patient/Stimulus bias differed between the no-pronoun condition and the pronoun-prompt condition. The pronoun prompt yielded a stronger bias to the Goal/Patient/Stimulus than was found in the no-pronoun condition, but only for the stimuli in which the Goal, Patient or Stimulus appeared in subject position (54b,d,f). The pronoun prompt yielded a weaker bias to the Goal/Patient/Stimulus than was found in the no-pronoun condition in the stimuli in which the Goal/Patient/Stimulus appeared in the non-subject position (54a,c,e).

The evidence that coreference varies depending on the presence or absence of a pronoun suggests that a model of coreference must decouple the choice of next mention from the choice to pronominalize. Stevenson et al. and Arnold construct pronoun interpretation models to account for this, positing a set of morphosyntactic and thematic-role heuristics that interact to yield the surface patterns. The model proposed here goes further by identifying the types of (coherence-driven) factors that

influence choice of next mention, distinguishing those factors from (grammatical role) factors that affect the choice to pronominalize. A Bayesian model is sketched below that establishes a relationship between the probability of a particular referent being mentioned again, $p(Ref)$, the probability of a pronoun being produced to refer to a particular referent, $p(Pro | Ref)$, and the probability of interpreting a pronoun to refer to a particular referent, $p(Ref | Pro)$. The TOP data from Experiment VIII with both no-pronoun and pronoun prompts make it possible to check whether the Bayesian model correctly captures the relationship between next-mention biases and pronoun-interpretation biases.

Bayesian Model of Pronoun Interpretation/Production

Given a context sentence, comprehenders are posited to generate expectations about the direction the discourse is likely to take and, from that, the probability that a particular referent will be mentioned again ($p(Ref)$). Bayes' rule can be used here to highlight the probabilities that come into play as comprehenders update that prior probability to derive a new probability ($p(Ref | Pro)$) when they encounter new information such as an ambiguous pronoun.

$$p(Ref | Pro) = \frac{p(Pro | Ref)p(Ref)}{p(Pro)} \quad (4.1)$$

In considering each of the terms in equation (4.1), the next-mention bias observed with the no-pronoun prompt corresponds to the prior probability $p(Ref)$. The pronoun production bias observed with the no-pronoun prompt corresponds to the conditional probability $p(Pro | Ref)$. The next-mention bias observed in the ambiguous pronoun-prompt condition corresponds to the posterior probability $p(Ref | Pro)$. Left implicit in (4.1) are those contextual cues that further condition the probabilities involved. One example of such a cue is verbal aspect, as the results from Experiments I, VII, and VIII showed. Equation (4.2) adapts equation (4.1) by conditioning all the probabilities on aspect.

$$p(\textit{Ref} | \textit{Pro}, \textit{Aspect}) = \frac{p(\textit{Pro} | \textit{Ref}, \textit{Aspect})p(\textit{Ref} | \textit{Aspect})}{p(\textit{Pro} | \textit{Aspect})} \quad (4.2)$$

The asymmetry that was noted here for TOP contexts and has been discussed in previous work concerns the pronoun production and interpretation terms. Given a pronoun, comprehenders have been shown to *interpret* it as coreferential with the non-subject Goal referent roughly 50% of the time, yet they choose to *produce* a pronoun far less than 50% of the time when they are re-mentioning the Goal. The asymmetry results from the fact that $p(\textit{Ref}=\textit{Goal} | \textit{Pro})$ cannot be equated directly with $p(\textit{Pro} | \textit{Ref}=\textit{Goal})$, since $p(\textit{Ref}=\textit{Goal} | \textit{Pro})$ depends not only on the probability that a Goal referent was pronominalized $p(\textit{Pro} | \textit{Ref}=\textit{Goal})$, but also on the probability that a Goal was mentioned at all $p(\textit{Ref}=\textit{Goal})$.

In fact, the pronouns in the continuations elicited with the no-pronoun prompt confirm that participants use pronouns in much the same way that they interpret pronouns in the pronoun-prompt condition. That is, when participants encountered a pronoun prompt in a context with a perfective verb, the results of Experiment VIII show that they interpreted it to refer to the Goal roughly half the time: $\hat{p}(\textit{Ref}=\textit{Goal} | \textit{Pro}, \textit{Aspect}=\textit{perf})=.498$, where \hat{p} represents the estimate of the true probability and is derived from observed counts. In the no-pronoun prompt condition, participants used pronouns in much the same way: $\hat{p}(\textit{Ref}=\textit{Goal} | \textit{Pro}, \textit{Aspect}=\textit{perf})=.463$. In the imperfective pronoun-prompt condition, participants were far less likely to interpret a pronoun prompt as coreferential with the Goal: $\hat{p}(\textit{Ref}=\textit{Goal} | \textit{Pro}, \textit{Aspect}=\textit{imp})=.260$. Likewise, participants were unlikely to produce a pronoun in the no-pronoun condition that referred back to the non-subject Goal: $\hat{p}(\textit{Ref}=\textit{Goal} | \textit{Pro}, \textit{Aspect}=\textit{imp})=.209$. These results are captured by the different probabilities of mentioning the Goal in these two conditions ($p(\textit{Ref}=\textit{Goal} | \textit{Aspect}=\textit{imp})$), whereas the rate of pronominalization was shown to be largely unaffected by aspect (no referent \times aspect interaction; see Section 4.4.2).

For this data, one can ask if the aggregate values just discussed are also representative of individual participants' behavior — are they using their knowledge of their behavior in production in their interpretation decisions? To examine this, one can compare $p(Ref|Pro, Aspect)$ in the pronoun-prompt condition and the no-pronoun condition across both verbal aspects at the level of individual participants. Indeed, across the two verbal aspects ($Aspect=\{perf, imp\}$), participants' interpretation and production biases are reliably correlated (adjusted R^2 by subjects and by items are as follows: $R_1^2=0.266$, $F_1(1,100)=37.65$, $p<0.001$; $R_2^2=0.404$, $F_2(1,71)=49.73$, $p<0.001$).⁹

In addition to the work of Stevenson et al. and Arnold and the work presented here, Miltsakaki (2007) reports evidence from a pair of story-continuation experiments in Greek that confirms that the probability of next mention is distinct from the probability of pronominalization. The Greek results show that a salient discourse entity is not necessarily re-mentioned with the most underspecified form (a null subject): In contexts in which the preferred referent for next mention is the non-subject, the most likely referring expression is a less reduced form (a strong pronoun). Without a distinction between the probability of next mention and the probability of pronominalization, the observed pattern of results is difficult to explain. The data presented here along with that of Miltsakaki is particularly troublesome for models that rely on so-called Accessibility hierarchies (Gundel, Hedberg, & Zacharski, 1993) which posit that a speaker's choice to produce a pronoun primarily reflects the relative salience of the antecedent.

The Bayesian approach provides a means of reconciling this apparent asymmetry between participants' pronoun interpretation and production. Under this approach, comprehenders rely on observed pronominalization rates in their input. By keeping track of these statistical regularities, comprehenders are in a position to reverse engineer what the speaker's intended referent may have been by considering two

⁹The correlations were calculated by taking a mean across observations for each of the n subjects (or n items) in the perfective and imperfective conditions for the probabilities in the pronoun-prompt and no-prompt conditions. The tests establish how well correlated the $2n$ means are in the pronoun-prompt condition with the $2n$ means in the no-prompt condition.

contributing factors — the probability of next mention and the probability of pronominalization given next mention. The proposed model eliminates the need for positing a subject assignment strategy in pronoun interpretation (e.g., as per Stevenson et al., 1994) while still explaining the common (but not ubiquitous) surface preference toward interpreting ambiguous pronouns to be coreferential with the subject of the preceding clause. This explanation for the facts becomes apparent only through the separation of next-mention and pronoun-interpretation biases in contexts in which those biases are not aligned.

Acknowledging that subject referents are re-mentioned with a pronoun more often than non-subjects, however, does not explain *why* this might be the case. The experiments presented here were not designed to shed light on that question, but one can speculate about why the entity mentioned in subject position may have a special status. Crosslinguistically, the subject position is associated with old, or topical, information (Foley & Valin, 1984; Givon, 1990; Mithun, 1991; Lambrecht, 1994, *inter alia*), and since pronouns refer to entities previously mentioned in the discourse, they are, by definition, old information. One could imagine that speakers use pronominal referring expressions to signal which referent is most topical (who the story is ‘about’). Comprehenders’ then generate expectations about whether the discourse will continue with the current most topical referent or whether there will be a shift toward a new referent. It is of course possible that the short two-sentence passages used in these experiments failed to provide participants with a sufficiently rich discourse context. Lacking more context, they used cues like subjecthood to extrapolate which referent was the intended topic of the discourse. Their choice to pronominalize a reference to the subject may therefore have signaled that they were continuing the discourse with the referent who they suspected was the central topic of the discourse; their choice to mention the non-subject (often with a name) may have signaled that they were shifting the discourse toward a new topic (see Grosz, Joshi, & Weinstein, 1995).

4.6 Conclusion

The results presented here point to several conclusions. First, the fact that patterns of coreference can influence coherence expectations extends the growing list of factors that appear to influence comprehenders' expectations about upcoming material. Whereas the factors that were shown to matter in Experiments I-IV were limited to properties of the preceding sentence (verb type, verbal aspect, object type) and the prior discourse context (a question posed in the instructions), the results in this chapter show that properties of the current sentence can also influence coherence expectations. The specific property that was shown to matter was the presence of a sentence-initial pronoun, even a fully ambiguous one. Secondly, this result, taken in combination with the results from the previous chapter, suggests that discourse processing is sensitive to bidirectional influences between coherence and coreference. Furthermore, the coreference results turn out to be relevant to a comprehension/production asymmetry that was previously reported in the pronoun literature by Stevenson et al. A Bayesian model was proposed in Section 4.5 to explain this asymmetry. This model crucially distinguishes between factors that influence choice of next mention (verb type and the operative coherence relation) and factors that influence choice of referring expression (grammatical role of the referent).

The results of the three experiments presented here suggest that comprehenders condition their biases about where the discourse is going and who is likely to be mentioned next on the prior context (e.g., the type of verb or the verbal aspect). During language processing, comprehenders then proceed to update their biases incrementally when they encounter a pronoun, unambiguous or not. The question that I take up in the next chapter is whether comprehenders also show evidence of coherence-driven incremental updates when they process ambiguous syntactic constructions. The focus of Chapter 5 is a syntactic ambiguity that emerges in contexts with ambiguously attached relative clauses.

4.7 Acknowledgments

Two of the experiments described in this chapter have been reported in earlier work and are discussed here with permission from the coauthors. The dissertation author was the primary investigator for the material presented here. Experiment VII was reported, in part, in Ferretti, Rohde, Kehler, and Crutchley (Submitted) [Verb Aspect, Event Structure, and Coreferential Processing. *Journal of Memory and Language*. Experiment VIII appeared in Rohde and Kehler (2008) [The Bidirectional Influence between Coherence Establishment and Pronoun Interpretation. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008].

5

Syntactic Ambiguity

The previous chapters have shown that a purely morphosyntactic model is untenable as an account of coreferential processing. The alternative model put forth is one which recognizes the importance of incorporating deeper biases conditioned on discourse coherence. Such a model contrasts with what I have referred to in previous chapters as a surface-cue model — a model which relies primarily on grammatical- and thematic-role biases to account for discourse-level processes such as coreference.

These types of surface-cue models that have characterized the discourse-processing literature have their parallels in the literature on sentence processing. In this chapter, I turn to a phenomenon that has been a hallmark of syntactic analyses and I consider whether deeper coherence-driven biases may also be relevant to the way that comprehenders process syntactic relationships within a sentence. The phenomenon I consider is relative clause (RC) attachment ambiguity, a phenomenon that I will show is influenced by the RC's simultaneous roles in both the syntactic structure and the discourse structure. An off-line story continuation experiment and an on-line reading-time study using RC contexts tests whether comprehenders bring expectations from the discourse level to bear on the resolution of local syntactic ambiguity. The reading-time study also tests whether expectations about upcoming coherence relations impact comprehenders' sentence-internal incremental processing.

5.1 Relative Clause Attachment

A great deal of work in syntactic processing has focused on the ambiguity that arises in sentences like (55).

(55) Someone shot the servant of the actress who was on the balcony.

NP1 (high) NP2 (low)

First discussed by Cuetos and Mitchell (1988), examples like (55) contain an RC (*who was on the balcony*) whose attachment site is ambiguous. The RC can be interpreted to modify one or the other of the two noun phrases (NPs) in the preceding complex NP (*the servant of the actress*). The RC is said to attach high if it is interpreted to modify NP1 (*the servant*) which occupies the higher position in the syntactic structure. It is likewise said to attach low if it modifies NP2 (*the actress*), the possessor NP in the complex NP. In English, a low attachment preference in a context like (55) is widely accepted as the default preference, having been confirmed in off-line studies with questionnaires and completion tasks and in most on-line studies (Frazier & Clifton, 1996; Carreiras & Clifton, 1999; Fernandez, 2003; but see also Traxler, Pickering, & Clifton, 1998).

This preference is not universal across languages, however, and the two decades following Cuetos and Mitchell's original paper have yielded a large body of research on RC attachment preferences within and across languages (see reviews in Cuetos, Mitchell, & Corley, 1996; Mitchell & Brysbaert, 1998; Desmet, Brysbaert, & De Baecke, 2002; and Papadopoulou & Clahsen, 2006). Languages like Spanish, French, German, and Dutch (among others) show evidence of a bias towards a high-attachment interpretation for cases like (55) (Cuetos & Mitchell, 1988; Zagar, Pynte, & Rativeau, 1997; Hemforth, Konieczny, & Scheepers, 2000; Brysbaert & Mitchell, 1996; *inter alia*). The dispreference for the local attachment site in these languages and the lack of a universal attachment preference have been problematic for theories of sentence processing that posit crosslinguistic syntactic constraints for the inter-

pretation of words and clauses (e.g., Frazier's (1978) principles of Late Closure and Minimal Attachment). Whereas most previous analyses have focused primarily on syntactically-driven biases to account for the reported preferences in RC attachment, the question I ask here is whether a larger discourse context can also affect RC attachment preferences.

An RC Role in Discourse: Restricting Referent Identity

The handful of investigations of discourse-driven biases in RC attachment have taken as their assumption that the pragmatic function of an RC is to modify or restrict the identify of a referent. This work has shown that an RC is more likely to attach to a host NP that has more than one potential referent in the preceding discourse context (Zagar et al., 1997; van Berkum et al., 1999; Desmet et al., 2002; Papadopoulou & Clahsen, 2006). These studies rely on the fact that an NP with a definite determiner is usually infelicitous in a context in which the referent is not uniquely identifiable: The use of the NP *the actress* from example (55) is typically not licensed in a context with more than one actress unless further modification is provided. If the comprehender assumes the speaker is obeying rules of cooperative discourse and will only use a definite determiner with an NP that has a unique referent, then the presence of a temporarily non-unique NP appears to bias the comprehender to anticipate further identifying material. The premise of the Zagar et al., Desmet et al., and Papadopoulou and Clahsen studies is that comprehenders are more likely to interpret an RC whose attachment site is ambiguous as modifying the NP that requires additional modification in order to be referentially unique. In the case of examples like (56), this bias has been shown to reinforce the low-attachment preference in a context with multiple actresses. In (57), the context with multiple servants induces a bias towards a high attachment.

- (56) There was a servant working for two actresses. Someone shot the servant of the actress who was on the balcony.
- (57) There were two servants working for a famous actress. Someone shot the servant of the actress who was on the balcony.

Studies have been conducted in French, Dutch, and Greek, respectively, and all report off-line effects of referential context: The assumed high-attachment preferences in these three languages can be reversed in contexts in which the preceding discourse context introduces multiple referents for the second NP. On-line effects have been more elusive, with the Greek study reporting significant effects using self-paced reading, but the French and Dutch studies finding no significant effects using eye-tracking.

In this chapter, I will be focusing on a class of examples in which the function of the RC goes beyond modification or restriction to a more general discourse function of contributing propositional content to the overall discourse structure.

Alternative RC Role: Providing an Explanation

Besides modifying or restricting, RCs can play an additional pragmatic role of providing an explanation of the event described in the matrix clause. Consider the example in (58).

- (58) The teenager detests the children who watch *Alvin and the Chipmunks* over and over at full volume.

The RC in (58) not only restricts which children the teenager detests, but it also carries a (cancelable) implicature that the repeated watching of *Alvin and the Chipmunks* is the reason for the teenager's ill will. Comprehending the sentence in (58) requires inferring a relationship between the event conveyed by the RC (the children watch *Alvin and the Chipmunks*) and the information conveyed by the matrix clause (the teenager detests the children). The implicature that the RC provides an explanation of the matrix clause event depends on the causal reasoning that watching animated chipmunks on full volume is a plausible reason for being detested.¹ As a point of comparison, the same RC in a different context need not induce the same kind of causal reasoning and merely modifies the preceding NP, as in (59).

¹The fact that this is an implicature and not an entailment is demonstrated by the fact that it is cancelable, as in (i).

(i) The teenager detests the children who watch *Alvin and the Chipmunks* over and over at full volume. The reason the teenager hates them is that they refuse to follow instructions, staying up past their bedtime and ignoring all warnings.

- (59) The teenager babysits the children who watch *Alvin and the Chipmunks* over and over at full volume.

In (59), the watching is not the cause of the babysitting; it merely serves to distinguish which children the teenager babysits.

This kind of intrasentential inferencing can be taken to mirror the inferencing that comprehenders use to establish coherence relations between sentences. The causal reasoning that links the clauses in (58) is similar to that used in establishing an Explanation relation between the two sentences in (60). Similarly, the lack of causal reasoning in (59) mirrors the non-causal Elaboration/Background relation that holds between the sentences in (61).

- (60) The teenager detests the children. The children watch *Alvin and the Chipmunks* over and over at full volume.
- (61) The teenager babysits the children. The children watch *Alvin and the Chipmunks* over and over at full volume.

The different relationship between the sentences in (60) and (61) stems from the fact that *Alvin and the Chipmunks* provides a plausible reason to detest someone but not a plausible reason to babysit someone. As was described in Section 2.3, the relationship between two sentences can sometimes be made explicit with the use of a connective; note that the connective *because* is appropriate in (60) but not in (61). In general, though, understanding the relationship between two sentences requires the use of real-world knowledge to determine what constitutes, for example, a plausible Explanation or a relevant Elaboration. I propose that the same type of knowledge and inferencing that is required for understanding intersentential coherence is also used to establish intrasentential coherence relations that hold between clauses within the same sentence. Although one could construct clauses that are related by any of the coherence relations defined by Kehler (2002) and laid out in Section 2.3, for the purposes of this chapter, I focus on Explanations.

In order to address the question posed at the beginning of the chapter—whether coherence-driven expectations impact RC attachment—I will bring together three observations related to discourse coherence. The first is the one noted above: that RCs can do more than simply modify or restrict; they can also serve to provide an explanation of the matrix clause event. The second and third observations center around the fact that certain contexts are biased towards upcoming Explanations and that these contexts have biases regarding which referent will be mentioned next, given an Explanation. These biasing contexts include sentences containing implicit causality verbs.

5.1.1 Implicit Causality Biases

Recall that Chapter 3 presented the results from a story continuation experiment (Experiment IV) that used context sentences which contained verbs like *detest* — a member of the class of so-called object-biased implicit causality (IC-2) verbs. These verbs are used to describe the attitudes or emotions of one individual regarding properties of another, and they differ from non-IC verbs in that they systematically attribute cause to one individual — for IC-2 verbs, it is the individual mentioned in the direct object position.² The results from Experiment IV showed that context sentences with IC-2 verbs like that in (62a) yielded significantly more Explanation continuations than did context sentences with non-IC verbs like that in (62b).

- (62) a. John detests_{IC-2} Mary. She is arrogant and rude.
 b. John babysits_{non-IC} Mary. Mary's mother is grateful.

Example (62a) shows a typical IC continuation (an Explanation), whereas (62b) shows a typical non-IC continuation (a Result). As noted, verbs like *detest* are object-biased, meaning that it is the direct object that comprehenders expect to find

²The class of IC verbs also includes subject-biased IC-1 verbs that attribute cause to the individual appearing in subject position. See Section 3.7 for a review of IC biases. For the purposes of this chapter, the critical examples require IC-2 and non-IC verbs.

mentioned again in subsequent discourse. The biases of non-IC verbs are reported to be less consistent. These next-mention biases have been measured using story completion tasks with prompts like those in (63), with sample continuations listed below each prompt (Garvey & Caramazza, 1974; Brown & Fish, 1983; Au, 1986; McKoon, Greene, & Ratcliff, 1993).

- (63) a. John detests_{IC-2} Mary because...
 ... she is rude and arrogant.
- b. John babysits_{non-IC} Mary because...
 ... he needs the money / ... she can't stay home alone

Experiment IV confirmed the next-mention preferences shown in (63) and also showed that a similar pattern of next-mention biases emerges for prompts like these even without the connective, but only if the analysis is restricted to the subset of continuations in which an Explanation relation is inferred to hold. That result showed that the previously reported next-mention preference for IC-2 verbs is driven by two distinct biases: an underlying bias for explanations following IC verbs and a set of next-mention biases conditioned on the operative coherence relation and the object-bias of the verb itself. While it makes sense that next-mention biases would be conditioned on coherence-driven factors since coreference is a discourse-level phenomenon, it is an open question whether such factors would also influence RC interpretation and the resolution of RC attachment ambiguity.

5.1.2 IC Biases in RC Attachment

To reiterate, the experiments presented in this chapter build on three observations: (i) RCs can provide an explanation of the event in the matrix clause; (ii) IC verbs create an expectation that an upcoming clause will contain an explanation; and (iii) certain IC verbs are biased such that an upcoming explanation is likely to mention the verb's direct object. With these in mind, the question I address is what

happens if the explanation that is anticipated following an IC verb is packaged as an RC. Consider examples (64) and (65).

(64) John babysits the children of the musician who ...

(65) John detests the children of the musician who ...

Under the standard result from the previous work on RC attachment in English, we would expect that with a neutral non-IC verb as in (64), a low-attaching RC that attaches to *the musician* would be preferred, and a high-attaching RC that attaches to *the children* would be dispreferred. However, consider now (65), in which the matrix clause contains an IC-2 verb. As we have seen, IC-2 verbs (henceforth IC verbs) like *detest* create an expectation for an upcoming Explanation, an Explanation that is likely to be about the direct object. In (65), the direct object, *the children*, is crucially in the high-attachment position. Therefore, if the RC serves the pragmatic function of providing an Explanation, comprehenders may be more likely to expect the RC to attach high in the IC context than in the non-IC context. In other words, the expectation for an Explanation and the expectation for a subsequent mention of the direct object may conspire to induce a high attachment for the RC, thereby reversing the default low-attachment preference.

That one might find a high-attachment bias in (65) may be intuitive, but bear in mind that this reasoning only goes through if comprehenders, having processed the initial part of a matrix clause, are implicitly aware of these three coherence-driven factors and use them when making a syntactic attachment decision mid-sentence. The null hypothesis is that verb type will have no effect on attachment preferences, whereas the discourse hypothesis contends that IC verbs will increase comprehenders' expectations for a high-attaching RC.

The remainder of this chapter presents a sentence-completion experiment and a reading-time experiment. If comprehenders are indeed using coherence-driven biases mid-sentence, then one might expect to see effects on the types of RC completions

they generate and in their RC attachment preferences following IC matrix-clause verbs (Experiment IX). If these biases also contribute to comprehender's on-line processing, then one might expect to see processing difficulty associated with RC attachments which violate the biases introduced by the matrix-clause verb (Experiment X).

5.2 Experiment IX: IC Biases

A sentence completion study was conducted to examine comprehenders' expectations about upcoming discourse continuations and to test whether such expectations influence RC attachment preferences. Participants were asked to write sentence completions following RC prompts. If expectations about upcoming coherence relations are generated mid-sentence, then RCs following an IC verb should be (i) more likely than RCs following a non-IC verbs to provide an explanation of the event described in the matrix clause, and, crucially, (ii) the RCs should be more likely to attach to the causally implicated direct-object referent, the higher NP. If, on the other hand, such expectations are not utilized during syntactic processing, RCs following either verb class should attach low.

In the story completion task, participants create a mental model of the event in the matrix clause and then write a completion that reflects their expectations about a likely continuation to the sentence. As such, the task involves both interpretation and production. The context sentence and relative pronoun prompt constrain the surface realization of their completion, but the completion depends on their expectations about how the discourse will proceed and which NP is most likely to be modified.

5.2.1 Methods

Participants

Fifty-two monolingual English-speaking undergraduates from UC San Diego participated in the experiment for credit in Linguistics courses.

Materials and Procedures

Stimuli consisted of twenty-one pairs of sentences that differed only in the matrix verb, as in (64) and (65), repeated here as (66) and (67).

(66) John *babysits* the children of the musician who ... [NON-IC]

(67) John *detests* the children of the musician who ... [IC]

The subject of the sentences was always a proper name with gender balanced across stimuli. The direct object consisted of a complex NP containing two NPs connected by the genitive marker *of*. Both NPs denoted human referents so that the relative pronoun *who* could plausibly be used to modify either NP. In order to make disambiguation easier for the judges, the complex NP consisted of a singular NP and a plural NP so that number agreement on the embedded verb could be used to assess the intended attachment site of the RC. The order of singular and plural was balanced across stimuli.

The verb in the matrix clause was either an object-biased IC verb or a non-IC verb. IC verbs were selected from two lexical semantic categories that Levin (1993) labels ‘psych’ and ‘judgment’ verbs. Non-IC verbs were adapted from those used by McKoon et al. (1993) in their study of implicit causality and pronoun interpretation. For these stimuli, psych verbs appeared in the present tense since they describe non-eventive states (e.g., *detest*, *adore*), whereas judgment verbs appeared in the simple past (e.g., *scolded*, *praised*). Each pair of IC and non-IC verbs was matched for tense as in (66-67). The full stimuli set can be found in the Appendix.

The experiment consisted of sixty-three items: twenty-one experimental items intermixed with twenty-one stimuli for an unrelated experiment (Experiment VIII) as well as a set of additional fillers. The stimuli for the interleaved experiment are described in Section 4.4 and consisted of transfer-of-possession context sentences followed by full-stop and ambiguous-pronoun prompts. The additional fillers consisted of sentences with non-transfer verbs and a variety of prompts as well as sentences

with complex NPs and unambiguous RC prompts (the ambiguity in those cases was resolved by the relative pronoun *who* in contexts in which only one of the two nouns in the complex NP was human). Half of the RC fillers enforced a low attachment and half enforced a high attachment.

Sentence completion were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their completion. The entire experiment took roughly thirty minutes, but participants were encouraged to have an hour available so that the experiment could be completed in one session. Participants could leave the website and return at a later time by identifying themselves with an ID number. They were instructed to imagine a natural sentence completion for each prompt, writing the first completion that came to mind and avoiding humor.

Evaluation and Analysis

Two trained judges—the author of this dissertation and a UCSD Linguistics undergraduate—annotated all responses for the type of RC (‘restriction/modification’ or ‘explanation’) and the intended attachment site (low or high). An RC was labeled ‘restriction/modification’ if it provided additional information about one of the nouns without providing additional information about the event in the matrix clause *per se*. Explanation RCs, on the other hand, established a causal link between the information conveyed by the matrix clause and the information in the RC. Disagreements were resolved through discussion. The intended attachment site was assessed in light of the matrix clause context and the elicited RC.

To provide a sense of how the RCs were annotated, the sample completions in (68) and (69) show examples of modification and Explanation RCs that attach either low and high. These examples also show that all the combinations of coherence and attachment height are possible. Note that the sample modification RCs listed do not differ between the non-IC and IC contexts (68a-b, 69a-b) because such RCs need not reflect information about the event described in the matrix clause. Explanation RCs,

on the other hand, provide an explanation of the matrix clause event and therefore are shown varying with the matrix clause (68c-d, 69c-d).

(68) NON-IC: John babysits the children of the musician who ...

- | | |
|---|--------------|
| a. ...is a singer at the club downtown. | [MOD - LOW] |
| b. ...are in elementary school. | [MOD - HIGH] |
| c. ...works a late shift every night. | [EXP - LOW] |
| d. ...are left home on Friday nights. | [EXP - HIGH] |

(69) IC: John detests the children of the musician who ...

- | | |
|---|--------------|
| a. ...is a singer at the club downtown. | [MOD - LOW] |
| b. ...are in elementary school. | [MOD - HIGH] |
| c. ...encourages their 3am drum solos. | [EXP - LOW] |
| d. ...are rude and arrogant. | [EXP - HIGH] |

As (68) and (69) show, both verb classes can be followed by either a modification RC or an explanation RC, and neither enforces a particular attachment. The hypothesis is that the combination of coherence biases and next-mention biases will render high-attaching explanation RCs more likely following IC verbs than non-IC verbs: Completions like (69d) will be more common than (68d). The low-attaching modification RCs, on the other hand, are predicted to be more likely following non-IC verbs than IC verbs: Completions like (68a) will be more common than (69a).

Analyses of variance were conducted on the assessed RC completion types and on the assessed attachment sites to test for a main effect of verb class. Verb class varied within subjects and within items. Because these measures involve examining proportions of binary outcomes, an arcsine transformation was first applied to the percentages of high/low attachments and modification/Explanation RCs. For clarity of presentation, I present means as raw proportions. For this experiment, t-tests were also conducted over raw means of attachment biases to measure the bias following IC and non-IC contexts, as compared to a hypothetical mean of 0.5.

5.2.2 Results

After setting aside completions whose attachment height was judged to be ambiguous by at least one coder (22.5% of the total number of RC completions), the remaining dataset contained 852 completions. Of those, only subject-extracted RCs were included in the analysis (N=846); object-extracted RCs made up less than 1% of the total completions and their inclusion did not influence the overall significance.

As predicted, RC completions were more likely to attach high in the IC context (50.6%) than in the non-IC context (36.5%; $F_1(1,51)=27.158$, $p<0.001$; $F_2(1,20)=6.8475$, $p<0.02$). The non-IC condition, the low-attachment bias replicates the default low-attachment preference that has been reported for English (significantly different from a hypothetical mean of 0.5: $t_1(51)=4.8454$, $p<0.001$; $t_2(20)=2.7294$, $p<0.02$). In the IC condition, the low-attachment bias disappears (not significantly different from a hypothetical mean of 0.5: $t_1(51)=0.3894$, $p<0.7$; $t_2(20)=0.21$, $p<0.84$). This pattern can be seen on the left-hand side of Figure 5.1.³




Figure 5.1: Proportions of high attachments and proportions of Explanation RCs

The shift towards high-attaching RCs occurs in the same context that exhibits a shift towards Explanation RCs: RC completions were more likely to provide an explanation of the event in the matrix clause in the IC context (63.9%) than in the non-IC context (11.0%; $F_1(1,51)=292.22$, $p<0.001$; $F_2(1,20)=87.665$, $p<0.001$). This

³For the experiments in this chapter, all means represent subject means, unless otherwise noted.

pattern can be seen on the right-hand side of Figure 5.1.⁴

As predicted, the prevalence of high-attaching RCs in IC contexts is driven by the prevalence of high-attaching explanation-providing RCs in those contexts: 66.3% of explanation-providing RCs attached high in IC contexts compared to 47.0% in non-IC contexts. Non-explanation-providing RCs were less likely to attach high — in both IC contexts (26.0%) and non-IC contexts (35.9%). See Figure 5.2.




Figure 5.2: Proportion of high attachments for Explanation and non-Explanation RCs

5.2.3 Discussion

The sentence completion study was designed to test a model of sentence processing that integrates coherence-driven biases regarding the larger discourse structure with strategies for resolving local ambiguity. As predicted, IC verbs yielded significantly more explanation-providing RCs, and significantly more high-attaching RCs compared with non-IC verbs. The results have implications both for RC processing and for general models of sentence comprehension.

The pattern of high attachments following IC verbs provides evidence that models that rely on a default low-attachment preference make incorrect predictions

⁴This analysis represents a conservative analysis in which an RC was excluded if at least one coder assessed it as ambiguous. The results remain significant if RCs are included when at least one coder assigned a non-ambiguous interpretation (Attachment height: $F_1(1,51)=53.52$, $p<0.001$; $F_2(1,20)=8.1197$, $p<0.01$; RC type: $F_1(1,51)=356.07$, $p<0.001$; $F_2(1,20)=96.407$, $p<0.001$).

(cf. Gilboy et al., 1995; Gibson et al., 1996). The fact that RC attachment is sensitive to the coherence relation that can be inferred to hold between the RC and the matrix clause suggests that the factors influencing RC processing are not limited to properties of the NPs themselves. A larger discourse context can influence how an RC is interpreted, and models of syntactic processing can thus not ignore coherence-driven biases.

The way in which the preceding context influences attachment preferences is complex, depending simultaneously on biases regarding upcoming coherence relations and next-mention biases conditioned on the matrix clause verb and the operative intrasentential coherence relation. Closer inspection of the elicited sentence completions reveals two subpatterns of behavior that are not apparent in the aggregate effects: one concerning the bias towards Explanations and one concerning the attachment biases. Regarding the coherence bias, some verbs that have been classified in the literature as non-IC were found to actually yield a larger number of Explanation-providing RCs than some IC verbs. For example, the verb *watch* which McKoon et al. (1993) included in a non-IC condition yielded 46.2% Explanations — more than IC verbs such as *like* (26.7% Explanation) and *value* (22.7% Explanation). This variation is in keeping with Caramazza et al.'s (1977) comment that IC biases lie along a continuum.

Regarding the attachment biases, the presence of both an IC verb and an Explanation-providing RC are not in themselves sufficient to yield a high-attachment preference; the relationship between the NPs in the complex NP also has an influence. The contrast between the two items in (70) shows how specific complex NPs can shift the bias to yield more high-attaching or more low-attaching RCs.

- (70) a. Alan punished the accountant of the businessmen who ...
 b. Bill congratulated the teacher of the second graders who ...

Example (70a) showed the predicted pattern: The IC matrix clause yielded a large proportion of Explanation-providing RCs (85.7%), and those RCs consistently attached to the higher NP (100%). Example (70b) showed a different pattern: The IC

matrix clause also yielded a large proportion of Explanation-providing RCs (81%), but in this case, the RCs tended to attach low (only 29% high attachment). Example (70b) differs from (70a) in that the lower NP in (70) refers to a set of individuals (NP1: *the second graders*) who are under the care or responsibility of the individual appearing as the direct object (NP2: *the teacher*). Because of the possibility of attributing responsibility to the NP1 referent for the NP2 referent's behavior, the Explanation-providing RCs could plausibly describe either the guardian or the guardian's wards, as in (71a) and (71b). (Similarly behaving items included *scold the landlady of the actors who...*, *detest the father of the students who...*; and *pity the bodyguards of the celebrity who...*)

- (71) a. Bill congratulated the teacher of the second graders who was the recipient of the Teacher-of-the-Year award.
- b. Bill congratulated the teacher of the second graders who were the winners of the annual spelling bee.

This subpattern in the data shows that is important to control for effects introduced by the relationship between the NPs in the complex NP because one set of complex NPs can yield results that support a high-attachment preference while another set can yield results that support a low-attachment preference. The aim of this experiment was to test whether discourse context can shift attachment preferences at all. Only under a model in which coherence-driven biases influence the resolution of syntactic ambiguity is a high-attachment bias predicted to be possible for this type of data. Since no other model predicts this result, the complex NPs in the stimuli that did induce high attachments provide evidence for the coherence-driven model. Nonetheless, the observed subpattern within the data suggests that, though the results were significant, a more consistently biased set of complex NPs could have yielded even stronger effects.

Details of the complex NPs aside, the overall pattern of results is also relevant to general models of sentence processing that aim to uncover the factors that guide

comprehenders' interpretation of a discourse structure and the scale of the linguistic elements over which such factors operate. The results confirm the conclusion from Chapter 3 that comprehenders are sensitive to two types of biases invoked by IC verbs: a clause-level coherence bias toward upcoming explanations and an entity-level next-mention bias conditioned on the presence of an explanation relation. These results go beyond that work by showing that coherence relations can be inferred to hold intrasententially and that RCs can embody such relations. If an RC can provide an explanation of an event in the matrix clause, then the process of linking together elements of a discourse into a larger coherent structure cannot be cast as a process of sentence-final integration.

If the biases that emerge in the off-line completion study are indeed deployed mid-sentence, then one would expect to see effects in comprehenders' incremental processing in a self-paced reading-time experiment. The goal of Experiment X is to test whether or not comprehenders construct discourse contexts dynamically during sentence processing, using pragmatic cues mid-sentence to generate expectations about the structural analysis of the remainder of the sentence.

5.3 Experiment X: On-line IC Biases

The reading time study was designed to test whether coherence-driven expectations have an impact on the immediate processing of RCs. The prediction is that high-attaching RCs will be read more quickly in the IC condition than the non-IC condition and that the reverse will be true for low-attaching RCs. The effects are predicted to emerge in the regions at or immediately following the point of disambiguation. In order for participants' reading times to be affected in these regions, several key assumptions must all hold: (i) The biases regarding upcoming coherence relations need to be conditioned on the IC context, (ii) the biases regarding next mention need to be conditioned on the probability of an Explanation; (iii) the knowledge that RCs can provide explanations of preceding material requires that comprehenders

have a working model of intrasentential discourse coherence. Since several different discourse cues contribute and their integration requires a fairly complicated chain of reasoning, an on-line effect would suggest that participants are invoking these coherence-driven biases mid-sentence and using them to generate expectations about upcoming syntactic structures.

5.3.1 Methods

A moving-window self-paced reading study was conducted to examine the time course of the biases found in the sentence completion study. The stimuli from Experiment IX were adapted for a 2×2 design, varying verb class and RC attachment height, as in (72-73).

(72) NON-IC MATRIX: John babysits the children of the musician ...

- a. [LOW ATTACHMENT] ...who *is* generally arrogant and rude.
- b. [HIGH ATTACHMENT] ...who *are* generally arrogant and rude.

(73) IC MATRIX: John detests the children of the musician ...

- a. [LOW ATTACHMENT] ...who *is* generally arrogant and rude.
- b. [HIGH ATTACHMENT] ...who *are* generally arrogant and rude.

In (72) and (73), the attachment height of the RC is signaled by the verb (e.g., the verb *is* in (72a) agrees in number with *the musician*, the NP at the low-attachment site). The low-attachment bias is expected to emerge for non-IC verbs (faster reading times for the critical/spillover region in (72a) than (72b)), but is predicted to be neutralized or even reversed following IC verbs (no speed advantage for (73a) over (73b)). Therefore, we are not looking for main effects of verb class or attachment height but rather for an interaction between the two. On the other hand, if integration into the larger discourse structure occurs only after the entire sentence has been processed, then the default low-attachment bias should hold across the board (faster reading times for the critical/spillover regions in (72a, 73a) than

for the critical/spillover regions in (72b, 73b)). The inclusion of a spillover region in the analysis is in keeping with work by Mitchell (1984) and others who report that it is common for reading time effects to appear in the immediate regions following a target disambiguation region, especially in cases in which the critical region is short, as is the case here (*is/are*).

Participants

Fifty-eight monolingual English-speaking undergraduates from UC San Diego participated in the experiment for credit in Linguistics and Psychology courses.

Materials

Each experimental item consisted of a matrix clause with a proper name, a verb, and a complex NP direct object, followed by a temporarily ambiguous RC, as in (72) and (73). Given the observations about the complex NPs in Experiment IX, stimuli were constructed with NP pairs for which an Explanation relation following an IC verb would be compatible with high attachment. The complex NP contained a singular NP and a plural NP so that number agreement on the embedded auxiliary verb served to disambiguate the attachment site of the RC. The order of singular and plural in the complex NP was balanced across stimuli so that high attachment was signaled with plural agreement for half the items and with singular agreement for the other half. After the disambiguation point, a semantically neutral adverb (intended not to affect the bias of the RC attachment level) appeared in the spillover region. The filler items were similar to the stimuli in that some included proper names and RCs or other subordinate clauses.

Given the observations from Experiment IX concerning the IC-like coherence biases of some non-IC verbs, verbs were selected from Experiment IX that had strong biases towards explanation RCs and then we supplemented those verbs with verbs taken from McKoon et al.'s (1993) study (the same verbs used in Kehler et al. (2008)). The verbs *jeer* and *dread* were replaced with *mock* and *fear* because of concerns that the original verbs did not sound natural when followed immediately by an animate

direct object. The 20 non-IC verbs consisted of a mix of verbs from McKoon et al. (1993), Levin (1993), and from Experiment IX. The full stimulus set can be found in the Appendix.

For the experiment, there were 10 practice items, followed by 20 experimental items mixed with 30 fillers, pseudo-randomized for each subject. Each participant saw a mix of IC and non-IC verbs paired with either a low-attaching or high-attaching RC.

Procedure

Items were presented in a moving-window self-paced reading paradigm, using DMDX experiment software (Forster & Forster, 2003). Sentences appeared in white letters on a dark background, left-justified on a 19" CRT screen, and no sentence was longer than one line of text. Sentences initially appeared as a series of dashes (— — —) obscuring the words, and participants pushed a button on a Logitech USB gamepad to reveal each region. The presentation was non-cumulative such that previous regions were replaced with dashes when the next region appeared. The critical region and the spillover regions were revealed one word at a time, but multi-word regions were used elsewhere to present short phrases such as a verb and a preposition (*stared at, stood near*) or a determiner and a noun (*the children*). Multi-word regions are indicated in the stimuli set in the Appendix. Participants pushed either a YES or NO button on the gamepad to answer a comprehension question after every sentence, and they received automatic feedback whenever they answered incorrectly. They were instructed to read as quickly and carefully as possible, making sure they understood the complete sentence and slowing down if they answered multiple questions incorrectly. Reading times were recorded for each region as well as the participant's response to the comprehension question.

5.3.2 Results

After excluding three participants whose comprehension-question accuracy was not significantly better than chance, the percentage of correct responses was 93.0% for fillers and 85.1% for experimental items (percentages over subject means), indicating that participants paid attention to the task. All critical trials in which the comprehension question was answered incorrectly were excluded from the analysis. Table 5.1 shows the raw reading times by condition for the critical region and the spillover regions, as well as the mean accuracy on comprehension questions for each condition.

Table 5.1: Raw RTs and question accuracy (subject means \pm standard error)

	Target	Spillover1	Spillover2	Accuracy
IC high	395.70 \pm 16.83	430.43 \pm 18.90	442.81 \pm 18.84	.873 \pm .02
IC low	398.83 \pm 16.71	474.16 \pm 23.26	477.19 \pm 26.34	.780 \pm .02
non-IC high	402.03 \pm 16.55	501.48 \pm 24.26	473.59 \pm 20.22	.862 \pm .02
non-IC low	403.96 \pm 13.83	462.63 \pm 20.03	437.50 \pm 15.91	.887 \pm .02

Figure 5.3 shows the residual reading times for each of the four conditions starting at the matrix verb. Comprehension-question accuracy and reading times were analyzed with 2×2 ANOVAs, by subjects and by items.

Comprehension question accuracy

All question-accuracy analyses were conducted on arcsine transformed proportions of correct answers. Accuracy was not uniform across conditions — in particular, low-attaching RCs following non-IC verbs yielded the best accuracy, whereas low-attaching RCs yielded the worst accuracy following IC verbs. There was a marginal main effect of attachment height favoring high-attaching RCs (marginal by subjects: $F_1(1,54)=3.889$, $p<0.06$; $F_2(1,19)=2.778$, $p<0.12$). There was also a main effect of




Figure 5.3: Residual reading times

verb class favoring non-IC verbs (significant by subjects: $F_1(1,54)=4.59$, $p<0.04$; $F_2(1,19)=2.206$, $p<0.16$). This main effect is driven by an interaction in which low-attaching RCs in the IC condition yielded lower accuracies than any other condition. The cross-over interaction that emerges is consistent with the predicted interaction for processing difficulty in cases in which the RC violates the expectations of the preceding context (significant verb-class \times RC-height by subjects and marginal by items: $F_1(1,54)=7.346$, $p<0.01$; $F_2(1,19)=3.89$, $p<0.07$).

Reading time results

Residual reading times were analyzed at the critical region and two spillover regions. Residual reading times adjust for overall differences in participants' reading rates as well as differences in readers' sensitivity to word length. They were calculated as the difference between the actual reading time on a word and the reading time predicted by a regression equation (computed separately for each participant, using all experimental and filler items) relating word length to reading time (F. Ferreira & Clifton, 1986). Residual reading times that were more than four standard deviations away from the mean, per region and per condition were removed (0.34% of the data).

At the disambiguating verb (*is/are*), there were no significant effects for verb class ($F_1(1,54)=1.442$, $p<0.24$; $F_2(1,19)=2.116$, $p<0.17$), attachment height ($F_s<1$),

or the attachment-height \times verb-class interaction ($F_s < 1$).

At the first spillover region (*generally*), there were again no main effects of verb class ($F_1(1,54)=2.646$, $p < 0.12$; $F_2(1,19)=1.701$, $p < 0.21$) or attachment height ($F_s < 1$). A significant cross-over interaction was observed in the predicted direction: High attachments were read more slowly than low attachments in non-IC conditions but faster in IC conditions ($F_1(1,54)=6.117$, $p=0.017$; $F_2(1,19)=5.216$, $p < 0.04$).

At the second spillover region (*arrogant*), there were again no main effects ($F_s < 1$), but the same cross-over interaction was significant ($F_1(1,54)=7.598$, $p < 0.01$; $F_2(1,19)=5.465$, $p < 0.04$).

Analyses of the raw reading times show a similar pattern, as do analyses of the residuals with incorrectly answered items included.⁵

5.3.3 Discussion

The results for the reading time experiment confirm that the discourse-level biases observed in the off-line completion experiment also influence comprehenders' on-line incremental processing. As predicted, high-attaching RCs were read more quickly than low-attaching RCs in the IC condition, whereas the reverse was true in the non-IC condition. The crossover interaction was observed in the region immediately following the disambiguation point. This suggests that expectations generated at the matrix verb can impact comprehension of an RC in the same sentence. Not only did the previously reported default low-attachment preference disappear in an IC

⁵Considering the raw reading times, there were no effects at the disambiguating verb ($F < 1$). At the first spillover region, there were effects of verb class by subjects and a significant cross-over interaction (verbclass: $F_1(1,54)=7.075$, $p < 0.02$; $F_2(1,19)=3.548$, $p < 0.08$; attachment: $F_s < 1$; interaction: $F_1(1,54)=6.853$, $p < 0.02$; $F_2(1,19)=5.434$, $p < 0.04$). At the second spillover region, there were no main effects ($F_s < 1$) and the same interaction was significant ($F_1(1,54)=6.705$, $p < 0.02$; $F_2(1,19)=6.078$, $p < 0.03$).

Considering the residual reading times over all items regardless of comprehension-question accuracy, there were no effects at the disambiguating verb (verbclass: $F_s < 1$; attachment: $F_1(1,54)=1.071$, $p < 0.31$; $F_2(1,19)=1.124$, $p < 0.31$; verbclass \times attachment: $F_s < 1$). At the first and second spillover regions there were also no main effects but significant cross-over interactions (Spillover1: verbclass: $F_s < 1$; attachment: $F_1(1,54)=1.295$, $p < 0.26$, $F_2 < 1$; verbclass \times attachment spillover1: $F_1(1,54)=5.522$, $p < 0.03$; $F_2(1,19)=6.167$, $p < 0.03$; Spillover2: verbclass: $F_s < 1$; attachment: $F_s < 1$; verbclass \times attachment $F_1(1,54)=6.588$, $p < 0.02$; $F_2(1,19)=4.967$, $p < 0.04$).

context, but the condition with the fastest reading times was the high-attaching RC in the IC context. In the analysis of comprehension-question accuracy, the condition with the worst accuracy was the low-attaching RC following an IC matrix clause.

5.4 General Discussion

As discussed in the introduction to this thesis, relative clause processing has been used as a testing ground for various phenomena, including ambiguity resolution (Cuetos & Mitchell, 1988), extraction (Kluender, 1992), optionality (V. S. Ferreira & Dell, 2000), expectation-based parsing (Levy, 2008) and thematic fit (McRae et al., 1998), among others. In this chapter, I have shown that RC processing also reveals the time course over which people construct a discourse context and the ways in which discourse-level information impacts syntactic disambiguation. The sentence completion experiment and the self-paced reading experiment yielded a reversal of the RC low-attachment preference in cases in which the causally implicated referent was located at the high attachment site. The on-line results from the self-paced reading experiment suggest that comprehenders track such expectations mid-sentence, and that these expectations influence syntactic processing.

A parallel can be drawn between RC processing and pronoun interpretation. This parallel has been proposed as a potential explanation for the crosslinguistic differences in RC attachment preferences. It has been noted that RCs in many high-attaching languages consistently use a relative pronoun, while English RCs vary between a relative pronoun *who* and an ambiguous, often optional, word *that* (Hemforth, Konieczny, Seelig, & Walter, 2000). This observation has led to a proposal that languages with obligatory relative pronouns treat the resolution of the RC attachment ambiguity as a problem of anaphora resolution: The most salient referent is identified as the best attachment site. Hemforth et al. suggested that languages with relative pronouns were more likely to show a high attachment preference: RCs were more likely to attach high to the syntactically and semantically prominent head of

the RC. In languages that lack a consistent relative pronoun, on the other hand, the low-attachment preference emerges due to recency.

In a study comparing RC attachment with pronoun resolution in German, Hemforth et al. (2000) report off-line results showing similar interpretations of RCs and pronouns in examples like (74) and (75).

(74) RELATIVE CLAUSE:

Der Student des Professors, der in dem neuen Labor war, las die Ergebnisse.

The student of the professor [who was in the lab] read the results.

NP1

NP2

(75) PRONOUN:

Der Student des Professors las die Ergebnisse, als er in dem neuen Labor war.

The student of the professor read the results when [he] was in the lab.

NP1

NP2

Hemforth et al. report that the RC in (74) is interpreted to attach to the higher NP (*the students*) roughly as often as the pronoun in (75) is interpreted to refer to the same NP (*the students*). This result highlights a shared component of coreference and RC attachment, namely that semantic and pragmatic biases influence which referent will be mentioned again as the discourse proceeds. Hemforth et al.'s approach was originally proposed to distinguish between those languages that favored high-attaching RCs and those that favored low-attaching RCs. German with its use of relative pronouns was posited to invoke strategies akin to anaphora resolution in the interpretation of RCs.⁶

Hemforth et al. did not pursue the anaphora analogy beyond its implications for differences between languages, but the results presented here show that factors that influence pronoun interpretation biases can also lead to different attachment preferences within a language. Recall from Chapters 3 and 4 that pronoun-interpretation

⁶As Carreiras and Clifton (1999) point out, however, the anaphora account cannot fully explain the crosslinguistic differences in RC attachment preferences because of evidence of low-attachment preferences in languages like Italian which consistently use a relative pronoun (cf. Vincenzi & Job, 1995).

and next-mention biases can be characterized largely as epiphenomena of the mechanisms that are used to establish intersentential discourse coherence. The results from this chapter follow a similar vein: RC attachment ambiguity can be better understood within a discourse-sensitive model of language processing because it is a phenomenon, like pronominal anaphora, whose complexity lies not only in the linguistic construction itself, but in its role in the composition of a coherent discourse.

5.5 Acknowledgments

The experiments described in this chapter have been reported in earlier work and are discussed here with permission from the coauthors. The dissertation author was the primary investigator for the material presented here. Both Experiments IX and X were reported in Rohde, Levy, and Kehler (2008) [Implicit Causality Biases Influence Relative Clause Attachment. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008].

6

Structuring a Discourse

The experiments in this chapter focus on another domain in which comprehenders may use contextual cues to generate expectations about upcoming units of discourse structure. Whereas the previous chapters analyzed the effects of contextual cues on upcoming coherence relations, this chapter considers whether contextual cues have an impact in another domain, the domain of questions. The previous chapters relied on a model of discourse structure in which adjacent clauses or discourse segments are related by intersentential coherence relations (Hobbs, 1979; Kehler, 2002); here, in order to test the utility of an expectation-driven model for questions, I will be appealing to an alternative model of discourse structure in which adjacent discourse segments are related by implicit and explicit questions.

Under the coherence relation approach, as we have seen, an example like (76) can be analyzed as an Explanation relation: Sentence (76b) describes the reason for the event expressed in (76a).

- (76) a. Hannah played countless games of online Scrabble.
b. She was avoiding writing her dissertation.¹

¹As of 8/01/08, the Scrabble scoreboard is at 27 wins, 31 losses. My best bingo was ‘RESETTLE’ for 122 points.

A competing account of discourse structure is the Question-Under-Discussion (QUD) model (Roberts, 1996; see also Van Kuppevelt, 1995; Büring, 2003; Larsson, 1998; Ginzburg & Sag, 2000), in which discourses are structured with respect to implicit questions: Roughly speaking, an utterance is coherent insofar as it provides an answer to an inferrable question that is relevant to the proceeding discourse. A QUD analysis of example (76) would posit an implicit question *Why?* as intervening between (76a-b). The QUD analysis applies not only to the implicit questions that must be inferred in monologue passages like (76), but to explicit questions in dialogs as well. This chapter asks whether a connection can be drawn between comprehenders' biases regarding coherence relations and their biases regarding questions.

The story-continuation experiments presented in the previous chapters found that different discourse contexts give rise to different distributions over coherence relations. For example, the results of Experiment IV showed that expectations about upcoming Explanation coherence relations depended in part on the verb type: whereas 60% of passages in implicit causality (IC) contexts like (77) (see Garvey & Caramazza, 1974; *inter alia*) were continued with Explanations, only 24% of non-IC contexts like (77) were.

- (77) a. Mary scolded_{IC} John. _____.
 b. Mary saw_{nonIC} John. _____.

The results of Experiment IV highlighted the pragmatic biases introduced by IC and non-IC verbs, and these biases were accounted for within a coherence-driven model. Although coherence models and QUD models have both played a role in the research on discourse structure and pragmatics, there has been limited discussion of the correspondence between these two types of models. One exception is Roberts' comment that "[Coherence] relations can often...be characterized in terms of questions and answers, e.g., the use of a why-question and its answer to characterize explanations" (Roberts, 1996, p. 50). Taken together with the story-continuation results from previous chapters, Roberts' comment gives rise to a novel prediction that

the biases people have regarding upcoming coherence relations in monologues will be mirrored in the biases they have towards certain explicit questions in dialogs. To test this, the experiments in this chapter compare elicited story and dialog continuations. Participants in the experiments were instructed to imagine a conversation with a friend and write natural continuations that either represented what the friend was likely to say next (monologue condition, as in (78)) or a question that they would be likely to pose to their friend (dialog continuation, as in (79)).

(78) Friend: Mary scolded_{IC} John. _____.

(79) a. Friend: Mary scolded_{IC} John.

b. You: _____.

The first experiment shows that, as predicted, the differing proportions of Explanations in the IC and non-IC monologue contexts are highly correlated with the proportions of explicit why-type questions that interlocutors ask in dialogs with the same contexts. To verify that the effects generalize beyond Explanation and ‘Why?’, the second experiment uses transfer-of-possession verbs that vary by aspect, a factor shown in Experiment I to yield different distributions of coherence relations. As predicted, the results of the second experiment show that, for both verbal aspects, the percentage of the questions falling into the ‘Why?’, ‘What next?’, and ‘Where/when/how?’ categories in the dialog condition was highly correlated with the percentage of Explanations, Occasions, and Elaborations in the story condition, respectively.

As I noted earlier in Chapters 3-5 on pronoun interpretation and relative clause attachment ambiguity, previous work has tended to focus on how comprehenders *recover* underlying structure from their input. The previous work on implicit questions makes similar assumptions about the role of the comprehender being largely to *recover* underlying discourse structure. In the following section, I discuss work showing that comprehenders are indeed capable of recovering implicit questions in a discourse.

Then, in keeping with the theme of this dissertation on expectation-driven processing, I turn to the idea that comprehenders also use contextual cues in a discourse to generate expectations about where the discourse is going and what questions are likely to be raised explicitly or implicitly. I then present the two story- and dialog-continuation experiments showing that the same factors that influence the distribution of upcoming coherence relations in monologues have a similar influence on the distribution of upcoming questions in dialogs.

6.1 Structuring a Discourse with Questions

In models built around implicit and explicit questions, discourse segments can be thought of as a series of comments which address some common topic, goal, or QUD. These comments are structured so that they fit together to form a discourse representation. An issue that has been raised in previous research is, if this structure exists meaningfully for speakers and hearers, what cues do comprehenders use to identify this underlying structure? In monologues, the underlying relationships between utterances may be signaled overtly with a connective, but much of the time, comprehenders rely on their real world knowledge and logical inferencing to establish such relationships. In conversations, overt questions provide cues to the direction the discourse is taking, but often discourse participants answer questions that were not explicitly asked. In certain contexts, monologues simulate aspects of a dialog by introducing the QUD explicitly through an embedded interrogative or other rhetorical device, but in both domains, comprehenders are often forced to rely on other cues to understand what question an utterance is answering. I discuss here two cues that have been argued to help comprehenders retroactively identify the QUD even when it has not been uttered explicitly.

In a study on intonation and questions, Most and Saltz (1979) showed that comprehenders are able to determine what *wh*- questions an intoned target answer would be an appropriate answer to (see Roberts (1996), Büring (2003), and Jackend-

off (1972) for semantic analyses of the contexts that license certain prosodic focus). Participants were told to imagine that the sentences they heard were replies to questions, and their task was to write down the most appropriate question to which the sentence might be an answer. The results showed that subjects distinguished between sentences like (80a-b), reliably generating questions in which the *wh*- element corresponded to the stressed element in the sentence they heard. (80a) with agent stress generated questions like (81a) 79.0% of the time. Sentences like (80b) with patient stress generated questions like (81b), though only 56.3% of the time.

(80) a. The PITCHER threw the ball.

b. The pitcher threw the BALL.

(81) a. Who threw the ball?

b. What did the pitcher throw?

Most and Saltz conducted a second experiment varying only active and passive voice as in (82).

(82) a. The pitcher threw the ball.

b. The ball was thrown by the pitcher.

Active sentences in English do not carry many information structural restrictions about which argument (subject or object) ought to be new or old in the discourse. But the patient-promoting passive construction has been associated with stricter constraints on information structure: A passive construction is appropriate in a context in which the patient in subject position is old information and the agent in the *by*-phrase is new. Therefore, Most and Saltz point out that passive sentences are good answers to questions about the agent of an event (especially in contexts in which the patient is topical). Their results show that the agent question in (81a) was more frequently selected for passive sentences like (82b) (70.8% of the time) than for active sentences like (82a) (53.8%).

Most and Saltz's results were designed to show that, even in the absence of a larger context, comprehenders could identify what the likely question was that might have generated the answer that they heard, based on cues from intonation and active/passive voice. In this light, Most and Saltz reinforce the treatment of comprehension as a retrieval problem — upon hearing an utterance, a listener searches the set of open QUDs to retrieve the relevant question; if a likely QUD is unavailable, then one must be created and accommodated. As in many of previous models of non-dialog sentence processing, the comprehender is portrayed as reacting rather than predicting. The alternative approach that I have advocated in this dissertation posits that comprehenders are anticipators, using contextual cues to generate expectations about where a discourse is going and how that discourse direction is likely to impact linguistic phenomena in the next sentence or the next clause.

Under expectation-driven accounts of processing, the advantage for the comprehender is that in cases in which the comprehender correctly anticipates subsequent material, the material is therefore easier to process. The expectation-driven approach has been used previously to model comprehension of sounds, words, and syntactic structures, and now in this dissertation, it has been extended to units of discourse, coherence relations. Now the issue is whether it can be further extended to model what comprehenders do in dialogs — do they anticipate upcoming discourse moves, specifically questions?

Van Kuppevelt (1995) proposes that a set of possible questions arise in a listener's mind during discourse comprehension and that these questions are entertained as ones that may be answered in subsequent discourse. He leaves open the question of which factors in a context contribute to the anticipation of upcoming questions:

An implicit question is a question which the speaker anticipates will arise in the listener's mind on interpreting preceding utterances (or some non-linguistic events occurring in the discourse). ... In this paper, however, we will largely leave undiscussed the way in which these questions arise as the result of the interaction of given contextual information and a given model of the addressee. (p. 117)

Van Kuppevelt identified ‘Feeder’ sentences that appear early in a discourse and give rise to the questions that follow, as in his example repeated in (83).

- (83) Feeder1: John is ill.
- a. Question1: What does he suffer from?
 - b. Question2: For how long already?
 - c. Question3: What is the reason?
 - d. Question4: When do you expect him to recover?

Consider the feeder sentence from (83) paired with possible answers to each of the questions, as in (84). The feeder sentence “John is ill” and the possible answer to each of the questions in (83) yield pairs of sentences for which the relationship that can be inferred to hold between the meanings of those two sentences corresponds to one of the coherence relations described in Kehler (2002).

- (84) a. John is ill. He is sick with the flu. [Elaboration]
 b. John is ill. He has been sick for a week. [Elaboration]
 c. John is ill. He forgot to wear a hat and scarf. [Explanation]
 d. John is ill. He will get better by next week. [Occasion]

If a feeder sentence yields a set of possible questions in a dialog, do all the questions have the same status or are some more likely than others? The work I have described so far in this dissertation suggests that a feeder sentence like (83) in a given monologue context would yield a probability distribution over the set of possible coherence relations. The analogy that I want to draw is that the same sentence in a dialog context also yields a probability distribution over a set of possible QUDs.

6.2 Experiment XI: Monologues/Dialogs in IC Contexts

To reiterate the predictions for Experiment XI, recall that Experiment IV, a story-continuation experiment, showed that the class of implicit causality (IC) verbs

yielded a distribution of coherence relations that was skewed toward Explanation relations when compared to the distribution following other types of verbs. By comparing monologue continuations and dialog continuations, Experiment XI tests whether the coherence relations generated in these contexts are correlated with questions posed following the same contexts.

The monologue and dialog continuations for this experiment come from two of the three conditions in Experiment IV. The continuations analyzed in Section 3.7 were elicited in monologue *because* and full-stop conditions, as in (85a-b). The third condition was a dialog continuation as in (85c).

(85) a. BECAUSE (MONOLOGUE) CONTINUATION:

Friend: John scolded Mary because _____.

b. FULL-STOP (MONOLOGUE) CONTINUATION:

Friend: John scolded Mary. _____.

c. DIALOG CONTINUATION:

Friend: John scolded Mary.

You: _____?

If coherence relations and QUDs are correlated, then the large proportion of Explanation coherence relations observed in the full-stop prompt following IC context sentences is predicted to be matched by a large proportion of ‘Why’-type questions in the dialog prompts following IC contexts; likewise, smaller proportions of Explanation coherence relations and ‘Why’-type questions are predicted following non-IC context sentences compared to IC contexts. Explanation coherence relations (and their ‘Why’ question counterparts) were selected because they were the most frequent relation in continuations elicited in Experiment IV with the IC and non-IC contexts.

6.2.1 Methodology

The participants, materials, task, and annotation for the monologue continuations are described in Section 3.7. The factors of verb type (IC-1/IC-2/non-IC

for the context sentence) and prompt type (full-stop/*because*/dialog for the elicited response) were both within-subjects variables. Responses to the *because* prompt are not discussed here.

For the dialog continuations, an annotation scheme was devised for categorizing QUDs. The categories were based on the set of coherence relations described in Section 2.3. Questions were annotated in terms of their possible answers such that a question like ‘Why’, which asks for an answer in the form of an Explanation, would be categorized as an Explanation QUD (other questions that correspond to Explanation QUDs would include ‘How come’, ‘What for’, or ‘For what reason’). Examples are given in (86a-e) and come directly from participants’ story continuations for the experiment. All of the examples use IC verbs, but the same annotation scheme was applied to non-IC contexts as well. *Wh-* and *yes/no* questions were treated the same (i.e., Friend: *John infuriates Mary.* / You: *Why?* / You: *Is Mary annoyed with John’s endless complaining?* — Both are used to request information about the Explanation of an event). No questions were annotated as Violated-Expectation because it was not clear what question one could pose to ask about an outcome that did not happen or was surprising.

(86) ANNOTATION EXAMPLES

- a. ELABORATION
Friend: Greg corrected Sally.
You: When did this happen ?
- b. EXPLANATION
Friend: Ryan hates Amy.
You: What had she done ?
- c. OCCASION
Friend: Craig reproached Kate.
You: What happened next ?
- d. PARALLEL
Friend: Laura values Luis.
You: Does Luis value Laura ?
- e. RESULT
Friend: James charmed Amber.
You: Did she blush ?

For the analysis, I tested how well correlated the proportion of Explanation story continuations was with the proportion of Explanation dialog continuations across IC and non-IC verbs. As in previous chapters, the coherence relations were coded as six binary outcomes. Because this measure involves examining proportions, I first applied an arcsine transformation to the percentages of pronoun interpretations before carrying out analyses of variance. All other analyses of binary outcomes that are presented in this chapter use arcsine transformed percentages. For clarity of presentation in this section and throughout the chapter, I present means in the form of raw proportions.

6.2.2 Results

After setting aside the *because* continuations and the responses in the dialog condition in which the participant misunderstood the task (interpreting the question prompt as a prompt to write a statement) or that were deemed to be ambiguous or non-information-seeking (rhetorical) (11.8% of the dialog responses), the remaining dataset contained 1788 dialog responses. The 1938 full-stop monologue responses analyzed in Experiment IV were used here. However, since there were no Violated Expectations in the dialog condition, Violated Expectation continuations in the monologue condition were excluded (7.3% of the monologue responses), leaving a combined monologue and dialog dataset of $N=3585$.

Correlation of Coherence and QUD Biases

As was reported in Section 3.7 for the monologue condition, IC verbs yielded significantly more Explanations compared to non-IC verbs. As predicted, the increase in the proportion of Explanation continuations is correlated with an increase in the proportion of ‘Why’-type questions (‘Why?’, ‘How come?’, ‘What for?’). Figure 6.1 shows the percentage of Explanations and ‘Why’-type questions in both the monologue condition and the dialog condition.²

²For both experiments in this chapter, all means represent subject means, unless otherwise noted.

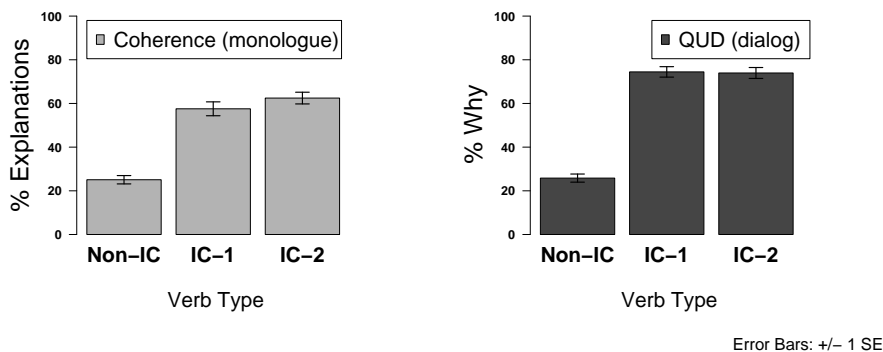


Figure 6.1: Proportion of Explanation/Why responses in Experiment XI

Across the three verb types, non-IC contexts yielded fewer Explanation-based responses (25.8% Explanations/26.1% ‘Why’-type questions) compared to IC-1 contexts (58.3% Explanations/75.4% ‘Why’-type questions) and IC-2 contexts (62.9% Explanations/73.6% ‘Why’-type questions). These biases are reliably correlated (adjusted R^2 by subjects and by items are as follows: $R_1^2=0.289$, $F_1(1,223)=92.23$, $p<0.001$; Adjusted $R_2^2=.488$, $F_2(1,78)=76.3$, $p<0.001$).³

6.2.3 Discussion

The goal of this experiment was to establish a connection between the implicit coherence relations that have been posited to underlie a monologue and the overt questions that structure a dialog. Contexts with IC and non-IC verbs were used because those had been shown previously to have strong biases regarding upcoming Explanation relations. As predicted, participants asked the question in their dialog responses that they were likely to answer in their monologue continuations: Compared to non-IC contexts, IC contexts gave rise to more Explanations and more ‘Why’-based questions, and the proportions of Explanation-based responses were reliably correlated across conditions.

³The correlations were calculated by taking a mean across observations for each of the n subjects (or n items) in the non-IC, IC-1, and IC-2 conditions for the monologue condition and the dialog condition; the tests establish how well correlated the $3n$ means are in the monologue condition with the $3n$ means in the dialog condition.

A question remains as to whether the effect shown here in Experiment XI is limited to contexts that are biased to Explanations and ‘Why’-type questions or whether this generalizes to other contexts. Experiment XII compares the distributions of a broader range of coherence relations and questions in contexts with transfer-of-possession verbs.

6.3 Experiment XII: Monologues/Dialogs in Transfer Contexts

Recall that Experiment I in Chapter 3 showed that manipulating the verbal aspect of a context sentence shifted the distribution of coherence relations. The verbal aspect manipulation was designed to alter the structure of the event evoked by the context sentence: Perfective aspect is used to describe a completed event, whereas imperfective aspect is used to describe an ongoing event. The experiment showed that perfective verbs were followed more frequently by Occasion relations than their imperfective counterparts, whereas imperfective verbs were followed more frequently by Elaborations and Explanations.

In comparing story continuations and dialog continuations across perfective and imperfective contexts, the goal is to test whether the correlation between coherence and questions that was noted in the previous section extends beyond Explanation and ‘Why’ questions. The experiment tests whether the proportion of the questions in the dialog condition that are categorized as ‘Why?’, ‘What next?’, and ‘Where/when/how?’ QUDs is reliably correlated with the proportion of continuations in the monologue condition that are categorized as Explanations, Occasions, and Elaborations, respectively. Violated Expectations were excluded because no questions were annotated as V-E. The three coherence types were selected because they were the most frequent relations in continuations elicited in transfer-of-possession contexts, excepting Violated Expectations.

6.3.1 Methods

Participants

Thirty-two monolingual English-speaking undergraduates from UC San Diego participated in the experiment for extra credit in Linguistics courses.

Materials and Procedures

Each context sentence contained one of forty-eight transfer-of-possession verbs from Experiment VII. All context sentences mentioned two possible referents, one male and one female with gender balanced across items. The prompts varied between a full-stop monologue prompt, a dialog prompt, and a *because* monologue prompt. (Again, the third condition with the *because* prompt will not be analyzed here.) The full stimulus set is given in the Appendix. The target items were interleaved with stimuli for an unrelated experiment with implicit-causality contexts (Experiment V described in Section 4.1) as well as forty-eight filler sentences that contained non-transfer verbs and were followed by various interclausal connectives or pronoun prompts (monologue continuation) or a dialog response that contained the beginning of a question (dialog continuation), for a total of one hundred thirty-two items.

The task was the same as that for Experiment XI: Participants were instructed to imagine a phone conversation with a friend and either write a continuation that continued what the friend said (monologue condition) or pose a question to the friend (dialog condition).

Again, story continuations were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their continuation.

Evaluation and Analysis

Two trained judges, the author of this dissertation and a UCSD Linguistics undergraduate, evaluated the participants' continuations. For the monologue condi-

tion, they assessed the coherence relation that could be inferred to hold between the context sentence and the continuation. For the dialog condition, they annotated the elicited question in terms of their possible answers (see Section 6.2.1 for examples of Elaboration, Explanation, Occasion, Parallel, and Result questions).

For the analysis, I tested how well correlated the proportions of Occasions, Elaborations, and Explanations were with the proportions of ‘What-next’, ‘How/Where/When’, and ‘Why’ questions across perfective and imperfective contexts.

6.3.2 Results

After setting aside the *because* continuations and the responses in the dialog condition that were deemed to be ambiguous (<1% of the dialog responses), the remaining dataset contained 510 dialog responses. Responses in the monologue condition that were deemed to be ambiguous were also set aside (<1%), leaving 508 monologue responses. Since there were no Violated Expectations in the dialog condition, Violated Expectation continuations in the monologue condition were excluded (14.6% of the monologue responses), leaving a combined monologue and dialog dataset of N=944. Due to an error in the presentation of stimuli, not all subjects saw all conditions. For this reason, I present only by-items analyses and item means in the text and graphs.

Replication of Previous Coherence Biases

The results for the monologue condition replicate the coherence biases observed in Experiment I from Chapter 3: Perfective verbs yielded more Occasions (31.5%) compared to imperfective verbs (15.8%; $F_2(1,47)=11.230$, $p<0.002$), whereas imperfective verbs yielded more Explanations (41.6%) and Elaborations (37.9%) compared to perfective verbs (*Elab*: 23.8%; $F_2(1,47)=6.032$, $p<0.02$; *Exp*: 23.7%; $F_2(1,47)=16.85$, $p<0.001$).

Correlation of Coherence and QUD Biases

As predicted, the increase in the proportion of each of these coherence relations is matched by an increase in the proportion of associated questions (‘What next’, ‘Why?’, ‘How?’). The effects for Occasions, Elaborations, and Explanations can be seen in Figures 6.2, 6.3, and 6.4, respectively.

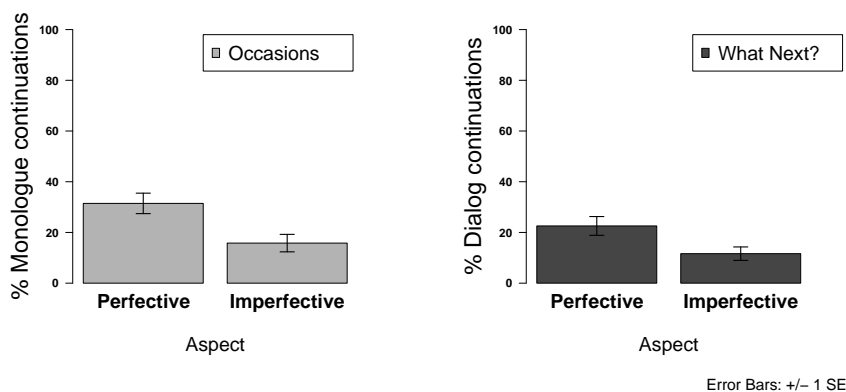


Figure 6.2: Effects of verbal aspect on Occasion/What-next responses in Experiment XII




Figure 6.3: Effects of verbal aspect on Elaboration/How responses in Experiment XII

Across perfective and imperfective aspect the two response types are reliably correlated. As Figure 6.2 shows, the larger proportion of Occasion relations following perfective context sentences (31.5%) than imperfective context sentences (15.8%) is correlated with the similar difference in proportions of ‘What next’ questions in the




Figure 6.4: Effects of verbal aspect on Explanation/Why responses in Experiment XII

same contexts (22.6% ‘What next’ QUDs following perfectives and 11.6% following imperfectives; $R^2=0.074$, $F_2(1,94)=8.638$, $p<0.005$). As Figure 6.3 shows, the larger proportion of Elaboration relations following imperfective context sentences (37.9%) than perfective context sentences (23.8%) is correlated with the similar difference in proportions of ‘How/when/where’ questions in the same contexts (43.2% ‘How’ QUDs following imperfectives and 37.4% following perfectives; $R^2=0.089$, $F_2(1,94)=10.24$, $p<0.002$). Lastly, as Figure 6.4 shows, the larger proportion of Explanation relations following imperfective context sentences (41.6%) than perfective context sentences (23.7%) is correlated with the similar difference in proportions of ‘Why’ questions in the same contexts (39.5% ‘Why’ QUDs following imperfectives and 30.3% following perfectives; $R^2=0.279$, $F_2(1,94)=37.82$, $p<0.001$).

6.3.3 Discussion

The goal of this experiment was to test whether the coherence~QUD connection observed in the implicit-causality contexts in Experiment XI extends beyond Explanation relations and ‘Why’-type questions to other coherence relations in other contexts. As predicted, the biases generated in transfer-of-possession contexts regarding other coherence relations (Occasions/Explanations/Elaborations) were reliably

correlated with related questions (‘What-next’, ‘Why’, ‘How/when/where’).

Although the direction of the coherence/QUD biases was correlated across perfective and imperfective contexts, the actual percentages differed between the story-continuation prompts and the dialog-continuation prompts. For example, there were more Occasion relations in the story continuations than there were ‘What-next’-type questions in the dialog continuations; there were more ‘How’-type questions in the dialog continuations than Elaborations in the story continuations. Certain questions and continuations appear to have different baseline rates — Occasion relations may simply have a higher prior probability in monologues than ‘What-next’-type questions do in dialogs, whereas Elaborations may have a lower prior probability in monologues than ‘How’-type questions do in dialogs.

Differences in the baseline rates of certain coherence relations and certain questions may correspond to differences in the type of information that is conveyed in statements and the type of information that is sought in questions. To the extent that speakers tend to communicate information about sequences of events, monologues may contain a large proportion of Occasion relations, whereas the types of questions that are most frequent may be biased towards clarification questions and requests for more detail. If this is the case, one could speculate that the different goals of conveying and clarifying information are what yield different base rates of coherence relations and questions.

6.4 General Discussion

This chapter extends the expectation-driven approach advocated in this dissertation from coherence relations in monologues to explicit questions in dialogs. The results show that participants’ expectations about upcoming coherence relations and about upcoming questions are both sensitive to some of the same contextual cues. If it is true that comprehenders use contextual cues to generate expectations about

upcoming material, then an important question is what the units are over which comprehenders themselves calculate such expectations — in other words, is there an underlying “discourse unit” and can a coherence or a QUD model better capture comprehenders’ behavior with regards to such a unit?

An advantage of postulating QUDs as underlying units of discourse is that the QUD model has been used to describe both single-speaker monologues and multi-speaker dialogs. However the coherence model could presumably be extended to similarly capture both monologues and dialogs. A question in a dialog can be considered a request for a proposition that would instantiate a particular coherence relation with the interlocutor’s previous utterance. Recall that Chapter 3 reported the results of an experiment that did just that by demonstrating that questions can be used to direct comprehenders’ expectations regarding the upcoming coherence relation. In Experiment III, participants were instructed to write either a story continuation that answered the question ‘Why?’ or one that answered the question ‘What happened next?’. The results showed that posing an explicit question in the instructions shifted the distribution of implicit coherence relations in participants’ continuations.

Note, however, that the inventory of coherence relations and the typology of *wh*- questions do not stand in a one-to-one relationship with each other, nor do the baseline rates of particular coherence relations and particular questions appear to align perfectly. For example, certain coherence relations lack correlate questions (e.g., Violated Expectations), and several questions appear to map onto the same coherence category (e.g., ‘how’/‘where’/‘when’ all correspond to an Elaboration relation). Furthermore, certain types of information are more representative of the information conveyed in monologues (descriptions of sequences of events, as in Occasion relations, which were the most frequent continuation in certain monologue contexts) whereas other types are more characteristic of dialogs (clarification and elaboration, as in ‘How’-type questions, which were the most frequent question in dialogs with similar contexts). These facts suggest that models of language processing need to condition

expectations about upcoming discourse continuations on other more general factors such as the goals and intentions of the speaker and their role as an interlocutor in a particular discourse context.

6.5 Acknowledgments

The experiments described in this chapter have been reported in earlier work and are discussed here with permission from the coauthors. The dissertation author was the primary investigator for the material presented here. Experiment XI and XII were reported in Rohde and Kehler (2008) [Coherence-Driven Expectations in Discourse and Dialog. Presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC, March 13-15, 2008].

7

Conclusion

The research presented in this dissertation fits within a larger context of work exploring expectation-driven effects in language comprehension and other areas of human cognition. Whereas previous psycholinguistic work has found evidence of expectation-driven effects at sentence-internal levels of processing (sound, word, syntactic structure), this dissertation set out to test whether comprehenders also show evidence of expectation-driven biases at the discourse level. The approach taken here centers around two phenomena, coreferential processing and the resolution of local syntactic ambiguity, using them as windows into comprehenders' discourse-level biases and expectations. The experiments aim to clarify our understanding of these two phenomena, while at the same time providing psycholinguistic evidence of expectation generation over the relationships that are inferred to hold between sentences. In this final chapter I review the findings, discuss their implications for theories of sentence and discourse processing, and consider directions for future work.

7.1 Contributions and Context

The work presented here lends support to a model of language comprehension in which pragmatic biases are integrated with biases at other levels of language

processing. Successful processing requires that comprehenders identify the structure and relationships between sounds, words, phrases, and sentences. One might expect that comprehenders wait until sentence-internal ambiguity has been resolved before considering the pragmatics of how entire clauses and larger discourse segments fit together to form the overall discourse structure. The results from the experiments described here, however, show that models of sentence-internal ambiguity resolution are incomplete without the incorporation of pragmatic biases concerning the relationships between sentences. This conclusion helps redefine our understanding of how comprehenders process two specific phenomena: coreference and relative clause attachment ambiguity. At a more general level, the results also provide novel evidence that coherence-driven pragmatic expectations are an integral part of sentence processing — thereby adding discourse structure to the list of linguistic levels at which comprehenders appear to make predictions.

In terms of the factors that influence sentence processing, much of the previous work on coreference and syntactic ambiguity resolution has relied on surface-level heuristics: Pronoun resolution is modeled with constraints such as a first-mention privilege, structural parallelism, and thematic-role biases; syntactic ambiguity resolution is described in terms of principles of adjacency and minimal structure building. Some existing models do appeal to deeper discourse-level processes and representations, but, as this dissertation has shown, many of these factors can still be better understood when one considers the relationships that are inferred to hold between sentences. Specifically, factors such as event structure and causal biases have previously been shown to influence pronoun interpretation (Stevenson et al., 1994; Arnold, 2001), but these factors are not relevant in all contexts. The coreference results from Chapter Three show that the influence of these factors is limited in predictable ways to particular coherence relations and that the effects disappear in contexts in which comprehenders have reason to expect different upcoming coherence relations. For discourse effects in syntactic processing, the existing work has focused on factors such as who has been mentioned in the preceding context and whether a referent is uniquely

identifiable (van Berkum et al., 1999). Taken together, this existing work calls out for a different framework that can begin to shed light on the puzzle of discourse-driven processing.

This dissertation has taken a step towards addressing this puzzle, and it does so by adopting inventories of intersentential coherence relations that could be inferred to hold between clauses (Hobbs, 1990; Kehler, 2002). These inventories of relations provide a unit over which to estimate comprehenders' discourse-driven expectations. For many researchers, the idea that discourse-level information should matter for processing may be plausible but finding a way to quantify that information and to manipulate it in controlled experiments had proven more difficult.

The approach advocated here is part of a growing shift in the psycholinguistic approach to both coreferential processing and syntactic ambiguity resolution. This shift concerns both the role of the comprehender and the types of factors that are posited to influence comprehension. Earlier approaches to these phenomena portrayed comprehension as a process of reaction rather than prediction: The presence of a pronoun was said to prompt the comprehender to launch a search for the most likely referent (Gordon & Scearce, 1995; Kaiser, 2003); the availability of more than one syntactic structure was said to require the comprehender to select a preferred interpretation (Frazier, 1978). This dissertation adopts an alternative approach, one in which comprehenders use prior context to generate expectations about upcoming material, whether it be upcoming patterns of coreference, syntactic structure, or discourse direction. This expectation-driven approach has gained momentum within several models of language comprehension (Arnold, 2001; Hale, 2001; Levy, 2008) as well as in more general models of cognition (Griffiths & Tenenbaum, 2000). Nonetheless, the model described here differs in important ways from other expectation-driven accounts.

The existing expectation-driven model of coreference (Arnold's Expectancy Hypothesis) posits that comprehenders generate expectations about upcoming mate-

rial based on contextual cues. Under this model, multiple contextual cues, ranging from the accessibility of entities mentioned in particular grammatical and thematic roles to the operative discourse continuation type, are all taken to influence comprehenders' expectations about which referent will be mentioned next, but the relationship between these cues remains unspecified. As such, the Expectancy Hypothesis does not capture the pattern of effects shown here — that particular biases (e.g., thematic role biases) do not have equal strength in all contexts but rather unfold along a causal chain of dependencies (e.g., the probability of a particular thematic role being re-mentioned depends on the probability that a particular coherence relation is operative which in turn depends on other contextual cues). Arnold did note that different continuation types yield different patterns of coreference, but she left open the question of how context guides the relative frequency of different continuation types. The coherence-driven model proposed here is in line with Arnold's but uses a theoretically motivated set of coherence relations, and the importance of these relations is emphasized in a different way. Furthermore, the model proposed here is one that captures more domain-general discourse expectations rather than expressly tracking referent accessibility in order to determine which individual is likely to be mentioned next. The coherence-driven model is not specific to a particular linguistic phenomenon; the insights into pronoun interpretation emerge as side effects of the deeper understanding of how comprehenders establish coherence generally in a discourse.

There are other expectation-driven models that have been proposed that do not target a particular phenomenon or ambiguity, but these models have thus far focused on domain-general expectations about syntactic outcomes (Levy, 2008; Hale, 2001) and semantic outcomes (Kamide et al., 2003; Bicknell et al., 2008). The predictability of pragmatic outcomes such as upcoming coherence relations, and thus the kind of phenomena this impacts, has remained largely unexplored.

7.2 Overview of Findings

As a simple summary, the results reported here show that models of language processing are incomplete if they fail to incorporate pragmatic, coherence-driven biases into their analysis of the comprehender's task. The chapters addressed two specific topics in language comprehension, showing that comprehenders' pragmatic biases about upcoming coherence relations influence both their coreferential processing and syntactic ambiguity resolution. Taken together, the results extend beyond these two phenomena to show more generally that comprehenders *have* expectations about outcomes beyond the next sound or the next word or the next phrase; comprehenders generate expectations about how upcoming clauses and discourse segments will be integrated into the larger discourse context.

The experimental work in this dissertation took as its starting point a surprising result from the pronoun interpretation literature showing that, in certain contexts, the preferred referent for an ambiguous pronoun is not the referent that appears in the syntactically parallel, first-mentioned, subject position. Rather, the non-subject, non-parallel referent is preferred as often or more often than the subject referent in contexts with particular verb types. The experiments in Chapter Three demonstrated that this result could be explained in terms of comprehenders' inferencing regarding event structure and causality. Furthermore, the pattern of inferencing, and its repercussions for coreference, differed depending on the operative coherence relation.

From that result sprang a set of experimental predictions that led to a set of new and surprising results about the role of discourse coherence in coreference more generally. The model put forth in Chapters Three and Four captures comprehenders' expectations regarding three different types of biases: the upcoming coherence relation ($p(CR | Context)$), the probability that a referent will be mentioned again ($p(Referent | CR)$), and the probability that a mention of a particular referent will be pronominalized ($p(Pronoun | Referent)$). From that model, a novel prediction emerged concerning relative clause processing.

Chapter Five adapted the coherence-driven model to test whether expectations about upcoming coherence relations influence processing of clausal relationships internal to the sentence, specifically for relative clause attachment ambiguity. The experiments addressed the question of whether or not comprehenders' resolution of local syntactic ambiguity reflects their coherence expectations. The results from the on-line self-paced reading-time experiment are of particular interest because they establish that comprehenders are constructing discourse contexts dynamically during sentence processing, using available coherence-driven biases mid-sentence to generate expectations about the structural analysis of the remainder of the sentence. Such a result echoes previous work showing that comprehenders use cues from different representational levels and that the integration of such cues is fully incremental (e.g., McRae, Spivey-Knowlton, and Tanenhaus' (1998) results for syntactic garden paths that are avoided or supported given the lexical semantics of the verb); in this case, however, the interaction is between a syntactic ambiguity and a set of pragmatic cues.

Having used the coherence-driven model to capture the way comprehenders understand the relationships between sentences and the relationships between clauses within a sentence, Chapter Six asked whether the coherence-driven model can also be applied to the domain of questions in dialogs. The results show that the same contextual cues that help comprehenders anticipate what direction a monologue will take also influence their expectations regarding upcoming questions in a dialog.

Each of these results stands on its own as a new insight in the well-studied domains of coreferential processing, syntactic disambiguation, and questions; however, taken together they shed light on language processing and the type of information sources that people use in comprehension. To understand language, comprehenders must not only fit coreferring elements together to track discourse entities and fit words together to establish the syntactic structure of a sentence, they must also fit clauses and sentences together to form a coherent discourse — and the experiments in this dissertation show that these processes are interdependent in a complex and

interesting way. The fact that comprehenders show sensitivity to an inventory of coherence relations provides important psycholinguistic evidence that comprehenders use contextual cues to generate expectations over units outside of the sentence.

7.3 Future Research

By opening the door to new research on expectation-driven approaches to discourse processing, this dissertation leads to several large questions about the nature of discourse-level structures and relationships, about their interaction with other levels of linguistic representation, and about the types of phenomena that may be sensitive to coherence-driven factors.

It is one thing to establish that discourse has an impact on processing, but this research program goes beyond the binary question of whether or not discourse matters. As has been discussed throughout the dissertation, the model put forth here represents more than just an addition to the list of factors that matter for processing; rather, I have shown that certain factors only have an effect in the context of certain other factors. One example that was already noted in Section 7.1 is that the strength of particular thematic role biases in pronoun interpretation is conditioned on the operative coherence relation. Likewise, the probability that a mention of a particular referent will be pronominalized is conditioned on grammatical role. The equations that were introduced here were used to highlight several contextual cues which the experimental evidence showed to be important; however the factors that appeared in those equations were not meant to be exhaustive. One avenue for future work is therefore to examine the role of other factors in a model in which coherence is already a part and to explore the space of possible ways that such factors might interact.

In a similar vein, it is one thing to establish that discourse matters for processing, but that does not address the open questions of what a unit of discourse really is and how comprehenders conceptualize discourse structure. Previous research re-

porting expectation-driven biases at other levels of linguistic structure has raised analogous questions before: What is the human sentence processor actually keeping track of? Do we have discrete representations for phonemes, syllables, morphemes, syntactic constructions, and coherence relations? What would such representations look like (clusters of observed tokens, idealized prototypes, a bundle of features) and how would an interactive model integrate different types of representations? Unsurprisingly, this dissertation is not able to provide answers to these questions. However, the experiments in Chapter Six did take a first step towards testing the nature of the discourse representation that was adopted in the earlier chapters.

Chapter Six considered two existing models of discourse structure: models of discourse coherence and models of Questions-Under-Discussion (QUD). Whereas coherence models propose that propositions in a single-speaker discourse are connected through a series of coherence relations, QUD models propose that single-speaker discourse consists of a series of utterances that answer inferrable questions and subquestions. The QUD models also have been used to describe multi-speaker discourses in which interlocutors may utter explicit questions as part of a dialog. The issue at stake in Chapter Six was whether coherence relations and QUDs could be linked in some way in order to clarify what properties comprise the underlying units of discourse structure. The results confirm that the distribution of possible coherence relations and the distribution of possible questions are both sensitive to similar contextual cues. If coherence relations and QUDs reflect similar behavior in comprehension, this invites a bridge between two discourse theories that have remained largely separate in the literature. The idea would be to interface a theory that appeals to general cognition and reasoning (coherence) with a theory that describes how underlying discourse structure is realized on the surface in speakers' questions and answers (QUD). Such a link would help further draw together two theories that are increasingly being used to address similar phenomena, for example, focus marking and accent placement (Roberts, 1996; Büring, 2003; Kehler, 2005) and pronoun and ellipsis interpretation (Kehler, 2002; Kertz, 2008; Kehler & Büring, 2007).

If discourse structure is to some degree predictable, that predictability has implications for production and the extent to which speakers signal the discourse structure they intend to convey. In the domain of phonetics, for example, it has been noted that less frequent versions of a homonym pair (*thyme/time*) are pronounced differently than the more frequent version (Jurafsky, Bell, Gregory, & Raymond, 2001; *inter alia*). In cases in which frequency influences production, the way a word is produced may be useful to the comprehender in determining whether the word is the more expected interpretation of those phonemes or a less frequent one. An analogy can be drawn between the way that speakers pronounce low frequency words and the way they may mark less frequent coherence relations, given contextual cues. Coherence relations can be marked overtly with an intersentential connective. Although analyses of the semantics of connectives exist, I know of no theories that predict when a connective is required, prohibited, or redundant. I suspect that the probability that a speaker will use a connective may depend in part on the predictability of the upcoming continuation type. For example, I would predict that contexts which favor ‘what-next’ continuations would yield fewer connectives like ‘next’ or ‘as a result’, whereas contexts that favor explanations would be more likely to use a connective to signal an unexpected ‘what-next’ continuation. Such a result would suggest that coherence-driven biases are influential in production as well as in comprehension.

7.4 Final Thoughts

Few people would deny that language processing must include mechanisms for understanding larger discourse contexts or that discourse factors affect the ultimate success or failure of communication, but it has been difficult to find concrete factors to manipulate in order to test these claims. Furthermore, the field of psycholinguistics has tended to focus on how comprehenders process sentence-internal units of linguistic structure—sounds, words, syntactic categories—but there has been little work on intrasentential units of structure, perhaps because it was not clear what

would constitute a unit of discourse for processing purposes.

This dissertation adopts an idea that grew out of research in theoretical linguistics and artificial intelligence, the idea that discourse can be characterized as a complex structured linguistic representation and that units of discourse can be identified and manipulated. What is unique about the work in this dissertation is that it uses controlled contexts in experimental settings to explore how comprehenders process relationships that are inferred to hold between sentences. I have presented a series of experiments to show that sentence comprehension and discourse comprehension reflect biases about intersentential coherence relations (a ‘unit of discourse’). The results serve as a reminder that speakers use language to describe ideas and events beyond those that fit within a single sentence. Speakers’ goals, assumptions, and expectations rely on linguistic representations that span multiple clauses and that must be integrated with a larger linguistic context in order for a discourse to make sense.

Appendix A

Experimental Stimuli

A.1 Experiment I

Pronoun prompts are marked in parentheses because Experiment VIII manipulated the presence/absence of the pronoun. The names were also changed for Experiment VIII so that each context sentence had two referents of the same gender.

1. John brought/was bringing a glass of water to Robert. (He) ...
2. Sarah carried/was carrying a tray to Brittany. (She) ...
3. Ben chucked/was chucking a wrench to Mark. (He) ...
4. Roger delivered/was delivering a subpoena to Joe. (He) ...
5. Brian faxed/was faxing a resume to Adam. (He) ...
6. Rebecca flung/was flinging a frisbee to Hannah. (She) ...
7. Angela forwarded/was forwarding a gossipy email to Kelly. (She) ...
8. Elizabeth gave/was giving a sweater to Ruth. (She) ...
9. Mike handed/was handing a book to Josh. (He) ...
10. Nick kicked/was kicking a soccer ball to Justin. (He) ...
11. Charles lobbed/was lobbing a football to Jacob. (He) ...
12. Heather mailed/was mailing a letter to Amy. (She) ...
13. Matt passed/was passing a sandwich to David. (He) ...

14. Peter rolled/was rolling a toy truck to Jeremy. (He) ...
15. Miriam sent/was sending a fruitcake to Rachel. (She) ...
16. Jessica served/was serving chili to Emily. (She) ...
17. Katherine shipped/was shipping a package to Laura. (She)
18. Jason threw/was throwing a hat to Andrew. (He) ...
19. Amanda tossed/was tossing a dish towel to Jenny. (She) ...
20. Jane transmitted/was transmitting a message to Nicole. (She) ...
21. Richard wired/was wiring money to Fred. (He) ...

A.2 Experiment II

1. John brought/was bringing a glass of water/a horse to Robert. He
2. Sarah carried/was carrying a tray/a 40-pound turnip to Brittany. She ...
3. Ben chucked/was chucking a wrench/an urn full of ashes to Mark. He ...
4. Roger delivered/was delivering a subpoena/four dozen lightbulbs to Joe. He ...
5. Brian faxed/was faxing a resume/a picture of a panda to Adam. He ...
6. Rebecca flung/was flinging a frisbee/a stolen wallet to Hannah. She ...
7. Angela forwarded/was forwarding a gossipy email/a Viagra SPAM email to Kelly. She ...
8. Elizabeth gave/was giving a sweater/a hand grenade to Ruth. She ...
9. Mike handed/was handing a book/a medieval to Josh. He ...
10. Nick kicked/was kicking a soccer ball/a pair of dirty sneakers to Justin. He ...
11. Charles lobbed/was lobbing a football/a wet dog to Jacob. He ...
12. Heather mailed/was mailing a letter/a bloody meat cleaver to Amy. She ...
13. Matt passed/was passing a sandwich/an empty jar of mayonnaise to David. He ...
...
14. Peter rolled/was rolling a toy truck/a barrel of rancid pickles to Jeremy. He ...
15. Miriam sent/was sending a fruitcake/a strand of dental floss to Rachel. She ...
16. Jessica served/was serving chili/stewed prunes to Emily. She ...

17. Katherine shipped/was shipping a package/a small lizard to Laura. She
18. Jason threw/was throwing a hat/a toilet plunger to Andrew. He ...
19. Amanda tossed/was tossing a dish towel/a condom to Jenny. She ...
20. Jane transmitted/was transmitting a message/a Shakespearean sonnet to Nicole. She ...
21. Richard wired/was wiring money/one million dollars to Fred. He ...

A.3 Experiment III

See Experiment I

A.4 Experiment IV

Pronoun prompts for Experiments VI and VIII are marked in parentheses. The names were changed for Experiment VIII so that each context sentence had two referents of the same gender.

IC-1 items

1. Mike aggravated Brittany. (He/She) ...
2. Amanda amazed Christopher. (He/She) ...
3. Josh amused Sarah. (He/She) ...
4. Stephanie annoyed David. (He/She) ...
5. Andrew apologized to Megan. (He/She) ...
6. Emily bored Justin. (He/She) ...
7. James charmed Amber. (He/She) ...
8. Heather offended Bob. (He/She) ...
9. Joe confessed to Tiffany. (He/She) ...
10. Rachel deceived Nick. (He/She) ...
11. Tony disappointed Courtney. (He/She) ...

12. Christina exasperated Brandon. (He/She) ...
13. Eric fascinated Rebecca. (He/She) ...
14. Steve frightened Chelsea. (He/She) ...
15. Crystal humiliated Brian. (He/She) ...
16. Zack infuriated Kristen. (He/She) ...
17. Kelly inspires Jacob. (He/She) ...
18. Ben intimidates Andrea. (He/She) ...
19. Angela scares Timothy. (He/She) ...
20. Tyler surprised Vanessa. (He/She) ...

IC-2 items

1. Paul assisted Melissa. (He/She) ...
2. Anna blames Frank. (He/She) ...
3. Beth comforted Simon. (He/She) ...
4. Jared congratulated Debbie. (He/She) ...
5. Greg corrected Sally. (He/She) ...
6. Jennifer detests Arnold. (He/She) ...
7. Bill fears Nicole. (He/She) ...
8. Candice envies John. (He/She) ...
9. Ryan hates Amy. (He/She) ...
10. Joel helped Claire. (He/She) ...
11. Ken mocked Jane. (He/She) ...
12. Sandra noticed Rob. (He/She) ...
13. Alice pacified Trevor. (He/She) ...
14. Tina praised Adam. (He/She) ...
15. Craig reproached Kate. (He/She) ...
16. Elizabeth scolded Alan. (He/She) ...
17. Carl stared at Veronica. (He/She) ...
18. Scott thanked Jessica. (He/She) ...

19. Ashley trusts George. (He/She) ...
20. Laura values Luis. (He/She) ...

Non-IC items

1. Ann chatted with Clinton. (He/She) ...
2. Arthur saw Wendy. (He/She) ...
3. Marcia read a funny story to Jim. (He/She) ...
4. Anthony went to visit Lucy. (He/She) ...
5. Rose threw a ball at Gary. (He/She) ...
6. Dennis poured a drink for Vicki. (He/She) ...
7. Ellen broke a window playing with Howard. (He/She) ...
8. Roy watched Pamela. (He/She) ...
9. Allison appreciated the flowers from William. (He/She) ...
10. Fred cooked dinner for Kim. (He/She) ...
11. Alyssa played the piano for Gabe. (He/She) ...
12. Miguel edited an essay for Molly. (He/She) ...
13. Caitlin made a sandwich for Logan. (He/She) ...
14. Neal waited to see Hannah. (He/She) ...
15. Eva repaired a bike for Keith. (He/She) ...
16. Jeff counted the money from Audrey. (He/She) ...
17. Kathleen was drawing a picture of Derek. (He/She) ...
18. Dustin fixed a broken printer for Alexa. (He/She) ...
19. Sylvia borrowed a bike from Kevin. (He/She) ...
20. Martin moved in next door to Ariana. (He/She) ...
21. Jasmine worked with Malcolm. (He/She) ...
22. Tom studied with Kristy. (He/She) ...
23. Dan sat next to Holly at lunch. (He/She) ...
24. Melanie did the dishes with Julia. (He/She) ...
25. Max went to the library with Tracy. (He/She) ...

26. Meredith cleaned the house with Todd. (He/She) ...
27. Ethan went jogging with Carrie. (He/She) ...
28. Teresa went to the beach with Rick. (He/She) ...
29. Walter bought a burger for Bridget. (He/She) ...
30. Justine enjoyed a book from Rodney. (He/She) ...
31. Darren waited in line behind Melinda. (He/She) ...
32. Virginia arrived at school before Ramon. (He/She) ...
33. Theo took a class with Gloria. (He/She) ...
34. Tanya went to the gym with Sebastian. (He/She) ...
35. Wayne got a ride home with Camille. (He/She) ...
36. Naomi was comparing grades with Trenton. (He/She) ...
37. Jerome took a pen from Carla. (He/She) ...
38. Charlotte stood next to Parker. (He/She) ...
39. Sonia split a pizza with Brady. (He/She) ...
40. Margaret ran into Mark. (He/She) ...

A.5 Experiment V

See Experiment IV

A.6 Experiment VI

See Experiment IV

A.7 Experiment VII

The * symbol indicates that the item only appeared in Experiment VII. Items without a * were used for Experiment XII as well.

1. *Philip administered/was administering medicine to Melissa. ...
2. Crystal batted/was batting a grounder to Jim. ...
3. Kyle blew/was blowing a kiss to Sasha. ...
4. Bruce bounced/was bouncing a basketball to Bridget. ...
5. Tom brought/was bringing a cup of tea to Stephanie. ...
6. *Paul carried/was carrying a tray to Julia. ...
7. Maria carted/was carting an old computer to Ryan. ...
8. Natalie chucked/was chucking a measuring tape to Jacob. ...
9. Shelly delivered/was delivering a letter to Kevin. ...
10. Brenda dispensed/was dispensing cough syrup to Brett. ...
11. *John donated/was donating a toy to Mary. ...
12. Eric dropped/was dropping a roof shingle to Nina. ...
13. Molly expedited/was expediting a financial report to Peter. ...
14. Diane faxed/was faxing a resume to Greg. ...
15. *Chad fed/was feeding a piece of chocolate to Nancy. ...
16. *Kristen flicked/was flicking a paperclip to Albert. ...
17. Alice flipped/was flipping a queen of spades to James. ...
18. Tim floated/was floating a life vest to Jessica. ...
19. *Rebecca flung/was flinging a coat to Charles. ...
20. Kara forwarded/was forwarding a gossipy email to Ken. ...
21. Adam gave/was giving a sweater to Laura. ...
22. *Sue handed/was handing a timecard to Fred. ...
23. Bill hauled/was hauling a wheelbarrow to Cindy. ...
24. *Larry heaved/was heaving a box to Jeanette. ...
25. Seth hit/was hitting an easy fly ball to Helen. ...
26. Claire hurried/was hurrying a transcript to Jeff. ...
27. Evan kicked/was kicking a soccer ball to Megan. ...
28. *Andrew hurled/was hurling a brick to Ashley. ...
29. Ethan lifted/was lifting a box to Katherine. ...

30. *Caitlin lobbed/was lobbing a football to Mitch. ...
31. *Sophia lowered/was lowering a first-aid kit to Joel. ...
32. *Janet lugged/was lugging a suitcase to Arthur. ...
33. Keith mailed/was mailing a fruitcake to Barbara. ...
34. Jason nudged/was nudging a microphone to Hannah. ...
35. *Chelsea passed/was passing a sandwich to Ben. ...
36. Caroline pitched/was pitching a ball to Jeremy. ...
37. Brian presented/was presenting an award to Emma. ...
38. *Luke propelled/was propelling a beer can to Colleen. ...
39. Rachel pushed/was pushing a shopping cart to Casey. ...
40. Kelly raised/was raising a ladder to Daniel. ...
41. Heather refunded/was refunding \$30 to Roger. ...
42. *Jerry returned/was returning a sweater to Christine. ...
43. *Naomi rolled/was rolling a toy truck to Steve. ...
44. *Anne rushed/was rushing a report to Sean. ...
45. *David sent/was sending a love letter to Gina. ...
46. Matt served/was serving chili to Jasmine. ...
47. Amanda shifted/was shifting some poker chips to Scott. ...
48. Angela shipped/was shipping a package to Cory. ...
49. *Henry shot/was shooting a puck to Brittany. ...
50. *Donald shoved/was shoving a mass of papers to Sharon. ...
51. Nick skipped/was skipping a beach ball to Courtney. ...
52. George slapped/was slapping a beachball to Sarah. ...
53. Monica slipped/was slipping \$50 to Ed. ...
54. Joe slugged/was slugging a line drive to Miranda. ...
55. Josh smacked/was smacking a ball to Katie. ...
56. Melanie snapped/was snapping a frisbee to Nathan. ...
57. Brandon socked/was socking a ball to Alyssa. ...
58. Karen spun/was spinning a top to Justin. ...

59. Robert supplied/was supplying food to Brianna. ...
60. Emily sold/was selling a postcard to Jack. ...
61. *Linda surrendered/was surrendering a toy to Grant. ...
62. Richard swatted/was swatting a tennis ball to Stacy. ...
63. Elizabeth took/was taking a meal to Frank. ...
64. Brad threw/was throwing a hat to Margaret. ...
65. Amy thrust/was thrusting a bouquet of roses to Mark. ...
66. *Jennifer tossed/was tossing a dishtowel to Colin. ...
67. Alan towed/was towing an old jalopy to Allison. ...
68. Diana transferred/was transferring \$1000 to Max. ...
69. *Mike transferred/was transmitting a message to Lisa. ...
70. *Beth whacked/was whacking a wiffleball to Dennis. ...
71. Kim wheeled/was wheeling a large bicycle to Todd. ...
72. Carl wired/was wiring money to Meredith. ...

A.8 Experiment VIII

See Experiment I

A.9 Experiment IX

1. Paul worships/listens to the coach of the cheerleaders who ...
2. Beth despises/babysits the children of the jazz musician who ...
3. Frank complimented/met the guests of the bride who ...
4. Jared blamed/noticed the friends of the athlete who ...
5. Greg adores/smiles at the secretaries of the lawyer who ...
6. Casey detests/looks like the father of the students who ...
7. Bill congratulated/visited the teacher of the second-graders who ...
8. Candice criticized/talked to the leader of the activists who ...

9. Ryan likes/resembles the captain of the old sailors who ...
10. Joel pities/hires the bodyguards of the celebrity who ...
11. Ken praised/videotaped the assistants of the CEO who ...
12. Sandra insulted/chatted with the gardeners of the millionaire who ...
13. Alice values/lives next to the surgeon of the soldiers who ...
14. Tina resents/knows the doctors of the supermodel who ...
15. Craig rewarded/inspected the servants of the dictator who ...
16. Alan punished/saw the accountant of the businessmen who ...
17. Carl admires/works with the agent of the rockstar who ...
18. Scott ridiculed/counted the fans of the singer who ...
19. George thanked/interviewed the representative of the employees who ...
20. Luis scolded/recognized the landlady of the actors who ...
21. Melissa dislikes/watches the little girls of the neighbor who ...

A.10 Experiment X

1. Anna scolded/studied_with the_chef of the_aristocrats who was routinely letting food go to_waste.
2. John stared_at/lived_next_to the_teacher of the_second_graders who was definitely smartest in the school.
3. Jenny assisted/joked_with the_maid of the_executives who was regularly late to work.
4. Nick trusted/stood_near the_captain of the_sailors who has consistently weathered big storms.
5. Angela corrected/gossiped_with the_secretary of the_lawyers who has occasionally made small mistakes.
6. Bob comforted/greeted the_leader of the_activists who was deeply disappointed by the_court's decision.

7. Laura envies/knows the_manager of the_cashiers who has supposedly received a_huge raise.
8. Zack valued/recognized the_daughter of the_shopkeepers who was usually willing to_spot him a_few_dollars.
9. Sarah fears/jogs_with the_uncle of the_toddlers who is often heard yelling and screaming.
10. Adam noticed/resembled the_representative of the_employees who was always wearing safety goggles.
11. Tina praised/met the_gardeners of the_millionaire who have recently installed a_solar powered sprinkler.
12. Justin hates/carpools_with the_cousins of the_accountant who are forever telling the_same tasteless jokes.
13. Emily blamed/waited_with the_nieces of the_florist who have repeatedly ruined expensive orchids.
14. Joe helped/ran_into the_brothers of the_athlete who are perpetually failing math class.
15. Jessica reproached/worked_with the_doctors of the_supermodel who were adamantly in_favor of plastic surgery.
16. Brian pacified/visited the_associates of the_businessman who were nearly bankrupted by the_new tax_policy.
17. Melissa detests/babysits the_children of the_musician who are generally arrogant and rude.
18. Frank thanked/talked_to the_servants of the_dictator who have lately been helping the_poor.
19. Tracy congratulated/chatted_with the_bodyguards of the_celebrity who were constantly fighting off the_paparazzi.
20. Kevin mocked/counted the_fans of the_singer who were continually stagediving and getting hurt.

A.11 Experiment XI

See Experiment IV

A.12 Experiment XII

See Experiment VII

Appendix B

Guidelines for Coherence Annotation

Any pair of sentences can be related by one or more coherence relations. The judges annotated the elicited passages with the goals of identifying the participant's intended meaning and selecting the most likely coherence relation that could be inferred to hold between the context sentence and the continuation. The set of coherence relations were taken from Kehler (2002) and consisted of Elaboration, Explanation, Occasion, Parallel, Result, and Violated Expectation, which together provided very good coverage of the elicited continuations. Biclausal continuations were annotated only in consideration of the first clause or the first event being described. Below I list some general guidelines, beyond Kehler's definitions, that were used in the annotation of the six coherence relations in the story-continuation passages.

Elaboration continuations were defined as those that provided more details of the event or state described in the context sentence. Such details included additional information about temporal, locative, instrumental, or benefactive properties of the event. Included in Elaborations were continuations about degree (*John hit Bob. He hit him hard.*) and those about details of the event arguments (*John made a sandwich for Mary. He made a tuna sandwich.*).

Explanations continuations were defined as those that established the cause of the event in the context sentence. Some difficulty arose in annotating passages with implicit causality verbs, especially subject-biased IC verbs, because it was difficult to distinguish between continuations that elaborated on an event (aggravating someone by pulling their hair) and those that described a reason for the event (aggravating someone for the pleasure of seeing them get annoyed). In most cases it was clear, but for tough cases, the judges agreed on an annotation scheme whereby continuations containing stative verbs or repeated/iterative events were generally characterized as Explanations (*John aggravated Mary. He was undoubtedly insensitive.*), whereas passages that described single events were treated as Elaborations of that main event (*John aggravated Mary. He told a terrible joke.*).

Occasions continuations were defined as those for which a temporal relation could be established between the events described by two sentences, whereby the event in the first sentence preceded the event in the second sentence (*John threw a ball to Bob. He caught it.*).

Result continuations were defined as those that appealed to causal inferencing to capture the relationship between a sequence of events (a temporally related sequence of events as in an Occasion relation in which the event described in the first sentence causes the event in the second sentence) or states of affairs (e.g., resulting emotional states). Results often described someone's reaction to an event (*John hit Bob. Bob got really mad.*). For cases in which more than one coherence relation held, the judges agreed that if a causal relation could be inferred, that relation trumped a non-causal relation such as Elaboration or Occasion.

Violated Expectations appealed to the type of real-world knowledge that was required for inferring a Result relation, but the continuation described an unexpected outcome given the knowledge of likely events and their consequences (*John insulted Mary. She was not offended.*).

Parallel relations are characterized by the description of two events or states

that share some common or contrasting property. Often these passages described the same referents but with different verbs (*John hates Mary. She adores him.*) or they used the same verb but varied the referents (*John hates Mary. Mary hates Bob.*).

Although the set of six coherence relations used in annotation provided good coverage, there were cases for which there was no suitable relation. For example, there was no obvious category for continuations that described events which could plausibly happen subsequent to or simultaneous to the event in the context sentence (*John was gardening. He was thinking about what to plant next year.*). Similarly, there was no obvious category for continuations that described events which happened prior to the event in the context sentence but did not provide information about the cause (*John frightened Mary. He had been planning to do so all day.*). Both of these types were labeled as Elaborations.

For several passages, more than one coherence relation could be inferred to hold. In those cases, a hierarchy for annotation was devised for passages with relations that maintained the meaning of the two sentences — e.g., a Result relation and an Occasion relation could both capture two temporally ordered events, but causal relations were given precedence. For cases in which the inferrable relations required very different interpretations of the passage (e.g., interpretations with opposite pronoun interpretations), the continuation was set aside as ‘Unknown’. In the Results section of each experiment for which judges annotated coherence relations, I report the percentage of the data that was excluded.

References

- Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review*, *106*, 748-765.
- Altmann, G. T. M. (1998). Ambiguity in sentence processing. *Trends in Cognitive Science*, *2*, 146-152.
- Arnold, J. E. (2001). The effects of thematic roles on pronoun use and frequency of reference. *Discourse Processes*, *31*, 137-162.
- Arnold, J. E., & Griffin, Z. M. (2007). The effect of additional characters on choice of referring expression: Everyone counts. *Journal of Memory and Language*, *56*, 521-536.
- Asher, N., & Lascarides, A. (2003). *Logics of Conversation*. Cambridge: Cambridge University Press.
- Au, T. K. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, *25*, 104-122.
- Bicknell, K., Elman, J. L., Hare, M., McRae, K., & Kutas, M. (2008). Online expectations for verbal arguments conditional on event knowledge. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*.
- Brennan, S. E., Friedman, M. W., & Pollard, C. J. (1987). A centering approach to pronouns. In *Proceedings of the 25th Meeting of the Association for Computational Linguistics*.
- Brown, R., & Fish, D. (1983). The psychological causality implicit in language. *Cognition*, *14*, 237-273.
- Brysbaert, M., & Mitchell, D. C. (1996). Modifier attachment in sentence parsing: Evidence from Dutch. *Quarterly Journal of Experimental Psychology*, *49A*, 664-695.
- Büring, D. (2003). On d-trees, beans, and b-accents. *Linguistics and Philosophy*, *26*,

511-545.

- Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behaviour*, *16*, 601-609.
- Carreiras, M., & Clifton, C. (1999). Another word on parsing relative clauses: Eye-tracking evidence from Spanish and English. *Memory and Cognition*, *27*, 826-833.
- Chafe, W. L. (1972). Discourse structure and human knowledge. In J. B. Carroll & R. O. Freedle (Eds.), *Language comprehension and the acquisition of knowledge* (p. 41-69). Washington, D.C.: Winston and Sons.
- Chambers, G. C., & Smyth, R. (1998). Structural parallelism and discourse coherence: A test of centering theory. *Journal of Memory and Language*, *39*, 593-608.
- Clark, H. H., & Sengal, C. J. (1979). In search of referents for nouns and pronouns. *Memory and Cognition*, *7*, 35-41.
- Crawley, R., Stevenson, R., & Kleinman, D. (1990). The use of heuristic strategies in the interpretation of pronoun. *Journal of Psycholinguistic Research*, *4*, 245-264.
- Cuetos, F., & Mitchell, D. (1988). Cross-linguistic differences in parsing: Restrictions on the use of the late closure strategy in Spanish. *Cognition*, *30*, 73-105.
- Cuetos, F., Mitchell, D., & Corley, M. (1996). Parsing in different languages. In M. Carreiras, J. García-Albea, & N. Sebastián-Gallés (Eds.), *Language processing in Spanish* (p. 145-187). Mahwah, New Jersey: Lawrence Erlbaum.
- DeLong, K., Urbach, T., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, *8*, 1117-1122.
- Desmet, T., Brysbaert, M., & De Baecke, C. (2002). The correspondence between sentence production and corpus frequencies in modifier attachment. *The Quarterly Journal of Experimental Psychology*, *55*, 879-896.
- Desmet, T., de Baecke, C., & Brysbaert, M. (2002). The influence of referential discourse context on modifier attachment in Dutch. *Memory and Cognition*, *30*, 150-157.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-211.
- Elman, J. L. (1993). Learning and development in neural networks: the importance of starting small. *Cognition*, *48*, 71-99.

- Fernandez, E. (2003). *Bilingual sentence processing: Relative clause attachment in bilinguals and monolinguals*. Amsterdam: John Benjamins.
- Ferreira, F., & Clifton, J. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*, 348-368.
- Ferreira, V. S., & Dell, G. S. (2000). The effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, *40*, 296-340.
- Ferretti, T., Rohde, H., Kehler, A., & Crutchley, M. (submitted). Verb Aspect, Event Structure, and Coreferential Processing. *Journal of Memory and Language*. (Under review)
- Foley, W. A., & Valin, R. D. V. (1984). *Functional syntax and universal grammar*. New York: Cambridge University Press.
- Forster, K. I., & Forster, J. C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments and Computers*, *35*, 1161-124.
- Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies*. Unpublished doctoral dissertation, University of Connecticut.
- Frazier, L., & Clifton, C. (1996). *Construal*. Cambridge, Mass: MIT Press.
- Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, *5*, 459-464.
- Gernsbacher, M. A., & Hargreaves, D. J. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, *27*, 699-717.
- Gibson, E., Pearlmutter, N., Canseco-Gonzalez, E., & Hickok, G. (1996). Recency preference in the human sentence processing mechanism. *Cognition*, *59*, 23-59.
- Gilboy, E., Sopena, J., Clifton, C., & Frazier, L. (1995). Argument structure and association preferences in Spanish and English complex NPs. *Cognition*, *54*, 131-167.
- Ginzburg, J., & Sag, I. (2000). *Interrogative investigations*. Stanford: CSLI Publications.
- Givon, T. (1990). *Syntax: a functional typological introduction*. Amsterdam: John Benjamins.
- Gordon, P., & Scarse, K. (1995). Pronominalization and discourse coherence, discourse structure and pronoun interpretation. *Memory and Cognition*, *23*, 313-323.

- Griffiths, T. L., & Tenenbaum, J. B. (2000). Teacakes, trains, toxins, and taxicabs: A Bayesian account of predicting the future. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*.
- Grosz, B. J., Joshi, A. K., & Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, *21*, 203-225.
- Gundel, J., Hedberg, H., & Zacharski, R. (1993). Referring expressions in discourse. *Language*, *69*, 274-307.
- Hale, J. (2001). A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics*. Pittsburgh, PA, June 2-7, 2001.
- Hemforth, B., Konieczny, L., & Scheepers, C. (2000). Syntactic attachment and anaphora resolution: The two sides of relative clause attachment. In M. Crocker, M. Pickering, & C. Clifton (Eds.), *Architectures and mechanisms for language processing* (pp. 259-281). Cambridge: Cambridge University Press.
- Hemforth, B., Konieczny, L., Seelig, H., & Walter, M. (2000). Case matching and relative clause attachment. *Journal of Psycholinguistic Research*, *29*, 81-88.
- Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, *3*, 67-90.
- Hobbs, J. R. (1990). *Literature and cognition, lecture notes 21*. Stanford, CA: CSLI.
- Hume, D. (1748). *An enquiry concerning human understanding*. New York: The Liberal Arts Press, 1955 edition.
- Hutchinson, B. (2005). *The automatic acquisition of knowledge about discourse connectives*. Unpublished doctoral dissertation, University of Edinburgh.
- Jackendoff, R. (1972). *Semantic interpretation in generative grammar*. Cambridge: MIT Press.
- Jurafsky, D., Bell, A., Gregory, M., & Raymond, W. D. (2001). Probabilistic relations between words: Evidence from reduction in lexical production. In J. Bybee & P. Hopper (Eds.), *Frequency and the emergence of linguistic structure*. Amsterdam: John Benjamins.
- Jusczyk, P. W., Luce, P. A., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language*, *33*, 630-645.
- Kadmon, N. (2001). *Formal pragmatics: Semantics, pragmatics, presupposition, and focus*. Malden, Mass.: Blackwell.

- Kaiser, E. (2003). *The quest for a referent: A crosslinguistic look at reference resolution (Finnish, Estonian, Dutch)*. Unpublished doctoral dissertation, University of Pennsylvania.
- Kamide, Y., Altmann, G. T. M., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, *49*, 133-156.
- Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. Stanford, CA: CSLI Publications.
- Kehler, A. (2005). Coherence-driven constraints on the placement of accent. In *Proceedings of the 15th Conference on Semantics and Linguistic Theory (SALT-15)*. Los Angeles, CA, March 25-27, 2005.
- Kehler, A., & Büring, D. (2007). Be bound or be disjoint! In *Proceedings of the 38th Meeting of the North East Linguistics Society (NELS-38)*. Ottawa, Canada, 2007.
- Kehler, A., Kertz, L., Rohde, H., & Elman, J. (2008). Coherence and coreference revisited. *Journal of Semantics*, *25*, 1-44.
- Kertz, L. (2008). Focus structure and acceptability in verb phrase ellipsis. In *Proceedings of the 27th West Coast Conference on Formal Linguistics Special Session on Experimental Syntax, Semantics, and Pragmatics*. Los Angeles, CA, May 16-18, 2008.
- Kertz, L., Kehler, A., & Elman, J. (2006). Grammatical and coherence-based factors in pronoun interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (p. 1605-1610). Vancouver, July 26-29, 2006.
- Kluender, R. (1992). Deriving island constraints from principles of predication. In H. Goodluck & M. Rochemont (Eds.), *Island constraints: Theory, acquisition and processing* (p. 223-258). Dordrecht: Kluwer.
- Knott, A. (1996). *A data-driven methodology for motivating a set of coherence relations*. Unpublished doctoral dissertation, University of Edinburgh.
- Koornneef, A. W., & van Berkum, J. J. A. (2006). On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language*, *54*, 445-465.
- Lambrecht, K. (1994). *Information structure and sentence form: Topic, focus, and the mental representations of discourse referents*. Cambridge: Cambridge University Press.
- Larsson, S. (1998). Questions Under Discussion and dialogue moves. In *Proceedings*

of TWLT 13/Twendial: Formal Semantics and Pragmatics of Dialogue.

- Levin, B. (1993). *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*. (In press)
- Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8, 243-281.
- Marcu, D. (2000). The rhetorical parsing of unrestricted texts: A surface-based approach. *Computational Linguistics*, 26, 395-448.
- McDonald, J. L., & MacWhinney, B. (1995). The time course of anaphora resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34, 543-566.
- McKoon, G., Greene, S., & Ratcliff, R. (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. *Journal of Experimental Psychology*, 19, 1040-1052.
- McRae, K., Hare, M., Elman, J. L., & Ferretti, T. (2005). A basis for generating expectancies for verbs from nouns. *Memory and Cognition*, 33, 1174-1184.
- McRae, K., Spivey-Knowlton, M., & Tanenhaus, M. K. (1998). Modeling the effect of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, 38, 283-312.
- Miltsakaki, E. (2007). A rethink of the relationship between salience and anaphora resolution. In *Proceedings of the 6th Discourse Anaphora and Anaphor Resolution Colloquium*.
- Mitchell, D. (1984). An evaluation of subject-paced reading tasks and other methods for investigating immediate processes in reading. In D. Kieras & M. Just (Eds.), *New Methods in Reading Comprehension Research* (p. 6989). Hillsdale, NJ: Earlbaum.
- Mitchell, D., & Brysbaert, M. (1998). Challenges to recent theories of crosslinguistic variation in parsing: Evidence from Dutch. *Syntax and Semantics*, 31, 313-335.
- Mithun, M. (1991). The role of motivation in the emergence of grammatical categories: The grammaticization of subjects. In E. C. Traugott & B. Heine (Eds.), *Approaches to grammaticalization* (Vol. 2, p. 160). Amsterdam: John Benjamins.
- Moens, M., & Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14, 15-28.

- Most, R. B., & Saltz, E. (1979). Information structure in sentences: New information. *Language and Speech*, 22, 89-95.
- Papadopoulou, D., & Clahsen, H. (2006). Ambiguity resolution in sentence processing: The role of lexical and contextual information. *Journal of Linguistics*, 42, 109-138.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *OSU Working Papers in Linguistics*, 49: *Papers in Semantics*.
- Rohde, H., Kehler, A., & Elman, J. (2006). Event structure and discourse coherence biases in pronoun interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (p. 697-702). Vancouver, July 26-29, 2006.
- Roland, D., O'Meara, C., Yun, H., & Mauner, G. (2007). *Processing object relative clauses: Discourse or frequency?* Poster presented at the CUNY sentence processing conference. La Jolla, CA, March 29-31, 2007.
- Saffran, J. R., Aslin, R. N., & Newport, E. (1996). Statistical learning by 8-month-old infants. *Science*, 274:5294, 1926-1928.
- Simner, J., & Pickering, M. J. (2005). Planning causes and consequences in discourse. *Journal of Memory and Language*, 52, 226-239.
- Smyth, R. J. (1994). Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, 23, 197-229.
- Stevenson, R., Crawley, R., & Kleinman, D. (1994). Thematic roles, focusing and the representation of events. *Language and Cognitive Processes*, 9, 519-548.
- Stevenson, R., Knott, A., Oberlander, J., & McDonald, S. (2000). Interpreting pronouns and connectives: Interactions among focusing, thematic roles, and coherence relations. *Language and Cognitive Processes*, 15, 225-262.
- Stewart, A. J., Pickering, M. J., & Sanford, A. J. (2000). The time course of the influence of implicit causality information: Focusing versus integration accounts. *Journal of Memory and Language*, 42, 423-443.
- Tabossi, P. (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 153-162.
- Traxler, M., Pickering, M., & Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, 39, 558-592.
- van Berkum, J. J. A., Brown, C. M., & Hagoort, P. (1999). Early referential context

- effects in sentence processing: Evidence from event-related brain potentials. *Journal of Memory and Language*, 41, 147-182.
- Van Kuppevelt, J. (1995). Discourse structure, topicality, and questioning. *Journal of Linguistics*, 31, 109-147.
- Vincenzi, M. D., & Job, R. (1995). An investigation of late closure: The role of syntax, thematic structure, and pragmatics in initial and final interpretation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1303-1321.
- Webber, B. (2006). Accounting for discourse relations: Constituency and dependency. In M. D. M. Butt & T. King (Eds.), *Intelligent linguistic architectures* (p. 339-360). Stanford: CSLI Publications.
- Webber, B., & Joshi, A. (1998). Anchoring a Lexicalized Tree-Adjoining Grammar for discourse. In *ACL/COLING Workshop on Discourse Relations and Discourse Markers*. Montreal, Canada.
- Wolf, F., Gibson, E., & Desmet, T. (2004). Coherence and pronoun resolution. *Language and Cognitive Processes*, 19, 665-675.
- Zagar, D., Pynte, J., & Rativeau, S. (1997). Evidence for early-closure attachment on first-pass reading times in French. *The Quarterly Journal of Experimental Psychology*, 50, 421-438.